Precision Luminosity Determination at the CMS Experiment

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Warusapperuma Don Nimmitha Karunarathna
August 2026

© by Warusapperuma Don Nimmitha Karunarathna, 2026
All Rights Reserved.

To my family, for their constant support and encouragement.

Acknowledgements

I sincerely thank my advisor, Dr. Stefan Spanier, for his guidance, support, and scientific mentorship throughout my doctoral research. His advice helped shape the direction of this dissertation and the way I approached my research. I am also grateful that he gave me the opportunity to participate in detector operations at CERN.

I would like to thank Dr. Andres Delannoy, Dr. Francesco Romeo, and the BRIL team for their support, technical guidance, and many useful discussions. Their expertise and collaboration were important to my understanding of luminosity measurements. I also thank Dr. Jesse Harris for his help and useful suggestions during the development of the rare-Higgs analysis.

I thank the members of the UTK-CMS group for their support, discussions, and shared experience throughout graduate school. I also thank my dissertation committee members, Dr. Sarah Cousineau, Dr. Robert Grzywacz, and Dr. Tova Holmes, for their time and service.

Finally, I am deeply thankful to my family, and especially to my wife, Ayesha, for their constant encouragement and support.

Abstract

The Compact Muon Solenoid (CMS) collaboration probes for new forces and particles at CERN's Large Hadron Collider (LHC). One approach searches for significant discrepancies in Higgs-boson properties relative to Standard Model predictions that arise from beyond-Standard Model processes.

Luminosity is essential for estimating background rates and for converting observed event yields into cross sections or decay branching ratios. This dissertation presents studies using the CMS detector aimed at achieving high-precision luminosity measurements during recent data-taking periods. It describes calibration studies for a dedicated luminosity detector, the Pixel Luminosity Telescope, including the van der Meer beam scan method, beam-related corrections, detector-response corrections, and nonlinearity studies.

Furthermore, an independent luminosity measurement was developed by counting reconstructed J/ψ [J/psi]-meson candidates decaying into muons in the large data set. Event yields are extracted from maximum-likelihood fits to dimuon invariant-mass distributions over short data-taking intervals. Event rates as a function of the LHC proton beam separation provide absolute calibration. This new method, together with corrections for pile-up and efficiency variations, yields a relative uncertainty of less than 2% with respect to reference measurements.

The dissertation presents a search for the rare Higgs boson decay into a Z boson and a J/ψ meson, with the Z decaying into a muon pair and the J/ψ into an electron or muon pair. A machine-learning (ML) discriminant is trained

to suppress background from misidentified event candidates. A two-dimensional unbinned maximum-likelihood fit using the 4-lepton invariant mass and the ML discriminant isolates the Higgs signal from background. According to Standard-Model predictions, no signal is expected in collision data taken during 2016–2018. Using the modified frequentist method, an expected upper limit on the decay branching ratio is estimated to be 10^{-3} at 95% confidence. This corresponds to about 400 times the Standard-Model value.

Table of Contents

1 Introduction 1

1.1 The Standard Model and Collider Measurements . . . 1

1.2 Rare Higgs Decays and Light-Quark Couplings . . . 3

1.3 Luminosity . . . 6

1.4 CMS Luminosity Instrumentation . . . 8

1.5 Dissertation Overview . . . 8

2 Detector Calibration and Luminosity Measurements at CMS 10

2.1 The CMS Detector . . . 11

2.1.1 The Large Hadron Collider . . . 11

2.1.2 CMS Overview . . . 12

2.1.3 Coordinate System . . . 13

2.1.4 Magnet . . . 14

2.1.5 Silicon Tracker . . . 14

2.1.6 Electromagnetic Calorimeter . . . 16

2.1.7 Hadron Calorimeter . . . 16

2.1.8 Muon System . . . 17

2.1.9 Trigger and Data Acquisition . . . 18

2.1.10 Particle-Flow Reconstruction . . . 18

2.1.11 Connection to the Dissertation Analyses . . . 19

2.2 Luminosity Observables . . . 20

2.3 Pixel Luminosity Telescope . . . 21
2.3.1 Telescope Geometry and Operational Role . . . 22
2.3.2 Triple-Coincidence Luminosity Observable . . . 23
2.3.3 Radiation Tolerance and Run 3 Operation . . . 24
2.4 Absolute Calibration . . . 25
2.4.1 Beam-Overlap Formalism . . . 26
2.4.2 Visible Cross Section from vdM Scans . . . 27
2.4.3 Scan Selection, Fit Models, and Averaging . . . 29
2.5 Beam-Related Corrections . . . 30
2.5.1 Background and Bunch-Current Corrections . . . 31
2.5.2 Length Scale and Orbit Drift . . . 32
2.5.3 Beam-Beam Corrections . . . 33
2.6 Detector-Related Corrections . . . 34
2.6.1 Efficiency Evolution . . . 35
2.6.2 Relative Linearity . . . 36
2.6.3 Independent Nonlinearity Extraction . . . 37
2.6.4 Cross-Detector Stability and Integration . . . 37
2.6.5 Summary of Corrections and Uncertainties . . . 38

3 Measurement of Integrated Luminosity in LHC Run 3 Using J/psi to Dimuons 41
3.1 Motivation . . . 41
3.1.1 Precision Luminosity in Run 3 . . . 41
3.1.2 Standard-Candle Counting . . . 42
3.1.3 Definitions and Analysis Strategy . . . 43
3.2 Experimental Setup and Data Selection . . . 46
3.2.1 CMS Detector Configuration for the Analysis . . . 46
3.2.2 Dataset and Trigger . . . 46
3.2.3 Event Selection and Reconstruction . . . 47

3.3 Signal Extraction Method . . . 49
3.3.1 Data Aggregation and Binning . . . 49
3.3.2 Invariant-Mass Fitting . . . 49
3.4 Calibration of the J/psi Luminosity Method in 2023 . . . 50
3.4.1 Pileup-Dependent Efficiency Correction . . . 51
3.4.2 Emittance-Scan Calibration . . . 55
3.4.3 Long-Term Stability Corrections in 2023 . . . 59
3.5 Application of the Calibrated Method to 2024 . . . 62
3.5.1 2024 Absolute Calibration . . . 64
3.5.2 2024 Long-Term Corrections . . . 66
3.6 Systematic Uncertainties . . . 67
3.6.1 Fit Model Uncertainty . . . 68
3.6.2 Visible Cross Section, Pileup, and Long-Term Uncertainties . . . 71
3.7 Results . . . 72
3.7.1 Lumisection Coverage . . . 72
3.7.2 Integrated Luminosity and Reference Comparison . . . 72
3.8 Summary . . . 73

4 Upper Limit on Rare Higgs Boson Decays to ZJ/psi 74
4.1 Theory Motivation and Polarization Convention . . . 75
4.1.1 Higgs Production and Analysis Scope . . . 75
4.1.2 Exclusive H to ZJ/psi Decays . . . 75
4.1.3 Polarization and Acceptance . . . 76
4.2 Data and Simulated Samples . . . 78
4.3 Analysis Strategy . . . 79
4.4 Event Reconstruction and Selection . . . 81
4.5 Machine-Learning Discriminant . . . 85
4.6 Signal and Background Modeling . . . 91
4.6.1 Invariant-Mass Modeling . . . 91

4.6.2 ML-Score Modeling 91
4.6.3 Two-Dimensional Diagnostic Configuration 94
4.7 Systematic Uncertainties 97
4.7.1 Impact of Systematic Uncertainties 99
4.8 Upper Limits 101
4.9 Summary 101

5 Conclusions 103

Bibliography 106

Vita 116

List of Tables

2.1 Corrections and uncertainties in the 2022 CMS luminosity calibration. 40

3.1 HLT filters associated with the `HLT_Mu0_L1DoubleMu` path used for the `ParkingDoubleMuonLowMass` dataset. 47

3.2 MC samples used to validate the pileup-dependent efficiency trend. . 55

3.3 Inputs to the 2023 J/ψ visible-cross-section calculation from emittance scans. 58

3.4 Stability regions used for the 2023 long-term slope correction and RAMSES residual normalization. Region 2023c contains the calibration period and is used as the reference region. The L_{int} range is the cumulative Run 3 integrated-luminosity coordinate used to locate bins along the data-taking timeline; it is not the final J/ψ-luminosity result for each region. 60

3.5 Normalized BPIX Layer 4 slopes used in the long-term efficiency correction. Slopes are expressed as fractional change per fb^{-1}. 61

3.6 Inputs to the 2024 J/ψ visible-cross-section calculation from emittance scans. 66

3.7 Stability regions used for the 2024 long-term slope correction and RAMSES residual normalization. Region 2024b2 contains the calibration period and is used as the reference region. The L_{int} range is the cumulative Run 3 integrated-luminosity coordinate; it is not the final J/ψ-luminosity result for each region. 67

3.8 Relative uncertainties in the J/ψ luminosity measurement. 68
3.9 Breakdown of the relative visible-cross-section uncertainty assigned to $\sigma_{\mathrm{vis}}^{J/\psi}$. 70
3.10 Final integrated luminosity results and retained `brilcalc` reference luminosities on the same lumisection coverage. 73

4.1 Reference-channel polarization acceptance benchmarks 78
4.2 Run 2 data eras and DAS dataset patterns used in the rare-Higgs analysis 78
4.3 Preselection requirements by final state. Masses are in GeV. 82
4.4 Soft-muon identification requirements used for selected muons. 82
4.5 Trigger paths used in the rare-Higgs analysis. 82
4.6 Final candidate-selection requirements used for the upper-limit inputs. Masses and transverse momenta are in GeV. 83
4.7 Selection yields after each selection step for the $Z(e^+e^-)J/\psi(\mu^+\mu^-)$ channel 84
4.8 Selection yields after each selection step for the $Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$ channel 85
4.9 Numerical denominator inputs for the signal normalization 85
4.10 Input variables used in the random forest classifiers. Both channels use the same set of physics observables. 86
4.11 Fitted signal mass-shape parameters for the Crystal Ball plus Gaussian model in each channel. Uncertainties reflect the fit precision from the signal Monte Carlo sample. 93
4.12 Fitted four-lepton mass background parameters. The background is modeled with an exponential function in each channel. 93
4.13 Fitted double-exponential signal ML-score parameters. The fit uses the random forest score over the interval $0 < s < 1$. 95
4.14 Fitted double-exponential background ML-score parameters. The fit uses the random forest score over the interval $0 < s < 1$. 95

4.15 Linear correlation coefficients between the four-lepton mass and the random forest score in the held-out signal test samples, quoted as percentages. 95
4.16 Experimental and analysis-model normalization nuisance parameters included in the statistical model. The magnitude column gives the fractional uncertainty applied as the log-normal width. 97
4.17 Theory uncertainties for the gluon-fusion signal normalization. 97
4.18 Expected 95% confidence-level upper limits on $\mathcal{B}(H \to ZJ/\psi)$ in the longitudinal-polarization interpretation. The 2D combined input gives the strongest expected sensitivity. 101
4.19 Expected 95% CL upper limits on $\mathcal{B}(H \to ZJ/\psi)$ for the combined ($Z \to e^+e^-$ plus $Z \to \mu^+\mu^-$) result in the longitudinal-polarization interpretation, with $\pm 1\sigma$ and $\pm 2\sigma$ quantile bands. Values are in units of 10^{-3}. 102
4.20 Luminosity scaling projection for the combined two-dimensional expected upper limit, assuming only $1/\sqrt{\mathcal{L}}$ scaling. 102

List of Figures

1.1 Elementary particles of the Standard Model 2
1.2 Standard Model Higgs-boson branching ratios 4
1.3 Direct and indirect amplitudes for $H \to ZJ/\psi$ 5

2.1 CERN accelerator complex. 12
2.2 Labelled cutaway drawing of the CMS detector. 13
2.3 Cross-section layout of the CMS Phase-1 pixel detector 15
2.4 CMS detector slice showing particle signatures. 19
2.5 Schematic of the PLT triple-coincidence geometry 21
2.6 Single-Gaussian fit of a Run 3 vdM scan curve. 28
2.7 PLT background measurements during super-separation scans in fill 8999. 32
2.8 Beam-beam deflection correction during a 2022 vdM scan. 34
2.9 Dynamic-β correction during a 2022 vdM scan. 34
2.10 PLT visible cross section from emittance scans versus accumulated luminosity in 2023. 35
2.11 Relative-linearity comparison of PLT and BCM1F in fill 8385. 36
2.12 Per-bunch scaled visible cross section versus SBIL. 38
2.13 Test-nonlinearity scan for a 2023 PLT emittance scan. 39
2.14 Extraction of the PLT nonlinearity from the zero-gradient point. . . . 40

3.1 Analysis chain for the J/psi-based luminosity measurement 45

3.2 Transverse-momentum distributions of the reconstructed muons used to form J/ψ candidates. The selection requires both muons to satisfy $p_{\mathrm{T}} > 5$ GeV. 48
3.3 Pseudorapidity distributions of the reconstructed muons used to form J/ψ candidates. Both muons are required to satisfy $|\eta| < 2.4$. 48
3.4 Representative dimuon invariant-mass fit in a 50-LS bin. The signal is modeled with a Crystal Ball plus Gaussian function, and the background is modeled with a first-order Chebyshev polynomial. . . . 51
3.5 Raw J/ψ rate and ratio to the reference luminosity in the 2023c calibration region before pileup, visible-cross-section, and long-term efficiency corrections. 52
3.6 Pileup-loss coefficients extracted from individual 2023 calibration runs. The dashed line indicates the weighted-average coefficient, and the shaded band represents its uncertainty. 53
3.7 Pileup-corrected J/ψ rate and ratio to the reference luminosity in the 2023c calibration region before visible-cross-section and long-term corrections. The relative standard deviation in the calibration region is reduced from 4.45% to 0.93% after the pileup correction. 54
3.8 Beam-separation values used in the selected 2023 nine-point emittance scans. 56
3.9 Time history of the J/ψ counting rate during a representative emittance scan. Highlighted intervals indicate stationary separation steps; unshaded intervals correspond to beam adjustment periods. 56
3.10 Per-scan J/ψ counting-rate estimate for the selected 2023 emittance scans as a function of beam separation. 57
3.11 Invariant-mass fit to the combined $m_{\mu\mu}$ distribution for the 2023 head-on separation step in the X scan. Events from all selected scans are combined before fitting. 57
3.12 J/psi fitted rates versus beam separation for the 2023 emittance scans 58

3.13 Distribution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ in the 2023c calibration region after the pileup correction and 2023 visible-cross-section calibration, before long-term corrections. 59
3.14 Time evolution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ in 2023 after pileup correction and visible-cross-section calibration. Stable running periods and the calibration-anchor region are indicated. 60
3.15 Observables motivating the long-term correction model. The BPIX Layer 4 cluster-charge evolution is used as the nominal proxy, and a trigger-ratio study tests the same trend independently. 62
3.16 RAMSES-based residual normalization in 2023. Region 2023c is fixed as the calibration reference while neighboring regions are adjusted for continuity. 63
3.17 Final 2023 long-term correction. The corrected $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ ratio has a relative standard deviation of 1.3%. 64
3.18 Raw J/psi rate and ratio in the 2024b2 calibration region 65
3.19 J/ψ fitted rates as a function of beam separation for the combined 2024 emittance scans. Gaussian fits determine the effective beam overlap widths and peak rates used in the 2024 visible-cross-section calculation. 66
3.20 Time evolution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ in 2024 after the pileup correction and 2024 visible-cross-section calibration. Region boundaries used for long-term correction are overlaid. 67
3.21 RAMSES-based residual normalization in 2024. Region 2024b2 is fixed as the calibration reference while neighboring regions are adjusted for continuity. 68
3.22 Final 2024 long-term correction. The corrected $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ ratio has a relative standard deviation of 2.2%. 69
3.23 Intrinsic fit bias as a function of the number of generated signal events in pseudo-experiments. 70

3.24 Validation of the fitted signal yield using pseudo-experiments for a representative sample size. 71

4.1 Normalized helicity-angle distributions 77

4.2 Ranking of the machine-learning input variables by separation power for the $Z(ee)J/\psi(\mu\mu)$ channel (top) and the $Z(\mu\mu)J/\psi(\mu\mu)$ channel (bottom). The rankings identify the variables most useful for discriminating the rare Higgs signal hypothesis from the background model in each channel. 88

4.3 ROC curves for the random forest classifiers evaluated on held-out test samples. The curves show the tradeoff between signal efficiency and background rejection for the mixed-flavor channel (a) and the four-muon channel (b). 89

4.4 Background-sample input-variable correlation matrices for the random forest classifiers. The matrices are used to check whether the training inputs contain strong internal correlations that could affect the stability of the discriminant shape. 90

4.5 Four-lepton invariant-mass signal and background parameterizations for the $Z(ee)J/\psi(\mu\mu)$ channel. The signal model is fitted to simulated signal MC, and the background model is fitted to data in the sideband-based background sample. 92

4.6 Four-lepton invariant-mass signal and background parameterizations for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The signal model is fitted to simulated signal MC, and the background model is fitted to data in the sideband-based background sample. 92

4.7 Illustrative pseudo-data four-lepton invariant-mass distributions for the blinded model. The displayed background spectrum is sampled from the fitted background model, while the signal shape is modeled with the Crystal Ball plus Gaussian parameterization. 93

4.8 ML discriminant parameterization for the $Z(ee)J/\psi(\mu\mu)$ channel. The signal distribution is shown in (a) using simulated signal MC events; the background distribution is shown in (b) using background candidates obtained with the loose selection criteria. Fitted double-exponential models are overlaid. 94
4.9 ML discriminant parameterization for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The signal distribution is shown in (a) using simulated signal MC events; the background distribution is shown in (b) using background candidates obtained with the loose selection criteria. Fitted double-exponential models are overlaid. 94
4.10 Diagnostic pseudo-data projections of the two-dimensional fit onto the four-lepton mass (a) and random forest score (b) axes for the $Z(ee)J/\psi(\mu\mu)$ channel. The fitted model is overlaid on the pseudo-data in both projections. 96
4.11 Diagnostic pseudo-data projections of the two-dimensional fit onto the four-lepton mass (a) and random forest score (b) axes for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The fitted model is overlaid on the pseudo-data in both projections. 96
4.12 First page of the combined impact plot for $H \to ZJ/\psi$. The nuisance parameters with the largest post-fit impact on the combined expected upper limit are ranked by magnitude and shown with their $\pm 1\sigma$ post-fit variations. 100

Chapter 1

Introduction

The Compact Muon Solenoid (CMS) experiment at the CERN Large Hadron Collider (LHC) records proton-proton collision data to test the Standard Model (SM) of particle physics and to search for phenomena beyond it. To interpret these data, CMS needs both reconstructed collision events and the integrated luminosity recorded by the experiment. The integrated luminosity converts an observed event yield into a production cross section and enters rate measurements, searches, and upper-limit interpretations based on CMS data.

This dissertation focuses primarily on luminosity measurement with CMS data. It includes detector calibration and luminosity work in the CMS Beam Radiation, Instrumentation, and Luminosity (BRIL) project and develops an independent $J/\psi \to \mu^+\mu^-$ luminosity measurement for Run 3. A secondary analysis studies a rare-Higgs upper-limit search that uses the calibrated Run 2 luminosity in the expected signal-yield normalization.

1.1 The Standard Model and Collider Measurements

The Standard Model describes the known elementary particles and their electromagnetic, weak, and strong interactions [1, 2]. Matter is represented by quarks and leptons, while interactions are mediated by gauge bosons. The Higgs mechanism

provides a field through which elementary particles acquire mass [3], and the associated Higgs boson was discovered by the ATLAS and CMS experiments in 2012 [4, 5]. Subsequent measurements established its spin-zero nature, measured its mass, and constrained its couplings to Standard Model particles [6–11]. A schematic summary of the particle content of the Standard Model is shown in Fig. 1.1.

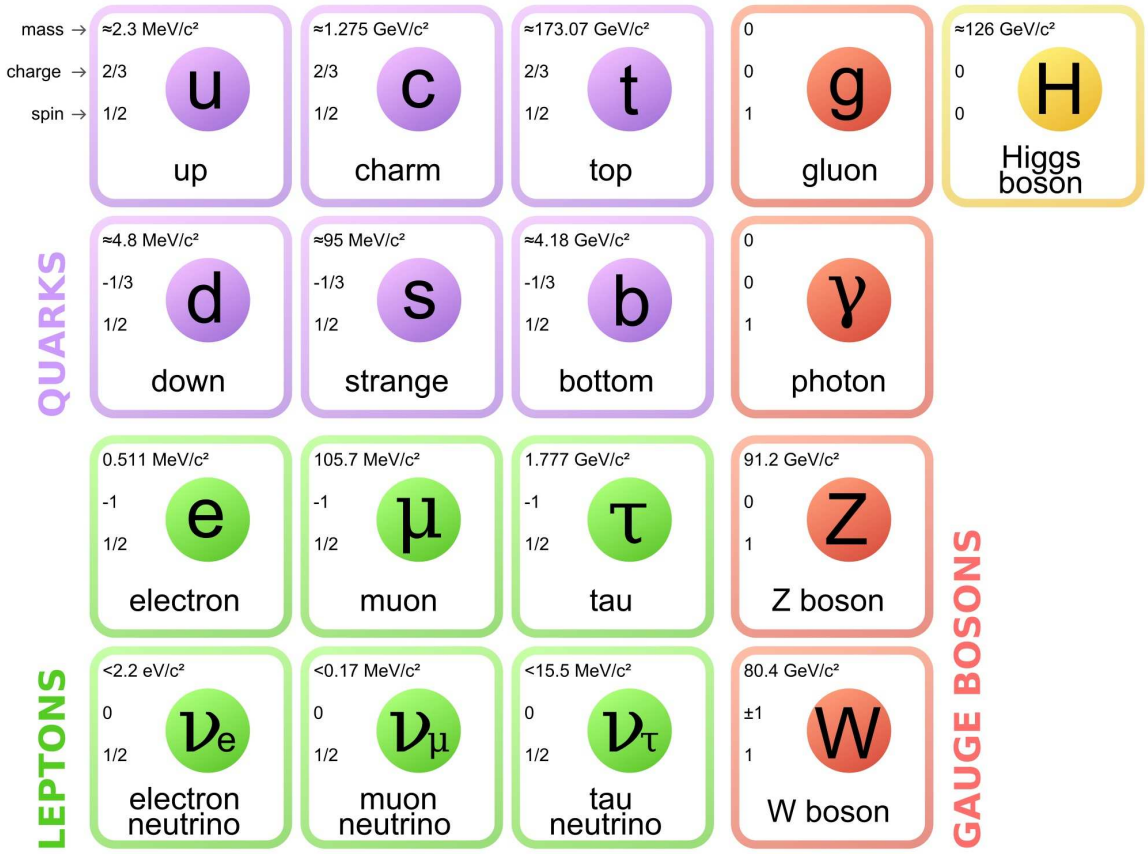


Figure 1.1: Elementary particles of the Standard Model, grouped into fermions and bosons [12].

Although the SM has been tested with high precision, it does not explain all observed phenomena. It does not include a particle candidate for dark matter, does not explain the origin of the observed hierarchy of fermion masses and mixings, and does not incorporate gravity. These limitations motivate precision measurements and searches for rare processes at the LHC [13–15].

The LHC accelerates counter-rotating proton beams and collides them at four main interaction points. During Run 3, proton-proton collisions are delivered at a center-of-mass energy of $\sqrt{s} = 13.6$ TeV. The CMS detector is located at one of these interaction points and is designed to reconstruct a broad range of final states using a silicon tracker, electromagnetic and hadronic calorimeters, a large superconducting solenoid, and muon detectors [16, 17].

Collider measurements of this kind share a common structure. A physics process is selected through reconstruction and identification requirements, the selected event yield is counted, and the yield is interpreted through the known efficiency and acceptance of the selection and through the integrated luminosity of the data sample. The first two ingredients are properties of the detector and the analysis; the third is a property of the delivered beams and of the luminosity calibration. Any improvement in the precision or reliability of the luminosity determination therefore benefits every rate measurement and search performed with the same data.

1.2 Rare Higgs Decays and Light-Quark Couplings

In the Standard Model, the Higgs boson couples to fermions through Yukawa interactions, which generate the fermion masses. The coupling of the physical Higgs boson to a fermion is proportional to that fermion's mass. Couplings to third-generation fermions have been established through the observations of $H \to \tau^+\tau^-$, $H \to b\bar{b}$, and Higgs-boson production in association with top quarks. Evidence has also been reported for $H \to \mu^+\mu^-$, which directly probes the Higgs coupling to second-generation leptons [8–10]. In contrast, the Higgs couplings to first- and second-generation quarks remain much less constrained by direct measurements, because the corresponding inclusive hadronic decays are overwhelmed by background at a hadron collider.

The Standard Model branching fractions for the dominant Higgs-boson decay modes near $m_H = 125$ GeV are shown in Fig. 1.2. The large branching fraction to

$b\bar{b}$ and the substantial decays to WW, gluons, and $\tau\tau$ set the scale of the established Higgs decay program. Rare exclusive modes such as $H \to ZJ/\psi$ are far below these inclusive decay rates, which is why they are treated as upper-limit searches rather than precision measurements with the present data sample.

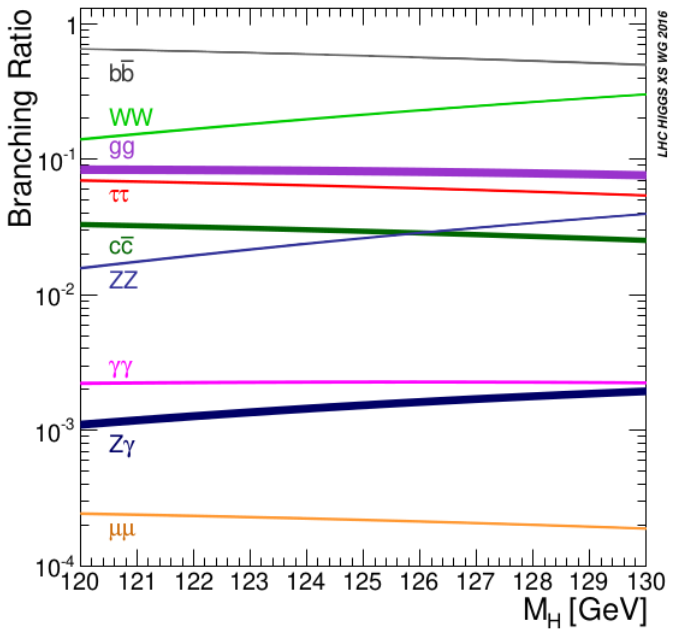
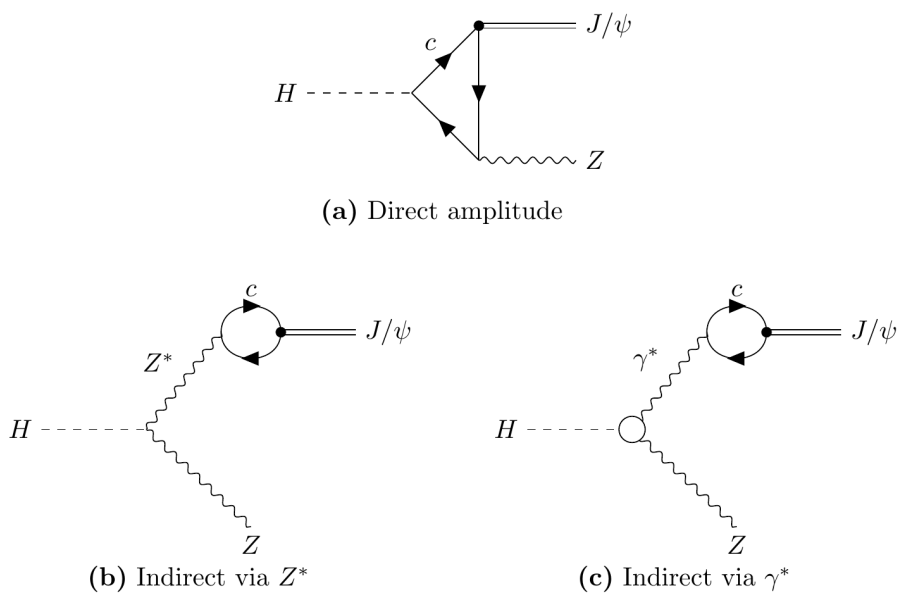


Figure 1.2: Standard Model Higgs-boson branching ratios near $m_H = 125$ GeV [18].

Rare exclusive Higgs decays to a gauge boson and a vector meson provide experimentally clean final states for probing these couplings, although the expected event yields are very small [19–25]. In decays such as $H \to ZJ/\psi$, the same final state can be reached through more than one amplitude. The direct amplitude depends on the Higgs coupling to the charm quarks that form the J/ψ, while the indirect amplitude proceeds through an intermediate gauge boson that subsequently produces the quarkonium state. Representative contributions are shown in Fig. 1.3. Because these amplitudes lead to the same final state, they interfere in the total decay rate.

The predicted SM branching fractions are of order 10^{-6}, so any observed signal at current sensitivities would indicate an enhancement from physics beyond the SM.



(a) Direct amplitude

(b) Indirect via Z^*

(c) Indirect via γ^*

Figure 1.3: Representative Feynman diagrams contributing to $H \to ZJ/\psi$. (a) Direct contribution. (b) Indirect contribution through $H \to ZZ^*$. (c) Indirect contribution through loop-induced $H \to Z\gamma^*$

Fully leptonic final states such as $H \to ZJ/\psi \to \ell^+\ell^-\mu^+\mu^-$ provide narrow resonance signatures for both intermediate particles, allowing low-rate signals to be searched for on top of smooth backgrounds. A CMS search in these final states with Run 2 data has set upper limits well above the SM expectation [26], and improved selections and discriminants are required to approach the SM sensitivity.

Because the predicted signal yields are so small, rare-decay searches of this kind depend directly on the normalization of the selected data sample. The integrated luminosity fixes the expected signal yield for a given branching fraction, and its uncertainty enters the statistical model as a normalization nuisance parameter. The rare-Higgs search therefore uses the luminosity scale developed for CMS rate measurements in a limit-setting calculation.

1.3 Luminosity

For a physics process with cross section σ, the expected event rate R is proportional to the instantaneous luminosity $\mathcal{L}$:

$$R = \mathcal{L}\sigma. \tag{1.1}$$

For detector-level observables, the corresponding rate relation is written in terms of a visible cross section, which includes the process cross section, detector acceptance, trigger and reconstruction efficiencies for the chosen observable. This form is introduced in Chapter 2, where it is used for luminometer calibration. The integrated luminosity,

$$L_{\text{int}} = \int \mathcal{L}(t)\, dt, \tag{1.2}$$

sets the normalization of cross-section measurements. A relative uncertainty in luminosity propagates directly into the normalization uncertainty of measured cross sections and into the expected signal yields used in upper-limit analyses.

CMS luminosity accounting distinguishes delivered, recorded, and certified luminosity. Delivered luminosity is provided by the accelerator to the interaction point. Recorded luminosity is the subset during which CMS is able to trigger on and store collision data. Certified luminosity is the portion that passes data-quality requirements for the relevant detector systems and physics objects. Physics results normally use recorded and certified luminosity, rather than the total delivered luminosity, because only that subset corresponds to usable reconstructed events. This distinction also matters for the J/ψ luminosity measurement in Chapter 3, where the comparison to the CMS reference luminosity is performed on the same retained lumisection coverage.

At a collider, the luminosity can also be related to beam parameters. For colliding bunches with populations N_1 and N_2, revolution frequency f_{rev}, and effective

transverse beam widths Σ_x and Σ_y, the idealized head-on luminosity is

$$\mathcal{L} = \frac{N_1 N_2 f_{\text{rev}}}{2\pi \Sigma_x \Sigma_y}. \tag{1.3}$$

In practice, CMS luminosity calibration combines accelerator information, detector rates, dedicated vdM scans, and several beam- and detector-related corrections [27, 28].

The required precision is set by the physics program. For many cross-section measurements at the LHC, the luminosity uncertainty is one of the largest individual systematic uncertainties, entering every measured rate as a common normalization factor. The CMS proton-proton luminosity has been determined with a combined Run 2 precision of 0.73% [29, 30] and a Run 3 (2022–2024) precision of 1.2% [31]. Reaching and maintaining percent-level precision requires three ingredients: an absolute calibration of each luminosity detector under controlled beam conditions, corrections for beam-related and detector-related effects that bias the calibration or its transfer to physics running, and continuous cross-checks between independent luminosity measurements. Chapters 2 and 3 of this dissertation contribute to all three ingredients.

Independent measurements are central to this strategy because each luminometer has its own acceptance, efficiency, background sensitivity, and nonlinearity. Agreement between detectors with different technologies and different systematic vulnerabilities is the main evidence that the luminosity scale is under control; disagreement localizes problems to a specific detector or period. Measurements based on reconstructed physics objects, such as the $J/\psi \to \mu^+\mu^-$ counting method developed in this dissertation, extend this redundancy beyond the dedicated luminometers, because their dominant systematic effects come from tracking, trigger, and reconstruction performance rather than from dedicated counting hardware.

1.4 CMS Luminosity Instrumentation

CMS uses several luminometers to provide real-time and offline luminosity measurements. The BRIL project operates dedicated systems that monitor collision rates independently of the primary physics trigger stream [32]. These measurements support LHC operations, detector monitoring, and the final normalization of CMS physics results.

The PLT is one of the dedicated CMS luminosity detectors. It uses small silicon pixel telescopes placed close to the beam pipe at high pseudorapidity. Charged particles originating at the CMS interaction point pass through three sensor planes in a telescope, producing a triple coincidence. The triple-coincidence rate is proportional to luminosity after calibration and corrections [33, 34]. The PLT operates alongside other BRIL luminometers with complementary technologies, including fast beam-condition monitors and luminosity algorithms based on the forward hadron calorimeter and on pixel-cluster counting; the redundancy among these systems underpins the stability and consistency checks used in the CMS luminosity determination [32].

Absolute calibration of a luminometer is obtained from special beam-separation scans. In a vdM scan, the two beams are displaced with respect to each other in the transverse directions while the detector rate is measured. The scan shape determines the effective beam overlap, and the visible cross section σ_{vis} for the luminometer is derived from the measured rate and bunch populations [35, 36]. Shorter emittance scans, which can be performed during physics running, provide additional information about detector stability, linearity, and time-dependent response changes [37].

1.5 Dissertation Overview

Chapter 2 describes the detector calibration and luminosity work performed within the CMS BRIL group. It introduces the PLT detector, the zero-counting method,

vdM and emittance-scan calibration, beam-related corrections, and detector-response corrections. The chapter also summarizes the author's contributions to Run 3 vdM scan-curve fits, super-separation background measurements, PLT emittance-scan monitoring, and PLT efficiency-evolution and nonlinearity studies.

Chapter 3 presents the central analysis of the dissertation: an independent integrated-luminosity measurement for the 2023 and 2024 data-taking periods using reconstructed $J/\psi \to \mu^+\mu^-$ decays in the Run 3 `ParkingDoubleMuonLowMass` data stream. The method extracts the J/ψ yield from invariant-mass fits, applies pileup-dependent and long-term efficiency corrections, calibrates the visible cross section with emittance scans, and compares the result with the CMS reference luminosity on the same lumisection coverage.

Chapter 4 presents a secondary analysis of the rare decay topology $H \to ZJ/\psi \to \ell^+\ell^-\mu^+\mu^-$. It uses the full Run 2 dataset, combines the electron and muon decay channels of the Z boson, and uses four-lepton invariant-mass and machine-learning discriminant inputs to set blinded 95% confidence-level upper limits on the direct $H \to ZJ/\psi$ channel. The calibrated Run 2 luminosity enters the signal-yield normalization and is propagated as a normalization uncertainty.

Chapter 2

Detector Calibration and Luminosity Measurements at CMS

The CMS luminosity program uses dedicated instrumentation, accelerator information, and offline calibration procedures to determine the luminosity delivered to the experiment. This chapter describes detector calibration and luminosity work performed within the CMS Beam Radiation, Instrumentation, and Luminosity (BRIL) project during Run 3, with emphasis on the Pixel Luminosity Telescope (PLT). The work includes PLT operation and monitoring, studies of PLT efficiency evolution and linearity, and calibration studies that support precision luminosity measurements with independent instruments.

The chapter first introduces the CMS detector systems relevant to luminosity measurements and to the reconstructed $J/\psi \to \mu^+\mu^-$ analysis used later in this dissertation. It then develops the luminosity formalism used for rate scaling, zero counting, vdM calibration, and emittance-scan monitoring. The final sections summarize Run 3 luminosity calibration results from the 2022 calibration campaign [28] and present analysis results obtained in this work from the 2022–2023 scan programs.

2.1 The CMS Detector

The Compact Muon Solenoid (CMS) experiment is a general-purpose detector located at one of the proton-proton collision points of the Large Hadron Collider (LHC). Its design is organized around precise charged-particle tracking, high-resolution electromagnetic calorimetry, hadronic calorimetry, and muon identification over a broad pseudorapidity range. These detector systems enter the luminosity studies in this dissertation in two ways: the central tracker and muon system reconstruct the low-momentum muons used in the J/ψ luminosity analysis, while the dedicated luminosity detectors measure collision rates that are converted to luminosity through beam-based calibration scans.

2.1.1 The Large Hadron Collider

The LHC is a superconducting proton and heavy-ion collider installed in a circular tunnel of approximately 27 km circumference at CERN. Two counter-rotating beams are accelerated and brought into collision at four main interaction regions, which host the ATLAS, CMS, ALICE, and LHCb experiments. During Run 3, CMS records proton-proton collisions at a center-of-mass energy of $\sqrt{s} = 13.6$ TeV. The beams are divided into bunches spaced by 25 ns in nominal operation, corresponding to a 40 MHz bunch-crossing frequency. Because only a subset of RF buckets are filled and because the same bunches circulate repeatedly, CMS luminosity is measured bunch by bunch and integrated fill by fill.

Figure 2.1 shows the CERN accelerator chain that injects beams into the LHC and the locations of the four main LHC experiments. For CMS luminosity measurements, this chain sets the beam preparation, bunch pattern, and collision conditions that determine the delivered luminosity in each LHC fill.

The instantaneous luminosity depends on the number of particles in the colliding bunches, the revolution frequency, and the transverse overlap of the two beams at the interaction point. The vdM and emittance-scan procedures described later in this

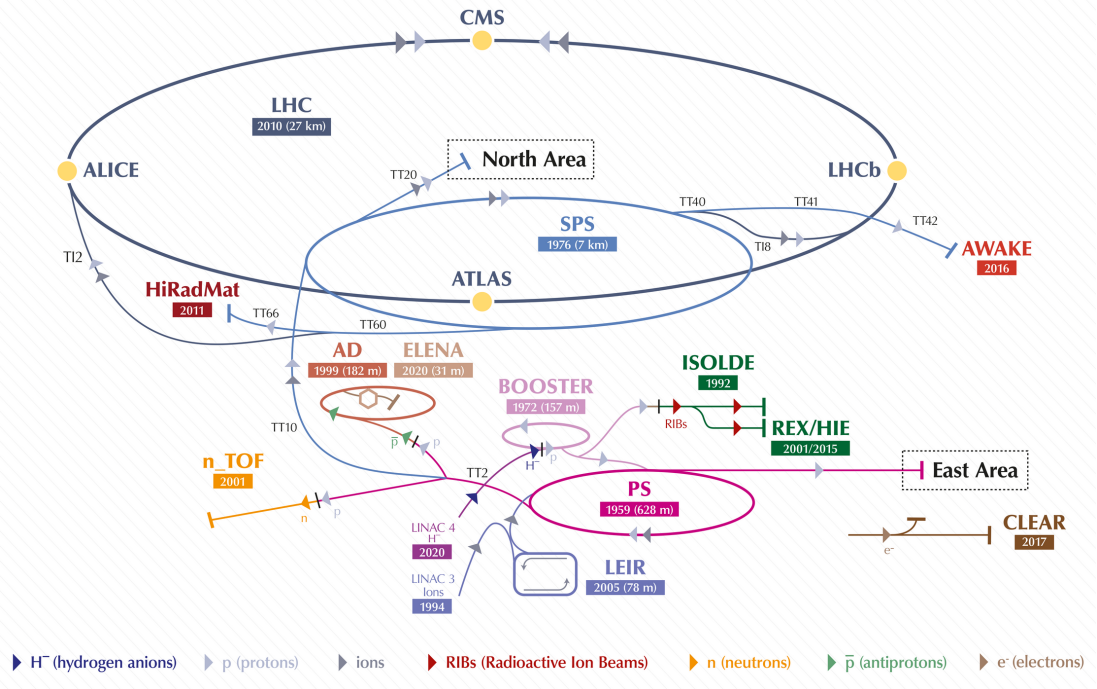


Figure 2.1: CERN accelerator complex, showing the sequence of accelerators that prepare beams for the LHC and the locations of the four main LHC experiments [38].

chapter use this dependence to relate measured detector rates to the beam conditions during calibration scans.

2.1.2 CMS Overview

CMS surrounds the interaction point with a nearly cylindrical, layered detector. Moving outward from the beam pipe, particles first traverse the silicon tracker, then the electromagnetic calorimeter (ECAL), the hadron calorimeter (HCAL), the superconducting solenoid, and finally the muon detectors embedded in the steel return yoke, as shown schematically in Fig. 2.2. The solenoid provides a magnetic field of 3.8 T, bending charged particles in the transverse plane so that their momenta can be measured from track curvature [16]. Run 3 operation included maintenance and upgrades to online reconstruction, trigger, data acquisition, and luminosity instrumentation [17].

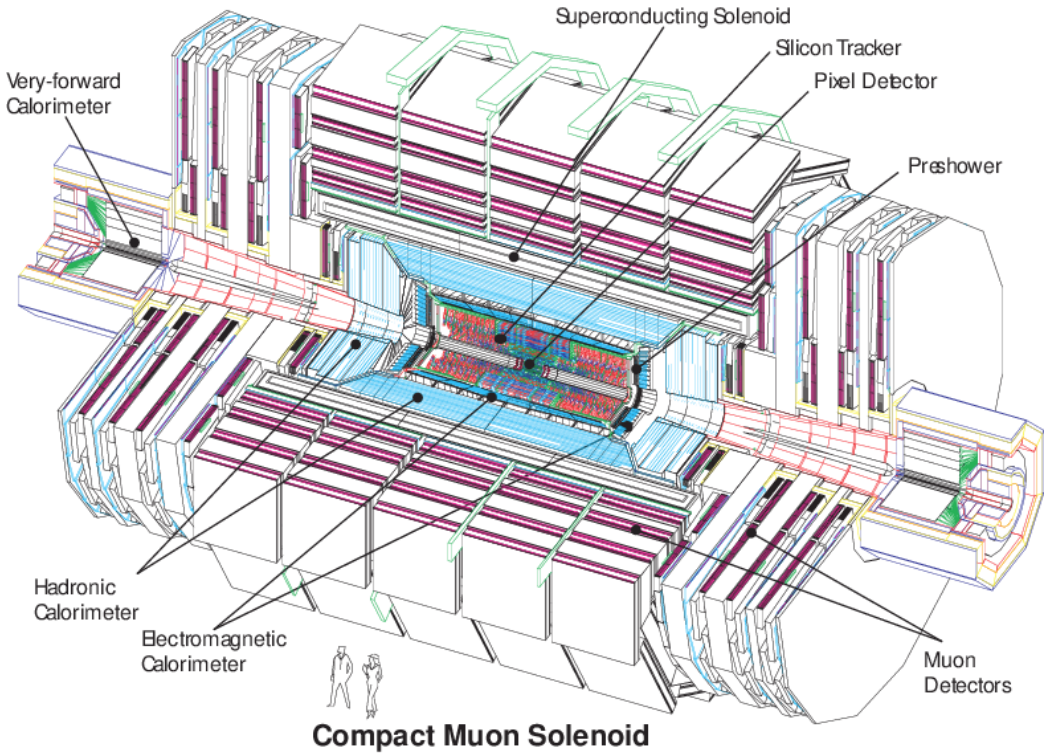


Figure 2.2: Labelled cutaway drawing of the CMS detector, showing the tracker, calorimeters, superconducting solenoid, forward calorimeters, and muon detectors [39].

The central design choice of CMS is the use of a large-bore high-field solenoid that encloses the tracker and calorimeters. This configuration gives precise momentum measurement for charged particles before they enter the calorimeters and keeps the muon system outside the coil, where the steel return yoke provides both magnetic flux return and absorber material. The arrangement is particularly effective for final states containing muons, photons, electrons, jets, missing transverse momentum, or combinations of these objects.

2.1.3 Coordinate System

CMS uses a right-handed coordinate system with its origin at the nominal interaction point. The x axis points radially inward toward the center of the LHC ring, the y axis points upward, and the z axis follows the counterclockwise beam direction.

The azimuthal angle ϕ is measured in the transverse x-y plane, and the polar angle θ is measured from the positive z axis. The pseudorapidity used throughout the dissertation is

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \tag{2.1}$$

Angular separations are commonly expressed as $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. Momenta and energies transverse to the beam line are denoted by p_{T} and E_{T}, respectively. This coordinate convention is used throughout the dissertation for detector acceptance, track selection, muon kinematics, and luminosity comparisons.

2.1.4 Magnet

The CMS magnet is a superconducting solenoid designed to provide a strong and uniform axial magnetic field across the inner tracking and calorimeter volumes. The nominal field of 3.8 T allows the tracker to measure the curvature of charged particles with high precision over a compact radial distance. Outside the solenoid, the magnetic flux returns through a massive steel yoke. The yoke is instrumented with muon chambers, so it also serves as absorber material that filters hadronic shower remnants before particles reach the outer muon system [16].

For the $J/\psi \to \mu^+\mu^-$ luminosity method, the magnetic field is essential because the invariant mass of the dimuon candidate depends directly on the reconstructed momenta of the two charged tracks. Stable field knowledge and alignment therefore support both the mass resolution of the resonance and the control of kinematic selection efficiencies.

2.1.5 Silicon Tracker

The silicon tracker is the innermost major CMS subsystem and provides the primary measurements used to reconstruct charged-particle trajectories and interaction vertices. It is built from pixel detectors close to the beam pipe and silicon strip

detectors at larger radii. The pixel detector supplies three-dimensional space points with high granularity near the interaction region, while the strip tracker extends the lever arm for momentum measurement. The tracker covers the central and forward regions relevant for most CMS physics analyses and is designed to operate in the high-occupancy environment produced by many simultaneous proton-proton interactions in the same bunch crossing.

The Phase-1 pixel detector, installed before Run 2 and used through Run 3, increased the number of barrel pixel (BPIX) layers and forward pixel (FPIX) disks relative to the original detector. Its four BPIX layers and three FPIX disks on each side improve hit efficiency, pattern recognition, impact-parameter resolution, and vertex reconstruction in high-pileup data [40]. For low-p_{T} muon tracks from J/ψ decays, precise hit information near the interaction point improves the primary-vertex association and the dimuon vertex fit.

Figure 2.3 shows the barrel and forward pixel geometry that is later used to define the BPIX/FPIX categories in the J/ψ luminosity analysis.

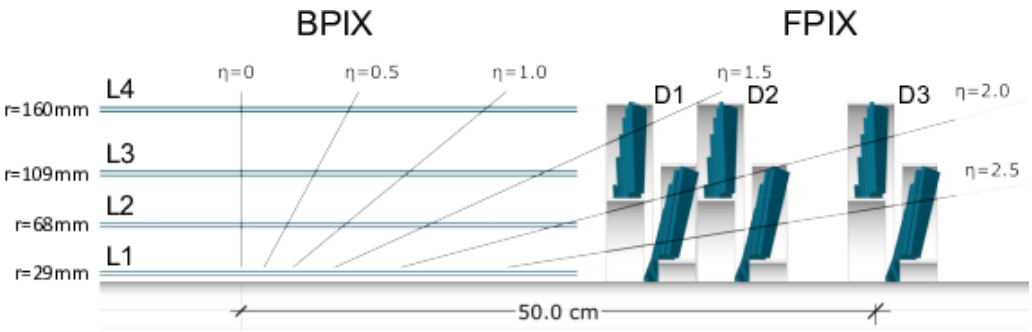


Figure 2.3: Cross-section layout of the CMS Phase-1 pixel detector showing the four BPIX layers (L1–L4) at radii of 29, 68, 109, and 160 mm and three FPIX disks (D1–D3) on each side of the interaction point, with the corresponding pseudorapidity coverage indicated [40].

Track reconstruction begins from compatible groups of hits, which are extended through the tracker using the measured magnetic field and material model. The fitted track parameters include the transverse momentum, direction, and impact parameters with respect to a chosen vertex. Vertices are reconstructed by clustering tracks that

are consistent with a common origin along the beam line. In a high-pileup event, the selected primary vertex is usually the vertex associated with the hardest reconstructed interaction.

2.1.6 Electromagnetic Calorimeter

The ECAL is located outside the tracker and measures the energy of electrons and photons. It is made from lead tungstate crystals, whose short radiation length and fast scintillation response allow compact shower containment and bunch-crossing-level timing. The ECAL consists of a barrel section and two endcaps, giving coverage over the central and forward detector regions. A preshower detector is installed in front of the endcap crystals to help distinguish photons from neutral meson decays in the forward region [16].

ECAL information is used in CMS trigger and event reconstruction, including electron and photon identification. In particle-flow reconstruction, ECAL clusters are combined with tracks and muon information to identify particles across subsystems [41]. The electron-channel component of the rare-Higgs analysis in Chapter 4 relies on this reconstruction chain.

2.1.7 Hadron Calorimeter

The HCAL measures the energy of hadrons and helps reconstruct jets, hadronic activity, and missing transverse momentum. In the barrel and endcap regions it is placed inside the solenoid, behind the ECAL, so that hadronic showers are sampled before particles reach the coil. Additional outer and forward calorimeters extend the hadronic energy measurement and improve hermeticity at large pseudorapidity. This angular coverage enters event-level quantities and helps separate beam-related and collision-related activity in different detector regions [16].

The HCAL and forward calorimeters also play a direct role in CMS luminosity operations. Dedicated luminosity algorithms use high-rate information from forward

calorimeter systems and related BRIL instrumentation to monitor bunch-by-bunch conditions during data taking. These measurements provide independent comparisons for the PLT calibration and linearity studies described later in this chapter.

2.1.8 Muon System

The muon system is installed outside the solenoid in the steel return yoke and is one of the defining features of CMS. Muons lose relatively little energy in the calorimeters and can penetrate to the outer detector layers, where gas-ionization chambers provide position and timing measurements. CMS uses drift tubes in the barrel region, cathode strip chambers in the endcaps, and resistive plate chambers in both barrel and endcap regions. Additional gas electron multiplier chambers have been added in the forward region as part of the CMS upgrade program for improved triggering and reconstruction at high rates [16, 17].

Muon reconstruction combines information from the silicon tracker and the muon chambers. A standalone muon can be reconstructed using only hits in the muon system, while a global muon fit combines muon-system information with a compatible inner-tracker track. Tracker muons are reconstructed by extrapolating inner tracks outward and matching them to muon detector segments. CMS analyses use combinations of these reconstruction categories and identification requirements to balance efficiency, momentum resolution, and background rejection [42].

Reconstructed $J/\psi \to \mu^+\mu^-$ candidates appear in both the luminosity measurement and the rare-Higgs study, so the efficiency and stability of low-p_{T} muon reconstruction are central to this dissertation. Because the J/ψ mass is narrow and the signal yield is extracted from the dimuon invariant-mass spectrum, tracker momentum resolution, muon identification, trigger efficiency, and vertex quality affect the reconstructed J/ψ yield. In Chapter 3, that yield is used as a luminosity observable; in Chapter 4, it is used to reconstruct the J/ψ candidate in the rare-Higgs final state.

2.1.9 Trigger and Data Acquisition

The LHC bunch-crossing rate is far above the rate that can be stored for offline analysis, so CMS uses a two-level trigger system. The Level-1 trigger is implemented in custom electronics and uses coarse information from the calorimeters and muon systems to select events within a fixed latency. The High-Level Trigger (HLT) is a software system running on a processor farm; it performs more detailed reconstruction and reduces the event rate to a level suitable for permanent storage [17, 43–45].

Trigger efficiency and stability enter analyses based on reconstructed physics objects. The $J/\psi \to \mu^+\mu^-$ luminosity measurement depends on the dimuon triggers used to collect the sample, or on the efficiency corrections used when more inclusive triggers are employed. The rare-Higgs study also relies on electron and muon trigger selections for the reconstructed four-lepton final state. Dedicated luminosity detectors follow a different data path, providing bunch-by-bunch rates for online monitoring and offline calibration.

2.1.10 Particle-Flow Reconstruction

CMS event reconstruction is organized around the particle-flow (PF) approach, which combines information from all subdetectors to identify and reconstruct individual stable particles in an event [41]. Tracks in the silicon tracker are linked to calorimeter clusters and muon-system segments. Charged hadrons are primarily measured using the tracker and calorimeters, photons and neutral hadrons are measured using calorimeter information, and muons are reconstructed using matched tracker and muon-system measurements. The resulting PF candidates provide the input for higher-level objects such as jets, missing transverse momentum, isolation sums, and reconstructed vertices.

Figure 2.4 illustrates how the layered CMS geometry gives different detector signatures for muons, electrons, charged hadrons, neutral hadrons, and photons.

This separation of signatures is the basis for PF reconstruction and for the object definitions used in CMS physics analyses.

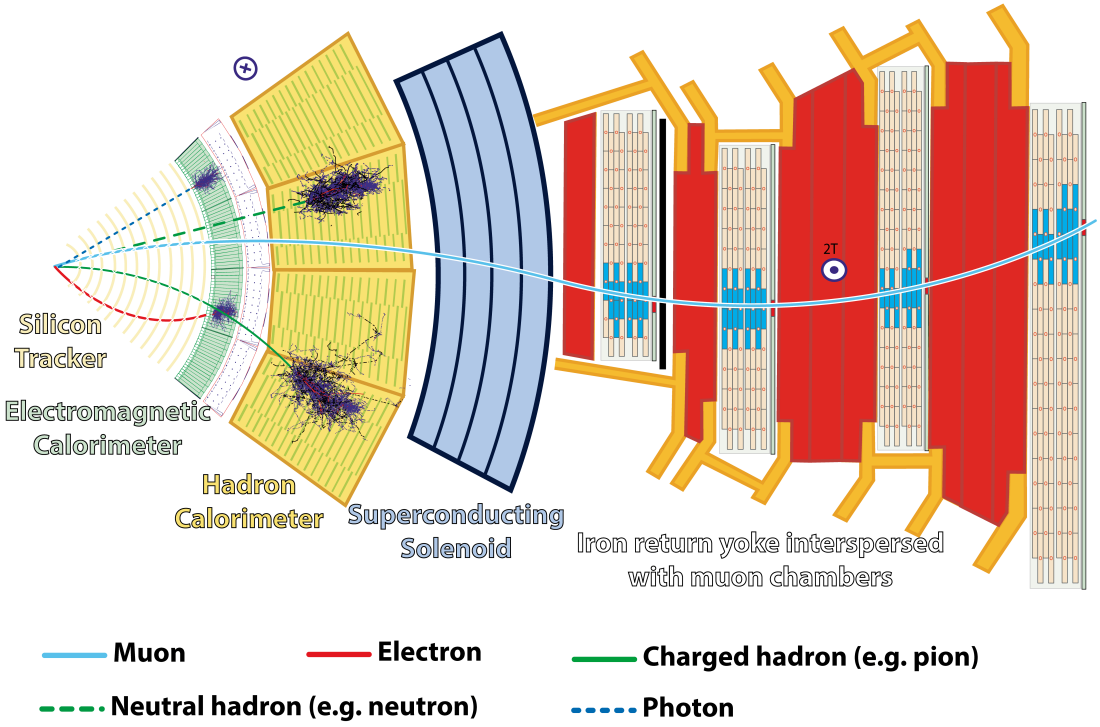


Figure 2.4: CMS detector slice illustrating the characteristic signatures of common reconstructed particles in the tracker, calorimeters, solenoid, and muon system [46].

The reconstructed-object analyses in this dissertation use the CMS reconstruction chain rather than only raw detector signals. The J/ψ luminosity measurement depends mainly on reconstructed muons and vertices, while the rare-Higgs study uses reconstructed muons, electrons, and four-lepton kinematic quantities. The detector alignment, calibration, tracking, and trigger infrastructure used in PF reconstruction therefore also controls the quality of the reconstructed objects used in the later analyses.

2.1.11 Connection to the Dissertation Analyses

The detector systems described above define the experimental inputs used in the later chapters. The PLT and other BRIL luminometers measure collision rates with

dedicated hardware for the calibration studies in this chapter. The central tracker and muon system reconstruct the $J/\psi \to \mu^+\mu^-$ candidates used in the luminosity measurement of Chapter 3. The tracker, ECAL, and muon system reconstruct the electrons, muons, and four-lepton kinematic quantities used in the rare-Higgs study of Chapter 4. In each case, the analysis depends on detector calibration: rate measurements require stability and linearity corrections, while reconstructed-object analyses require controlled trigger, reconstruction, and identification efficiencies.

2.2 Luminosity Observables

CMS luminosity measurements begin with detector observables that are proportional to the collision rate. The observable may be a direct count of reconstructed objects, an energy sum, a pixel-cluster count, an occupancy fraction, or a coincidence rate in a dedicated luminometer. In all cases, the observable can be used for luminosity only after its visible cross section, σ_{vis}, has been determined and after detector-response effects have been controlled. The visible cross section absorbs the detector acceptance, reconstruction or counting efficiency, trigger or readout response, and the probability that a proton-proton interaction produces the chosen visible signature [27, 47–49].

Two counting strategies are used throughout this dissertation. In a rate-scaling method, the measured mean observable is converted directly to luminosity using the visible-rate relation introduced in Sec. 2.4. Pixel-cluster counting and the J/ψ method in Chapter 3 follow this logic after their respective detector and efficiency corrections are applied. In a zero-counting method, the measured quantity is instead the fraction of bunch crossings with no visible object. Assuming Poisson-distributed visible counts, the empty-crossing fraction is converted to the mean visible response and then to luminosity using the same σ_{vis} formalism. The PLT triple-coincidence observable described below uses this second strategy.

The visible-cross-section calibration must be transferred from special scan conditions to normal physics running. Physics fills have higher pileup, larger rates, and

changing detector conditions, so the transfer requires stability and linearity checks before the calibrated rates are integrated.

2.3 Pixel Luminosity Telescope

The PLT is a dedicated forward luminosity monitor installed as part of the CMS BRIL instrumentation. It was deployed for Run 2 and rebuilt for Run 3 as part of the CMS luminosity and machine-induced-background monitoring program. The Run 2 configuration, which provides the reference geometry for the PLT operation discussed here, consisted of 48 silicon sensors arranged into sixteen telescopes, with eight telescopes installed around the beam pipe on each side of the CMS interaction point, approximately 1.75 m from the interaction point. Each telescope contains three silicon sensor planes positioned so that particles produced near the CMS interaction point can pass through all three planes [33, 34, 50, 51]. A sketch of the telescope geometry is shown in Fig. 2.5.

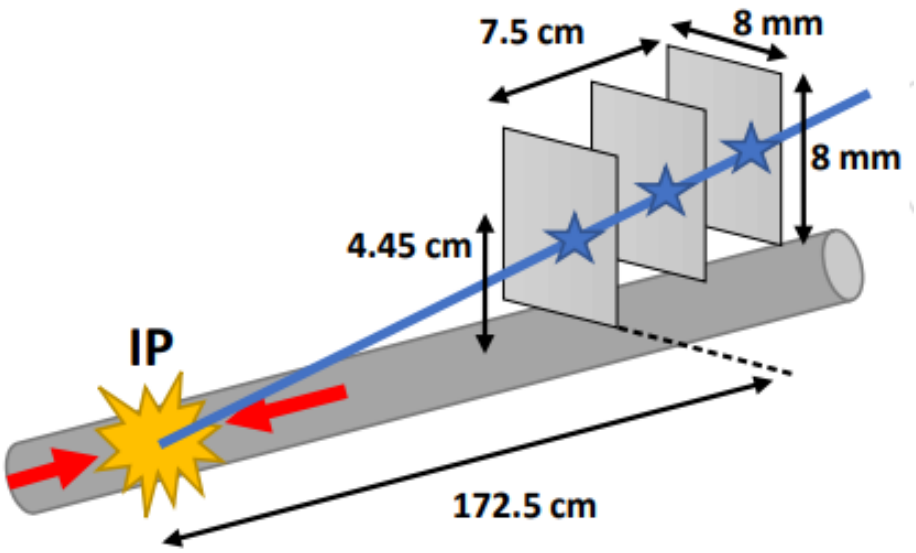


Figure 2.5: Schematic of the PLT triple-coincidence geometry [34]. A charged particle from the CMS interaction point crosses three sensor planes in one telescope.

The PLT uses two complementary data paths. The luminosity path is a fast-OR readout in which each sensor reports whether at least one pixel is above threshold in a 25 ns bunch crossing. A triple coincidence occurs when all three planes of a telescope record a fast-OR signal in the same bunch crossing. The fast-OR front-end driver histograms these coincidences by telescope and bunch crossing, allowing bunch-by-bunch luminosity monitoring at the 40 MHz LHC bunch-crossing frequency. A second, lower-rate full-pixel readout records hit positions and pulse-height information. These full pixel data are not the primary online luminosity observable, but they are used for alignment, active-area studies, beam-induced background studies, and detector-performance monitoring [51].

2.3.1 Telescope Geometry and Operational Role

The PLT geometry is optimized for particles emerging from the CMS interaction point at very forward pseudorapidity. The three planes in a telescope form a short lever arm, so a collision particle that crosses all three planes produces a geometrically consistent track-like signature. The telescopes are grouped into quadrants around the beam pipe, providing azimuthal coverage and allowing the luminosity measurement to be formed from multiple independent channels. This redundancy protects the luminosity measurement when individual channels have different efficiencies, noise conditions, or radiation histories.

Each PLT sensor has 80 rows and 52 columns of pixels, with a pixel pitch of 150 μm $\times$ 100 μm and a full active area of about 8×8 mm^2. In operation, a smaller central active region is used for the luminosity measurement. Restricting the active area reduces accidental coincidences from unrelated hits while preserving the region most relevant for particles originating from the interaction point [51].

In CMS operations, BRIL luminosity detectors deliver online information for LHC optimization and offline information for physics normalization. The PLT belongs to this instrumentation suite but differs from central CMS physics-object reconstruction:

its main observable is a high-rate bunch-by-bunch coincidence count, not a fully reconstructed event. The triple-coincidence requirement suppresses noise, beam-halo backgrounds, and backgrounds from activated material because uncorrelated hits are less likely to appear simultaneously in all three planes of a telescope. This separation allows PLT data to give immediate luminosity feedback with minimal dependence on the CMS trigger menu.

2.3.2 Triple-Coincidence Luminosity Observable

Because direct counting can underestimate rates when more than one charged particle crosses a telescope in the same bunch crossing, the PLT luminosity observable is based on zero counting. The fast-OR readout records whether a plane fired, but it does not distinguish all possible multiple-particle configurations in the same telescope and bunch crossing. The zero-counting method avoids this ambiguity by using the fraction of bunch crossings in which no visible object is observed [27].

The derivation begins from the Poisson probability for observing no visible PLT coincidences in a given telescope and bunch-crossing slot. If the mean number of visible coincidences is μ and the measured empty-crossing fraction is f_0, then Eqn. 2.2 gives

$$f_0 = P(0;\mu) = e^{-\mu}. \tag{2.2}$$

Solving Eqn. 2.2 for μ gives the zero-counting estimator in Eqn. 2.3:

$$\mu = -\ln(f_0). \tag{2.3}$$

For a bunch crossing i, the per-bunch rate is

$$R_i = \mu_i f_{\mathrm{rev}}, \tag{2.4}$$

where $f_{\mathrm{rev}} = 11245.6$ Hz is the LHC revolution frequency. The luminosity for that bunch crossing is the single-bunch instantaneous luminosity (SBIL):

$$L_i = \frac{R_i}{\sigma_{\mathrm{vis}}}, \tag{2.5}$$

after the visible cross section σ_{vis} has been calibrated. SBIL is therefore a per-bunch instantaneous quantity; integrated luminosity is obtained only after summing over bunch crossings and time.

The zero-counting formulation matches a detector with high bunch-crossing rate and limited per-crossing occupancy information. It converts the observed empty-crossing fraction into the mean number of visible coincidences under the assumption of Poisson-distributed counts, while absorbing the detector acceptance and the probability for an interaction to produce a visible PLT coincidence into σ_{vis}. The calculation is performed by telescope and bunch crossing before the channels are combined. At very high occupancy, the method would become statistically weak because the empty-crossing fraction approaches zero, a situation known as zero starvation. In normal PLT operation, however, the typical occupancy of about 0.2–0.3 triple coincidences per telescope and bunch crossing is low enough that zero counting provides a stable luminosity observable [51]. The resulting SBIL values can be summed over bunch crossings and time to obtain integrated luminosity after the absolute scale and detector-response corrections are applied.

2.3.3 Radiation Tolerance and Run 3 Operation

The PLT operates close to the LHC beams and is exposed to large particle fluence. Radiation damage can change sensor depletion voltage, charge collection, leakage current, timing, and effective detection efficiency. The PLT operating point must therefore be maintained with calibrations and occasional changes in high voltage or threshold settings. The Run 3 PLT rebuild and related pixel-sensor studies use

the same radiation-tolerant silicon-detector technologies developed for high-rate CMS operation.

For luminosity calibration, radiation tolerance is not only a hardware concern. Gradual response loss, per-channel efficiency changes, threshold changes, accidental coincidences, and nonlinear response at high instantaneous luminosity can all bias the calibrated rate. The detector-related corrections in Sec. 2.6 therefore include long-term efficiency evolution and linearity checks.

2.4 Absolute Calibration

The absolute scale of a luminometer is set using van der Meer (vdM) scans. The starting point is the detector-rate relation

$$R = \mathcal{L}\,\sigma_{\text{vis}}, \tag{2.6}$$

where R is the measured rate of a luminosity observable, $\mathcal{L}$ is the instantaneous luminosity, and σ_{vis} is the visible cross section for the detector and algorithm being calibrated. The visible cross section includes acceptance, efficiency, and algorithm-specific response, so it cannot be reliably predicted from simulation at the precision required for CMS luminosity measurements. Instead, CMS determines σ_{vis} experimentally using special beam-separation scans [27, 35].

Equation 2.6 is the rate-scaling form of luminosity measurement. For algorithms that count a mean number of visible objects per bunch crossing, the same relation can be written as Eqn. 2.7:

$$L_b = \langle N_{\text{obs}} \rangle \frac{f_{\text{rev}}}{\sigma_{\text{vis}}}, \tag{2.7}$$

where $\langle N_{\text{obs}} \rangle$ is the average detector observable in that bunch crossing. Rate-scaling algorithms, such as pixel-cluster or physics-object counting, rely on the average observable itself. Zero-counting algorithms, such as the PLT method in Sec. 2.3.2, first infer the mean visible response from the empty-crossing fraction and then use

the same visible-cross-section calibration. Chapter 3 applies the rate-scaling logic to reconstructed $J/\psi \to \mu^+\mu^-$ yields rather than to a dedicated hardware luminometer.

2.4.1 Beam-Overlap Formalism

For a single colliding bunch pair, the instantaneous luminosity can be written as the beam-overlap expression in Eqn. 2.8:

$$L_b = f_{\text{rev}} n_1 n_2 \iint \rho_1(x,y)\rho_2(x,y)\,dx\,dy, \tag{2.8}$$

where n_1 and n_2 are the bunch populations, f_{rev} is the LHC revolution frequency, and ρ_1 and ρ_2 are normalized transverse proton-density distributions for the two bunches. The double integral is the transverse overlap of the two bunches at the interaction point. The purpose of a vdM scan is to determine this overlap from measured rate curves as the transverse separation of the beams is varied [27].

If the bunch profiles factorize into independent horizontal and vertical components, as stated in Eqn. 2.9,

$$\rho_i(x,y) = f_i(x)g_i(y), \tag{2.9}$$

the overlap can be separated into horizontal and vertical terms. Scanning the beams across each other in x and y gives effective convolved beam widths, Σ_x and Σ_y, from the area under the rate-versus-separation curves. The scan-width definitions in Eqns. 2.10 and 2.11 use the measured luminosity, or any rate proportional to it, as a function of the horizontal separation w and vertical separation h:

$$\Sigma_x = \frac{\int L_b(w,0)\,dw}{L_b(0,0)}, \tag{2.10}$$

$$\Sigma_y = \frac{\int L_b(0,h)\,dh}{L_b(0,0)}. \tag{2.11}$$

For Gaussian scan curves this leads to Eqn. 2.12:

$$L_b = \frac{f_{\text{rev}} n_1 n_2}{2\pi \Sigma_x \Sigma_y}. \tag{2.12}$$

The Gaussian expression is a compact form, but the vdM method is more general: Eqns. 2.10 and 2.11 can be evaluated by numerical integration or by fitting non-Gaussian scan functions. In either case, Σ_x and Σ_y represent the effective transverse beam-overlap widths, not intrinsic widths of either individual beam alone [27]. Fig. 2.6 illustrates the corresponding fit to a rate-versus-separation curve.

2.4.2 Visible Cross Section from vdM Scans

The beam-overlap measurement is connected to the detector observable through the head-on rate. For an idealized factorized profile, the visible cross section is proportional to the product of the convolved beam widths and the peak detector rate:

$$\sigma_{\text{vis}} = \frac{2\pi \Sigma_x \Sigma_y R_0}{N_1 N_2 f_{\text{rev}}}, \tag{2.13}$$

where N_1 and N_2 are the bunch populations, Σ_x and Σ_y are the effective beam-overlap widths, and R_0 is the detector rate at zero nominal beam separation. Equivalently, for algorithms expressed as a mean visible response per bunch crossing, σ_{vis} can be written in terms of the calibrated head-on response and the measured beam overlap. Once σ_{vis} is determined, normal physics data can be converted from detector rates to luminosity using Eqn. 2.6.

Dedicated vdM fills are operated under conditions chosen to reduce systematic effects. The CMS 2015–2016 luminosity calibration used low pileup, special high-β^* optics, a limited number of colliding bunches, separated bunch trains to reduce afterglow and long-range effects, and repeated pairs of horizontal and vertical scans to check consistency. The target pileup in those scans was about $\mu \approx 0.6$, far below typical physics-running pileup, and the crossing angle at the CMS interaction point

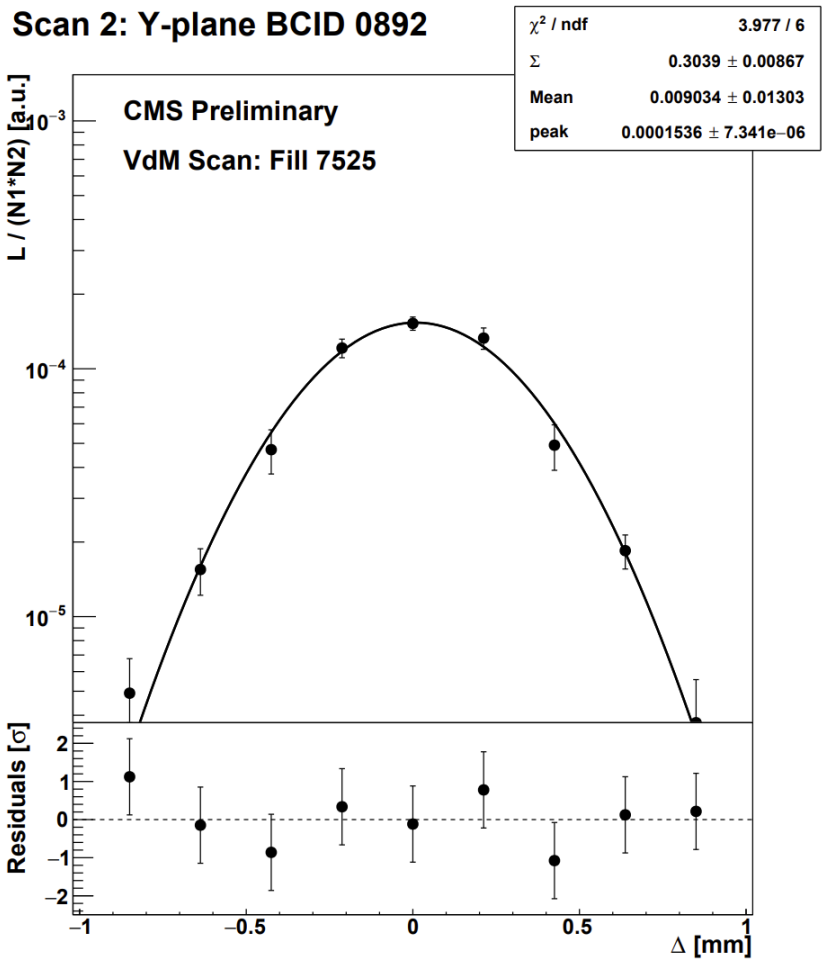


Figure 2.6: Single-Gaussian fit of a Run 3 vdM scan curve, produced in this work with the BRIL vdM framework. The measured rate, normalized to the beam-current product, is shown as a function of the beam separation in the vertical plane; the lower panel shows the difference between the measured rate and the fit divided by the statistical uncertainty.

was set to zero for the vdM scans [27]. These conditions are intentionally different from ordinary high-luminosity fills because the goal is to measure the absolute scale as cleanly as possible.

The practical calibration includes corrections for backgrounds, bunch currents, beam-beam effects, orbit drift, length scale, and possible factorization bias [27, 28, 36]. Beam current corrections account for the bunch-population product in Eqn. 2.13, including ghost and satellite charge. Orbit-drift and length-scale

corrections account for differences between nominal beam-separation settings and the actual beam displacement at the interaction point. Beam-beam corrections account for electromagnetic interactions between the two bunches, including coherent beam-beam deflection and the dynamic-β effect. Factorization uncertainties account for the fact that the real transverse bunch density may not separate exactly into independent x and y components.

Detector-specific corrections are also required when transferring the vdM calibration to physics running. Effects such as afterglow, radiation-induced efficiency evolution, threshold changes, and nonlinear response at high instantaneous luminosity can change the apparent σ_{vis} outside the special scan conditions. The PLT stability and linearity corrections described below are examples of this transfer problem.

Emittance scans are shorter vdM-like scans taken during physics fills. A dedicated vdM calibration scan typically lasts about 60 minutes, whereas an emittance scan typically lasts about 10 minutes. Emittance scans are not normally used as the primary absolute calibration of standard CMS luminometers, but they monitor detector stability and response changes as a function of time and instantaneous luminosity [37]. They also provide the calibration scans used for the J/ψ luminosity method in Chapter 3.

2.4.3 Scan Selection, Fit Models, and Averaging

A practical vdM calibration is built from a set of scan pairs rather than from a single horizontal and vertical scan. The scan program includes head-on periods, horizontal and vertical beam-separation scans, and auxiliary scans used to diagnose corrections. In a standard vdM scan both beams are moved in opposite directions to vary their relative separation, while beam-imaging-type scans move one beam at a time and provide complementary information about the beam profiles. The scan pairs are analyzed separately and then combined after applying the common correction sequence.

For each luminometer and bunch crossing, the rate is first normalized by the measured bunch-population product. The corrected rate-versus-separation curve is then fit in each transverse plane. Single-Gaussian fits give a compact description of the scan width and peak rate, while double-Gaussian fits allow non-Gaussian tails in the convolved beam profile to be described more flexibly. The fitted peak rates and effective widths are converted into σ_{vis} values for each bunch crossing. The final detector calibration is obtained by averaging over accepted bunch crossings and accepted scan pairs, with the bunch-to-bunch and scan-to-scan spreads used as consistency checks and uncertainty inputs [27, 28].

The vdM scan-curve fits for the Run 3 calibration campaigns were performed as part of this work, within the BRIL vdM analysis team. In the 2022 campaign, the calibration was based on six scan pairs recorded in a dedicated vdM fill: four standard vdM scan pairs with 25 separation steps per plane, in which both beams are moved in opposite directions, and two beam-imaging scan pairs with 19 steps, in which only one beam is moved [28]. The per-bunch scan curves were fitted with both the single-Gaussian and double-Gaussian models in the vdM framework. The double-Gaussian model generally gave the better fit quality and was used for the central calibration; its results agree with higher-order polynomial-times-Gaussian fit models within about 0.2%, while the single-Gaussian model differs by about -1%, and these model differences enter the fit-model systematic uncertainty. The bunch-to-bunch and scan-to-scan spreads serve as consistency checks on the calibration [28].

2.5 Beam-Related Corrections

Backgrounds, beam currents, length-scale effects, orbit drift, and beam-beam effects are major inputs to luminosity calibration. These corrections are applied before the final visible cross sections are averaged because each effect changes either the measured luminometer rate or the effective beam separation used in the scan fit.

2.5.1 Background and Bunch-Current Corrections

During special super-separation scans, the beams are separated by several transverse beam widths so that collision rates are strongly suppressed. The remaining rate measures beam-induced background and detector noise. The background can be estimated from noncolliding bunches or from super-separation periods. Super-separation periods sample bunches that are otherwise collidable while suppressing the true collision component. For luminometers whose online algorithm does not already subtract the background, the measured background is subtracted from the scan-step rates before fitting.

The super-separation background analysis for the Run 3 vdM scan programs was performed within the BRIL vdM analysis team. In these scans, the beams are displaced by about $5\sigma_b$ in each transverse direction, where σ_b is the transverse beam size, giving a total separation of about $7\sigma_b$, and the luminometer rates are recorded for about five minutes per scan. In the 2022 vdM fill, the PLT background level corresponds to a mean visible-response rate of $(3.11 \pm 0.47) \times 10^{-6}$ counts per bunch crossing [28]. Fig. 2.7 shows the corresponding PLT measurement made in this work during the 2023 vdM fill.

Accurate bunch-population measurements are required because Eqn. 2.13 depends directly on $N_1 N_2$. The LHC uses DC current transformers to measure total beam current and fast beam current transformers to measure bunch-by-bunch structure. The relative per-bunch pattern is normalized to the absolute total-current measurement, and the normalization is corrected for charge outside the nominal colliding bunch slots. Ghost charge contributes to the total current but is not assigned to nominal bunches, while satellite charge is displaced from the nominal RF bucket and changes the effective colliding population. The corrected bunch populations are propagated to the bunch-by-bunch rate normalization before the scan fits are performed [27, 52]. In the 2022 calibration, the normalization of the per-bunch pattern to the absolute current measurement corresponds to a correction of about

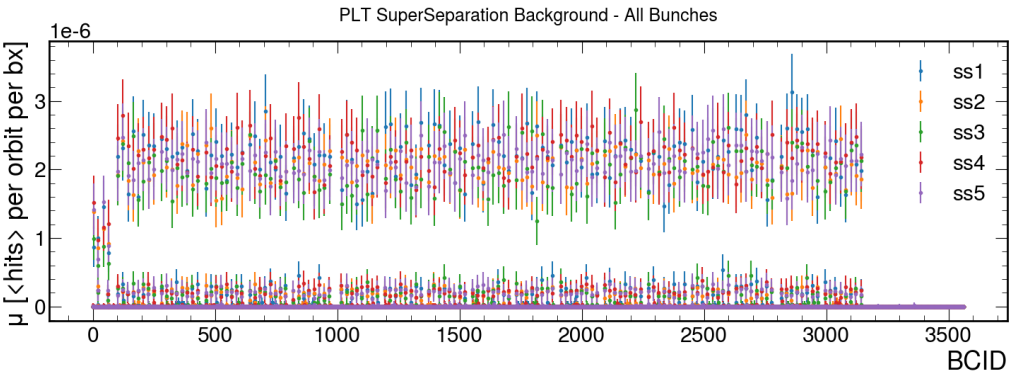


Figure 2.7: PLT background measurements during the five super-separation scans of the 2023 vdM fill 8999, produced in this work. Each point shows the average measured μ per bunch crossing during one scan period as a function of the bunch crossing identifier (BCID); the scans are labeled in the order they were performed. Colliding and noncolliding bunches separate into distinct populations, and only colliding-bunch BCIDs are used for the background estimate.

3% on the bunch-population product, the measured ghost and satellite fractions are between 0.1% and 0.4% per beam, and the combined bunch-current normalization uncertainty on σ_{vis} is about 0.25% [28].

2.5.2 Length Scale and Orbit Drift

The beam separation used in a scan fit is the realized separation at the interaction point, not simply the nominal displacement requested from the LHC magnets. A length-scale calibration measures the response between the nominal beam movement and the actual beam displacement. In constant-separation length-scale scans, both beams are moved together so that their separation is approximately fixed while the luminous region moves across the tracker. The reconstructed CMS beamspot position is then fit as a function of the nominal beam position to determine the effective displacement scale. Beam-position monitors such as DOROS add information on step-to-step nonuniformity and on slow drifts during the scan program.

Orbit drift is treated separately from the length scale. The head-on beam positions before and after a scan are compared, and a locally linear drift correction is interpolated across the scan steps. Residual nonlinear orbit effects can remain after this correction, for example from magnet hysteresis or from imperfect modeling of the beam displacement at the interaction point. These residuals are studied with beam-position monitor data and are propagated as corrections or uncertainties to the fitted scan widths. In the 2022 calibration, the length-scale correction changes σ_{vis} by about -1.0% with a residual uncertainty of 0.1%, the orbit-drift correction is about 0.1% with a comparable uncertainty [28].

2.5.3 Beam-Beam Corrections

Beam-beam interactions alter the nominal beam separation during a scan. The electromagnetic force between the two proton bunches deflects the beams and modifies the optics. These effects change both the effective separation and the scan rate profile. Fig. 2.8 shows the beam-beam deflection correction calculated for the first scan pair of the 2022 vdM fill, while Fig. 2.9 shows the corresponding dynamic-β correction derived from beam-beam simulations.

Two beam-beam effects are treated separately during the calibration. First, the coherent deflection changes the effective beam separation at each scan step, so the nominal scan coordinate is shifted by the calculated deflection. Second, the focusing interaction between the two beams changes the local optics during the scan. This dynamic-β effect is applied as a rate correction derived from beam-beam simulations and parameterizations of the vdM conditions [36, 53]. The corrected scan curves are then refit, and the change in σ_{vis} is included in the calibration result. In the 2022 calibration, the coherent deflection increases the fitted σ_{vis} by about $+2.4\%$ while the dynamic-β distortion contributes about -1.4%; because the two effects are anticorrelated, the net beam-beam correction is close to $+1\%$. [28].

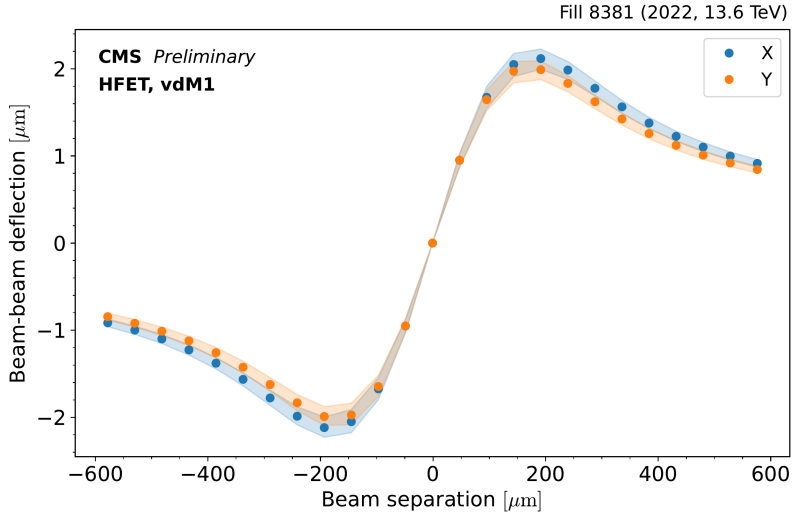


Figure 2.8: Calculated beam-beam deflection as a function of beam separation for the first vdM scan pair of the 2022 vdM fill 8381 [28]. The bands indicate the minimum and maximum values used for the per-bunch correction.

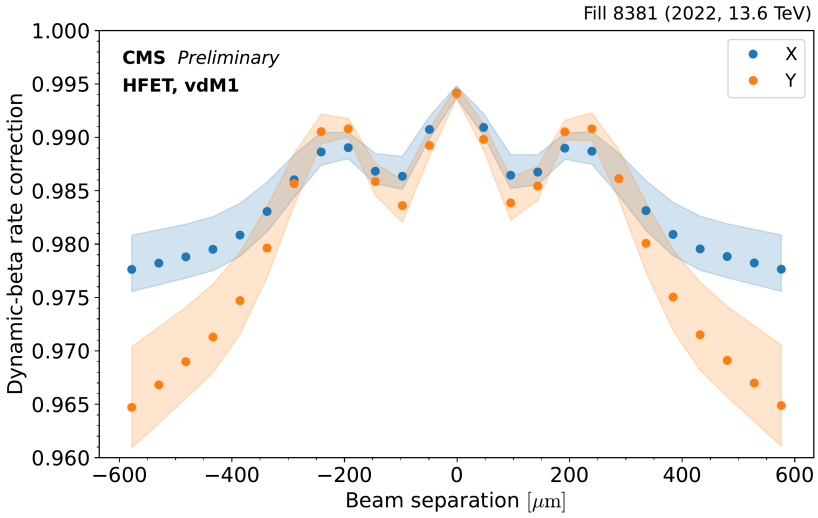


Figure 2.9: Dynamic-β rate-correction factor as a function of beam separation, calculated for the first vdM scan pair of the 2022 vdM fill 8381 [28, 53]. The bands indicate the minimum and maximum values used for the per-bunch correction.

2.6 Detector-Related Corrections

The response of a luminometer can change with operating conditions, radiation exposure, and instantaneous luminosity. The calibration must therefore include detector-related corrections in addition to beam-related corrections.

2.6.1 Efficiency Evolution

Radiation damage and operational changes can modify the detector efficiency over time. In the PLT, the effective σ_{vis} measured in emittance scans is monitored as a function of accumulated luminosity. This emittance-scan monitoring of the PLT response was performed in this work for the 2022–2023 data-taking periods. Fig. 2.10 shows the resulting trend for the 2023 data-taking period, with the per-scan σ_{vis} values extrapolated to zero SBIL. Four regions are identified from the observed trend: the first two reflect early adjustments of the detector and the LHC, the third corresponds to the nominal data-taking period, and the fourth begins after a calibration change. In the nominal region, the gradual downward slope is attributed primarily to radiation-induced changes in sensor response. Before the 2023 vdM fill, the reverse-bias voltage of the sensors was adjusted to bring all telescopes to their optimal operating points, illustrating how operating-point changes introduce discontinuities that must be handled separately from the gradual radiation-related trend.

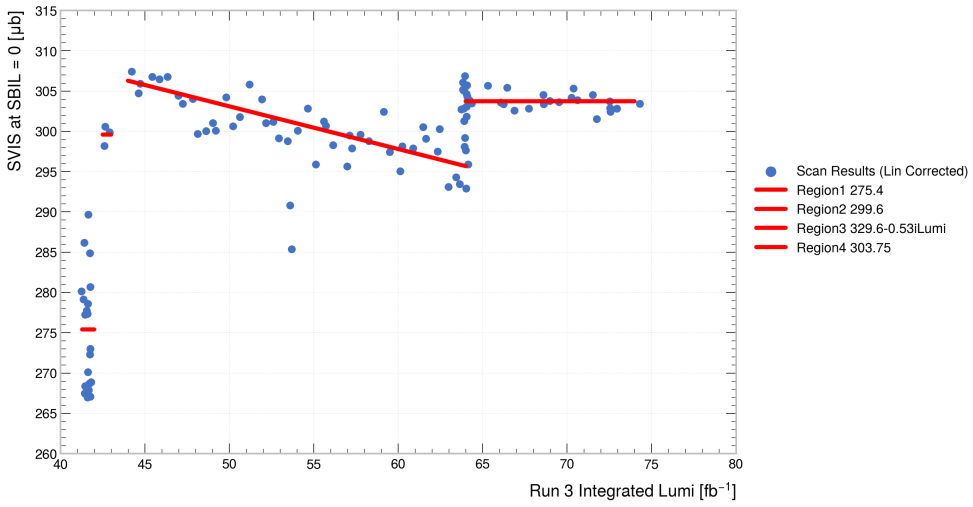


Figure 2.10: PLT visible cross section measured in emittance scans as a function of the accumulated integrated luminosity during the 2023 data-taking period, produced in this work. The values are extrapolated to zero SBIL, and straight lines are fitted to the trend in each region. Regions 1 and 2 reflect early detector and LHC adjustments, Region 3 is the nominal data-taking period with a slope attributed primarily to radiation damage, and Region 4 begins after a calibration change.

2.6.2 Relative Linearity

A luminometer should respond linearly to instantaneous luminosity after calibration. In practice, effects such as accidental coincidences, occupancy, and detector electronics response can introduce nonlinearity. In the PLT, the dominant effect is a positive nonlinearity from accidental triple coincidences, in which hits from different particles, beam-induced background, or activated material line up in the three planes without a single track from the interaction point; this over-reporting grows with per-bunch luminosity and is of the order of +1% per Hz/μb of SBIL [28]. The PLT linearity is studied by comparing the SBIL measured by PLT with the SBIL measured by an independent luminometer. The relative comparison in Fig. 2.11 shows one such measurement, in which the SBIL measured by a PLT channel is compared with the BCM1F luminometer and fitted with a quadratic function.

A purely linear detector response would follow a straight line after normalization; a nonzero quadratic term in the fit indicates a residual departure from linearity. The same comparison can be repeated fill by fill and then combined with luminosity weighting. For the PLT, accidental coincidences and occupancy effects motivate a correction that reduces the apparent PLT SBIL at high rates.

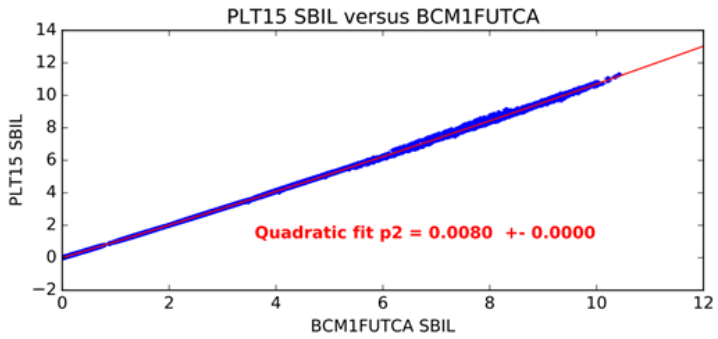


Figure 2.11: SBIL measured by PLT channel 15 compared with the BCM1F luminometer in LHC fill 8385, produced in this work. A quadratic function is fitted to the data points; the nonzero quadratic coefficient indicates a residual relative nonlinearity between the two detectors.

The correction can be parameterized as Eqn. 2.14:

$$\mathrm{SBIL}'_{\mathrm{PLT}} = \mathrm{SBIL}_{\mathrm{PLT}} - Q\,\mathrm{SBIL}^2_{\mathrm{PLT}}, \tag{2.14}$$

where Q is the nonlinearity parameter.

2.6.3 Independent Nonlinearity Extraction

The cross-detector comparison described above requires a second luminometer as a reference, introducing a dependence on its calibration stability. A complementary method developed in this work determines the PLT nonlinearity independently, without reference to another detector.

The method is based on the assumption that the visible cross section of a well-calibrated detector should not depend on SBIL within the short duration of an emittance scan. During each scan, the per-bunch σ_{vis} is measured as a function of SBIL for every bunch crossing. A non-zero slope in this distribution signals a residual nonlinearity (Fig. 2.12). A scan over candidate correction factors α is performed: for each value, the PLT rate is corrected as $\mathrm{SBIL}' = \mathrm{SBIL} - (\mathrm{SBIL})^2 \times \alpha$ and σ_{vis} is recomputed (Fig. 2.13); the value of α that yields a zero slope in σ_{vis} versus SBIL is taken as the extracted nonlinearity correction (Fig. 2.14). This procedure was applied to all PLT telescopes and all 2022 emittance-scan fills.

The independently extracted corrections reduce the nonlinearity effect by approximately a factor of 2.

2.6.4 Cross-Detector Stability and Integration

After detector-specific corrections are applied, calibrated luminometers are compared over common good-data periods. Ratios between detectors are formed in fixed luminosity-section intervals and studied as a function of time, accumulated luminosity, and SBIL. The mean of these ratios checks the relative normalization, while the spread

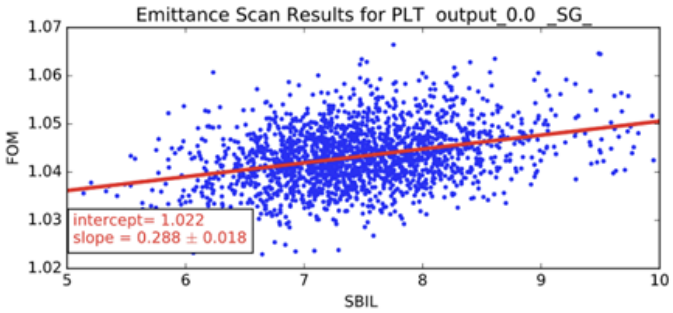


Figure 2.12: Per-bunch $\sigma_{\text{vis}}/\sigma_{\text{vdM}}$ as a function of the SBIL of each bunch in an emittance scan, produced in this work. The positive trend indicates a positive residual nonlinearity, which increases the measured σ_{vis} at higher SBIL.

and time dependence test residual stability. These cross-detector comparisons are also used to define which detector periods are certified for integration.

The final integrated luminosity delivered by CMS luminosity tools is controlled through calibration and certification inputs. An individual detector calibration is encoded in an interval-of-validity calibration tag, while a composite normtag specifies which detector calibration should be used for each run and luminosity-section range. The priority order between luminometers and the exclusion of known bad periods ensure that the integrated luminosity uses the best available calibrated detector for each time interval.

2.6.5 Summary of Corrections and Uncertainties

Table 2.1 summarizes the beam- and detector-related correction categories discussed in this chapter, with the magnitudes and systematic uncertainties documented for the 2022 CMS luminosity calibration [28]. The largest individual corrections are the bunch-current normalization, the anticorrelated beam-beam pair, the background subtraction, and the length scale, while the leading uncertainties come from the factorization bias, the scan-to-scan reproducibility, and the integration-period stability and linearity. Adding the contributions in quadrature gives a total luminosity

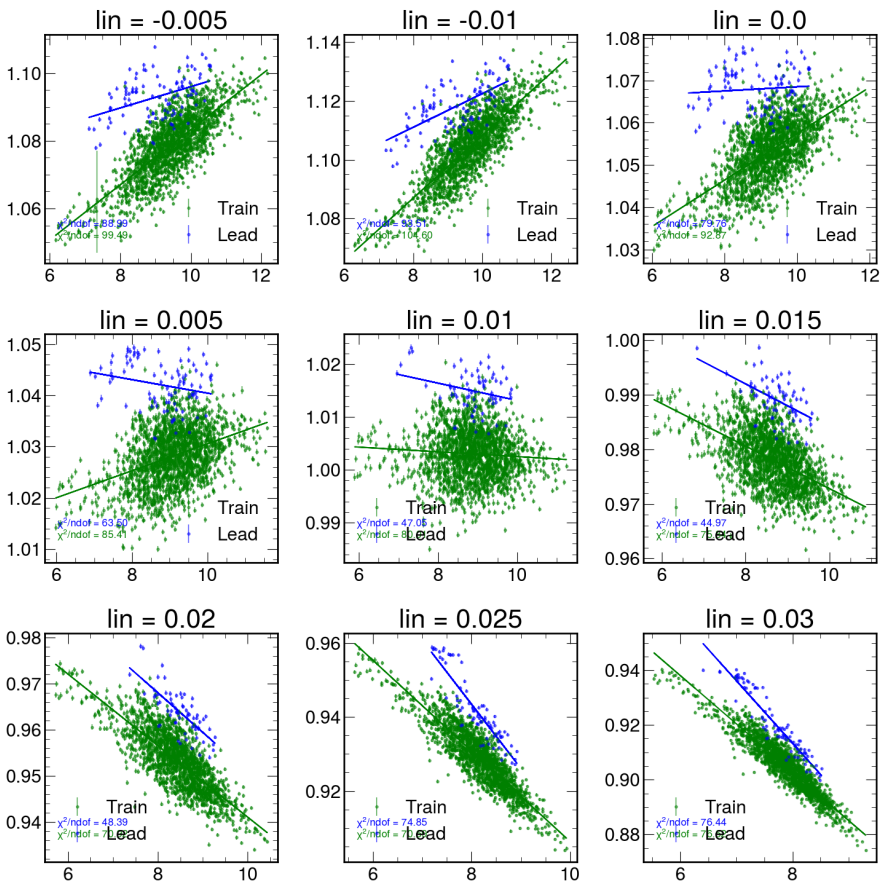


Figure 2.13: Scaled visible cross section $\sigma_{\mathrm{vis}}/\sigma_{\mathrm{vdM}}$ versus SBIL for a series of applied test nonlinearity corrections (indicated above each panel), from a PLT emittance scan in the 2023 fill 9023, produced in this work. Train bunches (green) and leading bunches (blue) are fitted separately with linear functions; the distinction is based on the position of the bunch crossing within a bunch train.

uncertainty of 1.4% for the 2022 data set, separated into a calibration part of 1.2% and an integration part of 0.8%.

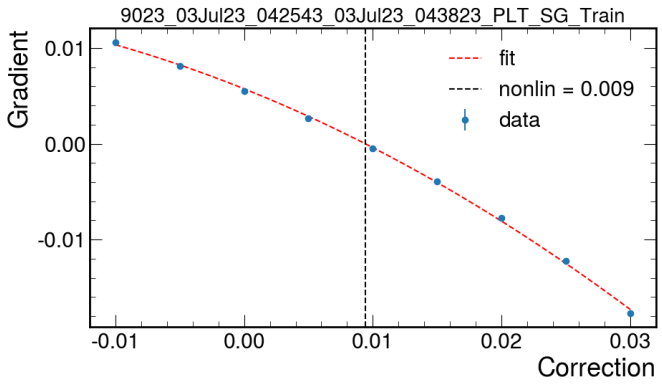


Figure 2.14: Fitted gradient of the scaled visible cross section versus the applied test nonlinearity correction for train bunches in the 2023 fill 9023, produced in this work. A quadratic fit describes the relationship, and the test value that yields a zero gradient defines the extracted nonlinearity correction.

Table 2.1: Corrections applied to the visible cross section and the associated systematic uncertainty contributions in the 2022 CMS luminosity calibration [28]. Entries without a correction value are consistency or stability terms that contribute only to the uncertainty. The total is obtained by adding the individual contributions in quadrature.

Source	Correction (%)	Uncertainty (%)
Calibration		
Bunch current	3.0	0.2
Ghost and satellite charge	0.43	0.15
Orbit drift	0.1	0.1
Residual beam positions	0.02	0.38
Beam-beam deflection / dynamic-β	$+2.36$ / -1.35	0.37
Length scale	-1.0	0.1
Factorization bias	1.0	0.8
Scan-to-scan variation	–	0.48
Bunch-to-bunch variation	–	0.06
Cross-detector consistency	–	0.4
Integration		
Afterglow corrections	–	0.2
Stability	–	0.5
Cross-detector linearity	–	0.5
Calibration total	–	1.2
Integration total	–	0.8
Total	–	1.4

Chapter 3

Measurement of Integrated Luminosity in LHC Run 3 Using $J/\psi \to \mu^+\mu^-$

This chapter presents the luminosity measurement based on counting reconstructed $J/\psi \to \mu^+\mu^-$ candidates in Run 3 CMS data. The analysis measures a corrected J/ψ rate in short time intervals and uses emittance scans to place that rate on an absolute luminosity scale. This analysis is the central work of this dissertation and is documented in CMS Analysis Note AN-2025/127 [54].

3.1 Motivation

3.1.1 Precision Luminosity in Run 3

Accurate determination of integrated luminosity, L_{int}, is a foundational requirement for the CMS physics program. In precision cross-section measurements, the luminosity uncertainty is often one of the dominant systematic uncertainties. In searches for physics beyond the Standard Model, the same uncertainty affects the interpretation of small deviations in event rates.

For colliding proton beams, the instantaneous luminosity can be written in terms of the bunch populations and transverse beam overlap as Eqn. 3.1:

$$\mathcal{L} = \frac{N_1 N_2 f_{\mathrm{LHC}} N_{\mathrm{coll}}}{2\pi \Sigma_x \Sigma_y}, \tag{3.1}$$

where N_1 and N_2 are the average bunch populations in the two beams, N_{coll} is the number of colliding bunch pairs, f_{LHC} is the LHC revolution frequency, and Σ_x and Σ_y are the effective convolved beam widths [27]. Accelerator parameters set the physical basis for luminosity calibration, but CMS luminometry relies on experimentally measured rates that are proportional to $\mathcal{L}$.

Run 3 operates at $\sqrt{s} = 13.6$ TeV and delivers high instantaneous luminosity with substantial pileup (PU). Pileup is the number of simultaneous proton-proton interactions in a bunch crossing. High pileup increases the available physics sample but also stresses detector hardware, trigger paths, reconstruction, and luminosity instrumentation. Online CMS luminometers such as the Pixel Luminosity Telescope (PLT), the Beam Conditions Monitor Fast (BCM1F), and algorithms based on the Hadron Forward (HF) calorimeter measure collision rates in real time. At high rates, these systems must be corrected for effects such as saturation, occupancy, afterglow, and radiation-induced efficiency changes. An offline luminosity method based on reconstructed physics objects tests the luminosity scale with different detector and reconstruction dependencies.

3.1.2 Standard-Candle Counting

For a standard-candle process, the reconstructed rate R is related to luminosity through the visible-cross-section relation in Eqn. 3.2:

$$R = \mathcal{L}\,\sigma_{\mathrm{vis}}. \tag{3.2}$$

The visible cross section includes the production cross section, the detector acceptance, trigger and reconstruction efficiencies, and the analysis selection efficiency. The goal of the luminosity method is not to calculate each factor separately from first principles, but to calibrate the product under controlled conditions and then monitor the corrected rate over time.

This rate-scaling structure was introduced in Sec. 2.4. In the PLT case, the detector rate is inferred from a zero-counting observable, while in the J/ψ method the observable is a fitted resonance yield in reconstructed CMS events. The two methods therefore have different detector dependencies, but both use a counted observable as the intermediate quantity between collision data and luminosity.

The process $Z \to \mu^+\mu^-$ is a standard candle for luminosity validation because it has a well-reconstructed dimuon signature and small background [29]. Its rate, however, is too low for fine-grained luminosity monitoring. The $J/\psi \to \mu^+\mu^-$ channel has a much larger yield. Its low-momentum muons make the method more sensitive to tracking, trigger, pileup, and detector-condition effects, but the high rate allows luminosity to be studied with short time bins. This tradeoff motivates the correction sequence developed in this chapter.

Physics-process luminosity methods can monitor relative luminosity stability by comparing corrected rates between time periods, detector conditions, or reference luminometers. If the visible cross section is calibrated in a well-defined reference period, the same corrected rate can also be placed on an absolute scale and integrated over the retained data-taking period.

3.1.3 Definitions and Analysis Strategy

The analysis uses fixed bins of 50 luminosity sections (LS). One LS corresponds to approximately 23.3 s, so each bin spans

$$\Delta t = 50\,\Delta t_{\mathrm{LS}}, \qquad \Delta t_{\mathrm{LS}} = 23.3\ \mathrm{s}. \tag{3.3}$$

The fitted J/ψ yield in bin i is denoted $N^{\text{fit}}_{J/\psi,i}$, and the raw rate is

$$R^0_i = \frac{N^{\text{fit}}_{J/\psi,i}}{\Delta t}. \tag{3.4}$$

The binned reference luminosity for the same LS interval is denoted $\mathcal{L}_{\text{ref},i}$ and is obtained from the CMS luminosity tools. The raw rate is corrected sequentially for pileup-dependent efficiency, long-term detector-response changes, and residual normalization offsets between stable running periods. The final corrected rate is converted to luminosity using the calibrated J/ψ visible cross section in Eqn. 3.5:

$$\mathcal{L}_{J/\psi,i} = \frac{R^{\text{corr}}_i}{\sigma^{J/\psi}_{\text{vis}}}. \tag{3.5}$$

The correction sequence is organized as a transfer from an observed resonance yield to an integrated luminosity. First, invariant-mass fits determine a raw J/ψ yield and rate in each 50-LS bin. Second, a pileup-dependent efficiency correction removes the dominant rate dependence on event occupancy. Third, emittance scans determine the visible cross section that sets the absolute scale for the corrected rate. Fourth, a long-term correction based on the barrel pixel detector response accounts for gradual detector-response evolution within stable running periods. Fifth, RAMSES supplies a relative monitor used to align residual offsets between neighboring stability regions. The final luminosity is then integrated over the retained certified lumisections and compared with the CMS reference luminosity on the same coverage. Fig. 3.1 summarizes this sequence.

The term stability region refers to a period of data taking with stable detector and beam conditions. The function $r(i)$ denotes the stability region containing bin i. In long-term correction formulas, L_{int} is also used as the cumulative Run 3 integrated-luminosity coordinate that locates bins along the data-taking timeline. The RAMSES radiation-monitoring system supplies a relative reference for region-to-region normalization, but it is not used to set the absolute luminosity scale.

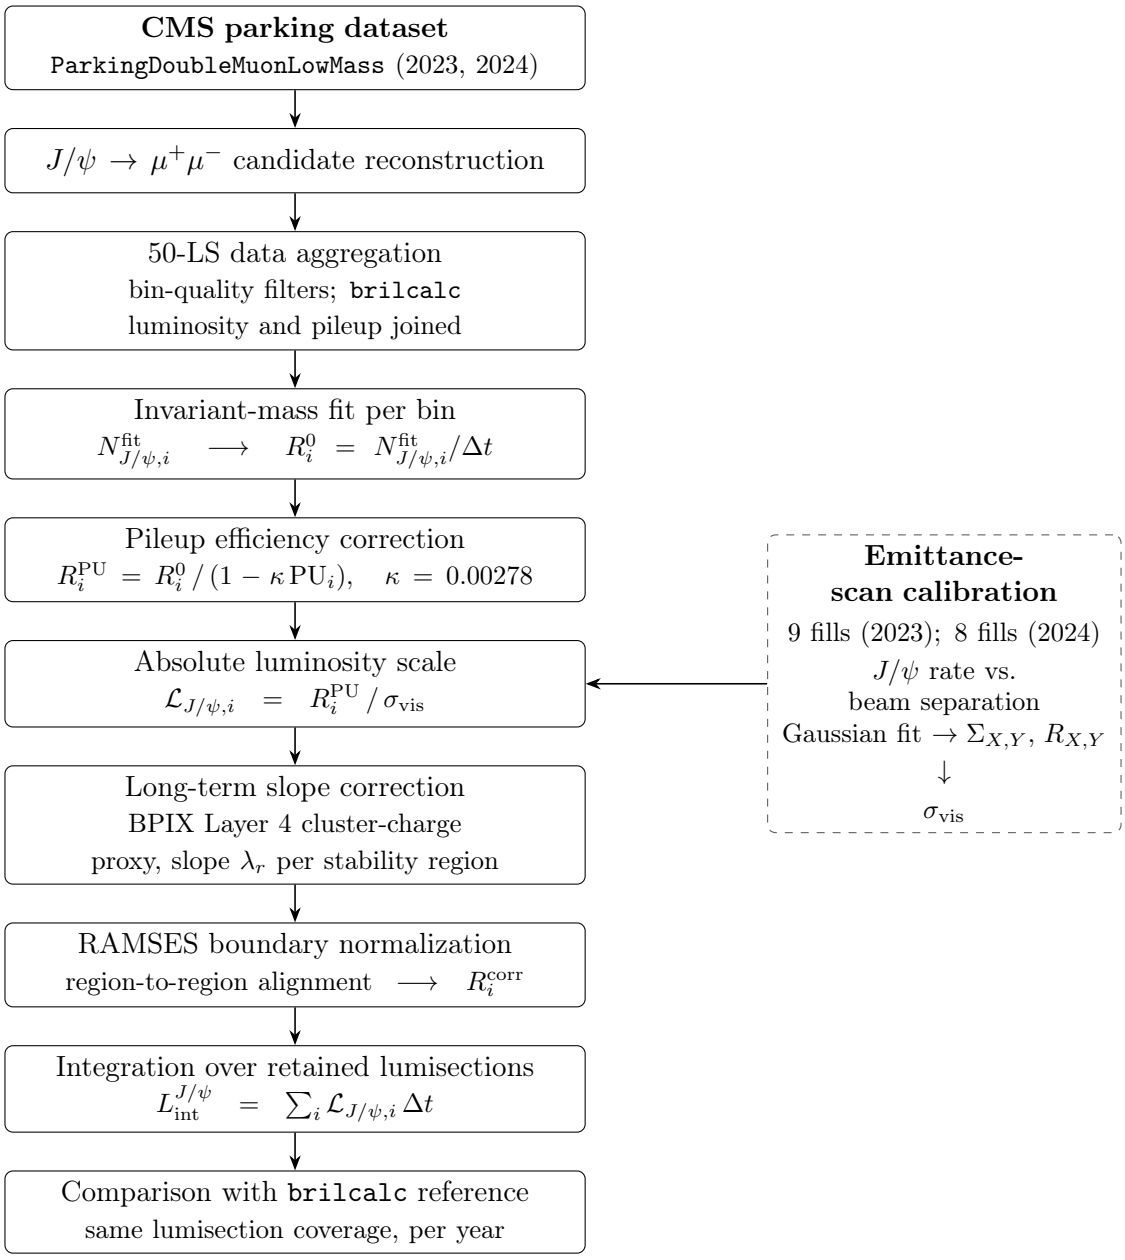


Figure 3.1: Analysis chain for the J/ψ-based luminosity measurement. The main correction sequence converts the raw reconstructed J/ψ candidate sample to an integrated luminosity result. The emittance-scan calibration (right, dashed) determines the visible cross section σ_{vis} that sets the absolute luminosity scale.

3.2 Experimental Setup and Data Selection

3.2.1 CMS Detector Configuration for the Analysis

The CMS detector configuration relevant to this analysis is described in Chapter 2. The tracking and muon systems are the central subsystems for the J/ψ measurement. During Run 3, the CMS tracker consists of the Phase-1 pixel detector and the silicon strip tracker. The pixel detector includes four barrel pixel (BPIX) layers and three forward pixel (FPIX) disks on each side of the interaction point [40]. BPIX Layer 4 is used in the long-term efficiency correction because its response tracks the gradual tracker-efficiency evolution relevant to the selected dimuon sample. Muon reconstruction and trigger performance are corresponding detector inputs for the selected dimuon resonance rate [42, 55].

RAMSES (RAdiation Monitoring System for the Environment and Safety) is installed in the CMS cavern and uses ionization chambers to monitor the radiation environment. Its comparatively low occupancy allows it to serve as a relative monitor across long periods. In this analysis, RAMSES is used only to align residual offsets between stability regions after the dominant within-region efficiency evolution has been corrected [56].

3.2.2 Dataset and Trigger

The analysis uses the `ParkingDoubleMuonLowMass` primary dataset from the 2023 and 2024 Run 3 data-taking periods. The parking strategy allows CMS to record high-rate event streams for later offline reconstruction, which is needed for low-mass dimuon studies.

The relevant high-level trigger path is `HLT_Mu0_L1DoubleMu_v*`. It is seeded by the logical OR of several Level-1 algorithms requiring two opposite-sign muon candidates. To maximize acceptance for low-mass resonances, the Level-1 seeds use combinations of low p_{T} thresholds, geometric requirements, and angular constraints.

A later hardware-level filter confirms the presence of at least two Level-1 muons within $|\eta| < 2.5$. At the HLT level, the path requires successful reconstruction of one muon candidate matched to a Level-1 muon. No explicit HLT-level requirement is imposed on the second Level-1 muon. This path is part of the CMS two-level trigger system described in Chapter 2 [43, 44]. The filters associated with this path are listed in Table 3.1.

Table 3.1: HLT filters associated with the `HLT_Mu0_L1DoubleMu` path used for the `ParkingDoubleMuonLowMass` dataset.

Quantity	Entry
Trigger path	`HLT_Mu0_L1DoubleMu`
Filter	`hltL1sDoubleMuForLowMassInclusive`
Filter	`hltPreMu0L1DoubleMu`
Filter	`hltL1fL1sL1DoubleMuL1Filtered0LowMassInclusive`
Filter	`hltL3fL1sDoubleMu0SQL1f0L2PreFilteres0L3Filtered0`

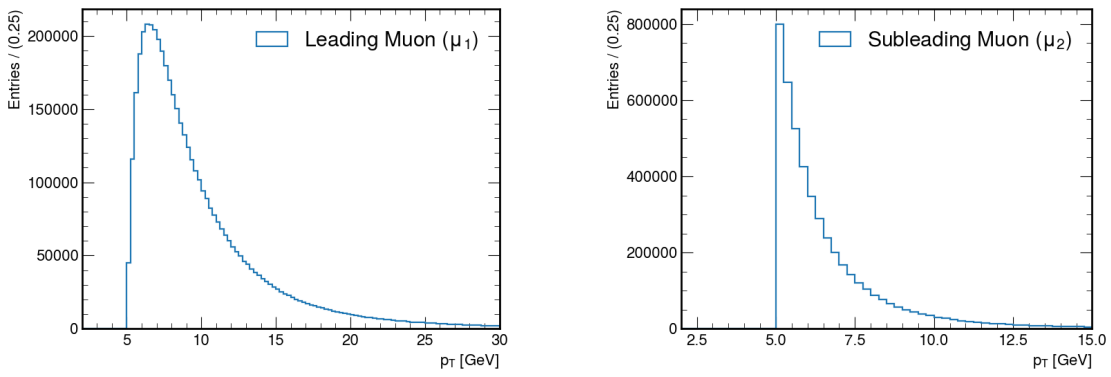
3.2.3 Event Selection and Reconstruction

Offline selection is applied to obtain high-quality J/ψ candidates and suppress combinatorial backgrounds. Muon candidates are selected from the standard CMS reconstructed-muon collection and are required to satisfy soft-muon identification criteria, following the standard CMS muon-identification framework [42, 55]. Both muons must be within $|\eta| < 2.4$, corresponding to the geometrical coverage of the CMS muon system, and each muon must satisfy $p_{\mathrm{T}} > 5$ GeV.

Dimuon candidates are formed from oppositely charged muon pairs. The two tracks are refitted to a common vertex using a Kalman vertex fit [57]. The vertex probability is required to be larger than 10%, and the invariant mass is required to lie in the range $2.7 < m_{\mu\mu} < 3.5$ GeV. This mass window contains the J/ψ peak and enough sideband region to constrain the smooth background.

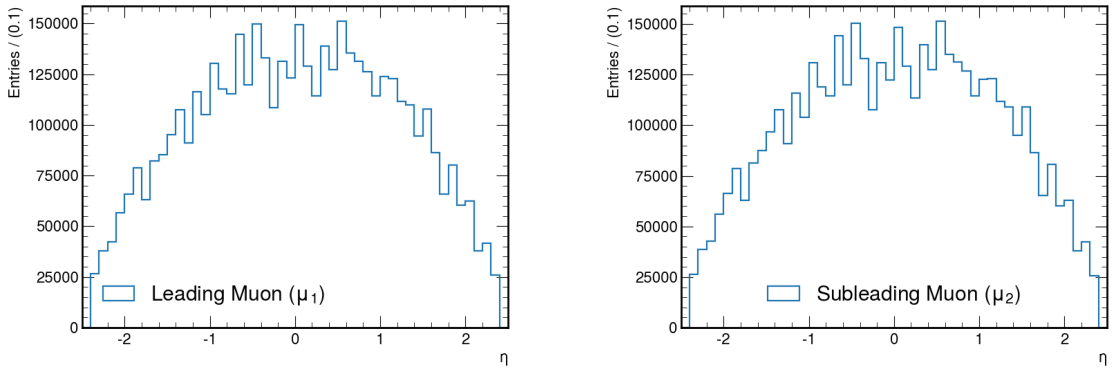
An event can contain more than one valid J/ψ candidate. To avoid double counting, candidates in the same event are sorted by vertex probability. The candidate with the highest vertex probability is retained first, and any remaining candidates

sharing a muon track with an already selected candidate are removed. This procedure keeps a unique set of dimuon candidates while favoring the highest-quality vertex. The selected muon transverse-momentum and pseudorapidity distributions are shown in Figs. 3.2 and 3.3.



(a) Leading muon **(b)** Subleading muon

Figure 3.2: Transverse-momentum distributions of the reconstructed muons used to form J/ψ candidates. The selection requires both muons to satisfy $p_{\mathrm{T}} > 5$ GeV.



(a) Leading muon **(b)** Subleading muon

Figure 3.3: Pseudorapidity distributions of the reconstructed muons used to form J/ψ candidates. Both muons are required to satisfy $|\eta| < 2.4$.

3.3 Signal Extraction Method

3.3.1 Data Aggregation and Binning

The data are grouped into 50-LS bins. One LS corresponds to approximately 23.3 s, so each 50-LS bin spans roughly 19 minutes of data taking. This bin width balances statistical precision, which requires an adequate J/ψ yield per bin, against time resolution, which must be fine enough to track short-term detector and beam variations.

For each 50-LS bin, the analysis records three primary quantities:

1. **Reference luminosity**: the average luminosity in the bin as reported by `brilcalc` using the recommended normtag for each year, where the normtag is the composite calibration record that selects the best available calibrated luminometer for each run and lumisection range (Sec. 2.6.4).

2. **Pileup**: the average number of simultaneous proton-proton interactions in the bin.

3. **Data-quality flag**: bins are excluded if the J/ψ sample contains fewer than 8,000 candidates, if the average pileup falls outside the calibrated range $0 < \mathrm{PU} < 80$, or if the lumisections correspond to non-physics detector states.

A small number of lumisections for which `brilcalc` reports a negative recorded luminosity are set to zero before constructing the binned reference. The reference integrated luminosity is computed over the same retained lumisections as the J/ψ result, ensuring that the final comparison is not affected by differences in dataset coverage.

3.3.2 Invariant-Mass Fitting

In each 50-LS bin, the signal yield $N^{\mathrm{fit}}_{J/\psi,i}$ is extracted from an unbinned maximum likelihood fit to the dimuon invariant-mass spectrum. For compactness, $m \equiv m_{\mu\mu}$

is used in the model equations. The fit is implemented with the RooFit statistical modeling toolkit [58]. The signal model combines a Crystal Ball function and a Gaussian function with a shared mean:

$$S(m) = f_{\mathrm{CB}}\,\mathrm{CB}(m;\mu,\sigma_{\mathrm{CB}},\alpha,n) + (1-f_{\mathrm{CB}})\,G(m;\mu,\sigma_G). \tag{3.6}$$

The Crystal Ball function describes the core resonance with a low-mass tail, while the Gaussian component absorbs additional mass-resolution structure. The background is modeled with a first-order Chebyshev polynomial,

$$B(m) = 1 + c_0 x(m), \tag{3.7}$$

where $x(m)$ maps the fitted mass range to the Chebyshev domain. The full model is

$$F(m) = N^{\mathrm{fit}}_{J/\psi,i} S(m) + N_{\mathrm{bkg}} B(m). \tag{3.8}$$

The signal and background shapes are normalized over the fitted mass interval, so the coefficients $N^{\mathrm{fit}}_{J/\psi,i}$ and N_{bkg} are fitted yields. The mean, widths, tail parameters, signal fraction, background shape coefficient, and yields are allowed to float in the fit. Fig. 3.4 shows a representative fit.

3.4 Calibration of the J/ψ Luminosity Method in 2023

The 2023 data are used to establish the correction and calibration strategy. The calibration-anchor period is in the 2023c stability region, where the uncorrected rate and reference ratio are shown in Fig. 3.5.

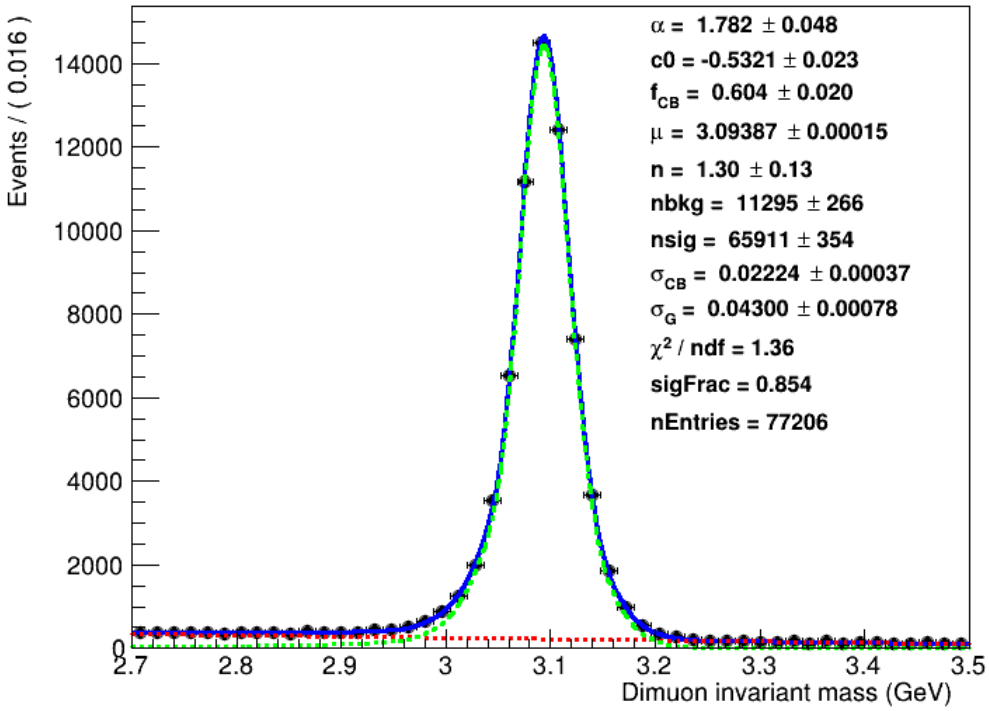


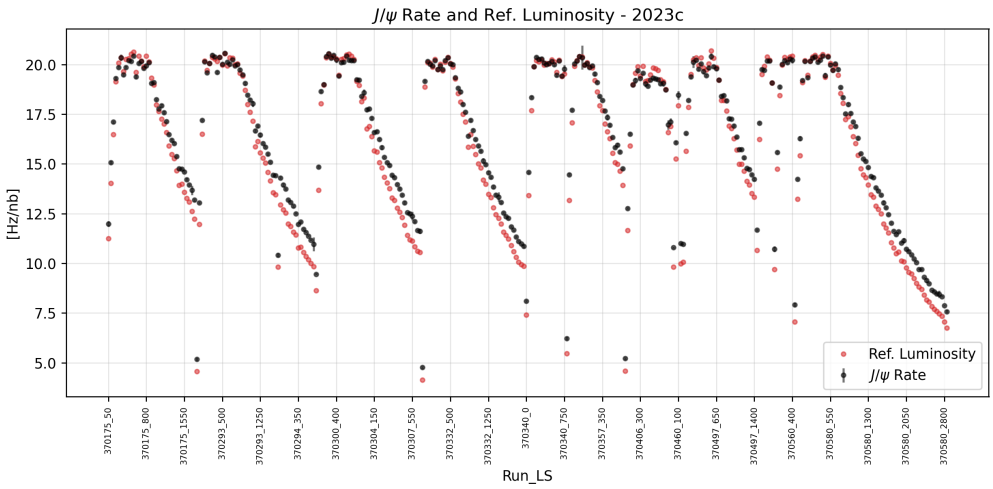
Figure 3.4: Representative dimuon invariant-mass fit in a 50-LS bin. The signal is modeled with a Crystal Ball plus Gaussian function, and the background is modeled with a first-order Chebyshev polynomial.

3.4.1 Pileup-Dependent Efficiency Correction

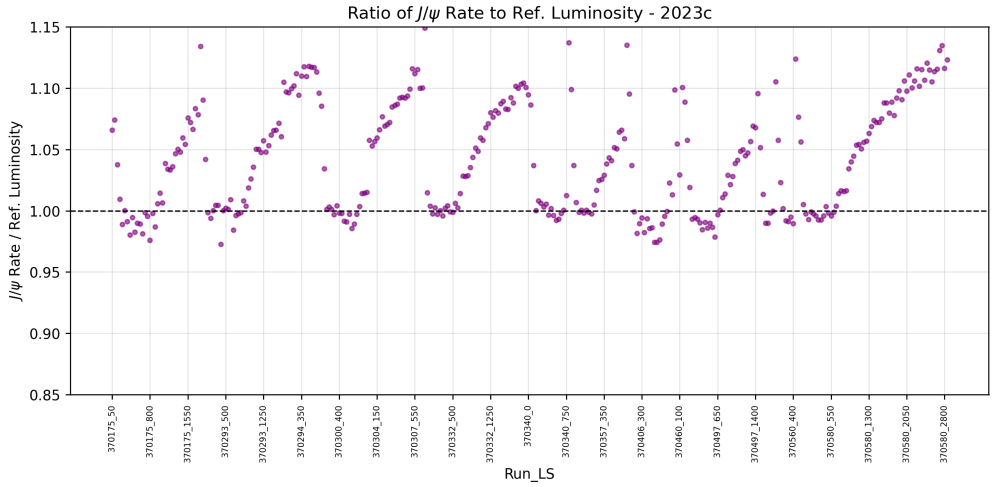
The selected J/ψ reconstruction efficiency depends on pileup because higher occupancy affects track reconstruction, vertexing, and trigger-related efficiencies. The correction is derived from the ratio in Eqn. 3.9:

$$\rho(\mathrm{PU}_i) = \frac{R_i^0}{\mathcal{L}_{\mathrm{ref},i}} \tag{3.9}$$

as a function of the average pileup in seven selected 2023 calibration runs. The same period is used for the 2023 emittance-scan calibration and later serves as the reference region for the RAMSES-based residual normalization.



(a) Raw rate



(b) Raw ratio

Figure 3.5: Raw J/ψ rate and ratio to the reference luminosity in the 2023c calibration region before pileup, visible-cross-section, and long-term efficiency corrections.

For each calibration run, $\rho(\mathrm{PU}_i)$ is fitted with a linear model $\rho = a + b\,\mathrm{PU}$. The pileup-loss coefficient is

$$\kappa = -\frac{b}{a}. \tag{3.10}$$

The weighted average is $\kappa = 0.00278 \pm 0.00006$, corresponding to an efficiency loss of about 0.28% per additional pileup interaction. The pileup efficiency factor and

corrected rate are

$$\epsilon_i^{\mathrm{PU}} = 1 - \kappa\,\mathrm{PU}_i, \qquad R_i^{\mathrm{PU}} = \frac{R_i^0}{\epsilon_i^{\mathrm{PU}}}. \tag{3.11}$$

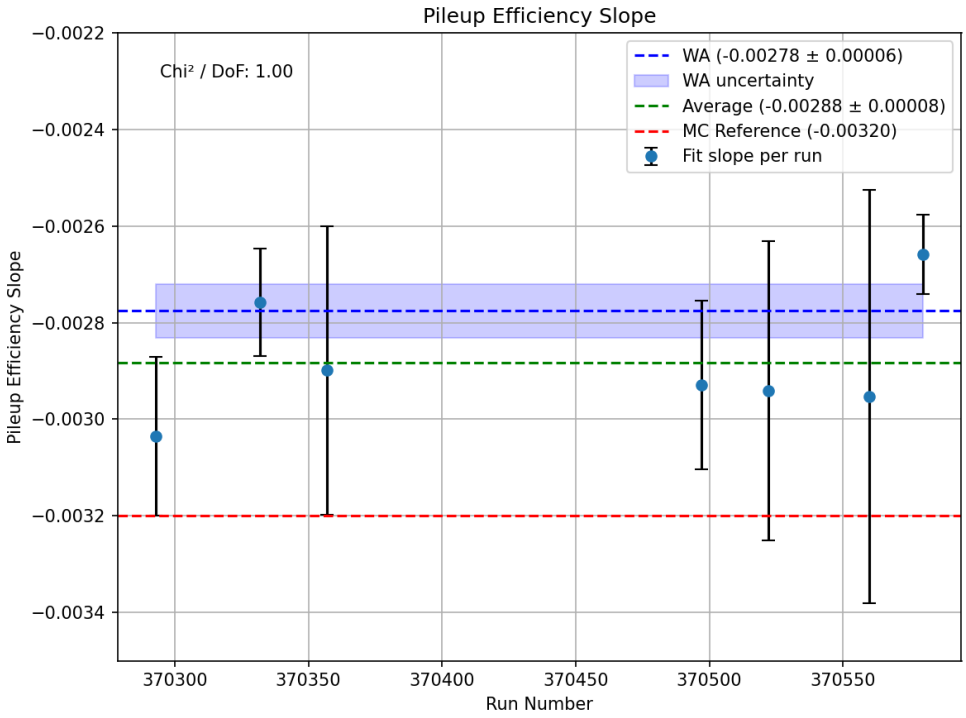
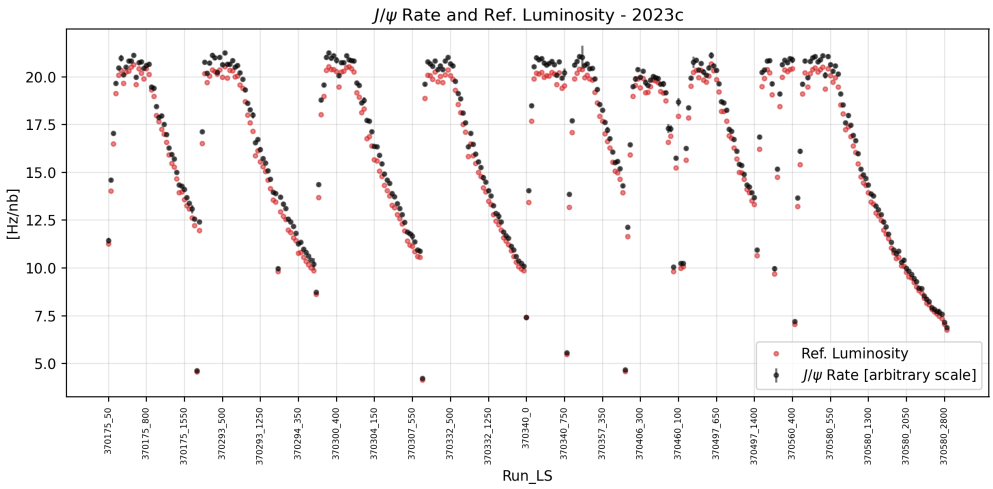


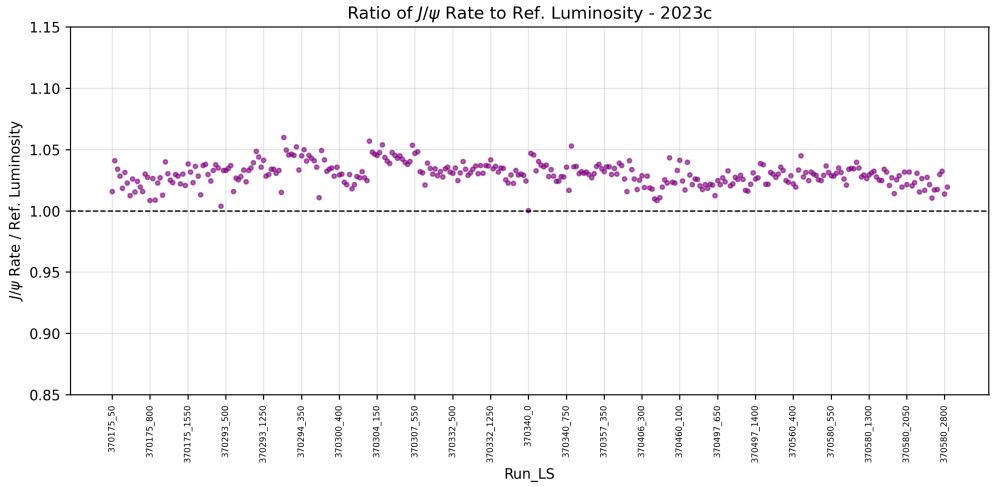
Figure 3.6: Pileup-loss coefficients extracted from individual 2023 calibration runs. The dashed line indicates the weighted-average coefficient, and the shaded band represents its uncertainty.

The pileup correction is summarized by the fitted slopes in Fig. 3.6 and by the corrected 2023c rate and ratio in Fig. 3.7.

The data-derived correction is cross-checked with separate Monte Carlo (MC) simulation samples for the 2024 and 2025 detector conditions. The MC samples are listed in Table 3.2. The simulation efficiency, ϵ_{MC}, is defined as the number of reconstructed J/ψ candidates divided by the number of generated J/ψ events in the fiducial region. The fitted positive pileup-loss coefficients from simulation are $\kappa_{\mathrm{MC}}^{2024} = 0.0032$ and $\kappa_{\mathrm{MC}}^{2025} = 0.0028$, consistent in scale with the data-derived value of $\kappa = 0.00278 \pm 0.00006$. These MC samples use PYTHIA 8 particle-gun production and therefore do not fully reproduce the detector occupancy and event activity present



(a) Pileup-corrected rate



(b) Pileup-corrected ratio

Figure 3.7: Pileup-corrected J/ψ rate and ratio to the reference luminosity in the 2023c calibration region before visible-cross-section and long-term corrections. The relative standard deviation in the calibration region is reduced from 4.45% to 0.93% after the pileup correction.

in collision data. The MC comparison confirms the expected trend of decreasing efficiency with increasing pileup, but the nominal correction remains the data-derived ϵ_i^{PU}.

Table 3.2: MC samples used to validate the pileup-dependent efficiency trend.

Quantity	Value
2024 dataset	`JPsiToMuMu_PT-0to100_pythia8-gun`
2024 conditions	`133X_mcRun3_2024_realistic_v8`
2025 dataset	`JPsiToMuMu_PT-0to100_pythia8-gun`
2025 conditions	`150X_mcRun3_2025_realistic_v6`

3.4.2 Emittance-Scan Calibration

The visible cross section of the J/ψ counting method is determined using emittance scans [37]. These scans are short vdM-like beam-separation scans taken during physics fills. Because each separation step is only about 7–9 s long, a single scan does not provide enough J/ψ statistics. The 2023 calibration therefore combines nine scans from fills 9043, 9045, 9046, 9049, 9050, 9055, 9059, 9062, and 9063. The selected scans have the same filling scheme, close temporal proximity, similar instantaneous luminosity, and complete nine-point scan patterns.

The scan timing information is taken from the BRIL vdM framework outputs [59]. Each selected event is assigned to a scan step using the step start and end timestamps. Events from all selected scans are combined at each separation step, and the J/ψ yield is extracted with the same invariant-mass fit model used for the 50-LS physics bins. The selected step pattern, scan timing, and resulting rate-versus-separation behavior are shown in Figs. 3.8, 3.9, and 3.10.

The step-wise fitted rates in the horizontal and vertical planes are fitted with Gaussian functions. Fig. 3.11 shows the head-on invariant-mass fit used to extract one of the scan-step yields, and Fig. 3.12 shows the Gaussian scan fits. The peak rates and convolved beam widths define the visible cross section in Eqn. 3.12:

$$\sigma_{\text{vis}} = \frac{0.5(R_X + R_Y)\, 2\pi\Sigma_X\Sigma_Y}{N_1 N_2 N_{\text{coll}} f_{\text{LHC}}}, \tag{3.12}$$

where R_X and R_Y are the fitted peak rates, Σ_X and Σ_Y are the effective beam-overlap widths, N_1 and N_2 are the average bunch populations, N_{coll} is the number of colliding

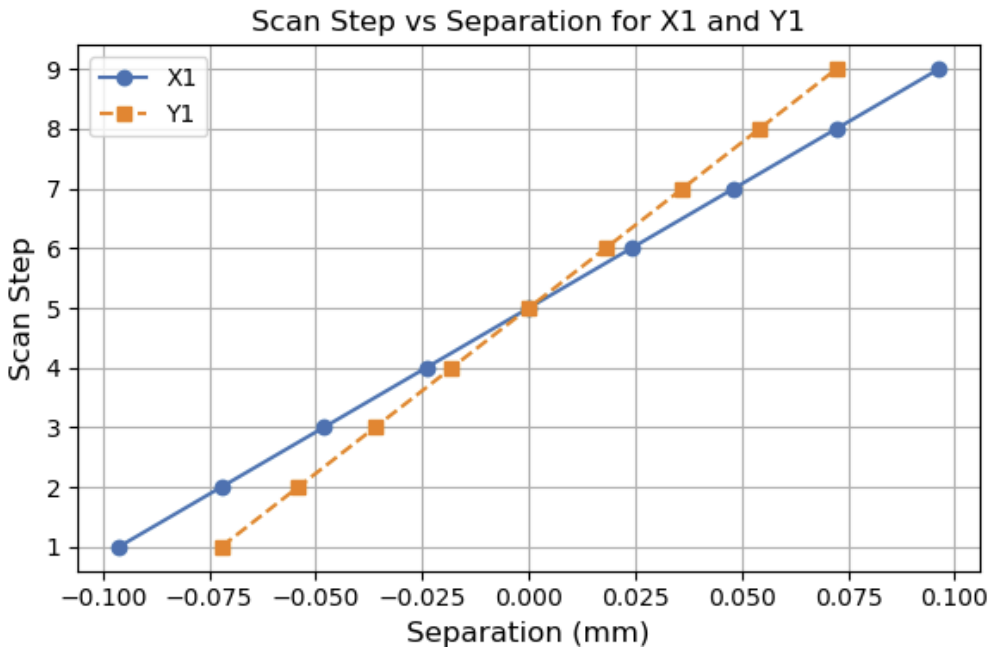


Figure 3.8: Beam-separation values used in the selected 2023 nine-point emittance scans.

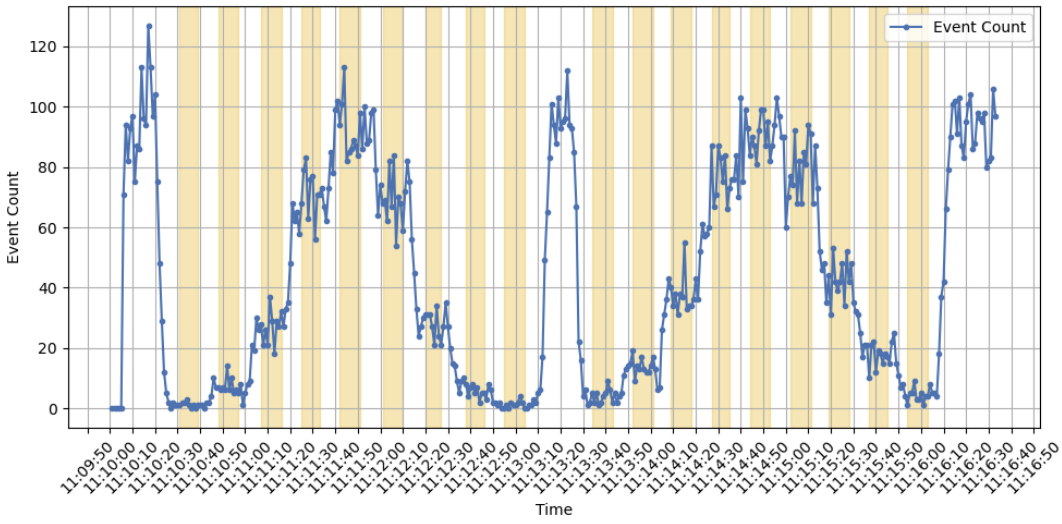


Figure 3.9: Time history of the J/ψ counting rate during a representative emittance scan. Highlighted intervals indicate stationary separation steps; unshaded intervals correspond to beam adjustment periods.

bunch pairs, and $f_{\mathrm{LHC}} = 11245$ Hz. The beam-population inputs are derived from LHC beam-current measurements, using DC current transformer and bunch-current information for the selected scan periods [52].

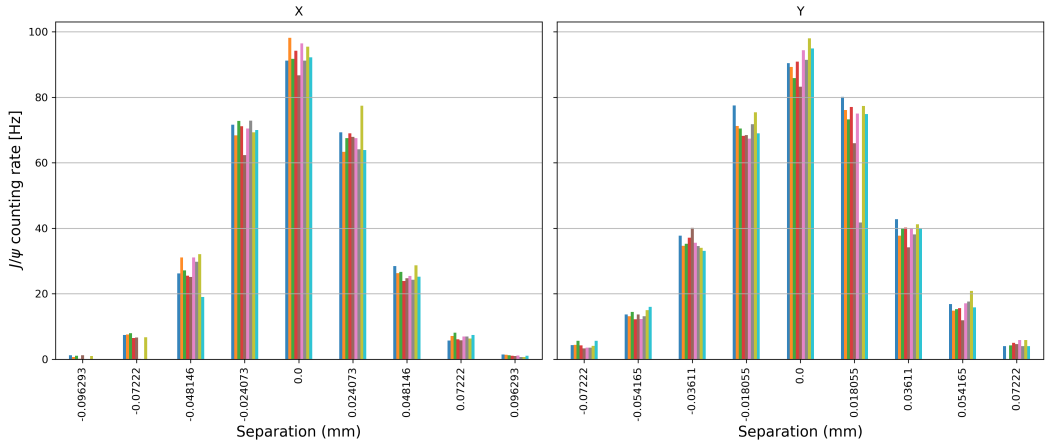


Figure 3.10: Per-scan J/ψ counting-rate estimate for the selected 2023 emittance scans as a function of beam separation.

Pileup also changes across the separation steps. The pileup at the head-on step is taken from the measurements immediately before and after the scan. For the nonzero separation steps, the pileup is estimated from the J/ψ counting rate relative to the head-on rate and then used to apply the pileup-efficiency correction step by step.

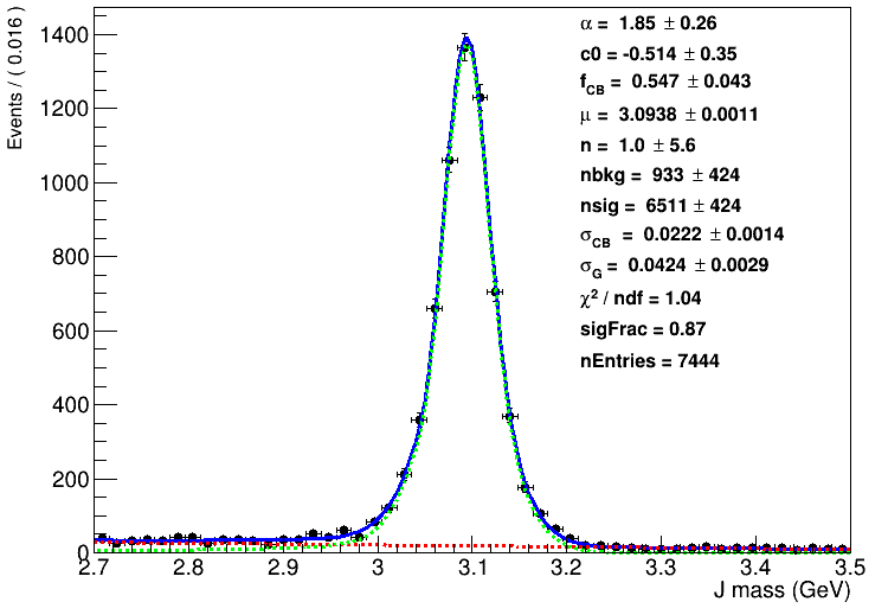
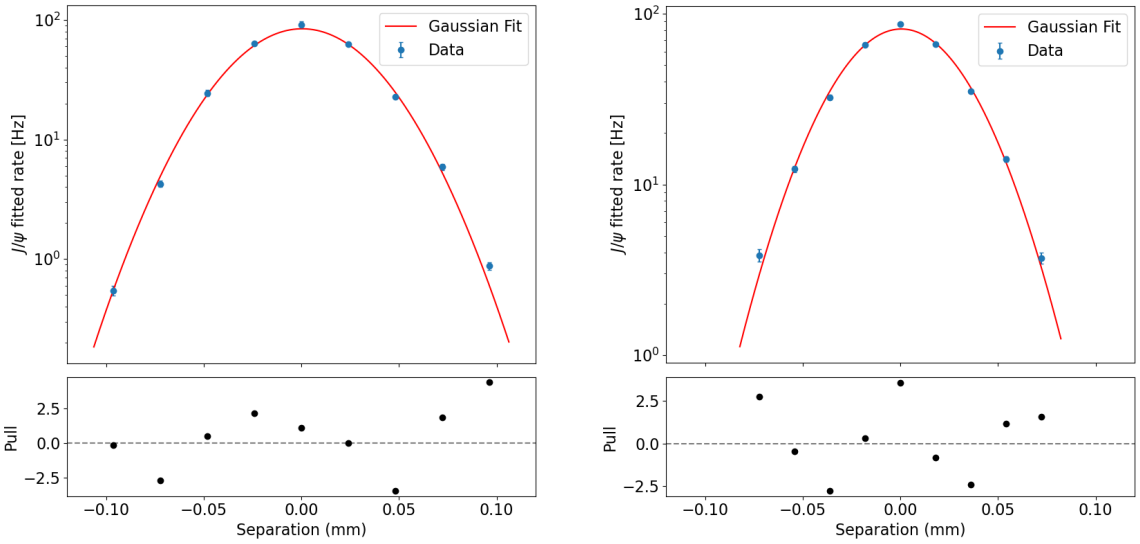


Figure 3.11: Invariant-mass fit to the combined $m_{\mu\mu}$ distribution for the 2023 head-on separation step in the X scan. Events from all selected scans are combined before fitting.



(a) x scan **(b)** y scan

Figure 3.12: J/ψ fitted rates as a function of beam separation for the combined 2023 emittance scans. Gaussian fits determine the effective beam overlap widths and peak rates used in Eqn. 3.12.

Table 3.3: Inputs to the 2023 J/ψ visible-cross-section calculation from emittance scans.

Quantity	Value
R_X	87.6 ± 1.2 Hz
R_Y	77.9 ± 1.3 Hz
Σ_X	0.0300 ± 0.0002 mm
Σ_Y	0.0288 ± 0.0002 mm
N_1	$(1.54 \pm 0.01) \times 10^{11}$
N_2	$(1.56 \pm 0.01) \times 10^{11}$
N_{coll}	2452
f_{LHC}	11245 Hz

The 2023 calibration inputs are summarized in Table 3.3. The resulting 2023 visible cross section is

$$\sigma_{\text{vis}}^{J/\psi,2023} = 6.8 \pm 0.1 \text{ nb}. \tag{3.13}$$

Using this visible cross section, the intermediate J/ψ luminosity is obtained from $R_i^{\text{PU}}/\sigma_{\text{vis}}$. The post-calibration ratio in the 2023c region is shown in Fig. 3.13.

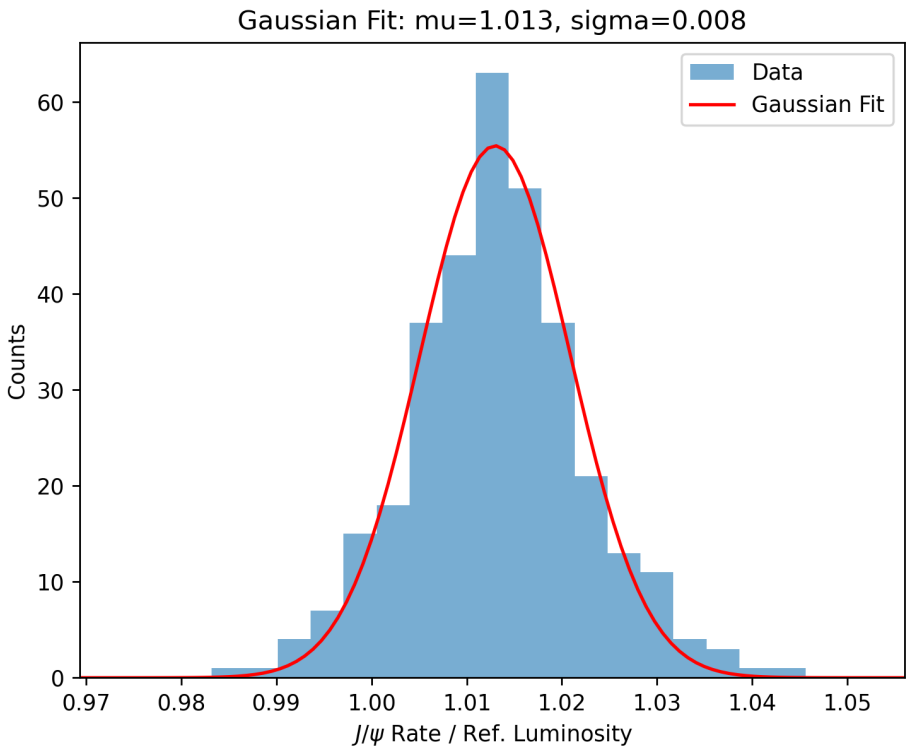


Figure 3.13: Distribution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ in the 2023c calibration region after the pileup correction and 2023 visible-cross-section calibration, before long-term corrections.

3.4.3 Long-Term Stability Corrections in 2023

After pileup correction and absolute calibration, the 2023 luminosity ratio still shows residual time dependence. The dominant effect is a gradual loss of efficiency within stable running periods, together with smaller discontinuities when detector or trigger conditions change. Changes in high voltage, gain calibration, technical stops, trigger menus, cluster-position-estimator calibrations, and alignment conditions are used to define stability regions. Fig. 3.14 shows these regions, and Table 3.4 lists the corresponding run and luminosity ranges.

The geometric distribution of the J/ψ candidates motivates the use of a BPIX-based efficiency proxy. Muons are categorized as BPIX-like or FPIX-like using an approximate $|\eta| = 1.5$ boundary. About 73% of candidates have both muons in the BPIX region, 17% have both muons in the FPIX region, and 10% are mixed. The

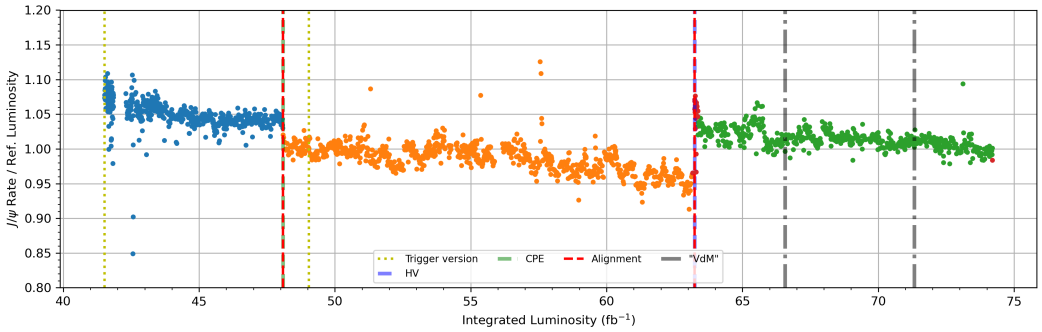


Figure 3.14: Time evolution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\mathrm{ref}}$ in 2023 after pileup correction and visible-cross-section calibration. Stable running periods and the calibration-anchor region are indicated.

Table 3.4: Stability regions used for the 2023 long-term slope correction and RAMSES residual normalization. Region 2023c contains the calibration period and is used as the reference region. The L_{int} range is the cumulative Run 3 integrated-luminosity coordinate used to locate bins along the data-taking timeline; it is not the final J/ψ-luminosity result for each region.

Region	First run	Last run	L_{int} range (fb^{-1})
2023a	366727	367516	41.52–48.08
2023b	367553	368822	48.09–63.15
2023c	369978	370776	63.36–74.20

dominance of BPIX acceptance supports the use of BPIX Layer 4 information for long-term corrections, with the relevant pixel geometry introduced in Fig. 2.3.

For each stability region r, the BPIX Layer 4 response is fitted with

$$M_r(L_{\mathrm{int}}) = p_r + q_r \left(L_{\mathrm{int}} - L^{\mathrm{min}}_{\mathrm{int},r}\right), \tag{3.14}$$

and the normalized slope is

$$\lambda_r = \frac{q_r}{p_r}. \tag{3.15}$$

The long-term efficiency factor is then

$$\epsilon^{\mathrm{LT}}_r(L_{\mathrm{int}}) = 1 + \lambda_r \left(L_{\mathrm{int}} - L^{0}_{\mathrm{int},r}\right), \tag{3.16}$$

where $L^0_{\mathrm{int},r}$ is the normalization point for region r. In the calibration-anchor region, this point is chosen at the center of the selected emittance-scan calibration period so that the correction preserves the absolute scale. The corrected rate after the long-term slope correction is

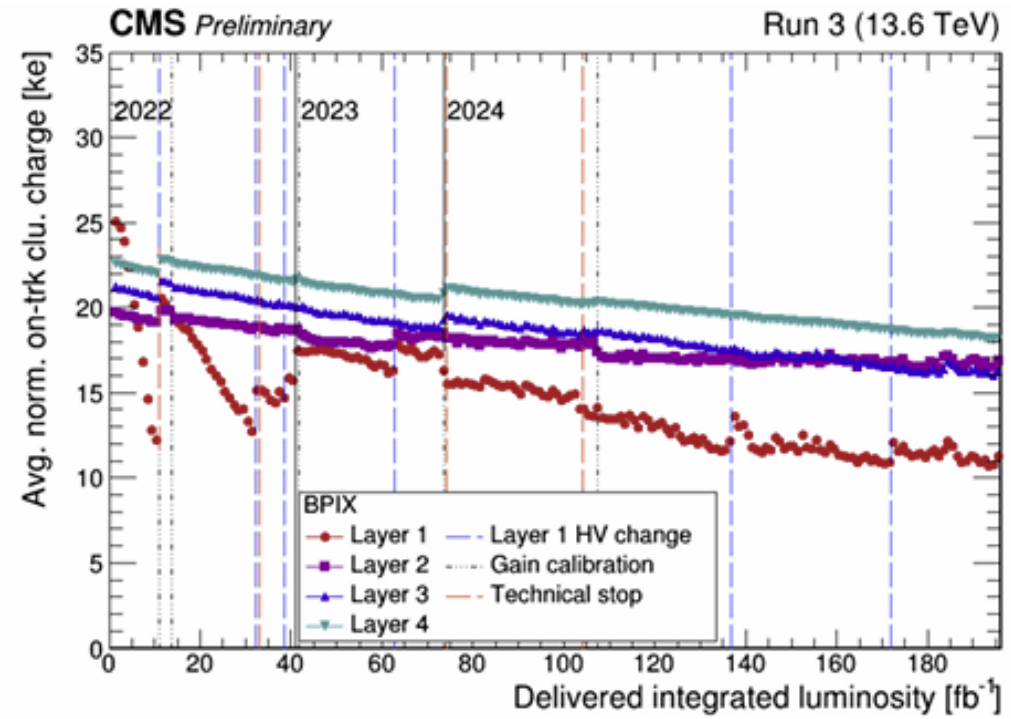
$$R_i^{\mathrm{LT}} = \frac{R_i^{\mathrm{PU}}}{\epsilon^{\mathrm{LT}}_{r(i)}(L_{\mathrm{int},i})}. \tag{3.17}$$

Table 3.5: Normalized BPIX Layer 4 slopes used in the long-term efficiency correction. Slopes are expressed as fractional change per fb^{-1}.

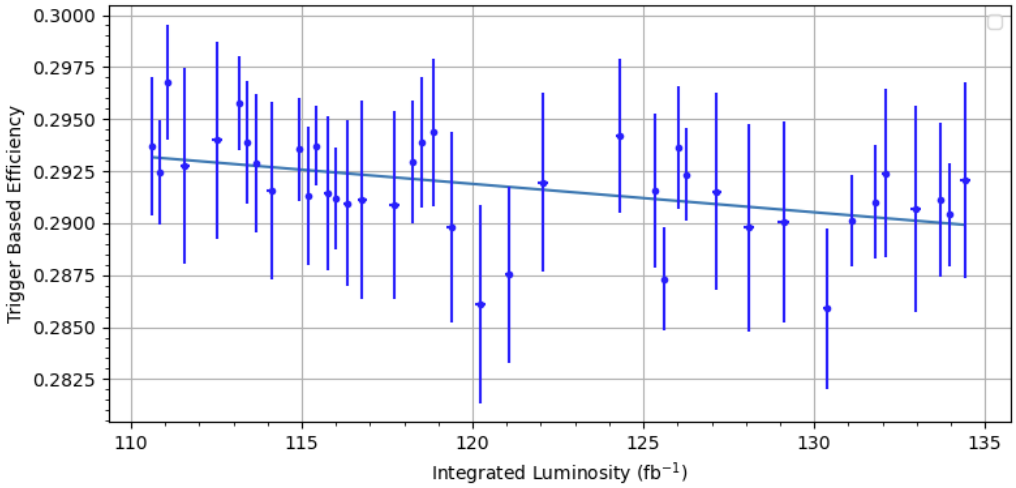
Year	Region	Layer 4 slope λ_r per fb^{-1}
2023	2023a	−0.00537
2023	2023b	−0.00306
2023	2023c	−0.00273
2024	2024a1	−0.00427
2024	2024a2	−0.00275
2024	2024a3	−0.00178
2024	2024a4	−0.00279
2024	2024b1	+0.00673
2024	2024b2	−0.00263
2024	2024c	−0.00237
2024	2024d	−0.00257
2024	2024e	−0.00271

The observables used to motivate the long-term correction and the fitted normalized slopes are shown in Fig. 3.15 and Table 3.5. After the region-wise slope correction, residual discontinuities between neighboring stability regions are removed using RAMSES as a relative monitor. The mean of the slope-corrected $\mathcal{L}_{J/\psi}/\mathcal{L}_{\mathrm{RAMSES}}$ ratio is evaluated in a 3 fb^{-1} boundary window on each side of each region boundary. The ratio of the two boundary means defines a multiplicative region-to-region adjustment, as shown in Fig. 3.16. The factors are propagated outward from the calibration-anchor region, which remains fixed.

After applying the slope and boundary corrections, the final 2023 comparison with the reference luminosity is summarized in Fig. 3.17.



(a) BPIX cluster charge

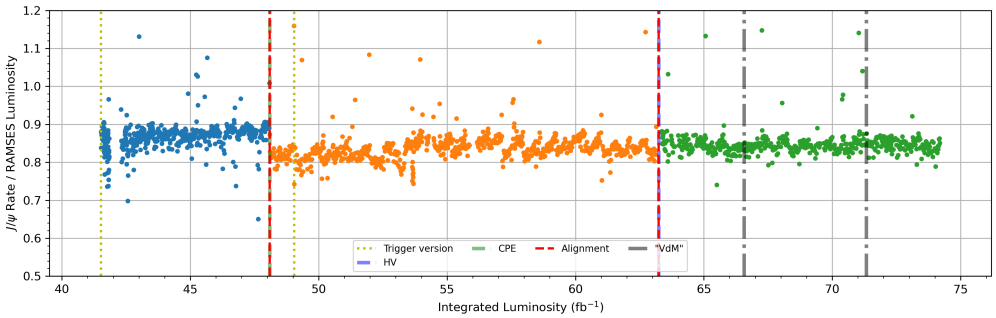


(b) Squared trigger-ratio probe

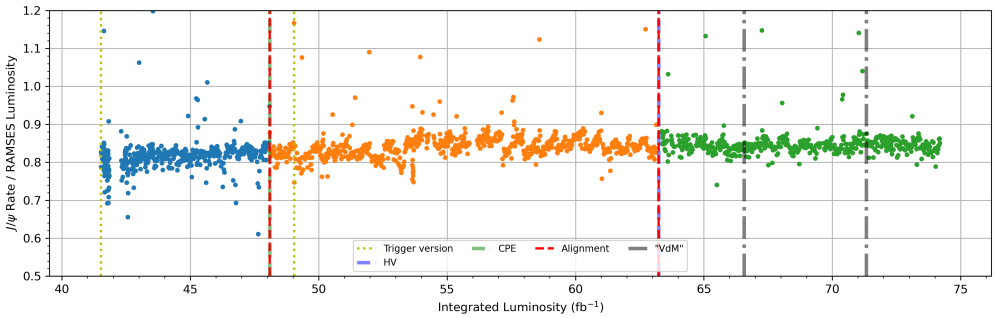
Figure 3.15: Observables motivating the long-term correction model. The BPIX Layer 4 cluster-charge evolution is used as the nominal proxy, and a trigger-ratio study tests the same trend independently.

3.5 Application of the Calibrated Method to 2024

The 2024 measurement follows the same method as 2023 but uses year-specific calibration and stability inputs. The detector response continues to evolve with



(a) Before boundary normalization

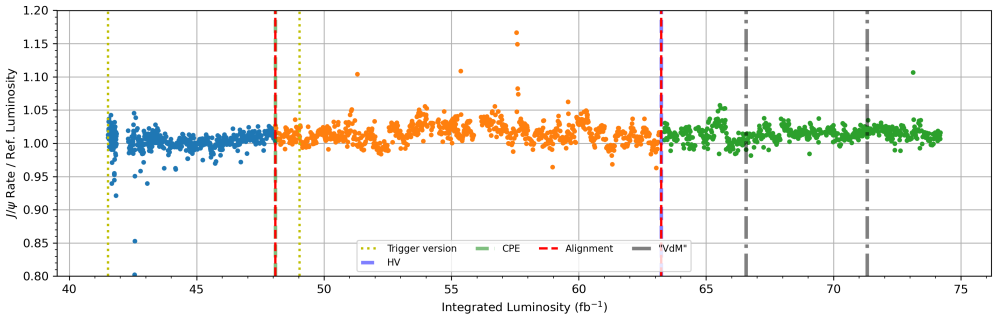


(b) After boundary normalization

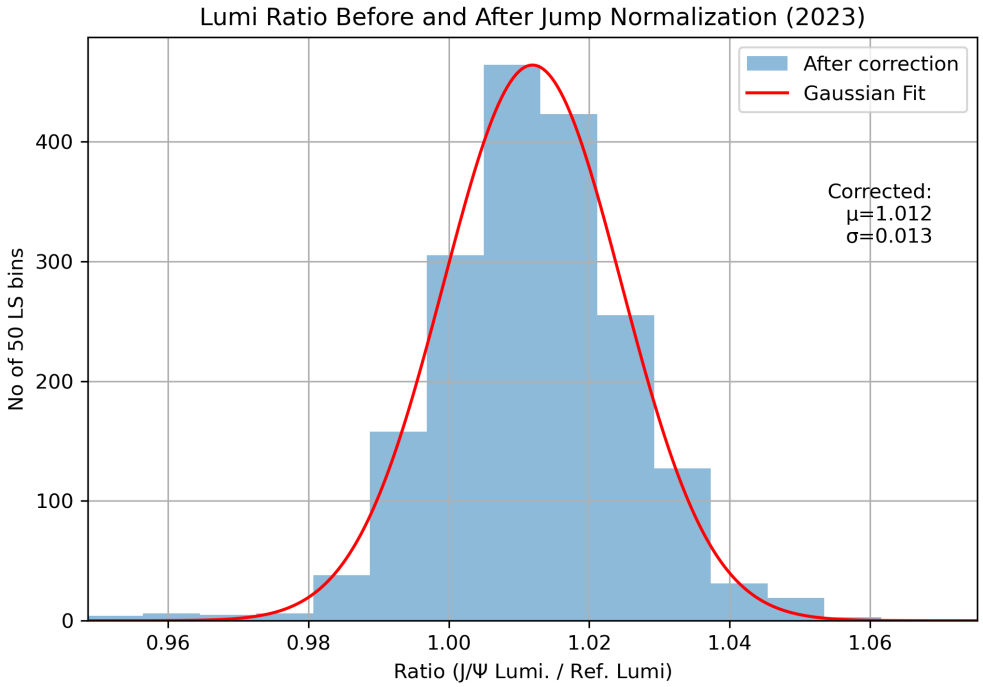
Figure 3.16: RAMSES-based residual normalization in 2023. Region 2023c is fixed as the calibration reference while neighboring regions are adjusted for continuity.

accumulated radiation dose, so the 2024 dataset requires its own stability-region definition, visible-cross-section calibration, and BPIX Layer 4 slope corrections. The trigger path, event selection, mass-fit model, and bin-quality requirements carry over unchanged from 2023.

The 2024 dataset spans 12 stability regions, compared to 3 regions in 2023, reflecting finer partitioning driven by detector hardware changes, technical stops, and greater total dose. The calibration period lies within region 2024b2, approximately $L_{\text{int}} = 129$–135 fb^{-1}. This region anchors the 2024 absolute scale and the RAMSES-based residual normalization, in the same role that 2023c plays for 2023. The uncorrected 2024b2 rate and reference ratio are shown in Fig. 3.18.



(a) Corrected ratio

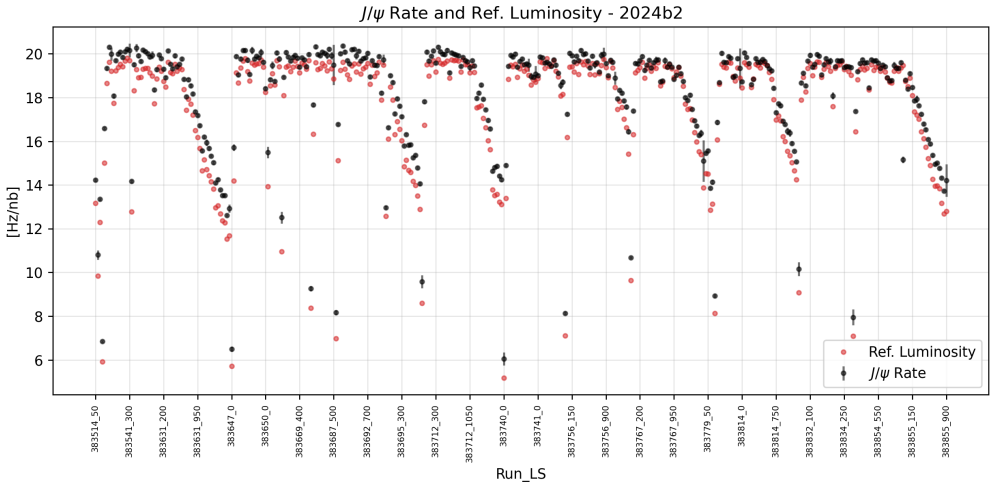


(b) Ratio distribution

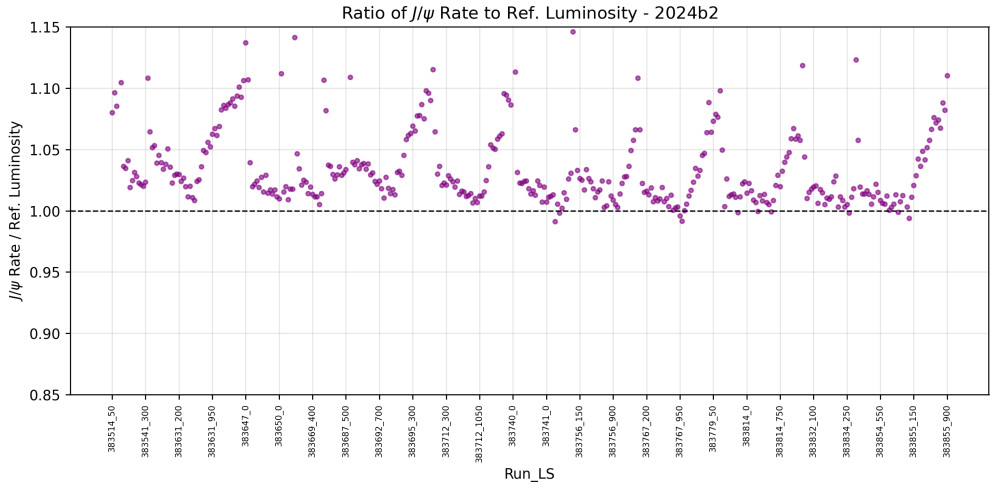
Figure 3.17: Final 2023 long-term correction. The corrected $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ ratio has a relative standard deviation of 1.3%.

3.5.1 2024 Absolute Calibration

The 2024 emittance-scan calibration uses fills 9931, 9932, 9934, 9935, 9936, 9937, 9939, and 9945. The scans are required to have complete nine-point X and Y patterns, closely matched beam conditions, and sufficient stability for the combined fit. The same yield extraction, pileup correction, Gaussian scan-curve fitting, and



(a) Raw rate



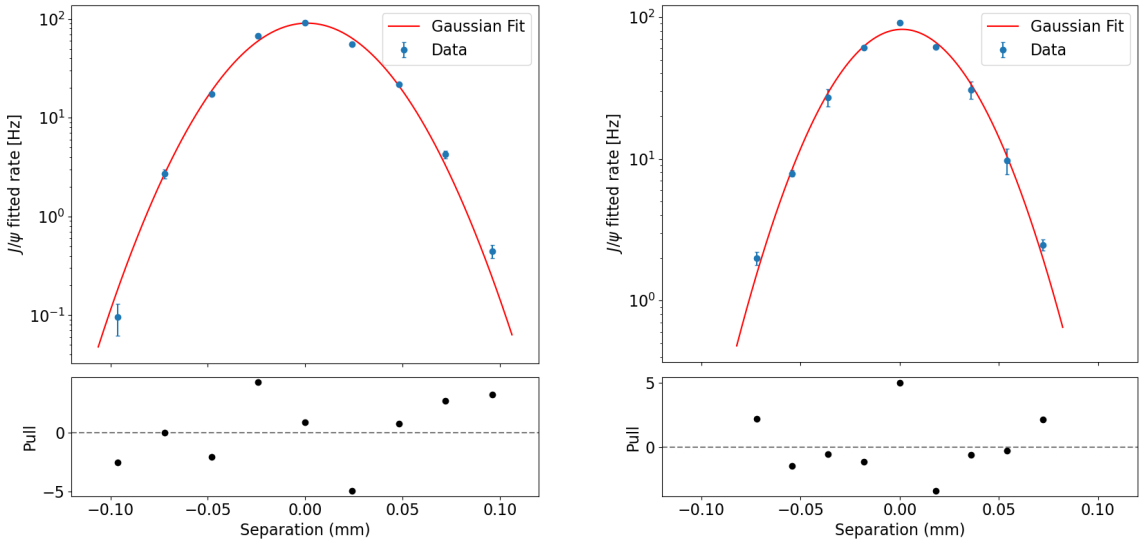
(b) Raw ratio

Figure 3.18: Raw J/ψ rate and ratio to the reference luminosity in the 2024b2 calibration region before pileup, visible-cross-section, and long-term efficiency corrections. This is the 2024 equivalent of the 2023c calibration-region starting point shown in Fig. 3.5.

beam-current treatment used in 2023 are applied to this scan set. Fig. 3.19 shows the 2024 scan fits, and Table 3.6 lists the inputs to the visible-cross-section calculation.

The 2024 visible cross section is

$$\sigma_{\text{vis}}^{J/\psi,2024} = 6.4 \pm 0.1 \text{ nb}. \tag{3.18}$$



(a) x scan **(b)** y scan

Figure 3.19: J/ψ fitted rates as a function of beam separation for the combined 2024 emittance scans. Gaussian fits determine the effective beam overlap widths and peak rates used in the 2024 visible-cross-section calculation.

Table 3.6: Inputs to the 2024 J/ψ visible-cross-section calculation from emittance scans.

Quantity	Value
R_X	90.7 ± 1.3 Hz
R_Y	82.0 ± 1.1 Hz
Σ_X	0.0276 ± 0.0002 mm
Σ_Y	0.0260 ± 0.0002 mm
N_1	$(1.52 \pm 0.01) \times 10^{11}$
N_2	$(1.52 \pm 0.01) \times 10^{11}$
N_{coll}	2340
f_{LHC}	11245 Hz

The ratio after the 2024 pileup correction and visible-cross-section calibration is shown in Fig. 3.20.

3.5.2 2024 Long-Term Corrections

The 2024 long-term correction uses the same two-stage strategy as 2023. First, a BPIX Layer 4 slope correction removes gradual within-region efficiency changes.

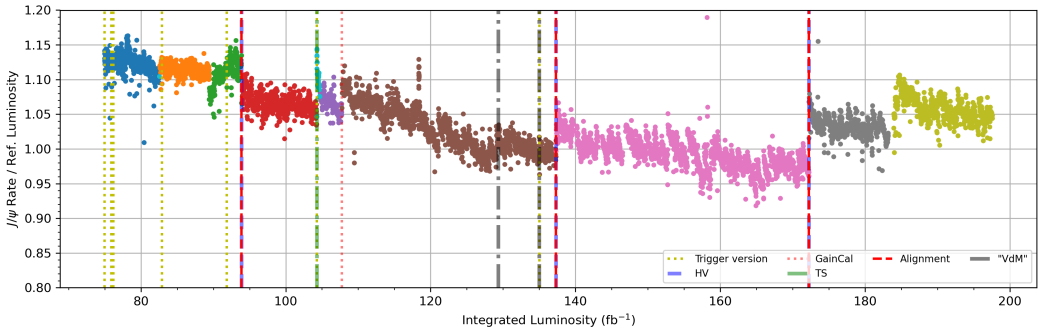


Figure 3.20: Time evolution of $\mathcal{L}_{J/\psi}/\mathcal{L}_{\mathrm{ref}}$ in 2024 after the pileup correction and 2024 visible-cross-section calibration. Region boundaries used for long-term correction are overlaid.

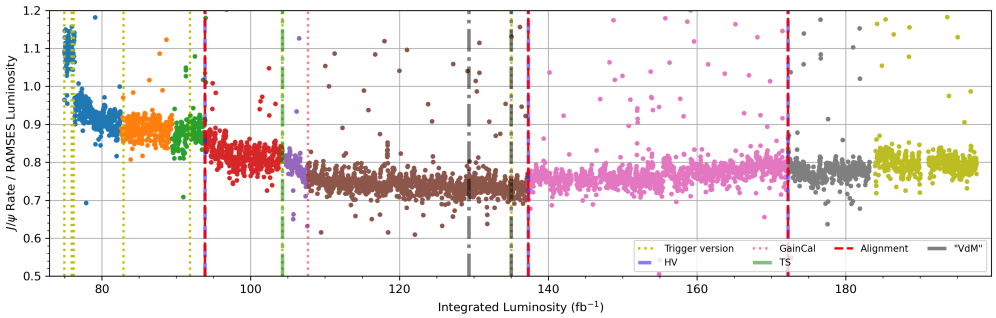
Second, RAMSES is used to align the residual normalization between neighboring stability regions while keeping the 2024 calibration-anchor region fixed. The 2024 stability regions are listed in Table 3.7, the RAMSES boundary normalization is shown in Fig. 3.21, and the final corrected comparison is shown in Fig. 3.22.

Table 3.7: Stability regions used for the 2024 long-term slope correction and RAMSES residual normalization. Region 2024b2 contains the calibration period and is used as the reference region. The L_{int} range is the cumulative Run 3 integrated-luminosity coordinate; it is not the final J/ψ-luminosity result for each region.

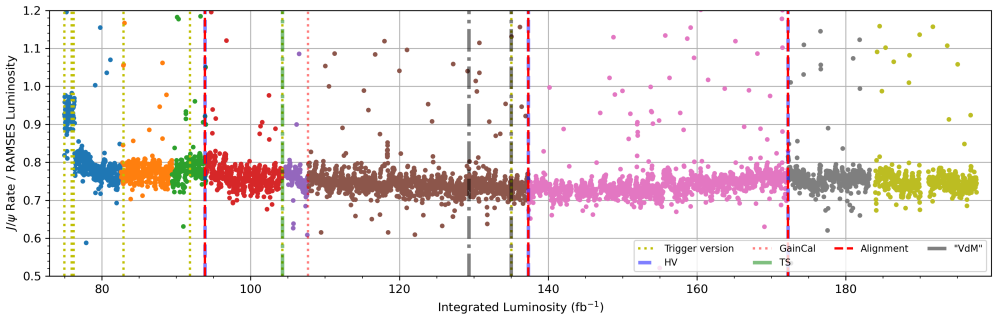
Region	First run	Last run	L_{int} range (fb^{-1})
2024a1	379416	380128	74.96–82.48
2024a2	380197	380538	82.58–89.41
2024a3	380539	381113	89.41–93.97
2024a4	381114	381544	93.99–104.23
2024b1	382258	382344	104.70–107.72
2024b2	382504	383855	107.73–137.30
2024c	383907	385515	137.31–172.24
2024d	385516	386025	172.24–183.19
2024e	386026	386951	183.94–197.58

3.6 Systematic Uncertainties

The systematic uncertainties considered in the measurement are summarized in Table 3.8. The dominant contribution is the visible-cross-section uncertainty from the



(a) Before boundary normalization



(b) After boundary normalization

Figure 3.21: RAMSES-based residual normalization in 2024. Region 2024b2 is fixed as the calibration reference while neighboring regions are adjusted for continuity.

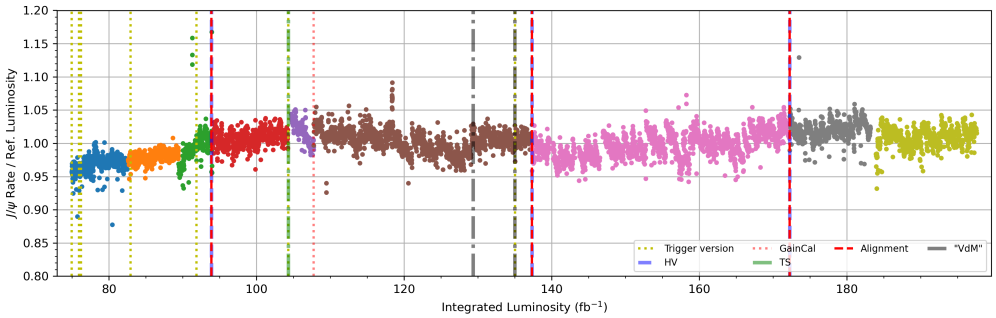
emittance-scan calibration, whose component breakdown is given in Table 3.9. The total uncertainty is obtained by adding the individual contributions in quadrature.

Table 3.8: Relative uncertainties in the J/ψ luminosity measurement.

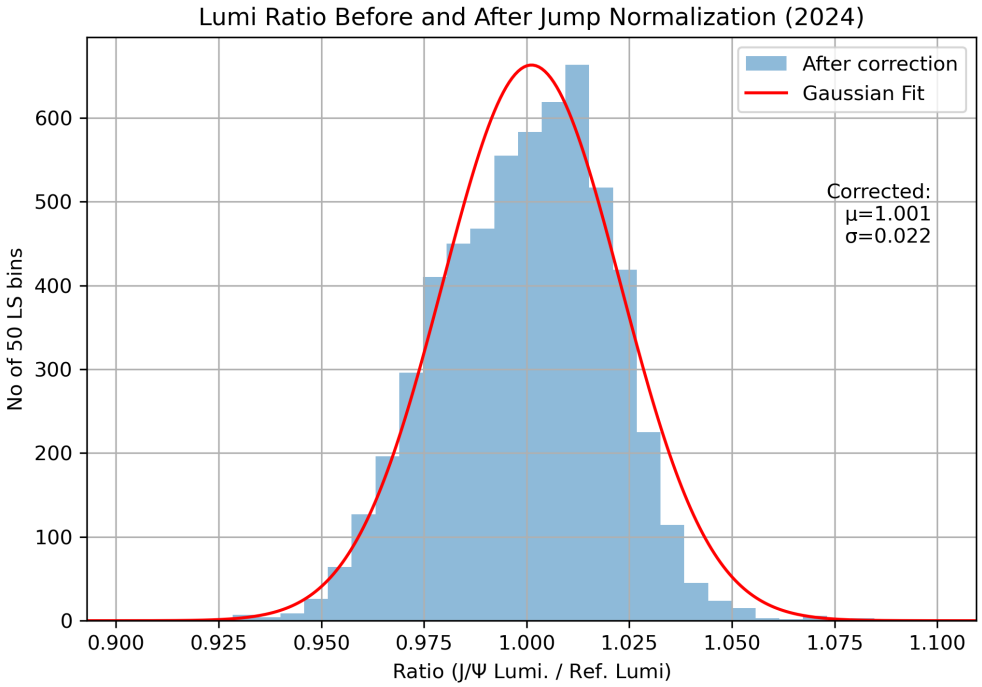
Source	2023 [%]	2024 [%]
Fit model bias	0.02	0.02
Long-term efficiency correction	0.01	0.01
Boundary normalization	0.23	0.61
Visible cross section	1.73	2.05
Total	1.74	2.14

3.6.1 Fit Model Uncertainty

The fit model uncertainty accounts for possible bias in the extracted J/ψ yield. Because the same signal-extraction model is used in the calibration and physics data,



(a) Corrected ratio



(b) Ratio distribution

Figure 3.22: Final 2024 long-term correction. The corrected $\mathcal{L}_{J/\psi}/\mathcal{L}_{\text{ref}}$ ratio has a relative standard deviation of 2.2%.

most model dependence is absorbed into the calibrated visible cross section. The remaining uncertainty is evaluated with 10,000 pseudo-experiments generated from nominal model parameters in a representative bin. Each pseudo-experiment is fitted with the same model used in data, and the study is repeated for sample sizes from 8,000 to 140,000 events. The resulting fractional-bias trend is shown in Fig. 3.23,

Table 3.9: Breakdown of the relative visible-cross-section uncertainty assigned to $\sigma_{\text{vis}}^{J/\psi}$.

Source	2023 [%]	2024 [%]
Beam current	0.25	0.25
Length scale	0.50	0.50
Beam position	0.39	0.39
Beam-beam effects	0.37	0.37
Factorization	0.80	0.80
Calibration component	1.11	1.11
Scan-statistical component	1.32	1.72
Total visible cross section	1.73	2.05

and the pull and fractional-bias distributions are shown in Fig. 3.24. The resulting fractional bias is $-0.021\% \pm 0.020\%$, and an uncertainty of 0.02% is assigned.

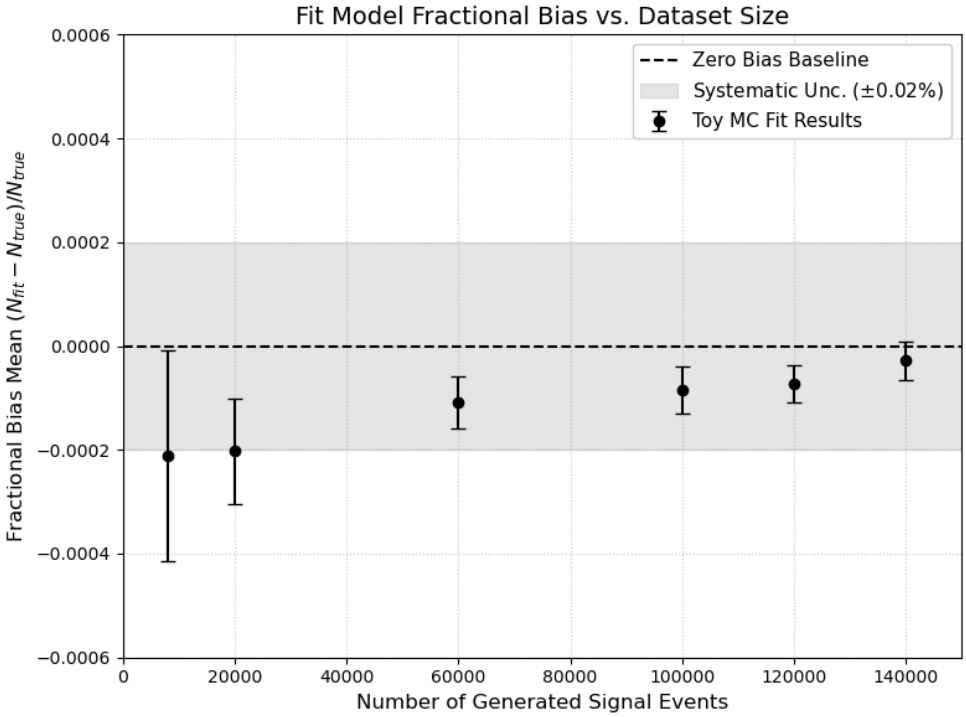
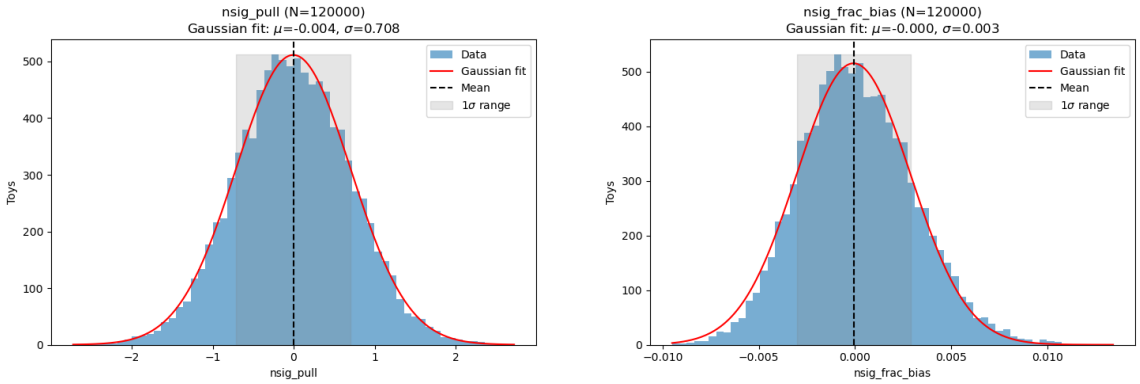


Figure 3.23: Intrinsic fit bias as a function of the number of generated signal events in pseudo-experiments.



(a) Pull distribution **(b)** Fractional-bias distribution

Figure 3.24: Validation of the fitted signal yield using pseudo-experiments for a representative sample size.

3.6.2 Visible Cross Section, Pileup, and Long-Term Uncertainties

The visible-cross-section uncertainty includes beam current, length scale, beam position, beam-beam effects, factorization, and scan-statistical components. Dedicated evaluations of the calibration terms are not feasible with the compact emittance-scan method used here, so the calibration terms are taken from corresponding CMS vdM-type luminosity-calibration studies and applied as common terms for the two years [27, 47–49]. The scan-statistical component is evaluated from the Gaussian fits to the emittance-scan rate curves and therefore differs between 2023 and 2024.

The pileup correction uncertainty is evaluated by varying the pileup-loss coefficient κ within the range observed across calibration runs. The resulting change in integrated luminosity is negligible, so no separate pileup-correction contribution is included in the final uncertainty table.

The long-term correction uncertainty has two components. The first corresponds to the BPIX Layer 4 slope correction and is estimated by varying the slope used in the nominal correction within the range observed across stability regions. The second corresponds to the RAMSES-based boundary normalization and is evaluated from

variations of the boundary-window width and from the propagated precision of the region-to-region adjustment factors.

3.7 Results

3.7.1 Lumisection Coverage

The Golden JSON defines certified lumisections suitable for physics analysis. In some runs, the certified lumisection range is wider than the part available in the processed `ParkingDoubleMuonLowMass` dataset. The J/ψ luminosity result is therefore evaluated only for lumisections that are both certified and present in the processed parking dataset.

The certified integrated luminosity missing from the processed parking dataset is 0.007 fb^{-1} in 2023 and 0.131 fb^{-1} in 2024. These lumisections are excluded from both the J/ψ result and the retained `brilcalc` comparison. A total of 101 bins in 2023 and 153 bins in 2024 are removed by quality requirements. The removed bins correspond to lumisections not present in the processed parking dataset, bins failing J/ψ yield or correction-quality requirements, and periods outside the stable region definitions.

3.7.2 Integrated Luminosity and Reference Comparison

The final integrated luminosity is

$$L_{\text{int}}^{J/\psi} = \sum_i \mathcal{L}_{J/\psi,i} \Delta t, \tag{3.19}$$

over the retained bins. The J/ψ result and the retained `brilcalc` reference are evaluated over the same LS coverage and are summarized in Table 3.10.

The final values agree with the retained reference luminosities within the assigned uncertainties. The J/ψ counting result gives an independent luminosity check based

Table 3.10: Final integrated luminosity results and retained `brilcalc` reference luminosities on the same lumisection coverage.

Year	$L_{\text{int}}^{J/\psi}$ [fb^{-1}]	$L_{\text{int}}^{\text{brilcalc}}$ [fb^{-1}]
2023	28.53 ± 0.50	28.19 ± 0.36
2024	107.57 ± 2.30	107.45 ± 1.73

on reconstructed physics objects. Its sensitivity to the tracking system, muon reconstruction quality, trigger path performance, and selection efficiency is different from that of dedicated hardware luminometers such as the PLT or BCM1F, which are based on fast particle-counting rates sensitive to occupancy, radiation damage, and bunch-crossing timing. The percent-level agreement indicates that the physics-object-based measurement remains stable over the full 2023–2024 Run 3 data-taking period and is consistent with the official CMS luminosity reference on the same lumisection coverage.

3.8 Summary

The J/ψ luminosity method uses reconstructed $J/\psi \to \mu^+\mu^-$ candidates from the `ParkingDoubleMuonLowMass` dataset to provide a physics-object-based luminosity measurement in Run 3. For the retained lumisections, the method gives $L_{\text{int}}^{J/\psi} = 28.53 \pm 0.50$ fb^{-1} in 2023 and 107.57 ± 2.30 fb^{-1} in 2024, consistent with the retained CMS reference luminosities for the same coverage.

Chapter 4

Upper Limit on Rare Higgs Boson Decays to ZJ/ψ

The upper-limit analysis in this chapter targets the rare Higgs boson decay topology in Eqn. 4.1:

$$H \to ZJ/\psi \to \ell^+\ell^-\mu^+\mu^-, \tag{4.1}$$

where the Z boson is reconstructed in either the e^+e^- or $\mu^+\mu^-$ final state and the J/ψ candidate is reconstructed through its dimuon decay. Rare exclusive Higgs decays are sensitive to small Standard Model amplitudes and to possible modifications of Higgs couplings or additional contributions from physics beyond the Standard Model [19–26]. The final state considered here contains four charged leptons, which gives a well-reconstructed topology, but the expected rate is small. The analysis uses event selection, shape information in the four-lepton mass and machine-learning discriminant, and a statistical combination of the two reconstructed Z channels.

The chapter documents the rare-decay search chain for the direct $H \to ZJ/\psi$ channel using a blinded analysis. The presentation follows the physics-analysis chain: reconstructed-object selection, candidate selection, discriminating observables, signal and background modeling, nuisance-parameter treatment, and the final 95% confidence-level upper limits.

4.1 Theory Motivation and Polarization Convention

4.1.1 Higgs Production and Analysis Scope

At the LHC, a Standard Model Higgs boson is produced mainly through gluon-gluon fusion (ggF), with smaller contributions from vector-boson fusion, associated production with a vector boson, and associated production with top quarks [8, 18]. The signal samples used in this dissertation model the ggF topology for the reconstructed signal shapes and efficiencies, while the yield normalization uses the total Standard Model Higgs production cross section listed in Sec. 4.3.

The nominal limit in this chapter is placed on the direct branching fraction $\mathcal{B}(H \to ZJ/\psi)$. Feed-down channels such as $H \to Z\psi(2S)$ followed by $\psi(2S) \to J/\psi + X$ populate the same reconstructed $J/\psi \to \mu^+\mu^-$ signature when the analysis uses the reconstructed Z and J/ψ leptons as the visible final state. The Particle Data Group reports a large inclusive branching fraction for $\psi(2S) \to J/\psi + X$ [60], and theoretical studies of exclusive Higgs decays include higher and excited quarkonium states [61, 62]. Rare-quarkonium searches include higher-quarkonium states and related exclusive final states in their broader search programs [26, 63–67]. Those channels are not included in the quoted limit because they require separate signal generation, acceptance corrections, and branching-fraction normalization.

The analysis described here is therefore a direct-channel implementation: it defines the selection, discriminating variables, likelihood inputs, and luminosity-normalized signal-yield calculation for the four-lepton topology. Extending the interpretation would require additional reconstructed Z channels, quarkonium states, feed-down contributions, and production-mode acceptance treatments.

4.1.2 Exclusive $H \to ZJ/\psi$ Decays

The rare-decay motivation and the direct and indirect amplitudes for $H \to ZJ/\psi$ were introduced in Sec. 1.2. For the statistical interpretation in this chapter, the

relevant theory input is the Standard Model branching-fraction prediction for the direct channel. Ref. [25] predicts $\mathcal{B}(H \to ZJ/\psi) = (2.30 \pm 0.11) \times 10^{-6}$.

4.1.3 Polarization and Acceptance

The polarization of the vector bosons changes the angular distributions of the final-state leptons and therefore changes the geometric and kinematic acceptance of the analysis. The CMS search for related rare Higgs decays presents limits under a longitudinal-polarization assumption and notes that alternative polarization scenarios change the signal acceptance [26]. The limit inputs in this dissertation are obtained from signal samples generated with an isotropic decay-angle treatment and are then converted to the longitudinal-polarization interpretation. The upper limits quoted in Sec. 4.8 are the longitudinal-polarized values.

The relevant angle is the helicity angle θ_h, defined here as the angle between the vector-boson direction in the Higgs rest frame and the positively charged daughter lepton in the vector-boson rest frame. For a scalar decay to two vector particles, the $H \to ZJ/\psi$ final state is expected to be dominantly longitudinal, so the longitudinal shape is used for the final interpretation. The isotropic and transverse shapes are retained as reference alternatives for acceptance studies. For a vector boson decaying to leptons, the polar-angle distribution can be parameterized by a coefficient multiplying $\cos^2\theta_h$. The limiting longitudinal and transverse cases give the shapes used below [68]. The normalized angular shapes used for the illustrative polarization comparison are given in Eqns. 4.2–4.4:

$$\frac{1}{N}\frac{dN_{\text{long}}}{d\cos\theta_h} = \frac{3}{4}\left(1 - \cos^2\theta_h\right), \tag{4.2}$$

$$\frac{1}{N}\frac{dN_{\text{trans}}}{d\cos\theta_h} = \frac{3}{8}\left(1 + \cos^2\theta_h\right), \tag{4.3}$$

$$\frac{1}{N}\frac{dN_{\text{iso}}}{d\cos\theta_h} = \frac{1}{2}. \tag{4.4}$$

The corresponding benchmark shapes are shown in Fig. 4.1. Longitudinal polarization favors leptons emitted perpendicular to the vector-boson flight direction, while transverse polarization gives more weight near $|\cos\theta_h| = 1$. Because the selected leptons must satisfy finite p_T and η requirements, these angular differences propagate into the efficiency times acceptance.

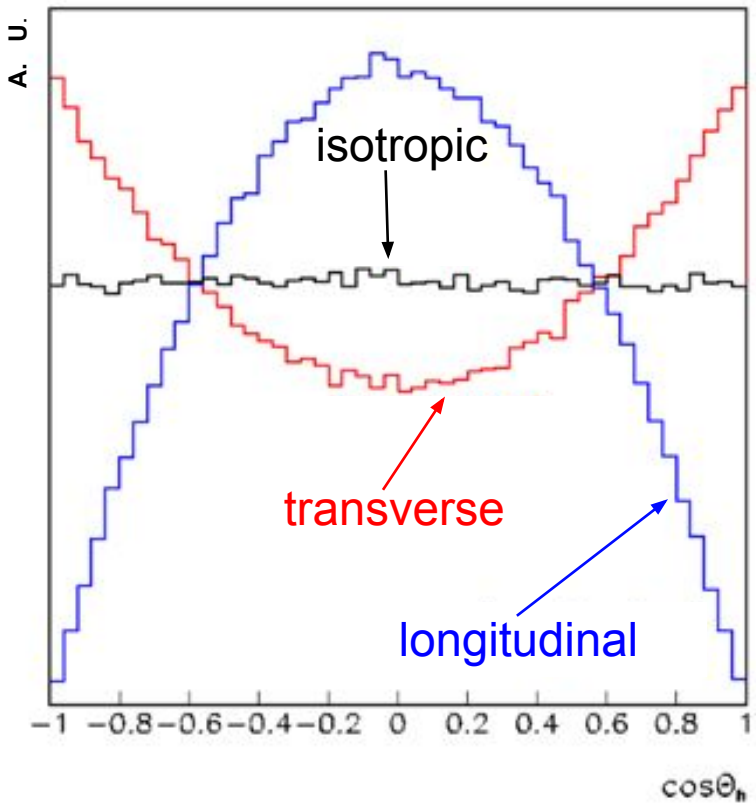


Figure 4.1: Normalized helicity-angle distributions for isotropic, transverse, and longitudinal vector-boson decay assumptions. The curves illustrate the polarization scenarios used to evaluate acceptance changes in the $H \to ZJ/\psi$ final state.

The reference same-channel study quotes acceptance changes under different polarization assumptions, taking the longitudinal scenario as the reference baseline; these benchmarks are shown in Table 4.1. For the $H \to Z(\ell\ell)J/\psi(\mu\mu)$ channel, the

isotropic acceptance is 20.4% lower than the longitudinal value. Therefore,

$$\varepsilon A_{\mathrm{iso}} = 0.796\,\varepsilon A_{\mathrm{long}}, \qquad B_{\mathrm{long}} = 0.796\,B_{\mathrm{iso}}. \tag{4.5}$$

This conversion is applied to the isotropic-limit outputs before quoting the final longitudinal-polarized limits in Sec. 4.8.

Table 4.1: Polarization acceptance benchmarks for $H \to Z(\ell\ell)J/\psi(\mu\mu)$ from the reference same-channel study. Changes are quoted relative to the longitudinal scenario.

	Longitudinal	Isotropic	Transverse
Acceptance change	0.0%	−20.4%	−30.2%

4.2 Data and Simulated Samples

The analysis uses proton-proton collision data collected by the CMS detector during Run 2 of the LHC at a center-of-mass energy of $\sqrt{s} = 13$ TeV. The selected samples are the CMS Ultra Legacy MiniAODv2 datasets summarized in Table 4.2. The mixed-flavor channel uses the electron primary datasets available in each year: `SingleElectron` in 2016 and 2017, and `EGamma` in 2018.

Table 4.2: Run 2 Ultra Legacy MiniAODv2 data eras and DAS dataset patterns used in the rare-Higgs analysis.

Data-taking period	Run eras	DAS patterns
2016 preVFP	Bver1, Bver2, C, D, E, F	`/SingleMuon/Run2016*-HIPM_UL2016_MiniAODv2*/MINIAOD`
2016 postVFP	F, G, H	`/SingleMuon/Run2016*-UL2016_MiniAODv2*/MINIAOD`
2017	B, C, D, E, F	`/SingleMuon/Run2017*-UL2017_MiniAODv2*/MINIAOD`
2018	A, B, C, D	`/SingleMuon/Run2018*-UL2018_MiniAODv2*/MINIAOD`
2016 preVFP	Bver1, Bver2, C, D, E, F	`/SingleElectron/Run2016*-HIPM_UL2016_MiniAODv2*/MINIAOD`
2016 postVFP	F, G, H	`/SingleElectron/Run2016*-UL2016_MiniAODv2*/MINIAOD`
2017	B, C, D, E, F	`/SingleElectron/Run2017*-UL2017_MiniAODv2*/MINIAOD`
2018	A, B, C, D	`/EGamma/Run2018*-UL2018_MiniAODv2*/MINIAOD`

The effective integrated luminosity is 140.33 fb^{-1} in the four-muon channel and 137.83 fb^{-1} in the mixed-flavor channel. These values are calculated using `brilcalc`

with the CMS Run 2 luminosity recommendations [30] and channel-appropriate data-certification JSON files from the Run 2 Legacy analysis recommendations [69]: the `SingleMuon` channel uses the Muon JSON, while the `SingleElectron`/`EGamma` channel uses the Golden JSON.

Signal samples are generated using PYTHIA 8 [70] with the gluon-fusion Higgs production mode and the decay chain $H \to ZJ/\psi \to \ell^+\ell^-\mu^+\mu^-$ forced at generator level. The Higgs boson mass is set to $m_H = 125$ GeV, and the signal mass window covering $112 < m_{4\ell} < 162$ GeV is used to retain both the core and tail of the reconstructed four-lepton distribution. The simulation uses the CMS Ultra Legacy campaign conditions, and the CMS detector simulation is based on Geant4 [71]. The generated sample defines the gluon-fusion signal acceptance and shape model used in the analysis.

Background estimation in this analysis uses reconstructed collision data rather than simulated background samples. The dominant background is non-peaking in the four-lepton invariant mass and is modeled from data sidebands after applying the same reconstruction and selection requirements used in the signal region. This data-driven approach treats the contributing physical processes inclusively and avoids dependence on limited simulated samples for rare multilepton backgrounds. A peaking background contribution could arise from $H \to ZZ^*$ decays where one dilepton pair falls in the J/ψ mass window; previous studies of the same $H \to ZJ/\psi$ final state found this contribution to be well below one event after full Run 2 reconstruction and selection [26], and it is therefore not included as a separate component in the statistical model.

4.3 Analysis Strategy

The analysis is divided into two final states. The four-muon channel reconstructs $Z \to \mu^+\mu^-$ with $J/\psi \to \mu^+\mu^-$, while the mixed-flavor channel reconstructs $Z \to e^+e^-$

with $J/\psi \to \mu^+\mu^-$. In both cases, the four-lepton invariant mass is the primary resonance observable for a Higgs boson near 125 GeV.

Three limit inputs are used. The first is a one-dimensional fit to the four-lepton invariant mass. The second is a one-dimensional fit to a machine-learning discriminant. The third combines the four-lepton mass and the machine-learning discriminant as a two-dimensional shape. These inputs give complementary sensitivity through the resonance mass and event-topology information.

The signal-strength parameter r is defined as the branching fraction $\mathcal{B}(H \to ZJ/\psi)$ directly. For a given fitted or tested signal yield, the corresponding branching fraction is

$$B_{H\to ZJ/\psi} = \frac{N_{\text{sig}}}{\sigma\,\mathcal{L}\,B_{Z\to\ell^+\ell^-}\,B_{J/\psi\to\mu^+\mu^-}\,(\varepsilon A)}, \tag{4.6}$$

where N_{sig} is the signal yield, $\sigma = 55.1$ pb is the total Standard Model Higgs production cross section at $\sqrt{s} = 13$ TeV [18], $\mathcal{L}$ is the channel-dependent integrated luminosity, $B_{Z\to\ell^+\ell^-}$ and $B_{J/\psi\to\mu^+\mu^-}$ are the measured branching fractions from Ref. [60], and εA is the product of selection efficiency and geometric acceptance evaluated from the signal simulation.

Upper limits on r are obtained at 95% confidence level using the modified frequentist CL_s construction [72–74]. The limit calculation uses a profile-likelihood ratio in which the signal-strength parameter is scanned and the nuisance parameters are profiled at each tested value. For a tested value of r, the CL_s ratio is written as

$$\text{CL}_\text{s}(r) = \frac{\text{CL}_{s+b}(r)}{\text{CL}_b(r)}, \tag{4.7}$$

where q_r is the profile-likelihood test statistic, CL_{s+b} is the signal-plus-background tail probability, and CL_b is the corresponding background-only tail probability. The 95% confidence-level upper limit is the value of r for which $\text{CL}_\text{s} = 0.05$. The statistical model is evaluated with the CMS Combine framework using asymptotic formulae for profile-likelihood test statistics [75, 76]. Nuisance parameters encode rate and shape uncertainties, including luminosity, branching-fraction inputs, lepton-efficiency

assignments, Higgs production theory uncertainties, and fitted mass and machine-learning shape parameters. The analysis is developed and validated with the Higgs signal window excluded from data-based checks. This blinding convention is distinct from the asymptotic approximation, which is the statistical method used to evaluate the profile-likelihood test statistic.

The likelihood inputs are derived from the reconstructed Run 2 samples and the signal simulation described above. The certified data selections define the integrated luminosity and the sideband events used for the background model in each channel. The signal simulation provides the acceptance, selection efficiency, and signal-shape parameters after the final candidate selection. The sideband data are fitted to obtain the nonresonant background shapes used in the blinded signal-region model. The one-dimensional inputs use either the four-lepton invariant mass or the random forest score, while the two-dimensional input uses their joint distribution. Systematic uncertainties are then assigned as normalization and shape nuisance parameters. The resulting blinded profile-likelihood model gives the expected upper limits reported in Sec. 4.8.

4.4 Event Reconstruction and Selection

The analysis begins from events containing well-reconstructed primary vertices and leptons compatible with CMS electron and muon reconstruction [16, 41, 42, 55, 77–79]. The preselection defines broad Z and J/ψ candidates before the final mass-window and discriminant inputs are formed. The main reconstructed-object requirements are summarized in Table 4.3.

The soft-muon identification requirement used in the final selection is summarized in Table 4.4.

Events are required to satisfy the relevant high-level trigger (HLT) path for the data-taking year and channel. The specific HLT paths used in each year are summarized in Table 4.5. In the mixed-flavor channel, the quoted electron transverse

Table 4.3: Preselection requirements by final state. Masses are in GeV.

Requirement	$Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$	$Z(e^+e^-)J/\psi(\mu^+\mu^-)$
Event and trigger	Muon trigger fired	Electron trigger fired
Z candidate	Opposite-sign dimuon pair with high-purity tracks; $60 < m_Z < 120$; vertex probability > 0.005	Opposite-sign electron pair with valid tracks; $60 < m_Z < 120$; vertex probability > 0.005
J/ψ candidate	Opposite-sign dimuon pair with high-purity tracks; $2 < m_{\mu\mu} < 4$; vertex probability > 0.005	Opposite-sign dimuon pair with high-purity tracks; $2 < m_{\mu\mu} < 4$; vertex probability > 0.005
Lepton kinematics	Four muons with $p_T \geq 2$ and $\lvert\eta\rvert \leq 2.4$; at least one muon with $p_T > 20$	Electrons with $p_T \geq 4$ and $\lvert\eta\rvert \leq 2.5$, leading electron $p_T \geq 20$; muons with $p_T \geq 2$ and $\lvert\eta\rvert \leq 2.4$
Four-lepton vertex	Four-muon vertex probability > 0.005	Four-lepton vertex probability > 0.005

Table 4.4: Soft-muon identification requirements used for selected muons.

Requirement	Selection
Tracker layers with hits	> 5
Pixel layers with hits	> 0
Muon chambers with matched segments	> 1
Hits in the muon system	> 0
Longitudinal impact parameter	$\lvert d_z\rvert < 20$ cm
Transverse impact parameter	$\lvert d_{xy}\rvert < 0.3$ cm

momenta use the EGamma post-correction values following the Ultra Legacy electron recommendations [80].

Table 4.5: Trigger paths used in the rare-Higgs analysis.

Year	$Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$	$Z(e^+e^-)J/\psi(\mu^+\mu^-)$
2016	`HLT_IsoMu24`	`HLT_Ele27_WPTight_Gsf`
2017	`HLT_IsoMu27`	`HLT_Ele32_WPTight_Gsf_L1DoubleEG` + `HLT_Ele35_WPTight_Gsf`
2018	`HLT_IsoMu24`	`HLT_Ele32_WPTight_Gsf`

After preselection, final selection tightens the resonance windows, lepton quality, transverse-momentum thresholds, and vertex requirements. These choices suppress combinatorial and nonresonant backgrounds while keeping the $H \to ZJ/\psi$ signal region in the four-lepton mass spectrum. The final candidate selections are summarized in Table 4.6.

Table 4.6: Final candidate-selection requirements used for the upper-limit inputs. Masses and transverse momenta are in GeV.

Requirement	$Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$	$Z(e^+e^-)J/\psi(\mu^+\mu^-)$
Trigger and identification	Muon trigger; all four muons pass soft identification	Electron trigger; both J/ψ muons pass soft identification; both electrons pass loose identification
Lepton kinematics	All muons have $p_T > 3$ and $\|\eta\| < 2.4$	J/ψ muons have $p_T > 3$ and $\|\eta\| < 2.4$; electrons have $\|\eta\| < 2.5$, with $p_T > 32$ and $p_T > 5$ for the leading and subleading electrons
Isolation	Included through the reconstructed variables used by the discriminant	Each J/ψ muon has absolute track isolation below 10 in a cone of $\Delta R = 0.3$
Candidate kinematics	$p_T(J/\psi) > 5$; $p_T(Z) > 5$; $p_T(4\ell) > 5$	$p_T(4\ell) > 5$
Vertex quality	J/ψ, Z, and four-lepton vertex probabilities > 0.01	J/ψ, Z, and four-lepton vertex probabilities > 0.01
Mass windows	$80 < m_Z < 100$; $3.0 < m_{J/\psi} < 3.2$; $112 < m_{4\ell} < 162$	$80 < m_Z < 100$; $3.0 < m_{J/\psi} < 3.2$; $112 < m_{4\ell} < 162$
ML background training sample	Relaxed windows, $70 < m_Z < 110$ and $2.8 < m_{J/\psi} < 3.4$, are used only for the ML background sample; the Higgs region $120 < m_{4\ell} < 130$ is excluded	Relaxed windows, $70 < m_Z < 110$ and $2.8 < m_{J/\psi} < 3.4$, are used only for the ML background sample; the Higgs region $120 < m_{4\ell} < 130$ is excluded

The sideband selection constrains the smooth nonresonant background from data while preserving blindness in the Higgs signal window. The relaxed Z and J/ψ mass windows in Table 4.6 are used only to increase the background sample available for ML training and validation. They are not a replacement for the nominal final signal-region selection. Simulated signal candidates are passed through the same final selection and are used to build the signal templates and signal-normalization inputs used in the limit calculation. Tables 4.7 and 4.8 summarize the event yields after each selection step for the two channels.

Table 4.7: Selection yields after each selection step for the $Z(e^+e^-)J/\psi(\mu^+\mu^-)$ channel. Signal counts are raw simulated-event yields from 10,000 generated PYTHIA signal events and are not normalized to the expected cross section here. Background counts correspond to the sideband-blinded data selection; the final row excludes events in the signal window $120 < m_{4\ell} < 130$ GeV, which is not applied to the signal simulation.

Selection step	Signal	Background
Preselection	3064	31943
Trigger	3064	31943
Muon identification	2982	9595
Electron identification	2847	4893
Muon p_T requirement	2747	2931
Electron p_T requirement	2746	2910
Detector acceptance	2746	2910
Z mass window	2395	2036
J/ψ mass window	2362	472
Dilepton vertex quality	2698	2846
4ℓ vertex quality	2664	2752
4ℓ p_T requirement	2584	2632
4ℓ mass window	2355	236
Signal window blinded	—	173

Table 4.9 lists the numerical denominator inputs used to quote the upper limits in branching-fraction units.

Table 4.8: Selection yields after each selection step for the $Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$ channel. Signal counts are raw simulated-event yields from 10,000 generated PYTHIA signal events and are not normalized to the expected cross section here. Background counts correspond to the sideband-blinded data selection; the final row excludes events in the signal window $120 < m_{4\ell} < 130$ GeV, which is not applied to the signal simulation.

Selection step	Signal	Background
Preselection	3701	9885
Trigger	3583	9885
Muon identification	3493	2678
Muon p_T requirement	3367	1678
Detector acceptance	3367	1678
Z mass window	3367	1311
J/ψ mass window	3367	803
Dilepton vertex quality	3311	794
4ℓ vertex quality	3272	760
J/ψ p_T requirement	3242	734
4ℓ p_T requirement	3114	707
4ℓ mass window	3103	347
Signal window blinded	—	268

Table 4.9: Numerical denominator inputs used with Eqn. 4.6.

Quantity	$Z(e^+e^-)J/\psi(\mu^+\mu^-)$	$Z(\mu^+\mu^-)J/\psi(\mu^+\mu^-)$
σ (pb)	55.1	55.1
$\mathcal{L}$ (fb^{-1})	137.83	140.33
$B_{Z\to\ell^+\ell^-}$ (%)	3.3632	3.3662
$B_{J/\psi\to\mu^+\mu^-}$ (%)	5.961	5.961
εA (%)	23.55	31.02

4.5 Machine-Learning Discriminant

A random forest is an ensemble of decision trees trained to discriminate signal from background. The classifier output is a continuous score in $[0, 1]$, where values near 1 indicate signal-like candidates [81]. Random forest classifiers are trained separately for the two reconstructed final states. The signal class is built from simulated $H \to ZJ/\psi$ samples, while the background class is built from data candidates selected outside the blinded Higgs mass window. The relaxed Z and J/ψ windows described

in Table 4.6 are used only for this ML background sample, where they increase the training statistics without using candidates in $120 < m_{4\ell} < 130$ GeV.

Each classifier uses 12 input variables, listed in Table 4.10. The variables describe the lepton kinematics, the isolation of the reconstructed J/ψ and Z systems, the J/ψ vertex quality, and the angular separation between the two reconstructed resonances. The relative isolation variables use the scalar sum of track transverse momenta in a cone of $\Delta R < 0.3$ around the resonance direction, excluding the lepton tracks assigned to that resonance when they fall inside the cone, divided by the corresponding resonance transverse momentum. The separation variable $\Delta R(Z, J/\psi)$ probes the event topology and helps distinguish the two-body Higgs decay hypothesis from continuum four-lepton production.

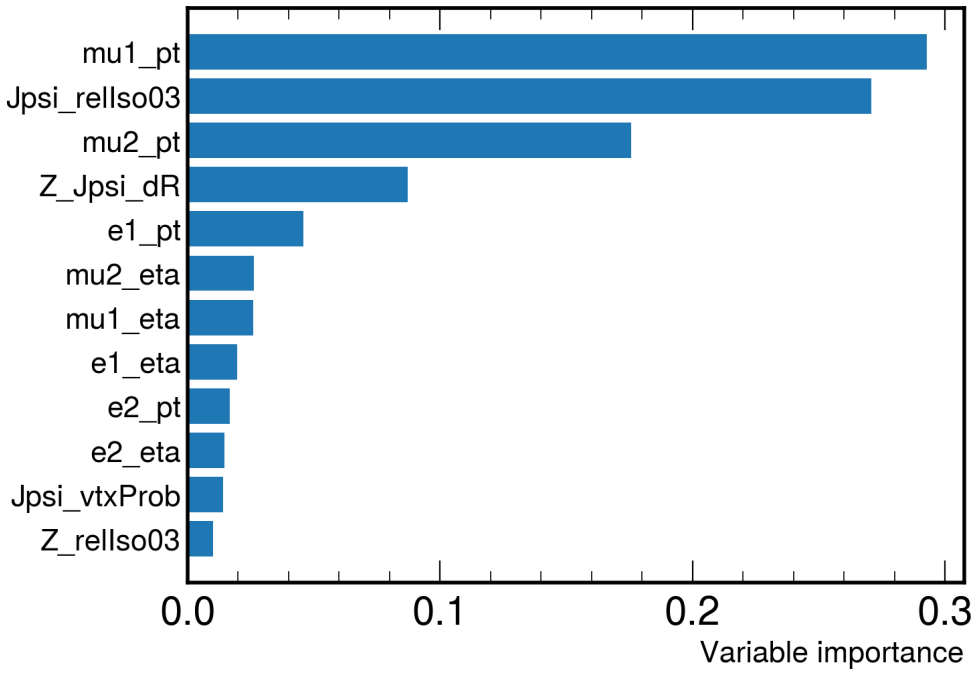
Table 4.10: Input variables used in the random forest classifiers. Both channels use the same set of physics observables.

Variable group	Variables
J/ψ daughter kinematics	$p_T(\mu_1)$, $p_T(\mu_2)$, $\eta(\mu_1)$, $\eta(\mu_2)$
Z daughter kinematics	$p_T(\ell_1)$, $p_T(\ell_2)$, $\eta(\ell_1)$, $\eta(\ell_2)$
Isolation and vertex quality	Relative J/ψ isolation, relative Z isolation, J/ψ vertex probability
Resonance separation	$\Delta R(Z, J/\psi)$

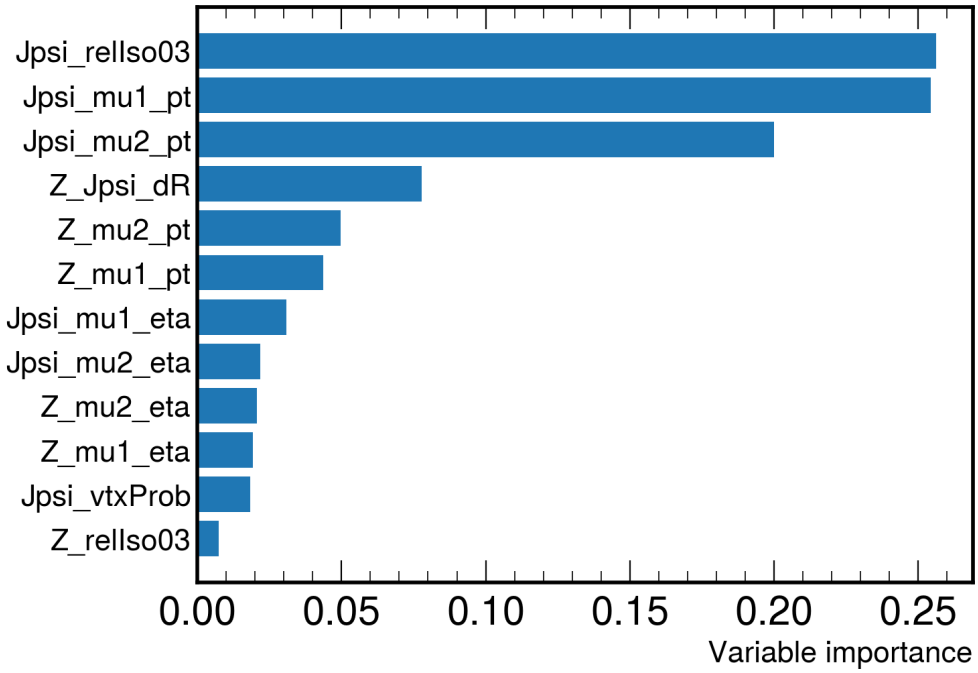
The classifiers are implemented in the Toolkit for Multivariate Analysis (TMVA) [82] as random forests with 200 decision trees and bagged training samples. Training uses 500 signal and 500 background candidates in the mixed-flavor channel and 350 signal and 350 background candidates in the four-muon channel. The held-out test samples contain 1855 signal and 295 background candidates for the mixed-flavor channel and 2753 signal and 1000 background candidates for the four-muon channel. The discriminant output is used directly in the interval $0 < \text{score} < 1$ for the shape analysis.

Figure 4.2 ranks the input variables by separation power in the trained random forest. The ranking is used as a diagnostic rather than as a variable-removal procedure: all variables in Table 4.10 remain in the nominal classifier. Figure 4.3 shows the ROC curves for the held-out test samples. The corresponding areas under the curve are 0.950 for the mixed-flavor channel and 0.955 for the four-muon channel. The high signal efficiency at moderate background rejection confirms that the classifier provides useful separation beyond the four-lepton mass. Figure 4.4 shows the linear correlations among background input variables. These correlations are monitored because strongly correlated inputs can reduce the interpretability of the ranking and can make the classifier more sensitive to mismodeling in small data control samples.

To assess model-to-model variation in the trained discriminant, the random forest is retrained 100 times with independent random seeds controlling the training/testing split of the input samples, while keeping the same training procedure. The resulting score distributions for signal and background test candidates remain consistent across retrainings, with no individual model producing a notably different shape. This test is treated conservatively, since it also includes possible stochastic effects from the training procedure. A 2% nuisance-parameter uncertainty is assigned to cover the resulting model-to-model variation.

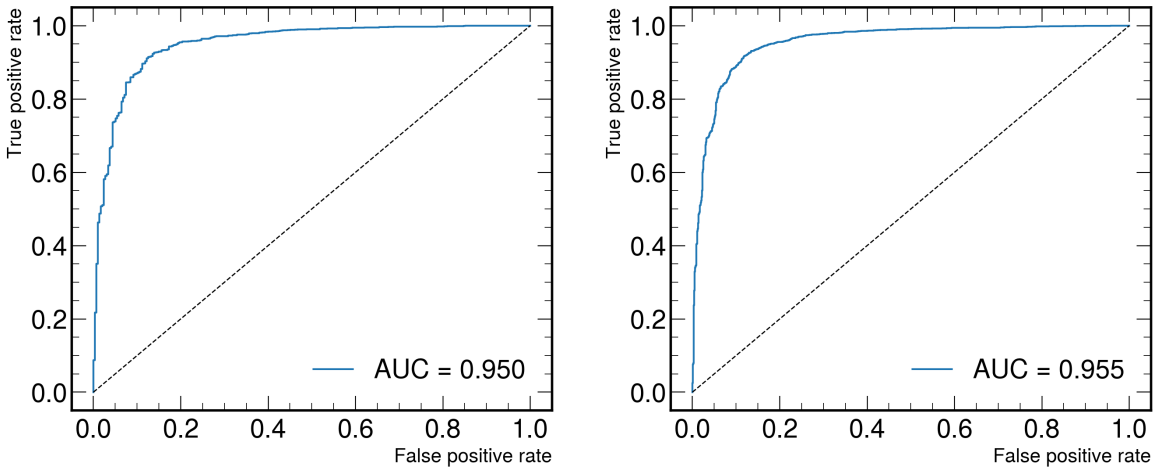


(a) $Z(ee)J/\psi(\mu\mu)$



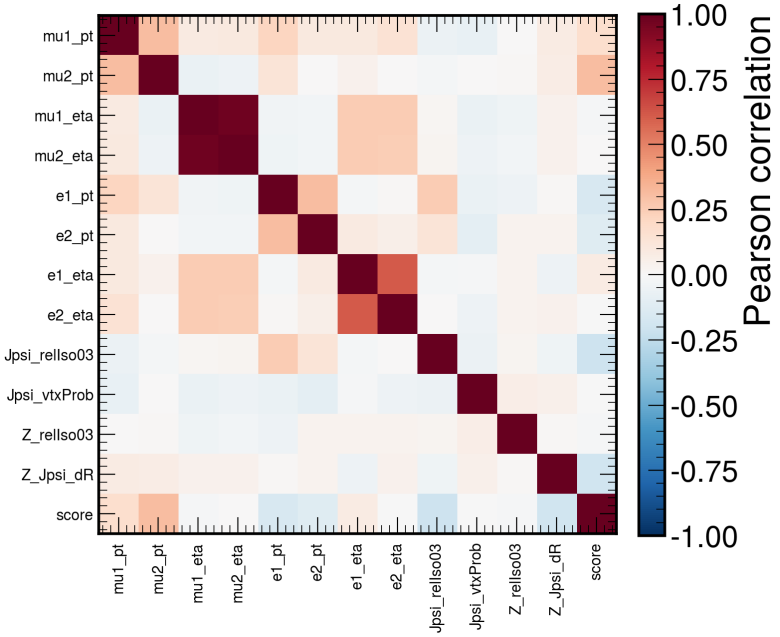
(b) $Z(\mu\mu)J/\psi(\mu\mu)$

Figure 4.2: Ranking of the machine-learning input variables by separation power for the $Z(ee)J/\psi(\mu\mu)$ channel (top) and the $Z(\mu\mu)J/\psi(\mu\mu)$ channel (bottom). The rankings identify the variables most useful for discriminating the rare Higgs signal hypothesis from the background model in each channel.

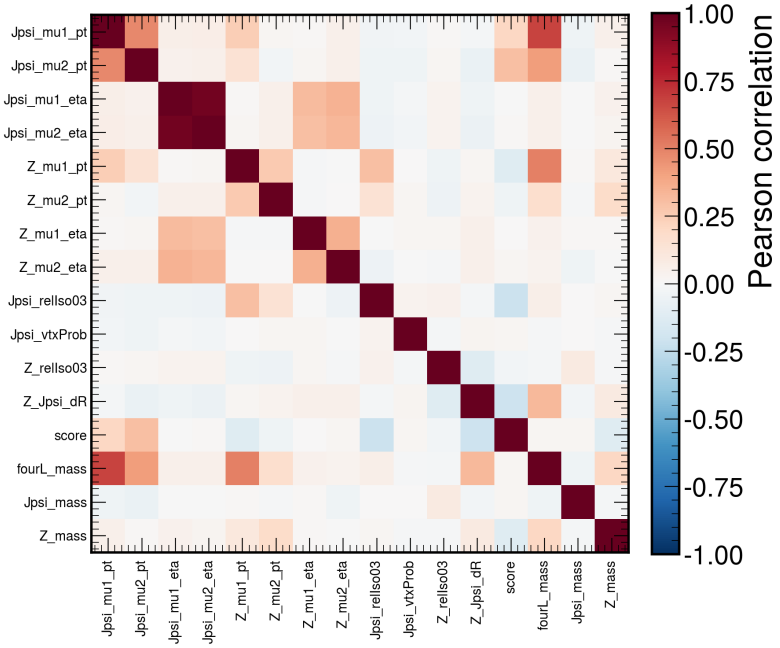


(a) $Z(ee)J/\psi(\mu\mu)$ **(b)** $Z(\mu\mu)J/\psi(\mu\mu)$

Figure 4.3: ROC curves for the random forest classifiers evaluated on held-out test samples. The curves show the tradeoff between signal efficiency and background rejection for the mixed-flavor channel (a) and the four-muon channel (b).



(a) $Z(ee)J/\psi(\mu\mu)$



(b) $Z(\mu\mu)J/\psi(\mu\mu)$

Figure 4.4: Background-sample input-variable correlation matrices for the random forest classifiers. The matrices are used to check whether the training inputs contain strong internal correlations that could affect the stability of the discriminant shape.

4.6 Signal and Background Modeling

4.6.1 Invariant-Mass Modeling

The invariant-mass model describes the Higgs signal peak near 125 GeV and a smooth nonresonant background. The mass fits are implemented with RooFit [58] and are used directly in the one-dimensional mass analysis and as one axis of the two-dimensional shape analysis. The signal shape in the four-lepton mass is modeled with a Crystal Ball plus Gaussian parameterization,

$$f_{\text{sig}}(m_{4\ell}) = f_{\text{CB}}\,\text{CB}(m_{4\ell};\mu,\sigma_{\text{CB}},\alpha_{\text{CB}},n) + (1-f_{\text{CB}})\,G(m_{4\ell};\mu,\sigma_G), \tag{4.8}$$

where the Crystal Ball component describes the narrow peak with a low- or high-side tail and the Gaussian component absorbs the broader detector-resolution contribution. The smooth background is modeled with

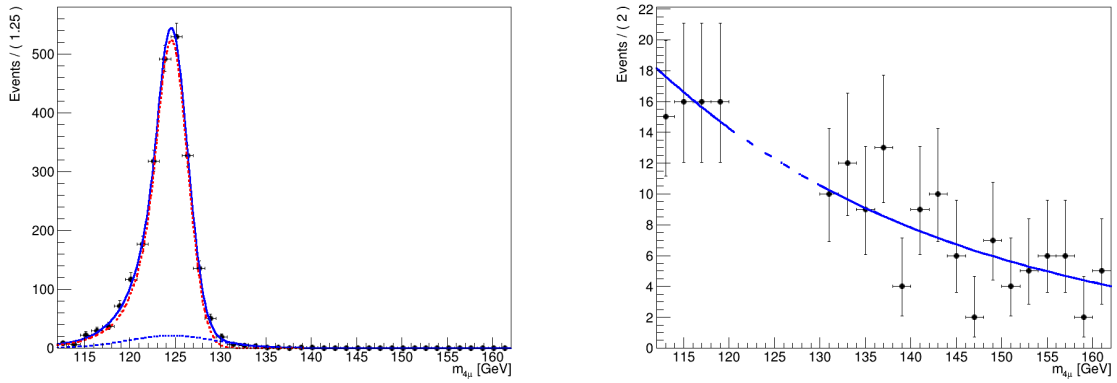
$$f_{\text{bkg}}(m_{4\ell}) \propto \exp(\alpha m_{4\ell}), \tag{4.9}$$

where the slope α is obtained from the sideband-based background sample. Figures 4.5 and 4.6 show the fitted signal and background mass parameterizations, and Fig. 4.7 shows an illustrative pseudo-data mass distribution used to visualize the blinded model. Tables 4.11 and 4.12 list the fitted mass-model parameters. The reduced chi-square values of the fitted models are close to unity.

4.6.2 ML-Score Modeling

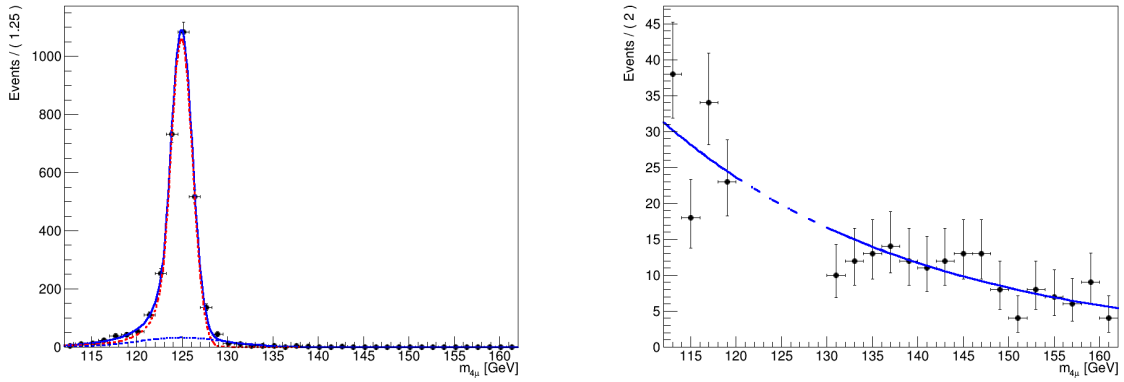
The machine-learning score distribution is modeled separately for signal and background. Double-exponential functions are used for the signal and background score shapes in both channels,

$$f(s) = f_1\exp(\alpha_1 s) + (1-f_1)\exp(\alpha_2 s), \qquad 0 < s < 1, \tag{4.10}$$



(a) Signal MC **(b)** Background data

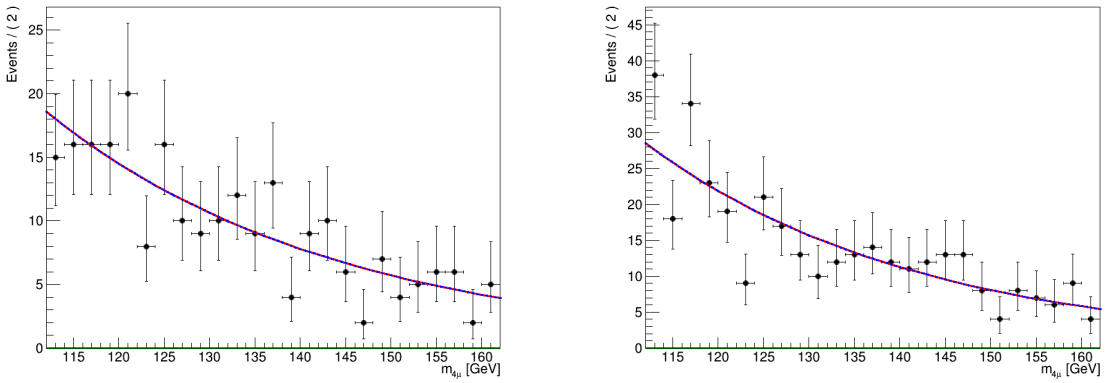
Figure 4.5: Four-lepton invariant-mass signal and background parameterizations for the $Z(ee)J/\psi(\mu\mu)$ channel. The signal model is fitted to simulated signal MC, and the background model is fitted to data in the sideband-based background sample.



(a) Signal MC **(b)** Background data

Figure 4.6: Four-lepton invariant-mass signal and background parameterizations for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The signal model is fitted to simulated signal MC, and the background model is fitted to data in the sideband-based background sample.

where s is the random forest score, α_1 and α_2 are the two score-shape slopes, and f_1 is the fitted mixture fraction. This form gives enough flexibility to describe both the signal distribution concentrated near high scores and the background distribution concentrated near low scores without introducing a high-order polynomial shape. Figures 4.8 and 4.9 show the fitted signal and background score parameterizations in



(a) $Z(ee)J/\psi(\mu\mu)$ **(b)** $Z(\mu\mu)J/\psi(\mu\mu)$

Figure 4.7: Illustrative pseudo-data four-lepton invariant-mass distributions for the blinded model. The displayed background spectrum is sampled from the fitted background model, while the signal shape is modeled with the Crystal Ball plus Gaussian parameterization.

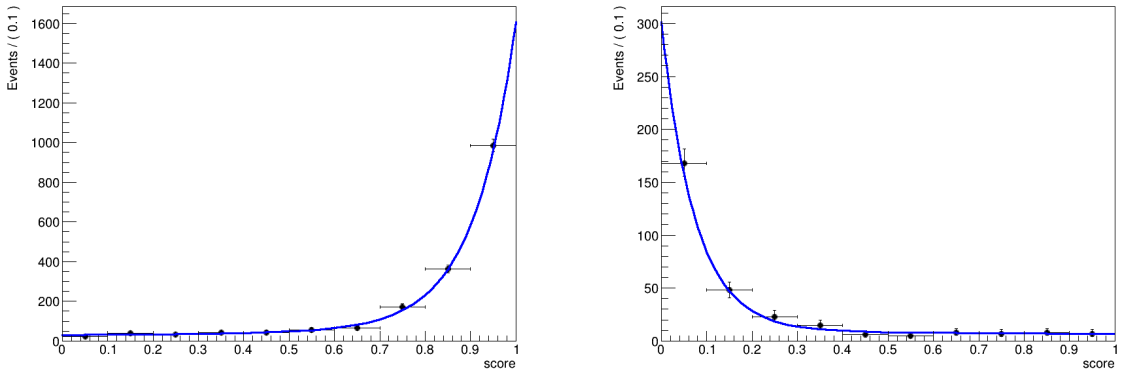
Table 4.11: Fitted signal mass-shape parameters for the Crystal Ball plus Gaussian model in each channel. Uncertainties reflect the fit precision from the signal Monte Carlo sample.

Parameter	$Z(\mu\mu)J/\psi(\mu\mu)$	$Z(ee)J/\psi(\mu\mu)$
Mean	124.913 ± 0.032	124.553 ± 0.088
σ_{CB}	1.210 ± 0.034	1.794 ± 0.097
α_{CB}	1.465 ± 0.119	0.913 ± 0.140
n	2.110 ± 0.416	7.867 ± 5.643
f_{CB}	0.899 ± 0.016	0.912 ± 0.032
σ_G	4.784 ± 0.361	4.873 ± 0.636
N_{fit}	3102.7 ± 55.7	2355.1 ± 48.5

Table 4.12: Fitted four-lepton mass background parameters. The background is modeled with an exponential function in each channel.

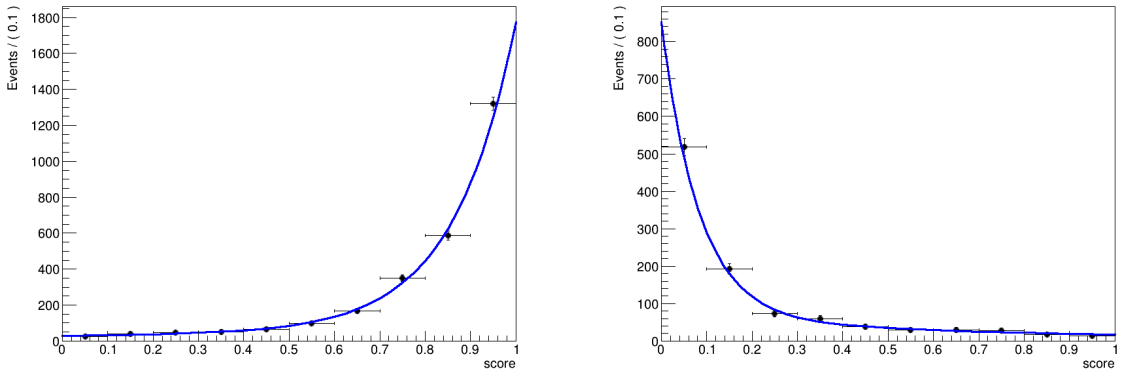
Parameter	$Z(\mu\mu)J/\psi(\mu\mu)$	$Z(ee)J/\psi(\mu\mu)$
Exponential slope α	-0.03509 ± 0.00403	-0.03012 ± 0.00499
N_{fit}	269.07 ± 16.40	172.99 ± 13.15

the two channels. Tables 4.13 and 4.14 list the fitted signal and background score-shape parameters. The reduced chi-square values of the fitted models are close to unity.



(a) Signal Monte Carlo (MC) **(b)** Background data (loose selection)

Figure 4.8: ML discriminant parameterization for the $Z(ee)J/\psi(\mu\mu)$ channel. The signal distribution is shown in (a) using simulated signal MC events; the background distribution is shown in (b) using background candidates obtained with the loose selection criteria. Fitted double-exponential models are overlaid.



(a) Signal Monte Carlo (MC) **(b)** Background data (loose selection)

Figure 4.9: ML discriminant parameterization for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The signal distribution is shown in (a) using simulated signal MC events; the background distribution is shown in (b) using background candidates obtained with the loose selection criteria. Fitted double-exponential models are overlaid.

4.6.3 Two-Dimensional Diagnostic Configuration

The two-dimensional analysis combines the four-lepton mass and machine-learning score in a joint likelihood, using the two-dimensional signal and background templates built from the parameterizations described above. For visualization, diagnostic

Table 4.13: Fitted double-exponential signal ML-score parameters. The fit uses the random forest score over the interval $0 < s < 1$.

Parameter	$Z(ee)J/\psi(\mu\mu)$	$Z(\mu\mu)J/\psi(\mu\mu)$
α_1	14.064 ± 2.466	6.859 ± 0.514
α_2	2.440 ± 0.877	-0.392 ± 1.811
Fraction	0.638 ± 0.104	0.925 ± 0.050
N_{fit}	295.0 ± 17.2	1000.0 ± 31.6

Table 4.14: Fitted double-exponential background ML-score parameters. The fit uses the random forest score over the interval $0 < s < 1$.

Parameter	$Z(ee)J/\psi(\mu\mu)$	$Z(\mu\mu)J/\psi(\mu\mu)$
α_1	-13.673 ± 2.236	-12.056 ± 1.094
α_2	-0.498 ± 0.975	-1.309 ± 0.465
Fraction	0.722 ± 0.077	0.655 ± 0.051
N_{fit}	295.0 ± 17.2	999.9 ± 31.6

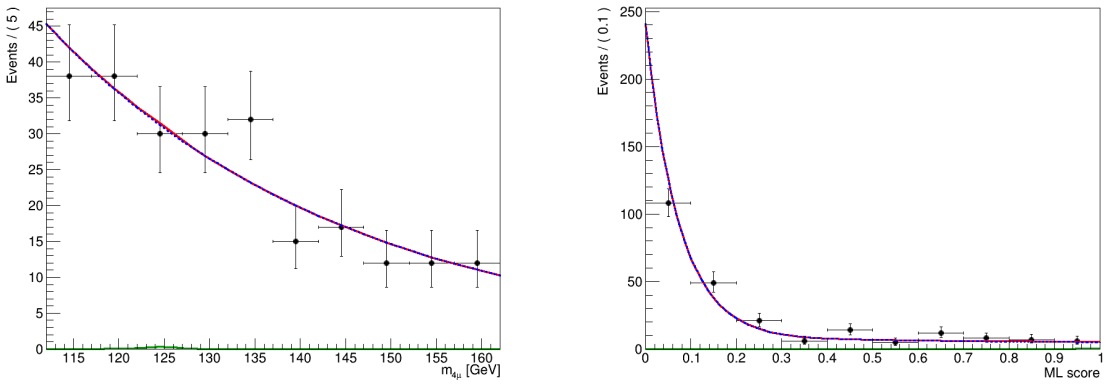
pseudo-data samples are drawn from the fitted background model in the blinded signal window.

Table 4.15 summarizes the linear correlation coefficients between the four-lepton mass and the random forest score in the held-out signal test samples. In both channels the correlation is small, supporting the use of the two observables together without a large residual dependence between the mass resonance and ML discriminant axes.

Table 4.15: Linear correlation coefficients between the four-lepton mass and the random forest score in the held-out signal test samples, quoted as percentages.

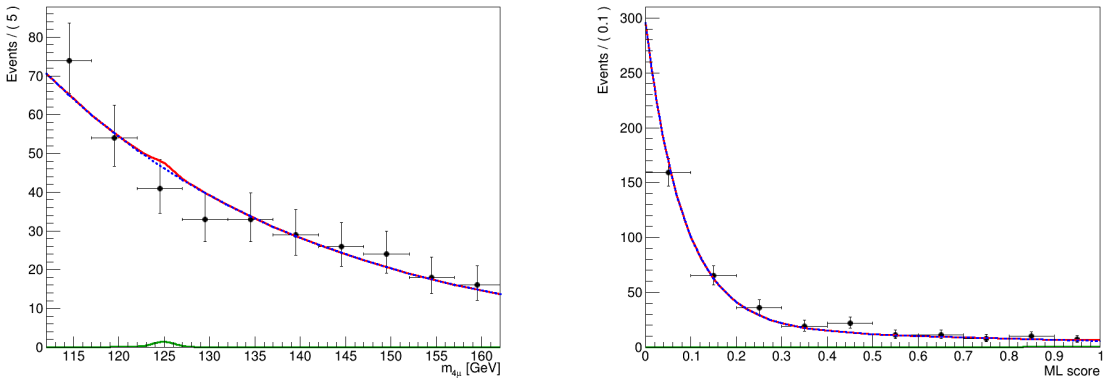
Channel	Signal test events	Correlation (%)
$Z(ee)J/\psi(\mu\mu)$	2005	0.09
$Z(\mu\mu)J/\psi(\mu\mu)$	2753	2.86

Figures 4.10 and 4.11 show the corresponding mass and ML-score projections for the two channels.



(a) Four-lepton mass projection **(b)** Random forest score projection

Figure 4.10: Diagnostic pseudo-data projections of the two-dimensional fit onto the four-lepton mass (a) and random forest score (b) axes for the $Z(ee)J/\psi(\mu\mu)$ channel. The fitted model is overlaid on the pseudo-data in both projections.



(a) Four-lepton mass projection **(b)** Random forest score projection

Figure 4.11: Diagnostic pseudo-data projections of the two-dimensional fit onto the four-lepton mass (a) and random forest score (b) axes for the $Z(\mu\mu)J/\psi(\mu\mu)$ channel. The fitted model is overlaid on the pseudo-data in both projections.

4.7 Systematic Uncertainties

The statistical model includes normalization nuisance parameters, encoded as log-normal constraints, and fitted shape parameters encoded as Gaussian-constrained parameters in the Combine framework. The normalization uncertainties account for the limited knowledge of the integrated luminosity, decay branching fractions, detector response, lepton efficiencies, and signal-modeling inputs. The fitted shape parameters describe the mass and ML-score probability density functions and are constrained by the corresponding signal and sideband fits.

Table 4.16: Experimental and analysis-model normalization nuisance parameters included in the statistical model. The magnitude column gives the fractional uncertainty applied as the log-normal width.

Source	Magnitude
Integrated luminosity	0.73%
$\mathcal{B}(J/\psi \to \mu^+\mu^-)$	0.6%
$\mathcal{B}(Z \to e^+e^-)$	0.1%
$\mathcal{B}(Z \to \mu^+\mu^-)$	0.2%
Trigger efficiency (muon)	1.0%
Trigger efficiency (electron)	1.0%
Muon reconstruction and identification	2.0%
Muon isolation	0.5%
Four-lepton vertex probability	2.0%
Signal and background modeling	2.0%

Table 4.17: Theory uncertainties for the gluon-fusion signal normalization.

Source	Magnitude
QCD scale, $gg \to H$	$+4.6\%/-6.7\%$
PDF uncertainty, Higgs gg initial state	3.2%

The entries in Table 4.16 include both externally recommended inputs and conservative analysis assignments:

- **Integrated luminosity:** the 0.73% uncertainty is the combined Run 2 CMS luminosity uncertainty [29, 30].

- **Branching fractions:** the uncertainties for $J/\psi \to \mu^+\mu^-$, $Z \to e^+e^-$, and $Z \to \mu^+\mu^-$ are taken from Ref. [60].

- **Trigger efficiency:** 1% uncertainties are assigned for the muon and electron trigger categories. These are conservative normalization terms, with the scale motivated by CMS lepton-trigger performance studies based on tag-and-probe methods [42, 79].

- **Muon reconstruction and identification:** a 2% uncertainty is assigned to cover residual muon reconstruction and identification efficiency effects. The size is a conservative event-level assignment guided by CMS muon performance measurements [42].

- **Muon isolation:** a 0.5% uncertainty is assigned for residual isolation-efficiency effects.

- **Four-lepton vertex probability:** a 2% uncertainty is assigned to cover possible mismodeling of the common four-lepton vertex requirement.

- **Signal and background modeling:** the mass and machine-learning score shapes are fitted from signal simulation and sideband-selected background data. Their fitted parameters enter the likelihood as Gaussian-constrained nuisance parameters; the additional 2% term conservatively covers residual dependence on the chosen parameterizations.

The entries in Table 4.17 are applied to the gluon-fusion signal normalization:

- **QCD scale, $gg \to H$:** the scale uncertainty reflects the theory uncertainty on the gluon-fusion Higgs production cross section used in the signal-yield normalization.

- **PDF uncertainty, Higgs gg initial state:** the PDF uncertainty is treated as a separate production-rate uncertainty for the same gluon-fusion signal component.

The fitted mass and ML-shape parameters listed in Tables 4.11, 4.12, 4.13, and 4.14 are included in the likelihood as Gaussian-constrained nuisance parameters. The smooth background normalization is constrained by the sideband data rather than fixed from simulation.

4.7.1 Impact of Systematic Uncertainties

Figure 4.12 shows the first page of the impact ranking for the combined two-dimensional expected-limit model. The largest impacts come from the gluon-fusion theory uncertainties and from the fitted ML-score shape parameters, with additional contributions from muon-efficiency, integrated-luminosity, and residual modeling uncertainties.

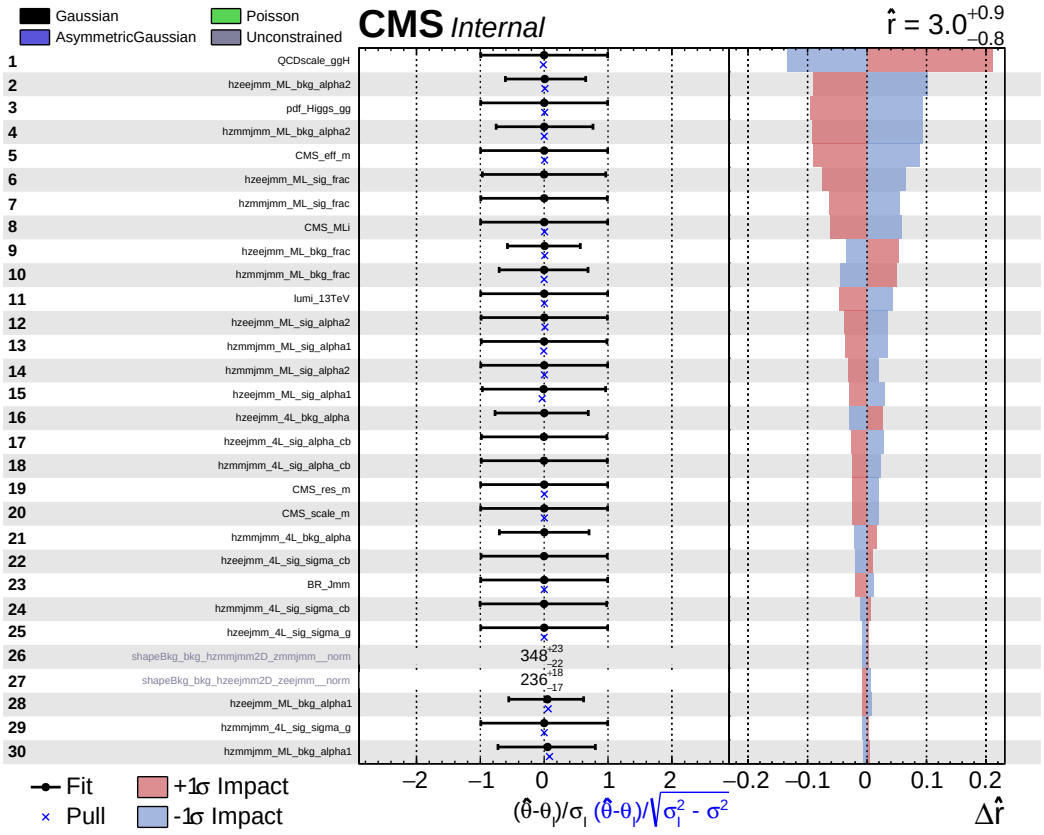


Figure 4.12: First page of the combined impact plot for $H \to ZJ/\psi$. The nuisance parameters with the largest post-fit impact on the combined expected upper limit are ranked by magnitude and shown with their $\pm 1\sigma$ post-fit variations.

4.8 Upper Limits

The limit calculation is performed separately for each final state and each analysis input, and then repeated after combining the $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ channels. The values quoted here are converted to the longitudinal-polarization interpretation described in Sec. 4.1.3. Table 4.18 gives the expected 95% confidence-level upper limits on $\mathcal{B}(H \to ZJ/\psi)$.

Table 4.18: Expected 95% confidence-level upper limits on $\mathcal{B}(H \to ZJ/\psi)$ in the longitudinal-polarization interpretation. The 2D combined input gives the strongest expected sensitivity.

Category	4L ($\times 10^{-3}$)	ML ($\times 10^{-3}$)	2D ($\times 10^{-3}$)
$Z \to e^+e^-$	4.11	2.62	1.47
$Z \to \mu^+\mu^-$	2.75	2.23	1.07
Combined	2.22	1.62	0.78

The expected limit improves when the two reconstructed Z channels are combined. In the combined result, the two-dimensional input gives an expected limit of 0.78×10^{-3} on $\mathcal{B}(H \to ZJ/\psi)$.

The Standard Model prediction used for context in Sec. 4.1.2 is $\mathcal{B}(H \to ZJ/\psi) = 2.30 \times 10^{-6}$ [25]. The combined two-dimensional expected limit is therefore about 340 times larger than the Standard Model expectation.

Table 4.19 gives the full set of expected quantile bands for the combined channel result. The $\pm 1\sigma$ and $\pm 2\sigma$ bands are the expected quantiles from the asymptotic $\mathrm{CL_s}$ calculation and quantify the statistical uncertainty in the expected sensitivity. The two-dimensional input has both the strongest median expected sensitivity and the narrowest bands.

4.9 Summary

The upper-limit analysis searches for the rare Higgs boson decay topology $H \to ZJ/\psi \to \ell^+\ell^-\mu^+\mu^-$ in the $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ channels. The Run 2 dataset

Table 4.19: Expected 95% CL upper limits on $\mathcal{B}(H \to ZJ/\psi)$ for the combined ($Z \to e^+e^-$ plus $Z \to \mu^+\mu^-$) result in the longitudinal-polarization interpretation, with $\pm1\sigma$ and $\pm2\sigma$ quantile bands. Values are in units of 10^{-3}.

Input	-2σ	-1σ	Median exp.	$+1\sigma$	$+2\sigma$
4L (mass only)	1.14	1.55	2.22	3.24	4.55
ML (score only)	0.82	1.12	1.62	2.37	3.38
2D (mass + ML)	0.37	0.52	0.78	1.22	1.82

used in the analysis corresponds to 137.83 fb^{-1} in the mixed-flavor channel and 140.33 fb^{-1} in the four-muon channel at $\sqrt{s} = 13$ TeV. The blinded analysis gives a combined expected upper limit of 0.78×10^{-3} on $\mathcal{B}(H \to ZJ/\psi)$ in the longitudinal-polarization interpretation, as summarized in Tables 4.18 and 4.19.

The HL-LHC program is projected to deliver approximately 3000 fb^{-1} of integrated luminosity at $\sqrt{s} = 14$ TeV [83], representing an increase by a factor of approximately twenty relative to the Run 2 dataset used here. For a search limited by signal-region statistics, the expected upper limit scales approximately as the inverse square root of the integrated luminosity. The corresponding simple luminosity-scaling projection is given in Table 4.20.

Table 4.20: Luminosity scaling projection for the combined two-dimensional expected upper limit, assuming only $1/\sqrt{\mathcal{L}}$ scaling.

Scenario	Luminosity (fb^{-1})	Expected limit ($\times10^{-3}$)
Run 2 (this analysis)	140.33 / 137.83	0.78
HL-LHC projection	3000	≈ 0.17

Chapter 5

Conclusions

This dissertation uses CMS collision-rate information in three contexts. Dedicated luminosity detectors convert calibrated detector rates into the luminosity scale used by CMS. Reconstructed $J/\psi \to \mu^+\mu^-$ candidates provide an independent rate-based luminosity measurement. In the rare-Higgs search, the selected Run 2 luminosity normalizes the expected signal yield used in the upper-limit calculation.

The detector-calibration component described luminosity work within the CMS Beam Radiation, Instrumentation, and Luminosity program. The Pixel Luminosity Telescope, van der Meer and emittance-scan analysis, beam-related corrections, detector-response corrections, and vdM analysis software all contribute to the calibration chain that connects detector rates to an absolute luminosity scale. This work contributed the vdM scan-curve fits, the super-separation background measurements, the emittance-scan response monitoring, and the PLT linearity and efficiency corrections for the Run 3 calibration campaigns. In the 2022 campaign, this calibration chain gave a PLT visible cross section of 300.5 ± 0.2 μb and contributed to a total luminosity uncertainty of 1.4% for the 2022 data set. The correction categories discussed in this dissertation include background and bunch-current effects, length-scale and orbit corrections, beam-beam deflection, dynamic-β optics changes, detector

efficiency evolution, and nonlinearity. These studies support the stability and cross-detector consistency required for CMS luminosity measurements.

The second part developed a luminosity measurement based on $J/\psi \to \mu^+\mu^-$ counting in the `ParkingDoubleMuonLowMass` dataset. The analysis extracts the J/ψ signal yield in fixed 50-LS time bins, corrects the rate for pileup-dependent efficiency effects, calibrates the visible cross section with emittance scans, and applies long-term detector-response corrections based on the barrel pixel detector response and RAMSES region-to-region normalization. For the retained parking-dataset coverage, the resulting integrated luminosities are 28.53 ± 0.50 fb^{-1} in 2023 and 107.57 ± 2.30 fb^{-1} in 2024, consistent with the retained CMS reference luminosities on the same lumisection coverage.

The dissertation also presented a blinded upper-limit analysis in the $Z(\ell\ell)J/\psi(\mu\mu)$ final state. The expected signal yield is proportional to the integrated luminosity of the selected dataset, and the luminosity uncertainty enters the profile-likelihood model as a normalization uncertainty. The analysis combines the $Z \to e^+e^-$ and $Z \to \mu^+\mu^-$ channels using invariant-mass and machine-learning shape information. Under the longitudinal-polarization interpretation, the combined two-dimensional input gives an expected 95% confidence-level upper limit of 0.78×10^{-3} on the branching fraction $\mathcal{B}(H \to ZJ/\psi)$.

The percent-level agreement between the J/ψ luminosity measurement and the CMS reference across the 2023 and 2024 running periods indicates that the luminosity scale is consistent between methods with fundamentally different instrumental dependencies. Because the dominant uncertainties of the J/ψ method — tracking efficiency, muon reconstruction, trigger behavior, and resonance-fit stability — differ from those of the dedicated hardware luminometers, this agreement supports the use of the J/ψ method as a complementary cross-check of the CMS luminosity program.

As Run 3 data accumulates and the High-Luminosity LHC program approaches, the calibration and cross-check strategies developed in this dissertation remain important. The J/ψ method can be extended to larger parking datasets and to

additional running periods, preserving an independent luminosity cross-check based on reconstructed resonances rather than dedicated luminometer rates alone. For rare-Higgs searches in the same final state, improved luminosity control and multivariate discrimination can strengthen the expected sensitivity established in this blinded study.

The Run 3 PLT luminosity measurement was presented in the ICHEP 2022 proceedings [84], and broader CMS Run 3 luminosity measurements were presented at the 2023 and 2024 APS April Meetings [85, 86]. The CMS luminosity calibration was documented through the CMS LUM-22-001 approval process and the corresponding public CMS Physics Analysis Summary [28, 87]. The $J/\psi \to \mu^+\mu^-$ luminosity measurement is described in CMS Analysis Note AN-25-127 [54].

Bibliography

[1] Sheldon L Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579–588, 1961. ISSN 0029-5582. doi: 10.1016/0029-5582(61)90469-2. 1

[2] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19(21):1264–1266, November 1967. doi: 10.1103/PhysRevLett.19.1264. 1

[3] Peter W Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13(16):508–509, October 1964. doi: 10.1103/PhysRevLett.13.508. 2

[4] G. Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1): 1–29, September 2012. ISSN 03702693. doi: 10.1016/j.physletb.2012.08.020. 2

[5] S. Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, September 2012. ISSN 03702693. doi: 10.1016/j.physletb.2012.08.021. 2

[6] Georges Aad et al. Evidence for the spin-0 nature of the Higgs boson using ATLAS data. *Phys. Lett. B*, 726:120, 2013. doi: 10.1016/j.physletb.2013.08.026. 2

[7] ATLAS and CMS Collaborations. Combined measurement of the Higgs boson mass in pp collisions at $\sqrt{s}$=7 and 8 TeV with the ATLAS and CMS experiments. *Phys. Rev. Lett.*, 114(19):191803, 2015. doi: 10.1103/PhysRevLett.114.191803.

[8] A M Sirunyan et al. Combined measurements of the Higgs boson's couplings at s=13 TeV. *Eur. Phys. J. C*, 79:421, 2019. doi: 10.1140/epjc/s10052-019-6909-y. 3, 75

[9] CMS Collaboration. Observation of Higgs Boson Decay to Bottom Quarks. *Phys. Rev. Lett.*, 121(12):121801, September 2018. doi: 10.1103/PhysRevLett.121.121801.

[10] CMS Collaboration. Evidence for Higgs boson decay to a pair of muons. *Journal of High Energy Physics*, 2021(1):148, 2021. doi: 10.1007/JHEP01(2021)148. 3

[11] CMS Collaboration. A search for the standard model Higgs boson decaying to charm quarks. *Journal of High Energy Physics*, 2020(3):131, March 2020. doi: 10.1007/JHEP03(2020)131. 2

[12] CERN Document Server. Standard model of elementary particles. CERN Document Server, 2018. URL https://cds.cern.ch/record/2315477/files/2000px-Standard_Model_of_Elementary_Particles.png. 2

[13] Gordon Kane. *Modern Elementary Particle Physics*. Cambridge University Press, 2017. doi: 10.1017/9781316691434. 2

[14] John Ellis. Limits of the Standard Model, November 2002.

[15] C Grojean. Higgs Physics. *arXiv e-prints*, page arXiv:1708.00794, August 2017. 2

[16] CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, February 2008. ISSN 1748-0221. doi: 10.1088/1748-0221/3/08/S08004. 3, 12, 14, 16, 17, 81

[17] C. M. S. Collaboration. Development of the CMS detector for the CERN LHC Run 3. *J. Inst.*, 19(05):P05064, May 2024. ISSN 1748-0221. doi: 10.1088/1748-0221/19/05/P05064. 3, 12, 17, 18

[18] D. de Florian et al. Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector. Technical Report CERN-2017-002-M, CERN, 2017. 4, 75, 80

[19] M Bander and A Soni. Decays of Higgs scalars into vector mesons and photons. *Phys. Lett. B*, 82:411, 1979. ISSN 0370-2693. doi: 10.1016/0370-2693(79)90255-7. 4, 74

[20] W.-Y. Keung. Decay of the Higgs boson into heavy-quarkonium states. *Phys. Rev. D*, 27:2762, 1983. doi: 10.1103/PhysRevD.27.2762.

[21] G Bodwin et al. Higgs boson decays to quarkonia and the Hcc coupling. *Phys. Rev. D*, 88:53003, 2013. doi: 10.1103/PhysRevD.88.053003.

[22] Alexander L Kagan et al. Exclusive Window onto Higgs Yukawa Couplings. *Phys. Rev. Lett.*, 114:101802, 2015. doi: 10.1103/PhysRevLett.114.101802.

[23] Gino Isidori, Aneesh V. Manohar, and Michael Trott. Probing the nature of the Higgs-like Boson via h -> VF decays. *Physics Letters B*, 728:131–135, January 2014. ISSN 03702693. doi: 10.1016/j.physletb.2013.11.054.

[24] Dao-Neng Gao. A note on Higgs decays into Z boson and $J/\Psi(\Upsilon)$. *Physics Letters B*, 737:366–368, October 2014. ISSN 03702693. doi: 10.1016/j.physletb.2014.09.019.

[25] Stefan Alte, Matthias König, and Matthias Neubert. Exclusive Weak Radiative Higgs Decays in the Standard Model and Beyond. *J. High Energ. Phys.*, 2016 (12):37, December 2016. ISSN 1029-8479. doi: 10.1007/JHEP12(2016)037. 4, 76, 101

[26] A. Tumasyan et al. Search for Higgs boson decays into Z and J/ψ and for Higgs and Z boson decays into J/ψ or Y pairs in pp collisions at s=13 TeV. *Physics Letters B*, 842:137534, July 2023. ISSN 0370-2693. doi: 10.1016/J.PHYSLETB.2022.137534. 5, 74, 75, 76, 79

[27] A. M. Sirunyan et al. Precision luminosity measurement in proton–proton collisions at $$\sqrt{s} = 13\,\hbox {TeV}$$in 2015 and 2016 at CMS. *Eur. Phys. J. C*, 81(9):800, September 2021. ISSN 1434-6052. doi: 10.1140/epjc/s10052-021-09538-2. 7, 20, 23, 25, 26, 27, 28, 30, 31, 42, 71

[28] CMS Collaboration. Luminosity measurement in proton-proton collisions at 13.6 TeV in 2022 at CMS. Technical Report CMS-PAS-LUM-22-001, CMS Collaboration, 2023. 7, 10, 28, 30, 31, 32, 33, 34, 36, 38, 40, 105

[29] CMS Collaboration. Luminosity determination using Z boson production at the CMS experiment. *Eur. Phys. J. C*, 84(1):26, 2024. doi: 10.1140/epjc/s10052-023-12268-2. 7, 43, 97

[30] CMS Collaboration. CMS luminosity recommendations for run 2. CMS TWiki, 2024. URL https://twiki.cern.ch/twiki/bin/view/CMS/LumiRecommendationsRun2. 7, 79, 97

[31] CMS Collaboration. CMS luminosity recommendations for run 3. CMS TWiki, 2026. URL https://twiki.cern.ch/twiki/bin/view/CMS/LumiRecommendationsRun3. 7

[32] Olena Karacheban and on behalf of the CMS Collaboration. Performance of the BRIL luminometers at CMS in Run 2. *Proceedings of Science*, 364:194, November 2020. ISSN 18248039. doi: 10.22323/1.364.0194. 8

[33] A. Kornmayer. The CMS pixel luminosity telescope. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 824:304–306, July 2016. ISSN 01689002. doi: 10.1016/j.nima.2015.09.104. 8, 21

[34] Edy Ayala et al. The Pixel Luminosity Telescope: A detector for luminosity measurement at CMS using silicon pixel sensors. *Eur. Phys. J. C*, 83(7):673, 2023. ISSN 1434-6052. doi: 10.1140/epjc/s10052-023-11713-6. 8, 21

[35] S. van der Meer. Calibration of the Effective Beam Height in the ISR. Technical Report CERN-ISR-PO-68-31, CERN, June 1968. 8, 25

[36] Vladislav Balagura. Van der Meer scan luminosity measurement and beam–beam correction. *The European Physical Journal C*, 81(1), 2021. ISSN 1434-6044. doi: 10.1140/epjc/s10052-021-08837-y. 8, 28, 33

[37] Olena Karacheban and Peter Tsrunchev. Emittance scans for CMS luminosity calibration. *EPJ Web Conf.*, 201:04001, 2019. ISSN 2100-014X. doi: 10.1051/epjconf/201920104001. 8, 29, 55

[38] CERN. The CERN accelerator complex. CERN Document Server, Poster-2019-858, 2019. URL https://cds.cern.ch/record/2684277. 12

[39] CMS Collaboration. CMS detector drawings, 2012. URL https://twiki.cern.ch/twiki/bin/view/CMSPublic/SketchUpCMS. 13

[40] W. Adam et al. The CMS Phase-1 pixel detector upgrade. *J. Inst.*, 16(02): P02027, February 2021. ISSN 1748-0221. doi: 10.1088/1748-0221/16/02/P02027. 15, 46

[41] CMS Collaboration. Particle-flow reconstruction and global event description with the CMS detector. *JINST*, 12(10):P10003, 2017. doi: 10.1088/1748-0221/12/10/P10003. 16, 18, 81

[42] A. M. Sirunyan et al. Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s}$=13 TeV. *JINST*, 13(06): P06015, 2018. doi: 10.1088/1748-0221/13/06/P06015. 17, 46, 47, 81, 98

[43] W Adam et al. The CMS high level trigger. *Eur. Phys. J. C*, 46:605–667, 2006. doi: 10.1140/epjc/s2006-02495-8. 18, 47

[44] V Khachatryan et al. The CMS trigger system. *JINST*, 12:P01020, 2017. doi: 10.1088/1748-0221/12/01/P01020. 47

[45] A. Tapper and Darin Acosta. CMS Technical Design Report for the Level-1 Trigger Upgrade. Technical Report CERN-LHCC-2013-011, CERN, 2013. URL https://cds.cern.ch/record/1556311. 18

[46] David Barney. CMS detector slice. CMS public detector drawing, 2016. URL https://cms-docdb.cern.ch/cgi-bin/PublicDocDB/ShowDocument?docid=4172. 19

[47] CMS Collaboration. CMS luminosity measurements for the 2016 data-taking period. Technical Report CMS-PAS-LUM-17-001, CERN, 2017. URL https://cds.cern.ch/record/2257069. 20, 71

[48] CMS Collaboration. CMS luminosity measurements for the 2017 data-taking period at $\sqrt{s}$=13 TeV. Technical Report CMS-PAS-LUM-17-004, CERN, 2018. URL https://cds.cern.ch/record/2621960.

[49] Joanna Wanczyk et al. CMS luminosity measurement for the 2022 proton-proton collisions data-taking period at $\sqrt{s}$=13.6 TeV. Technical Report CMS-PAS-LUM-22-001, CERN, 2023. 20, 71

[50] Anne Dabrowski. Upgrade of the CMS Instrumentation for luminosity and machine induced background measurements. *Nucl. Part. Phys. Proc.*, 273–275: 1147, 2016. doi: 10.1016/j.nuclphysbps.2015.09.180. 21

[51] Himal Acharya et al. Detector Performance and Operational Experience of the Pixel Luminosity Telescope. Technical Report CMS-DN-19-028, CMS Collaboration, 2019. 21, 22, 24

[52] Michal Krupa and Marek Gasior. A New Wall Current Transformer for Accurate Beam Intensity Measurements in the Large Hadron Collider. *Energies 2023, Vol. 16, Page 7442*, 16(21):7442, November 2023. ISSN 1996-1073. doi: 10.3390/EN16217442. 31, 56

[53] A. Babaev et al. Impact of beam–beam effects on absolute luminosity calibrations at the CERN Large Hadron Collider. *Eur.Phys.J.C*, 84(1):17, January 2024. ISSN 14346052. doi: 10.1140/EPJC/S10052-023-12192-5. 33, 34

[54] Nimmitha Karunarathna and Stefan Spanier. Luminosity determination with J/ψ counting. CMS Analysis Note AN-2025/127, CERN, 2025. 41, 105

[55] A M Sirunyan et al. Performance of the CMS muon trigger system in proton-proton collisions at (s) = 13. *Journal of Instrumentation*, 16(07):P07001, July 2021. doi: 10.1088/1748-0221/16/07/p07001. 46, 47, 81

[56] L. Scibile et al. The LHC radiation monitoring system for the environment and safety: From design to operation. In *Proceedings of EPAC08*, page MOPD030, Genoa, Italy, 2008. JACoW. 46

[57] R. Frühwirth. Application of Kalman filtering to track and vertex fitting. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 262(2-3):444–450, December 1987. ISSN 01689002. doi: 10.1016/0168-9002(87)90887-4. 47

[58] W Verkerke and D P Kirkby. The RooFit Toolkit for Data Modeling. In *Proceedings, 13th Int'l Conf. on Computing in High-Energy and Nuclear Physics (CHEP 2003)*, La Jolla, California, US, March 2003. 50, 91

[59] CMS BRIL. BRIL / VdMFramework. GitLab repository, 2026. URL https://gitlab.cern.ch/bril/VdMFramework. 55

[60] R. L. Workman et al. Review of particle physics. *Progress of Theoretical and Experimental Physics*, 2022(8):083C01, 2022. doi: 10.1093/ptep/ptac097. 75, 80, 98

[61] V Kartvelishvili, A V Luchinsky, and A A Novoselov. Double vector quarkonia production in exclusive Higgs boson decays. *Phys. Rev. D*, 79:114015, 2009. doi: 10.1103/PhysRevD.79.114015. 75

[62] Qi-Li Liao and Jun Jiang. Excited heavy quarkonium production in Higgs boson decays. *Phys. Rev. D*, 100(5):53002, September 2019. doi: 10.1103/PhysRevD.100.053002. 75

[63] Morad Aaboud et al. Searches for exclusive Higgs and Z boson decays into J/$\psi\gamma$, ψ(2S)γ, and Υ(nS)γ at s=13 TeV with the ATLAS detector. *Phys. Lett. B*, 786: 134, 2018. doi: 10.1016/j.physletb.2018.09.024. 75

[64] Albert M Sirunyan et al. Search for rare decays of Z and Higgs bosons to J/ψ and a photon in proton-proton collisions at s = 13 TeV. *Eur. Phys. J. C*, 79(2): 94, 2019. doi: 10.1140/epjc/s10052-019-6562-5.

[65] CMS Collaboration. Search for Higgs and Z boson decays to J/ψ or Υ pairs in the four-muon final state in proton-proton collisions at s = 13 TeV. *Physics Letters B*, 797:134811, 2019. doi: 10.1016/j.physletb.2019.134811.

[66] CMS Collaboration. Search for decays of the 125 GeV Higgs boson into a Z boson and a ρ or ϕ meson. *Journal of High Energy Physics*, 2020(11), July 2020. doi: 10.1007/JHEP11(2020)039.

[67] ATLAS Collaboration. Search for exclusive Higgs and Z boson decays to $\omega\gamma$ and Higgs boson decays to K$*\gamma$ with the ATLAS detector. *Physics Letters B*, 847: 138292, January 2023. doi: 10.1016/j.physletb.2023.138292. 75

[68] Seong Youl Choi. Z-boson polarization as a model-discrimination analyzer. *Phys. Rev. D*, 98(11):115037, December 2018. doi: 10.1103/PhysRevD.98.115037. 76

[69] CMS Collaboration. PdmV run 2 legacy analysis recommendations. CMS TWiki, 2024. URL https://twiki.cern.ch/twiki/bin/view/CMS/PdmVRun2LegacyAnalysis. 79

[70] Torbjörn Sjöstrand et al. An introduction to 8.2. *Comput. Phys. Commun.*, 191: 159, 2015. doi: 10.1016/j.cpc.2015.01.024. 79

[71] S. Agostinelli et al. Geant4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506 (3):250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8. 79

[72] Thomas Junk. Confidence level computation for combining searches with small statistics. *Nucl. Instrum. Meth. A*, 434:435, 1999. doi: 10.1016/S0168-9002(99)00498-2. 80

[73] Alexander L Read. Presentation of search results: The CL_s technique. *J. Phys. G*, 28(10):2693, 2002. doi: 10.1088/0954-3899/28/10/313.

[74] ATLAS and CMS Collaborations and the LHC Higgs Combination Group. Procedure for the LHC Higgs boson search combination in summer 2011. Technical Report CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11, CERN, 2011. URL https://cds.cern.ch/record/1379837. 80

[75] Glen Cowan et al. Asymptotic formulae for likelihood-based tests of new physics. *European Physical Journal C*, 71:1554, February 2011. doi: 10.1140/epjc/s10052-011-1554-0. 80

[76] CMS Collaboration. HiggsAnalysis-CombinedLimit documentation. CMS Combine documentation, 2026. URL https://cms-analysis.github.io/HiggsAnalysis-CombinedLimit/latest/. 80

[77] Wolfgang Adam et al. Reconstruction of Electrons with the Gaussian-Sum Filter in the CMS Tracker at the LHC. Technical Report CMS-NOTE-2005-001, CERN, Geneva, 2005. URL https://cds.cern.ch/record/815410. 81

[78] Jonas Rembser. CMS Electron and Photon Performance at 13 TeV. *Journal of Physics: Conference Series*, 1162:12008, 2019. doi: 10.1088/1742-6596/1162/1/012008.

[79] CMS Collaboration. Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at s = 8 TeV. *Journal of*

Instrumentation, 10(6):P06005, June 2015. doi: 10.1088/1748-0221/10/06/P06005. 81, 98

[80] CMS Collaboration. EGamma UL 2016 to 2018. CMS TWiki, 2024. URL https://twiki.cern.ch/twiki/bin/view/CMS/EgammaUL2016To2018. 82

[81] Leo Breiman. Random forests. *Machine Learning*, 45:5–32, 2001. doi: 10.1023/A:1010933404324. 85

[82] A. Hoecker et al. TMVA: Toolkit for multivariate data analysis. *PoS*, ACAT:040, 2007. doi: 10.22323/1.050.0040. 86

[83] M. Cepeda et al. Physics at the HL-LHC, and perspectives for the HE-LHC. *CERN Yellow Reports: Monographs*, 7:221, 2019. doi: 10.23731/CYRM-2019-007.221. 102

[84] Nimmitha Karunarathna. Run 3 luminosity measurements with the Pixel Luminosity Telescope. In *Proceedings of Science*, volume ICHEP2022 of *Proceedings of Science*, page 936, 2022. doi: 10.22323/1.414.0936. 105

[85] Warusapperuma Don Nimmitha Karunarathna. Online luminosity measurements at CMS during LHC run-3. Bulletin of the American Physical Society, 2023 APS April Meeting, Session WW01.00001, 2023. URL https://meetings.aps.org/Meeting/APR23/Session/WW01.1. 105

[86] Warusapperuma Don Nimmitha Karunarathna. Luminosity measurements with the CMS detector. Bulletin of the American Physical Society, 2024 APS April Meeting, Session GG03.00007, 2024. URL https://meetings.aps.org/Meeting/APR24/Session/GG03.7. 105

[87] CMS BRIL. Approval of LUM-22-001: Offline luminosity measurement for the 2022 pp collisions at 13.6 TeV data set. CMS luminosity approval meeting, 2023. URL https://indico.cern.ch/event/1253089/. 105

Vita

Nimmitha Karunarathna received his Bachelor of Science degree in Physics from the University of Colombo, Sri Lanka in 2017. He began his doctoral studies in Physics at the University of Tennessee, Knoxville.

His research focuses on experimental high-energy physics within the CMS experiment at CERN. His work includes luminosity calibration studies and detector performance evaluation. Throughout his doctoral studies, he has participated in collaborative research efforts, contributing to analyses related to proton–proton collision data from the Large Hadron Collider.