

PhD Dissertation

**PROBING NEW PHYSICS:
SEARCH FOR SUPERSYMMETRY WITH HIGGS PARTICLES
AND HIGH-PRECISION LUMINOSITY DETERMINATION AT
THE CMS EXPERIMENT**

Péter Major

EÖTVÖS LORÁND UNIVERSITY
FACULTY OF SCIENCE
INSTITUTE OF PHYSICS AND ASTRONOMY

DOCTORAL SCHOOL OF PHYSICS
PARTICLE AND NUCLEAR PHYSICS PROGRAM

Supervisor:

Dr. Gabriella Pásztor

Assistant Professor, PhD, Habil
Institute of Physics and Astronomy
Department of Atomic Physics

Head of doctoral school:

Dr. Palla Gergely

Professor, DSc

Head of doctoral program:

Dr. Zoltán Trócsányi

Professor, DSc, Member of
the Hungarian Academy of Sciences



Budapest, 2024

DOI 10.15476/ELTE.2024.190

Acknowledgments

Not only completing but merely surviving a PhD can be a challenging task. This is true even if the journey is rewarding and holds a lot of joy. Completion is not only to the credit of the candidate, at least not in my case, but also to the support of their colleagues, friends, and family. I wish I had the time and talent to aptly express my gratitude and appreciation to all of you, but the deadline is near.

I am immensely grateful to Gabriella Pásztor, my supervisor, for sharing her in-depth knowledge, guiding my way from a BSc student to an almost PhD with unparalleled patience, the friendly atmosphere she created as a work environment, and giving me a large freedom to develop new techniques for luminosity, supporting my work with her expert insight.

I am thankful to my immediate colleagues at ELTE and CERN who have welcomed me into their circles and involved me in the lengthy but fruitful collaborative research that this thesis is based on.

I would like to thank some of my favorite people in this world for their friendship and for putting up with my peculiarities and periods of unavailability. Without hope for completeness, Zsolt Szabó, Péter Lővei, and my dearest brother, Gábor, are definitely on this list. Thank you for existing.

Lastly, but most importantly, my warmest gratitude goes out to my parents, Csilla and Antal, for their love and unconditional support that enabled me to achieve everything I have.

Abstract

The standard model (SM) of particle physics, while extremely successful in describing the phenomena observed in high-energy particle interactions, is under extreme scrutiny as it leaves many unexplained observations, such as the presence of dark matter in the universe. The Large Hadron Collider (LHC) with its record energy and high statistics data is an ideal place to look for phenomena that lie beyond the SM. Using the Compact Muon Solenoid (CMS) detector, I contributed to the efforts to test the validity of the SM and attempted to find direct evidence for the production of supersymmetric (SUSY) partner particles or constrain the parameter space of the models that predict them.

High-precision cross-section measurements provide stringent constraints on the parameters of the SM. An essential ingredient to these studies is the accurate determination of luminosity, a measure of the analyzed data volume. I played a prominent role in the CMS luminosity calibration effort for the Run2 and early Run3 data. In particular, I have developed two new methods to derive corrections for the so-called non-factorization bias which is related to the bunch particle density distributions in the accelerator. This effect constitutes one of the dominant sources of uncertainty for the absolute calibration of the luminosity instrumentation. I also introduced various other new ideas and tools for the van der Meer calibration method as well as for monitoring the performance of the luminometers. These had a significant contribution to achieving a twofold improvement in the precision achieved in Run2, reaching percent level uncertainties for high energy proton - proton, as well as lead ion collisions.

The discovery of the Higgs boson in 2012 at LHC has opened the way to use this unique scalar particle as a probe of new physics, which I exploit in my attempt to find direct evidence for the manifestation of low-energy supersymmetry in the framework of theories with gauge mediated supersymmetry breaking. I perform a search for rare events containing a Higgs boson and a photon that exhibit significant missing transverse momentum due to undetected gravitinos which serve as dark matter candidates. I developed a semi-data-driven estimation of the dominant background in this SUSY search both via a classical method and by implementing it into the complex statistical model I used for hypothesis testing to gauge the sensitivity of the analysis in terms of upper limits on the production cross sections and constraints on the masses of SUSY particles.

Kivonat

Noha a részecskefizika standard modellje (SM) hihetetlen sikerrel írja le a nagyenergiás részecskék kölcsönhatásait, mégis komoly bírálatok érik, amiért megannyi kérdést, mint például a sötét anyag létezésének rejtélyét, megválaszolatlanul hagyja. A rekord energián működő és nagy adatmennyiséget szolgáltató Nagy Hadronütköztető (LHC) ideális eszköz a standard modellen túli fizikai jelenségek kutatására. A Kompakt Műion Szolenoid (CMS) detektor segítségével hozzájárultam a SM érvényességének ellenőrzésére tett erőfeszítésekhez, illetve megkíséreltem közvetlen bizonyítékot találni szuperszimmetrikus (SUSY) részecskék létezésére, vagy megszorítani az ezeket jóvendőlő elméletek paraméterterét.

A nagy pontosságú hatáskeresztmetszet-mérések szigorú megkövetéseket rónak a SM paramétereire. Az egyik alapvető feltétele ezen méréseknek a feldolgozott adatmennyiség pontos ismerete, amelyet a luminozítás számszerűsít. Munkám döntő szerepet játszott a CMS Run2-es és korai Run3-as luminozítás kalibrációiban. Két módszert is kifejlesztettem az úgynevezett nem-faktorizálhatósági torzítás mérésére. Ez az effektus a gyorsítóban keringő protoncsomagok alakjával kapcsolatos, és a luminozításmérés egyik vezető szisztematikus hibája. További újításokat vezettem be a kalibrációs folyamat kapcsán, munkám érintette a luminométerek stabilitását és linearitását, amellyel jelentős mértékben hozzájárultam a Run2-es luminozításbizonytalanság kevesebb, mint felére csökkentéséhez, amivel végül mind a proton-proton, mind az ólom-ólom ütközések esetén százalékos nagyságrendű hibákat értünk el.

A Higgs-bozon 2012-es felfedezésével lehetőség nyílt ezen különleges skalár részecskén keresztül kutatni az új fizikai jelenségeket. Ezt kihasználva kísértem meg közvetlen bizonyítékot találni alacsonyenergiás-szuperszimmetria (SUSY) nyomaira két olyan egyszerűsített modellben, amelyek a mértékkölcsönhatás-közvetített SUSY-sértés mechanizmusát (GMSB) használják. Olyan ritka események után kutattam, amelyeket a Higgs-bozon mellett egy nagy transzverzális lendületű foton és nagy hiányzó transzverzális lendület jellemez, amit a SUSY sötétanyag-jelölt gravitínók nem észlelt távozása okoz. Kidolgoztam egy részben adatvezérelt becslést a mérésben megjelenő domináns SM háttér hozzájárulására mind a klasszikus hibaterjedés módszereit használva, mind a becslést egy statisztikus modellbe ágyazva. A statisztikus modell segítségével hipotézisvizsgálatnak vettem alá az egyes SUSY modelleket a kutatás várható érzékenységének meghatározására, amit a jel keletkezési hatáskeresztmetszet felső korlátjával számszerűsíték a model paramétereire függvényében. Ezek alapján a SUSY részecskék megengedett tömegetartománya megszorítható.

Contents

1	Modern physics and new physics:	
	A theoretical overview	1
1.1	The standard model of particle physics	1
1.1.1	Matter fields	3
1.1.2	Force carriers	5
1.1.3	The Higgs boson	5
1.1.4	Formalizing the standard model	6
1.1.5	The Higgs field	8
1.2	Limitations of the standard model	12
1.2.1	The hierarchy problem	13
1.2.2	Dark matter	16
1.3	Supersymmetry as a possible solution	17
1.3.1	The spectrum of the MSSM	19
1.3.2	Final states for colliders in the MSSM	26
1.4	Alternative theories	28
1.4.1	The anthropic approach	28
1.4.2	Compact extra dimensions	28
1.4.3	Conformal solution	31
1.4.4	Technicolor	31
1.4.5	Modified Newtonian dynamics	32
2	Where the magic happens	33
2.1	The Large Hadron Collider at CERN	33
2.1.1	The CERN accelerator complex	35
2.1.2	The LHC collider	38
2.2	The CMS detector	40
2.2.1	The solenoid magnet	45
2.2.2	The beryllium beam pipe	46
2.2.3	The silicon tracker	46
2.2.4	The electromagnetic calorimeter	48
2.2.5	The hadron calorimeter	49

2.2.6	The muon system	50
2.2.7	The trigger system	51
2.2.8	Object reconstruction	53
3	Precision luminosity at the CMS experiment	57
3.1	Cross sections and luminosity	58
3.2	Luminometers in CMS	62
3.3	Data quality monitoring for PLT	68
3.4	The van der Meer method	74
3.5	Systematics affecting the vdM calibration	78
3.6	Emittance change	81
3.7	Non-factorization bias	84
3.8	The 2D rate fit method	85
3.8.1	Data consistency and rate matching	86
3.8.2	Fit functions	88
3.8.3	Deriving the correction	92
3.8.4	Uncertainties	93
3.9	The luminous region method	95
3.9.1	Input data	96
3.9.2	Fit models	100
3.9.3	From observables to bunch proton densities	104
3.9.4	Sequential fitting	106
3.9.5	Fit results	108
3.9.6	Deriving the vdM correction factors	108
3.9.7	Closure test	108
3.9.8	Results and uncertainties	113
3.10	Normalization results	121
3.11	Linearity and stability	122
3.11.1	Emittance scans	122
3.11.2	The relative non-linearity of PLT	123
3.11.3	The mean luminosity approach	125
3.12	Summary	128
4	In pursuit of new physics	131
4.1	Analysis strategy	131
4.2	Data samples	132
4.3	Signals	133
4.4	Object definitions	136
4.5	Event selection and binning	139

4.5.1	The signal and validation regions	141
4.6	Background processes	142
4.6.1	The $Z(\nu\nu)$ background	144
4.6.2	The misidentified electron background	144
4.6.3	The multijet background	145
4.7	The lost lepton background	147
4.7.1	Selecting a hadronic tau definition	149
4.7.2	Estimating the LL background	150
4.7.3	Validation of the background estimates	155
4.8	Statistical interpretation	155
4.9	Prospects	164
Conclusion		165
Appendix		167
A Normalization and CapSigma parametrization		167
A.1	Double Gaussian with shared means	167
A.2	Super Gaussian	168
A.2.1	Normalization in N dimensions	168
A.2.2	Normalization for product constructions	169
A.3	Polynomial-modulated Gaussian	169
A.3.1	Radial polynomial-modulated Gaussian normalization	170
A.3.2	2D polynomial-modulated Gaussian normalization	171
A.4	q-Gaussian	173
A.4.1	2D q-Gaussian normalization	175
B Full derivation of the luminous region parameters		177
B.1	The zero crossing angle case	177
B.1.1	The dynamic β^* correction	184
B.2	The non-zero crossing angle case	184
C Simplified derivation of the luminous region parameters		187
D Derivation of the emittance scan based linearity modification factor		193
D.1	Effect on the beam overlap width	197
Bibliography		199

Chapter 1

Modern physics and new physics: A theoretical overview

To provide a better understanding of the particle physics landscape and the motivation that drives the research for exotic phenomena such as particles predicted by supersymmetric theories, this section introduces the currently most widely accepted framework of high energy physics, the standard model (SM) [1–4].

Despite the broad applicability of this model and its consequent enormous success, it is well known, that the standard model cannot be the fundamental theory of the micro world. For this reason, a race for deeper insight into the fabric of reality is fueling the contemporary research scene.

This section also introduces some of the shortcomings of the Standard Model that motivate the search for its successor, as well as a handful of extensions that all contend for this title. The chapter puts an emphasis on supersymmetry (SUSY) with its phenomenological consequences in focus.

1.1 The standard model of particle physics

The Standard Model of particle physics (SM) is a spontaneously broken

$$SU(3)_C \times SU(2)_L \times U(1)_Y$$

gauge theory which describes all currently known elementary particles and three out of the four elementary forces, these being the strong, the weak and the

electromagnetic interactions. The gravitational interaction is not incorporated into the Standard Model.

The model was set up in the 20th century as a collective effort of many researchers. It was brought to its final form in the '70s after experimental evidence was found to support the quark hypothesis of Gell-Mann and Zweig [5] between 1967 and 1973¹. Concurrently, in '71 a large open question regarding the mathematical consistency of the theory was addressed² in [7]. The discovery of weak neutral currents [8, 9] in 1973 at CERN, followed by the direct observation of the W^\pm [10, 11] and Z bosons [12, 13] in 1983³ increased the trust in the SM which became a widely accepted theory. Ever since the predictions of SM (augmented with the neutrino masses) show an exceptionally good agreement with the experimental data of the last 50 years.

In the followings, we familiarize ourselves with the minimum vocabulary of the Standard Model. First, the particle content of the SM is discussed. The standard Model is then formalized while constructing its Lagrangian. In the process, the Higgs field is introduced at the point when we are faced with the problem of the disallowed explicit mass terms.

¹Using the Stanford Linear Accelerator, at SLAC, in deep inelastic electron-proton scattering (DIS) experiments [6].

²The proof of renormalizability earned Gerardus 't Hooft and Martinus J.G. Veltman the Nobel prize of 1999.

³In proton-antiproton collisions on the Super Proton Synchrotron using two movable detectors, the UA1 and UA2. For this breakthrough, a Nobel Prize was awarded in 1984.

According to our findings, the world consists of some non-further-dividable, so-called fundamental or elementary particles. These are represented as excitations of either bosonic or fermionic fields in quantum field theory (QFT). In the SM model, all fermions are of spin $1/2$ while bosons may have a 1 or a 0 spin quantum number.

1.1.1 Matter fields

In the SM framework, the matter in our universe is built up exclusively by particles of the fermionic sector. These particles are divided into two categories depending on whether they take part in the strong interaction or not. The ones that do are called quarks, the rest are called leptons. To engage with the strong force, quarks carry a so-called color charge and are only observable as color-neutral (singlet) bound states, called hadrons, as a result of a phenomenon called color confinement the reason of which is still an active field of research⁴. There are in total six different kinds of quarks differentiated by their charges and mass, the latter of which can be traced back to different coupling strengths in the SM. These subcategories are called flavors.

Leptons can be subdivided into two large groups depending on their electric charge. Neutral leptons are called neutrinos and are very special in the sense that they only interact via the weak nuclear force and hence are extremely elusive, with a mean free path length of approximately a light year in solid lead.

The fermionic sector can alternatively be divided into three families (or generations) each containing two quark flavors, a charged lepton, and a neutralino. Within each generation the strong and electromagnetic interactions act identically, however, differences in the electroweak interaction give rise to a split in the mass spectrum. The structure just described is best understood by examining Figure 1.1.

⁴At a simplistic level, confinement is due to the strength-distance characteristics of the strong force's potential which describe the interaction between color charged particles. The resulting force can be well approximated with a constant for any distance and hence if the quarks' separation is adequately great it is more favorable for a new quark-antiquark pair to form from this accumulated energy and, as a result, we get two separate pairs of quarks at a distance. The phenomenon could be described with an analogy of the non-separability of the ends of a rubber band. If one tries then the band splits into two, resulting in four ends. The distance characteristics are explained by the self-interaction of the particles mediating the strong force. As this phenomenon belongs to the non-perturbative regime of QCD, it is not very well understood, and creating rigorous proof for it is one of the Millennium Prize Problems.

Each of the fermions also possesses an antiparticle pair that has all of its charges flipped with respect to the original particle. While for electrically charged particles the difference between the particle and antiparticle partner is clear, it is not currently known whether neutrinos could be their own antiparticles.

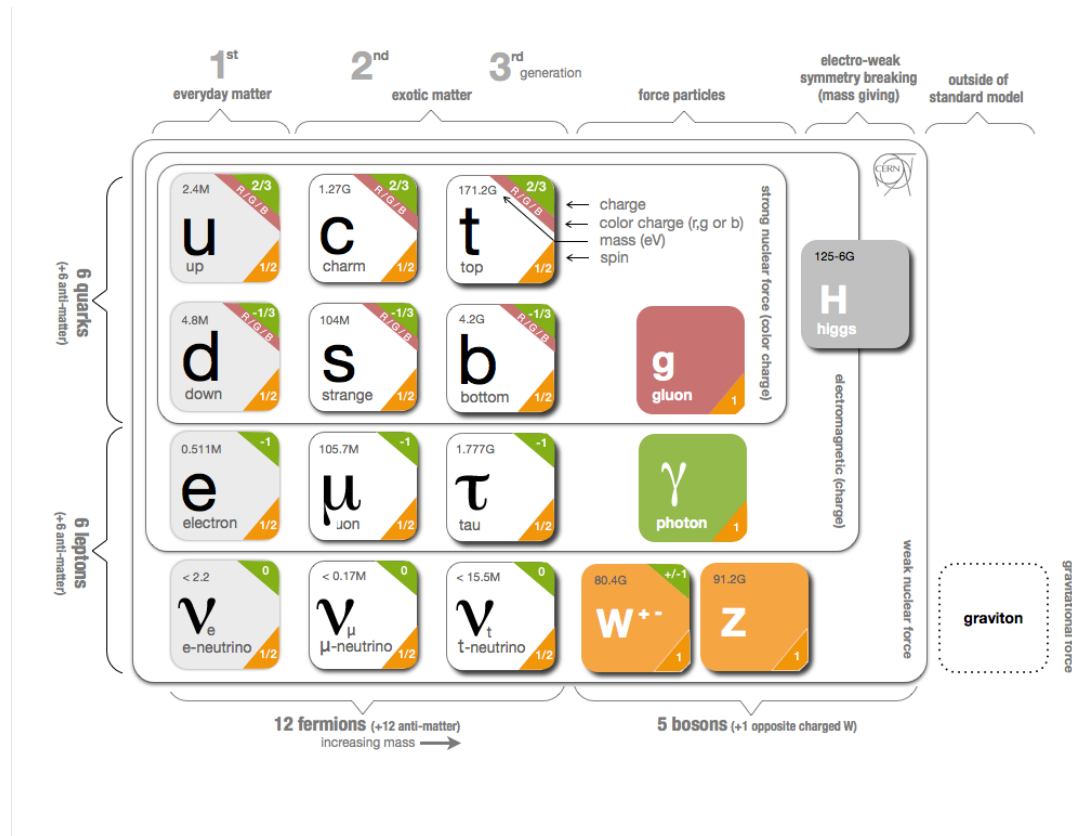


Figure 1.1: The standard model with three generations of matter particles (first three columns), the force carrier bosons (in colored boxes), and the Higgs boson (in a dark grey box). The graviton, the hypothetical particle mediating the gravitational force, is not described by the model. Source⁵.

All stable matter in the universe is built up of particles of the first generation, as particles of the second and third generations are heavier and rather unstable⁶. For example, atoms consist of a nucleus and electrons which are leptons from the first generation. The nucleus holds protons and neutrons which are types of hadrons, and contain up and down valence quarks that are also from the first

⁵<https://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>

⁶The neutrinos are of course exceptions from the previous statement as they are all massless in the SM. Moreover, in the SM, neutrinos cannot further decay due to kinematic constraints and the conservation of the lepton number in each generation. We will see, however, that experimental data contradict the hypothesis of massless neutrinos.

generation⁷.

1.1.2 Force carriers

Matter particles interact with each other through four fundamental forces of different strength-distance characteristics. These are gravity, and the strong, weak, and electromagnetic forces. Gravity, although having an infinite range⁸, is much weaker than the other three forces on the distance scale which is relevant in describing the micro-world that is accessible to us. Moreover, it is not incorporated into the SM due to technical difficulties emerging in the theories attempting to unify quantum field theory and general relativity.

The other three interactions come to life in the SM through the exchange of spin 1 force carrier bosons that transfer energy, momentum and charges between two interacting particles. These are the massless gluon (g) and photon (γ), and the massive W^\pm and Z^0 bosons.

Not every particle is capable of participating in all interactions. This is well portrayed by Figure 1.1. Strong interaction only affects particles carrying color charge, which means that its force carriers, the gluons, interact with the quarks of all the matter fields as well as themselves as gluons carry a color-anticolor charge pair. This self-interaction of massless force fields is unique to the gluons and is a source of many technical difficulties in computation.

Particles with an electric charge take part in the electromagnetic interaction, mediated by the photon. This means the W^\pm bosons and all fermions except for the neutrinos. Finally, all known matter particles take part in the weak interaction which is mediated by the massive vector bosons W^\pm and Z^0 .

1.1.3 The Higgs boson

The final puzzle piece in the standard model is the Higgs boson, an excitation of the Higgs field, which is special in a sense that it assumes a non-zero value everywhere, not only when a particle is present. This field interacts with all

⁷Of course, like all hadrons, protons, and neutrons contain not only their on-shell (valence) quarks but also a sea of virtual (off-shell) quark-antiquark pairs (a part of which are heavy quarks) and gluons. The composition of the sea is described by the parton distribution functions (PDF).

⁸The range of an interaction is the distance over which its influence is decreased e fold. Since gravity follows an inverse square law, it eventually outperforms exponential trends.

massive particles in the SM and is responsible for their mass through the so-called Brout–Englert–Higgs mechanism to be discussed in the next section.

1.1.4 Formalizing the standard model

The mathematical framework of the SM is relativistic quantum field theory. More specifically it is a non-abelian gauge theory in which the fermionic fields are predefined and the bosonic fields emerge as the gauge fields. The SM was, however, preceded by two theories that describe the strong and the electroweak interactions separately. These are quantum chromodynamics (QCD) and the Salam-Weinberg or electroweak (EW) model, respectively, the fusion of which led to the standard model.

QCD is an $SU(3)_C$ (C standing for chromo or color) gauge theory. The only fermions of QCD, the quarks are represented as Dirac fields. In this theory, it is postulated that each quark flavor exists in three different copies that span a \mathbb{C}^3 space and they transform under the $SU(3)$ group (the quarks are in the fundamental representation of $SU(3)$). The vectors of an orthonormal base on this space are labeled by the colors red, green, and blue (or 1,2,3):

$$\Psi = (\Psi_1, \Psi_2, \Psi_3). \quad (1.1)$$

Based on this, quarks have a color, a flavor, and a Dirac index. The latter two we will omit in our equations.

From here one can build the Lagrangian using the gauge principle. Let T^a ($a = 1, \dots, 8$) be a set of infinitesimal generators of $SU(3)$ which are most often chosen to be the Gell-Mann matrices. Let f^{abc} be the structure constant of the $\mathfrak{su}(3)$ Lie-algebra:

$$[T^a, T^b] = if^{abc}T^c. \quad (1.2)$$

Then we introduce the gauge field A_μ^a ($a = 1, \dots, 8$) which represents the gluons, the covariant derivative, and the gluon field strength tensor:

$$D_\mu = \partial_\mu - igT^a A_\mu^a, \quad (1.3)$$

$$G_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc}A_\mu^b A_\nu^c, \quad (1.4)$$

where g is the strong interaction's coupling constant.

Finally, the Lagrangian in the Einstein notation is

$$\mathcal{L}_{\text{QCD}} = \bar{\Psi}(i\gamma_\mu D^\mu)\Psi - \frac{1}{4}G_{\mu\nu}^a G^{\mu\nu a}, \quad (1.5)$$

where the first part is a Dirac term for all the quarks, with the interaction terms hidden in the covariant derivative, while the second term gives the dynamics of the gluons (Yang-Mills action). $\gamma_\mu \in \mathbb{C}^{4 \times 4}$ for $\mu \in \{0, 1, 2, 3\}$ are the Gamma matrices acting on the Dirac indices of the spinors. With these choices \mathcal{L}_{QCD} is invariant under gauge transformations.

Deriving the Salam-Weinberg Lagrangian is somewhat more complicated due to its hidden symmetries⁹. The minimum required symmetry group capable of accommodating the photon and the weak force bosons is $G = \text{SU}(2) \times \text{U}(1)$, which is also the scenario realized by nature.

The associated gauge bosons are usually labeled W_μ^a ($a = 1, 2, 3$) and B_μ . Similarly to the previous definitions:

$$D_\mu = \partial_\mu - ig_1 \frac{\tau^a}{2} W_\mu^a - ig_2 \frac{Y}{2} B_\mu, \quad (1.6)$$

where g_1 and g_2 are the coupling constants associated to $\text{SU}(2)$ and $\text{U}(1)$ respectively, $\tau^a/2$ and $Y/2$ are infinitesimal generators of these groups in the appropriate representation¹⁰. The field strength tensors are:

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + g_1 \frac{\epsilon^{abc}}{2} W_\mu^b W_\nu^c, \quad (1.7)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu. \quad (1.8)$$

Several experiments have demonstrated the parity symmetry (P-symmetry) violating nature of the weak interactions. Based on this, it is deduced, that only particles with a left-handed chirality couple to the W bosons of the spontaneously unbroken regime. The chiral states can be obtained from the spinors by means of projection

$$\Psi_L = \frac{1 - \gamma_5}{2} \Psi \quad \text{and} \quad \Psi_R = \frac{1 + \gamma_5}{2} \Psi, \quad (1.9)$$

where $\gamma_5 = i\gamma_0\gamma_1\gamma_2\gamma_3$. This suggests that the left-handed fields are in a doublet $\text{SU}(2)$ representation, while the right-handed fields are in the singlet one. For this reason, this group is often denoted as $\text{SU}(2)_L$. Using the notations of Figure 1.1 to reference the Dirac fields of each fermion, we get that the doublets are

$$Q_1 = (u_L, d_L), \quad Q_2 = (c_L, s_L), \quad Q_3 = (t_L, b_L),$$

⁹A symmetry of the theory which no ground state solution respects (also sometimes called a spontaneously broken symmetry). A good example is a ferromagnetic material below the Currie temperature. The whole material and thus the equations are symmetric to rotations but the magnetization takes a certain direction below the Currie temperature breaking this symmetry.

¹⁰For example for fields in the $\text{SU}(2)$ doublet representation τ^a is the Pauli matrix σ^a , but for singlets $\tau^a = 0$.

$$L_1 = (\nu_e, e_L), \quad L_2 = (\nu_\mu, \mu_L), \quad L_3 = (\nu_\tau, \tau_L),$$

while the singlets are:

$$u_R, d_R, c_R, s_R, t_R, b_R, \quad \text{and} \quad e_R, \mu_R, \tau_R.$$

The lack of the right-handed neutrino in this list accounts for the fact, that as it is now not interacting via any of the three forces with the rest of the SM, it is completely decoupled.

The charge associated with the U(1) symmetry is called weak hypercharge and its generator value is found to be

$$Y[Q] = \frac{1}{3}, \quad Y[L] = -1$$

and

$$Y[e_R] = -2, \quad Y[u_R] = \frac{4}{3}, \quad Y[d_R] = -\frac{2}{3},$$

where in the last line e_R stands for all right-handed charged leptons just like u_R stands for all up-type right-handed quarks, d_R for all down-type right-handed quarks and Q and L represent Q_i and L_i ($i = 1, 2, 3$) in the previous line. We will use this notation from now on to shorten our formulae.

These distinctions are critical as they determine how the covariant derivative acts on the fields. Using the Dirac slash notation ($\gamma_\mu D^\mu = \not{D}$) we expect a Lagrangian of the form

$$\begin{aligned} \mathcal{L} = & i\bar{Q}\not{D}Q + i\bar{L}\not{D}L + i\bar{e}_R\not{D}e_R + i\bar{u}_R\not{D}u_R + i\bar{d}_R\not{D}d_R - \\ & - \frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu} \end{aligned} \quad (1.10)$$

but this is incomplete still. To finish our derivation we need to introduce the last component of the SM we have not yet talked about, the one responsible for the mass of the weak gauge bosons and all other elementary masses, the Higgs boson.

1.1.5 The Higgs field

Our experiences show that most of the standard model particles have a mass, however, an explicit fermion or gauge boson mass term in the standard model Lagrangian would not respect the gauge symmetry and therefore is forbidden.

The Brout–Englert–Higgs mechanism [1, 14, 15] was proposed to account for the mass we experience through a spontaneous breaking of the local symmetry.

In the standard model only the Salam-Weinberg part is spontaneously broken, the color symmetry is exact.

The spontaneous symmetry breaking

$$\text{SU}(2)_L \times \text{U}(1)_Y \rightarrow \text{U}(1)_{e.m.}$$

is achieved by introducing an $\text{SU}(2)_L$ doublet complex scalar field, the Higgs field¹¹, Φ ($Y[\Phi] = 1$) which has a potential that guarantees a non-zero vacuum expectation value. This is accomplished by a so-called sombrero potential which takes on its global minimum at some non-zero field value. A cross-section of such a potential is shown in Figure 1.2. This potential is usually parametrized in the following way:

$$V(\Phi) = -\mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (1.11)$$

where $\mu^2 > 0$, and $\lambda > 0$ are needed for the sombrero to take shape. With the previous V the Higgs Lagrangian takes the form

$$\mathcal{L}_H = |D_\mu \Phi|^2 + \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2 \quad (1.12)$$

from where we can already see that it couples to the electroweak vector bosons through the covariant derivative. The Higgs field also interacts with the fermionic sector and it does so through the so-called Yukawa couplings. These are terms, invariant under the gauge transformations, of the form $\bar{\chi}_L \Phi \Psi_R + (\text{h.c.})$ where χ_L is an $\text{SU}(2)_L$ doublet, Ψ_R is a singlet and h.c. stands for the Hermitian conjugate. In the most general form the Yukawa interaction Lagrangian is

$$\mathcal{L}_{\text{Yuk.}} = f_{AB}^{(e)} \bar{L}_A \Phi e_{RB} + f_{AB}^{(d)} \bar{Q}_A \Phi d_{RB} + f_{AB}^{(u)} \bar{Q}_A \tilde{\Phi} u_{RB} + \text{h.c.} \quad (1.13)$$

where the $f^{(\cdot)}$ are 3×3 matrices containing the coupling constants, $A, B \in \{1, 2, 3\}$ can be thought of as generation indices and $\tilde{\Phi} = i\tau_2 \Phi^*$.

With this, the final SM Lagrange density function is written as

$$\mathcal{L}_{SM} = \mathcal{L}_{\text{Dirac}} + \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yuk.}} \quad (1.14)$$

¹¹To underline the difference between the observable Higgs boson and the underlying Higgs field, I will use Φ for the latter one and η for the former one.

Where

$$\begin{aligned}
\mathcal{L}_{\text{Dirac}} &= \mathcal{L}_{\text{quark}} + \mathcal{L}_{\text{lepton}}, & (\text{All fermion dynamical terms.}) \\
\mathcal{L}_{\text{quark}} &= \bar{Q}_A i \not{D} Q_A + \bar{u}_{RA} i \not{D} u_{RA} + \bar{d}_{RA} i \not{D} d_{RA}, & (\text{Quark dynamical terms.}) \\
\mathcal{L}_{\text{lepton}} &= \bar{L}_A i \not{D} L_A + \bar{e}_{RA} i \not{D} e_{RA}, & (\text{Lepton dynamical terms.}) \\
\mathcal{L}_{\text{gauge}} &= -\frac{1}{4} W_{\mu\nu}^a W^{\mu\nu a} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} G_{\mu\nu}^b G^{b\mu\nu}, & (\text{Gauge boson dynamical terms.}) \\
\mathcal{L}_{\text{Higgs}} &= |D_\mu \Phi|^2 + \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2, & (\text{Higgs dyn. and self interaction terms.}) \\
\mathcal{L}_{\text{Yuk}} &= f_{AB}^{(e)} \bar{L}_A \Phi e_{RB} + f_{AB}^{(d)} \bar{Q}_A \Phi d_{RB} + & (\text{Fermion - Higgs interaction terms.}) \\
&\quad + f_{AB}^{(u)} \bar{Q}_A \tilde{\Phi} u_{RB} + \text{h.c.}
\end{aligned}$$

Here care needs to be taken that for quarks the covariant derivative is

$$D_\mu = \partial_\mu - ig_1 \frac{\tau^a}{2} W_\mu^a - ig_2 \frac{Y}{2} B_\mu - ig T^a A_\mu^a \quad (1.15)$$

with all the generators in the appropriate representation. In all other cases, D is the EW covariant derivative seen in Equation 1.6. It is also worth noting that the interaction between the fermions and the gauge bosons is encoded into the dynamical terms via the covariant derivative.

In the following, the most important consequences of the Higgs mechanism in the SM are introduced in a few words. This mechanism can be described and used in more general frameworks as well [1] which is out of the scope of this chapter. For a visual representation of the result see Figure 1.2.

Near the ground state of the Higgs field, we may use the

$$\Phi(x) = \exp\left(-i \frac{\tau^a \Theta^a(x)}{2v}\right) \begin{pmatrix} 0 \\ \frac{v+\eta(x)}{\sqrt{2}} \end{pmatrix} \xrightarrow{\text{unitary gauge}} \begin{pmatrix} 0 \\ \frac{v+\eta(x)}{\sqrt{2}} \end{pmatrix} \quad (1.16)$$

parametrization, where $v = \sqrt{\mu^2/\lambda}$. Choosing the unitary gauge¹³ we get the second form.

Substituting this into the Yukawa term we get

$$\begin{aligned}
\mathcal{L}_{\text{Yuk}} &= \bar{e}_{LA} \left(\frac{v}{\sqrt{2}} f_{AB}^{(e)} \right) e_{RB} + \bar{d}_{AL} \left(\frac{v}{\sqrt{2}} f_{AB}^{(d)} \right) d_{RB} + \bar{u}_{LA} \left(\frac{v}{\sqrt{2}} f_{AB}^{(u)} \right) u_{RB} + \\
&\quad + \frac{\eta(x)}{\sqrt{2}} \left[\bar{e}_{LA} (f_{AB}^{(e)}) e_{RB} + \bar{d}_{LA} (f_{AB}^{(d)}) d_{RB} + \bar{u}_{LA} (f_{AB}^{(u)}) u_{RB} \right] + \text{h.c.} \quad (1.17)
\end{aligned}$$

¹²https://commons.wikimedia.org/wiki/File:Standard_Model_Of_Particle_Physics--Most_Complete_Diagram.png

¹³The unitary gauge is the result of a particular gauge fixing condition which eliminates the Goldstone boson components, in this case with an $\exp\left(i \frac{\tau^a \Theta^a(x)}{2v}\right)$ transformation.




Figure 1.2: The SM fields before and after the EW symmetry breaking. Note that this figure uses $Y_{\text{fig}} = Y/2$ as the weak hypercharge going against the conventions. Also, the gluon fields are represented with g and the EM field with γ . Source¹².

Hence we see, that fermions (with the exception of the neutrinos) obtained a mass term that mixes the particles of different chirality. It is important to note, that the mass matrix $M_{AB}^{(\cdot)} = \frac{v}{\sqrt{2}} f_{AB}^{(\cdot)}$ is not necessarily diagonal in the flavours. This fact is responsible for the mixing of the charged quark currents as described by the Cabibbo–Kobayashi–Maskawa (CKM) matrix.

Inserting Equation 1.16 into $\mathcal{L}_{\text{Higgs}}$ and collecting the terms containing only the gauge fields we get

$$\frac{v^2}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix}^\dagger \left[\frac{g_1}{2} \tau^a W_\mu^a + \frac{g_2}{2} B_\mu \right]^\dagger \left[\frac{g_1}{2} \tau^a W_\mu^a + \frac{g_2}{2} B_\mu \right] \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (1.18)$$

which is a mass term for the gauge bosons, but again it is not diagonal in the flavors. After calculating the mass eigenstates we get:

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_\mu^1 \pm iW_\mu^2) \quad (M_W = \frac{vg_1}{2}) \quad (1.19)$$

$$Z_\mu = W_\mu^3 \cos \Theta_W - B_\mu \sin \Theta_W \quad (M_Z = \frac{M_W}{\cos \Theta_W}) \quad (1.20)$$

$$A_\mu = W_\mu^3 \sin \Theta_W + B_\mu \cos \Theta_W \quad (M_A = 0) \quad (1.21)$$

where we used the weak mixing angle which is defined by $\tan \Theta_W = g_2/g_1$.

These are the massive vector bosons and the photon that we often see in our experiments.

A massless vector boson has two degrees of freedom while one with a non-zero mass has three hence it seems like the W^\pm and Z^0 gained three extra degrees of freedom in total in the process. As, by choosing the unitary gauge, the Higgs field lost three degrees of freedom, theorists often say that these degrees were eaten by the new massive vector bosons.

Collecting some other terms from $\mathcal{L}_{\text{Higgs}}$, we get

$$\frac{1}{2}(\partial_\mu\eta)^2 - \mu^2\eta^2 \tag{1.22}$$

which is the Lagrangian of a scalar field of mass $m_H = \mu/\sqrt{2}\mu$.

This single scalar field, the physically observable Higgs boson, was the puzzle piece predicted by the standard model that was the last one to be found in 2012 by the CMS and ATLAS collaborations [16, 17]. The confirmation of the discovery came in March 2013 when they found that the results favor a spin 0, even parity particle with couplings that are consistent with the theory. Francois Englert and Peter Higgs were awarded the Nobel prize in physics the same year.

1.2 Limitations of the standard model

According to our present knowledge the standard model, despite its success, is not a complete theory. One argument of aesthetic motivation is that a fundamental theory cannot possibly have as many free parameters as the SM (19 in total).

However, there are more pressing reasons to believe that the SM is just an effective theory underlaid by a more fundamental one, many of which are discussed below.

Clearly, a new theory is needed on the energy scales where the gravitational effects become significant, but we are far from this regime experimentally and this is not expected to change in the foreseeable future.

The first substantial experimental deviation from the SM was the existence of neutrino oscillations. These show that the neutrinos are not all massless [18]. This contradicts the SM, which can be augmented to describe, but not explain, this phenomenon.

Currently, it is not clear through what mechanism neutrinos gain mass. Although the SM does not have right-handed neutrinos, it can be augmented to have them. This would allow them to gain mass through the Higgs mechanism, but in this case, the coupling constant of this term would need to be very small compared to all the others, which is theoretically unfavorable. The right-handed neutrino would also remain sterile, meaning that it would not take part in any of the three fundamental interactions. In general sterile neutrinos are disfavoured by many theorists for lack of the simplicity that other aspects of the SM show. Another possibility is that neutrinos are Majorana particles, meaning, that they are their own antiparticles. The discussion of this option is beyond the scope of this chapter.

Another very apparent and puzzling problem is the matter-antimatter asymmetry [19]. Astronomical observations indicate that the universe is mostly made out of ordinary matter and hardly any antimatter, however the CP violation encoded into the CKM matrix does not explain this huge difference.

There are several other discrepancies, like the (lack of) strong CP violation [20], or the anomalous magnetic moment of the muon¹⁴ [21, 22]. Although, the magnitude of the latter tension was brought in question after the BMW collaboration suggested that the theoretical prediction may be biased by accounting for the hadronic vacuum polarization inaccurately [23]. Their results based on lattice QCD calculations moved the prediction closer to the SM value, but the validity of the approach is still actively discussed.

In the following, two additional anomalies are introduced that can be mitigated by a supersymmetric theory and therefore are of particular interest to us. These are the hierarchy and the dark matter problem.

1.2.1 The hierarchy problem

The hierarchy problem [24] relates to naturalness which is an important consideration as it has been the guiding principle, together with simplicity, in the creation of the standard model [25], although some question the validity of such arguments. In order to qualify as a natural solution, the parameters of the theory should not have to be selected with a large precision for them to conform to

¹⁴This phenomenon suggests the invalidity of the lepton universality assumption which states that charged leptons are distinguished only by their masses (and their generational lepton numbers) which is a key component of the SM.

reality. This means that large cancellations between unrelated terms without a mechanism to guide this procedure are disfavoured.

In this setting, the SM is treated as an effective theory that describes the behavior of the "theory of everything" in a low-energy regime that is available to our experiments. It is well understood that the standard model loses its validity at higher energies at some point, at the latest at the Planck scale, which is the point where gravitational effects become unavoidable and the behavior of the micro world is expected to change substantially. The Planck scale is defined as

$$m_p = \sqrt{\frac{1}{G_N}} \approx 1.2 \cdot 10^{19} \text{ GeV}, \quad (1.23)$$

where G_N is the Newton constant. It is expected, that the high energy regimes hide new particles that although couple to the SM, are not directly available to us due to their immense mass.

In quantum field theory, the measured quantities are not the same as the parameters set in the Lagrangian. The latter, called bare quantities, are modified by quantum corrections. The Higgs mass obtains first-order corrections due to all particles that couple to it. For a fermion that couples to the Higgs field via the Yukawa term

$$-f_\Psi \bar{\Psi}_L \eta \Psi_R \quad (1.24)$$

using perturbative one-loop calculations we get that the mass correction associated to the first diagram of Figure 1.3 is

$$\Delta m_H^2 = -\frac{|f_\Psi|^2}{8\pi^2} [\Lambda_{UV}^2 - 2m_\Psi^2 \ln(\Lambda_{UV}/m_\Psi) + \dots] \quad (1.25)$$

where the ellipses stand for at most logarithmic terms in Λ_{UV} , called the ultraviolet cutoff energy. It is the scale up to which the SM is considered valid [26]. The interpretation of this value is, that for momenta $k^2 > \Lambda_{UV}^2$ the loop contribution is compensated for by new degrees of freedom. The standard model declares its validity up to the magnitude of the Planck scale, $\Lambda_{UV} \sim m_p$, since it predicts a region devoid of new particles up to this regime.

Even if one rejects this interpretation of Λ_{UV} and treats it as a pure regularization artifact, the correction also contains a term proportional to m_Ψ^2 , which means that the existence of any heavy fermion on the new physics scale would result in a comparable effect. Similar corrections arise for bosons coupling to the Higgs field as well, and even in case the new, heavy particles interact with the Higgs field only indirectly via the other SM particles [26].

Such an enormous correction would prompt us to expect the Higgs mass to be situated much closer to the new physics scale, m_p according to the SM, than we measured it to be. Finally, the problem can be stated in its first form: Why is the Higgs boson so unexpectedly light?

Then, since the mass of the Higgs boson is close to the electroweak scale defined by the Higgs vacuum expectation value

$$v = \sqrt{\frac{1}{\sqrt{2}G_F}} \approx 250 \text{ GeV}, \quad (1.26)$$

where G_F is the Fermi constant. As $G_N \ll G_F$, one may rephrase this problem as: Why is gravity so much weaker than even the aptly named weak nuclear force? The difference is 33 magnitudes.

The fact that m_H is so tiny compared to m_p indicates that there could be an incredibly fine-tuned counter term in the bare mass m_H that almost annihilates the quantum corrections not only due to the regularization term but also due to the new massive particles. This however is not favorable due to naturalness considerations.

An alternative possibility is a striking cancellation of the quantum correction terms by some sort of extra symmetry of the spectrum.

Now suppose that in the yet unseen high energy sector of the spectrum, there are scalar fields (Σ) that couple to the Higgs field according to

$$-f_\Sigma |\eta|^2 |\Sigma|^2. \quad (1.27)$$

In this case, the quantum correction we get for the Higgs mass using the bottom diagram of Figure 1.3 is

$$\Delta m_H^2 = \frac{f_\Sigma}{16\pi^2} [\Lambda_{UV}^2 - 2m_\Sigma^2 \ln(\Lambda_{UV}/m_\Sigma) + \dots]. \quad (1.28)$$

Notice, that if there are two scalars with couplings such that $f_{\Sigma_{1,2}} = |f_\Psi|^2$ for all fermions then the correction terms quadratic in Λ_{UV} cancel out. If the corresponding particles have the same mass, the second term in the correction also disappears. We shall later see that this is exactly what the minimal supersymmetric standard model proposes.

Note that the corrections arising through indirect interactions and many-loop corrections are also automatically canceled once supersymmetry is imposed on the theory.




Figure 1.3: The Feynman graphs of two single loop quantum corrections contributing to the mass of the Higgs boson. η is the Higgs boson, Ψ is a fermion and Σ is a scalar particle.

1.2.2 Dark matter

Turning our attention from the smallest scales to the largest ones, one still can not escape the puzzles of the microcosm. Astronomical observations, such as the angular speed profiles of galaxies [27], and gravitational lensing [28], show that the gravitational fields are not consistent with the amount of matter we see in the universe. It is widely accepted, that there must be extra non-luminous matter that accounts for these observations, which is called dark matter (DM), however alternative theories do exist (e.g., MOND, see Sec. 1.4.5). It is not yet known whether the nature of DM is baryonic, or non-baryonic but astronomical observations disfavor the former scenario.

A number of non-baryonic DM candidates were proposed, such as weakly interacting massive particles (WIMP), which behave similarly to neutrinos but are much heavier, or axions, which are also motivated by the lack of strong CP violation in the SM. Unfortunately, the SM does not provide any good candidates for DM while its supersymmetric extensions, in most cases, contain suitable particles. Phenomena necessitating DM is yet to be observed in collider experiments.

Precision cosmology measurements conducted on SN-1a type supernovae, using the Λ CDM model of cosmology, indicate that in terms of energy, our

universe consists of approximately 5% ordinary (or luminous) matter, 27% dark matter and 68% dark energy, an unknown form of energy hypothesized to drive the accelerating expansion of the universe. [29].

1.3 Supersymmetry as a possible solution

In a broad sense, SUSY is a symmetry that relates particles of different spins. In this interpretation, Wigner's $SU(4)$ theory relating nuclei of different spins was a supersymmetric model [30], although many sources recognize Miyazawa's work [31] of relating baryons and mesons as the first attempt in general supersymmetry, which however was widely disregarded at its time.

In the modern day context, by supersymmetry we mean a spacetime symmetry (meaning that the Poincaré group¹⁵ is locally isomorphic to a subgroup of it) that relates fermions and bosons. This hybrid structure is a result of efforts aimed at uniting internal and spacetime symmetries. A no-go theorem by Coleman and Mandula [32] states that under some quite general, physically justifiable assumptions all such unifications forming a Lie-group break up into direct products of the Poincaré symmetry and the internal symmetries, so in some sense they are boring.

SUSY can be thought of as a possible loophole in the no-go theorem as the generators of the SUSY group do not form a Lie-algebra, but a graded Lie-algebra, meaning that some generators, in fact, all but the Poincaré generators, admit anticommutator relations instead of commutator relations [33].

Another reason for the popularity of SUSY is that it is a prediction of string theory, an even broader framework that might finally let us combine all four fundamental forces in a single theory. In a string theory, gravity is also mediated by particles called gravitons which are massless spin-2 particles. Thus, a supersymmetric theory that acts as an effective description of a string theory also has them.

¹⁵The Poincaré group is a 10-dimensional Lie-group generated by the group of translations, spatial rotations, and boosts. It is the central symmetry group of special relativity.

For a quick dive in graded Lie-algebras, let us use B for the generators of bosonic properties (they admit a commutator relation), and F for the generators of fermionic properties (these are anti-commuting). We use a \mathbb{Z}_2 graded Lie-algebra, also called superalgebra. The elements of the superalgebra satisfy

$$\begin{aligned} [B_a, B_b] &= ic_{abc} B_c \\ [F_\alpha, B_b] &= s_{\alpha b\beta} F_\beta \\ \{F_\alpha, F_\beta\} &= v_{\alpha\beta\gamma} F_\gamma \end{aligned} \tag{1.29}$$

where c_{abc} , $s_{\alpha b\beta}$ and $v_{\alpha\beta\gamma}$ are the structure constants.

The earliest suggestions of anticommutators appear in a paper of Golfand and Likhtman in 1971 [34] and the first field theory incorporating SUSY was developed by Volkov and Akulov [35]. It was quickly recognized that SUSY removes a lot of divergent results in field theory and greatly simplifies renormalization, making SUSY an attractive possibility.

The first realistic supersymmetric extension of the SM was proposed in 1977 by Pierre Fayet [36] and by 1981 the minimal supersymmetric standard model (MSSM) was born which is still in use today. This theory gives a natural answer to the hierarchy problem and provides a possible dark matter candidate as well. Its predictions (new particles near the EW scale) also give hope for experimental verification of this theory with the current accelerator technology.

One surprising fact is, that while in the SM there is no energy at which the three described elementary forces unite (their running coupling constants become equal), the MSSM very closely guarantees such a point [26] at around 10^{16} GeV called the grand unification scale (GUT) (in case the SUSY masses are around 1 TeV), as illustrated in Figure 1.4.




Figure 1.4: A sketch of the inverse of the running coupling constants in the SM and MSSM assuming the SUSY masses are around 1 TeV. α_1 , α_2 and α_3 correspond to the electromagnetic, the weak and the strong force respectively. Source¹⁶.

Since the very first attempts, a lot of models were created that are consistent with the SM and employ the supersymmetry principle. Today a big family of SUSY theories is considered to be possible beyond the standard model (BSM) descriptions of the physics of elementary particles. The great number of different models span an enormous parameter space, large part of which is open for exploration at the LHC. However, not the entirety of the allowed SUSY phase space can be surveyed, hence even if SUSY is a true symmetry of the universe, it might remain hidden.

1.3.1 The spectrum of the MSSM

The infinitesimal generators of the Poincaré group form a regular Lie-algebra. Following the literature, we use P , J , and K for the generators of the translations, rotations, and boosts respectively. The simplest extension of this algebra to a superalgebra is the super-Poincaré algebra, where the ten generators of the Poincaré group are augmented with N Weyl spinor generators that obey the anticommutator relations

$$\begin{aligned}
 \{F_{\alpha}^n; F_{\beta}^{m\dagger}\} &= 2\sigma_{\alpha\beta}^{\mu} P_{\mu} \delta^{nm}, \\
 \{F_{\alpha}^n; F_{\beta}^{m\dagger}\} &= \epsilon_{\alpha\beta} Z^{nm}, \\
 \{F_{\dot{\alpha}}^I; F_{\dot{\beta}}^{m\dagger}\} &= \epsilon_{\dot{\alpha}\dot{\beta}} (Z^{nm})^*,
 \end{aligned} \tag{1.30}$$

where σ^{μ} is the four vector of Pauli matrices, $n, m \in \{1, \dots, N\}$, $\alpha, \beta \in \{1, 2\}$ and Z , called the central charge, is antisymmetric. The dot above the spinor indices

¹⁶https://www.nobelprize.org/nobel_prizes/physics/laureates/2004/popular.html

signifies that it belongs to the conjugate spinor representation.

All commutators between F and P vanish. The commutation relation with the rotation and boost generators is not trivial, however. Let us rewrite the J and K generators as

$$M_{0i} = -M_{i0} = K_i \quad \text{and} \quad M_{ij} = \epsilon_{ijk} J_k. \quad (1.31)$$

With these, the commutators can be expressed as

$$\begin{aligned} [M_{\mu\nu}; F_\alpha^n] &= i(\sigma_{\mu\nu})_\alpha^\beta F_\beta^n, \\ [M_{\mu\nu}; F_\alpha^{n\dagger}] &= i(\sigma_{\mu\nu}^\dagger)_\alpha^\beta F_\beta^{n\dagger}, \end{aligned} \quad (1.32)$$

where $\sigma_{\mu\nu} = \frac{i}{2}[\gamma_\mu, \gamma_\nu]$.

In the case of MSSM, $N = 1$. This family of models assumes the simplest supersymmetric group and thus it predicts the least possible new particles, exactly one for each chiral SM fermion and vector boson (as well as an additional Higgs doublet). Of course, MSSM theories that contain a graviton will also have its corresponding supersymmetric partner.

The Poincaré group has two Casimir operators $m^2 = P_\mu P^\mu$ and $W_\mu W^\mu$ where

$$W_\mu = \frac{1}{2} \epsilon_{\mu\nu\rho\sigma} M^{\nu\rho} P^\sigma \quad (1.33)$$

is the Pauli-Lubansky pseudovector. In the superalgebra extension however, unlike m^2 , $W_\mu W^\mu$ does not commute with F. In fact it can be shown that F changes the eigenvalue of the eigenstates of $W_\mu W^\mu$ by one half. This indicates that the generator F associates fermionic and bosonic states. We call a fermion and a boson linked by F superpartners.

It can be shown that in the MSSM each left- and right-handed SM fermion has a unique bosonic superpartner (spin zero) and each SM boson has a unique fermionic superpartner (spin one half). In case this applies to the given theory, the supersymmetric mirror particle of the spin-2 graviton is the spin- $\frac{3}{2}$ gravitino. In the analysis introduced in this thesis, we use a model which contains a graviton.

The fact that m^2 commutes with all elements of the superalgebra implies that superpartners should have the same mass if SUSY is an intact symmetry.

This means, for example, that there should be a superpartner of the electron with the same mass and electric charge (as $U(1)_{e.m.}$ is an internal symmetry

that creates a direct product with the Super-Poincaré group) which we will call selectron. But the selectron should be bosonic. As a selectron interacts with matter through the electromagnetic interaction we should have discovered it long ago in our experiments as it is the lightest charged particle, besides the electron, and hence it does not decay.

The failure to observe superpartners so far indicates that supersymmetry is hidden, or in other words spontaneously broken, causing superpartners to have different masses. The breaking occurs at some energy higher than the spontaneous breaking of the Salam-Weinberg model [37] so when we address the fields of the exact SUSY it makes sense to talk about SM fields before the EW symmetry breaking.

Several possible mechanisms were proposed for the spontaneous supersymmetry breaking that splits the SUSY and SM masses. We have briefly discussed how SUSY solves the hierarchy problem in 1.2.1. We wish to preserve this feature in the case of a broken SUSY as well. This is achieved in "soft breaking". In this case, the Lagrangian can be written in the form

$$\mathcal{L} = \mathcal{L}_{\text{SUSY}} + \mathcal{L}_{\text{soft}}, \quad (1.34)$$

where $\mathcal{L}_{\text{SUSY}}$ is the SUSY Lagrangian containing the gauge and Yukawa interactions and $\mathcal{L}_{\text{soft}}$ contains only mass terms and possibly some couplings with coefficients of positive mass dimension. $\mathcal{L}_{\text{soft}}$ of course does not respect supersymmetry. These masses can be independent parameters or they can be linked by a mechanism that generates them. As the previously proposed symmetry removed the quadratic divergences, the remaining mass correction from these mass terms is of the order

$$\Delta m_{\text{H}}^2 = m_{\text{soft}}^2 \frac{f}{16\pi^2} \log(\Lambda_{\text{UV}}/m_{\text{soft}}) \quad (1.35)$$

where m_{soft} is the largest mass scale in $\mathcal{L}_{\text{soft}}$. For the Higgs mass to stay small these soft masses must not be much bigger than the EW scale which means a few TeVs at most.

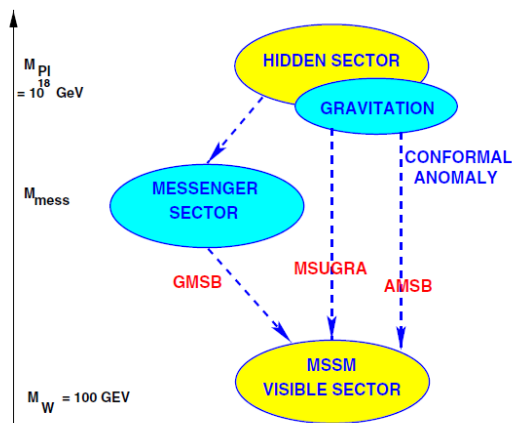


Figure 1.5: A schematic figure of the interaction structure of the most popular SUSY breaking mechanisms. The hidden sector refers to yet undiscovered fields that do not couple to SM gauge bosons and hence they do not directly interact with SM particles through the three SM forces but might interact through gravity for example. The hidden sector and the SM are linked through some massive or for other reason suppressed set of fields which constitute the messenger sector. [37].

Some of the most popular mechanisms generating the soft masses are the gravity-, gauge-, and anomaly-mediated supersymmetry breaking (SURGA, GMSB, AMSB respectively). The schematic of these is presented in Figure 1.5. In our research we use the MSSM model augmented with GMSB [38]. This model contains a new field with a non-zero vacuum expectation value that does not directly couple to any of the MSSM fields only to some messengers which are new supermultiplets charged under $SU(3)_C \times SU(2)_L \times U(1)_Y$ [37]. This provides an indirect connection to the MSSM to generate the soft masses in two steps.

In practice, however, we use the general gauge mediation (GGM) model [39], which allows for a theoretically well-grounded, yet model-independent exploration of GMSB phenomenology. The parameters of this model also fit the phenomenological description better, hence it is more convenient to use.

Finding the lightest supersymmetric particle (LSP) is a crucial part of the spectrum analysis of a theory. This, in GMSB models, is always the gravitino.

In general, the naming convention for superpartners of SM fermions is using an "s" (for scalar) prefix added to the name of the SM partner. This is also valid for the groups of particles. For example squarks (\tilde{q}) for quarks (q), sleptons (\tilde{l}) for leptons (l), selectron (\tilde{e}) for electron (e), and stop (\tilde{t}) for top (t). For the third generation of sfermions, the mass difference between the left- and right-handed partners can be significant but for all other generations, they are considered

degenerate in mass.

For the partners of the bosonic SM particles and the graviton we use an "ino" suffix. This rule gives wino (\widetilde{W}) for the W , bino (\widetilde{B}) for the B gauge boson, gluino (\widetilde{g}) for the gluon (g), gravitino (\widetilde{G}) for graviton (G) and higgsino ($\widetilde{\Phi}$) for the Higgs bosons (Φ). The collective name for these predicted particles is sparticles. As seen above sparticles are noted with a tilde. Figure 1.6 summarizes the MSSM particle spectrum.

One notable novelty of the spectrum, besides the presence of a graviton, is that in the MSSM one needs to introduce a second Higgs doublet field to avoid anomalies in the theory. One of these two doublets couples to the up- and the other to the down-type fermions. This raises the degrees of freedom of the Higgs fields to eight. After the EW symmetry breaking however only five degrees of freedom remain. This means that there are five Higgs fields in the MSSM of which two are charged. These are the h^0 , H^0 , A^0 and H^\pm if CP is conserved in the Higgs sector. The Higgs boson found at the LHC in 2012 is assumed to be the lightest of the neutral ones (h^0) as this possesses similar properties to the SM Higgs in many versions of the MSSM. An important parameter of the theory is the ratio of the vacuum expectation values of the two doublets

$$\text{tg}\beta = \frac{\langle 0|\Phi_u|0\rangle}{\langle 0|\Phi_d|0\rangle}. \quad (1.36)$$

After the electroweak symmetry breaking, the SM particles act the same way as they do according to the SM. The wino, bino, and some higgsino degrees of freedom mix to form the $\widetilde{\chi}_1^0, \widetilde{\chi}_2^0, \widetilde{\chi}_3^0, \widetilde{\chi}_4^0$ majorana fermions called neutralinos and the $\widetilde{\chi}_1^\pm, \widetilde{\chi}_2^\pm$ charged particles called charginos which are the new mass eigenstates. These are labeled in ascending order in terms of mass.




Figure 1.6: The MSSM spectrum below the EW energy scale. The different colors represent the big groups of particles. The SM particles are shown in front of a green background for example. The figure also shows the $SU(2)_L$ representation of the fields. Figure by Marton Bartok.

The experimental signature of MSSM models is also largely determined by the next to lightest supersymmetric particle (NLSP). We assume a $\tilde{\chi}_1^0$ NLSP in our analysis. In GGM almost any superpartner can be an NLSP which makes it a versatile signature generator model.

The neutralino mass matrix, which describes the composition of the NLSP in this case, depends on the weak mixing angle (Θ_W), β , the supersymmetric Higgs mass (μ), and the soft masses of bino and wino (M_1 and M_2). Depending on the matrix, the lightest neutralino might be bino-, wino- or higgsino-like. The last option is favored by naturalness. Depending on its composition, the $\tilde{\chi}_1^0$ favors different channels of decay.

These parameters also determine the mass spectra of the charginos and neutralinos at tree level. Three examples of characteristic spectra are shown in Figure 1.7.




Figure 1.7: The typical spectra of the charginos and neutralinos depend on the values of μ , M_1 and M_2 . [40].

The most general MSSM has over a hundred free parameters besides the SM parameters (105+45 to be exact) which gives rise to a great number of possible mass spectra even when constrained to be consistent with our experimental observations. Such a huge parameter space is hopeless to be exhaustively searched. By making some reasonable assumptions, like banning flavour-changing neutral currents, new sources of CP violation, and requiring the degeneracy of the first and second generations, the number of additional parameters is reduced to 19. This model is called the phenomenological MSSM (pMSSM-19).

1.3.2 Final states for colliders in the MSSM

Baryon and lepton number conservation is not hardwired into the MSSM under all renormalizable couplings we could fit in its Lagrangian. In the case when the coupling constants for the offending terms are of order 1, the theory predicts proton decay with a $\sim 10^{-2}$ second half-life. Analysis of the data collected by super-Kamiokande verified that the proton half-life is greater than $1.6 \cdot 10^{34}$ years, a result from 2016, which makes this prediction not theoretically favorable, to say the least [41].

To mitigate this problem the concept of R-parity is introduced which is a discrete \mathbb{Z}_2 extra symmetry forbidding the non-conserving couplings. The R-parity operator is defined as

$$P_R = (-1)^{3(B-L)+2s} \quad (1.37)$$

where B, L, and s are the baryon, lepton, and spin quantum numbers respectively. It is easy to check that SM particles are eigenstates of this operator with a (+1) eigenvalue, while their supersymmetric partners are eigenstates with a (-1) eigenvalue. Note that proton decay violates both baryon and lepton number conservation thus in theory it would be possible to prohibit it by only banning the violation of one of these quantum numbers.

If R-parity is in fact conserved (RPC) in interactions, then the decay products of non-SM particles should always contain an odd number of sparticles. Consequently, the lightest supersymmetric particle (LSP) is stable. While in models incorporating gauge-mediated supersymmetry breaking the LSP is always the gravitino, in other SUSY theories the lightest neutralino is a popular candidate.

The LSP is a potential candidate for the missing mass observed in astronomy, the previously discussed dark matter. Note that if the LSP was charged, it would have likely been discovered at the LHC already.

The searches conducted now are aiming to observe the decay of possible supersymmetric particles created in collisions and find traces of SUSY decay chains in the process.

The LHC is a proton-proton collider and as such it produces strongly interacting particles in the initial hard scattering with the highest probability, which has the most energy available to create new particles. The particles charged under $SU(3)_C$ are the quarks, gluons, and their superpartners the squarks and gluinos. If however, charginos are significantly lighter than gluinos they can take over the

leading role in production.

The many MSSMs variants provide a wide spectrum of possible final states to which SUSY decay chains can contribute and many analysis teams work simultaneously to cover all the possible interesting scenarios with the most advanced techniques available.

Let's examine, as an example, a possible decay chain of a gluino to build a collision event, assuming that \tilde{G} is the LSP and $\tilde{\chi}_1^0$ is the NLSP.

Gluinos can only decay either into an anti-quark and a squark or a quark and an anti-squark. These processes are represented as

$$\tilde{g} \rightarrow \bar{q}\tilde{q}^{(*)} \quad \text{and} \quad \tilde{g} \rightarrow q\tilde{q}^{(*)}. \quad (1.38)$$

The (*) symbol highlights, that squarks might be heavier than the gluino, and thus they might only emerge in this decay as virtual particles. Squarks may decay to a quark and a neutralino or a quark and a chargino

$$\tilde{q} \rightarrow q\tilde{\chi}^0 \quad \text{and} \quad \tilde{q} \rightarrow q'\tilde{\chi}^\pm. \quad (1.39)$$

Neutralinos and charginos may decay through many channels but in the final state phenomenology of interest the radiative decays of

$$\tilde{\chi}_1^0 \rightarrow \tilde{G}\gamma, \quad \tilde{\chi}_1^0 \rightarrow \tilde{G}H, \quad \tilde{\chi}_1^0 \rightarrow \tilde{G}Z^0, \quad (1.40)$$

$$\tilde{\chi}^\pm \rightarrow W^\pm\tilde{\chi}^0 \quad (1.41)$$

are of special importance. The decay channel the $\tilde{\chi}_1^0$ favors depends on its composition. Here H stands for one of the neutral Higgs bosons.

Using these building blocks one may construct decay chains with a characteristic final state, for example, if one of the neutralinos decays to $\gamma\tilde{G}$ while the other to $H\tilde{G}$:

$$\tilde{g}\tilde{g} \rightarrow (\bar{q}\tilde{q}^{(*)})(\bar{q}\tilde{q}^{(*)}) \rightarrow (\bar{q}q\tilde{\chi}^0)(\bar{q}q\tilde{\chi}^0) \rightarrow \bar{q}\bar{q}qq\gamma H\tilde{G}\tilde{G}. \quad (1.42)$$

Such a final state is characterized by a big missing transverse energy, a high energy photon, and signs of a Higgs boson which has many decay channels. The Higgs can decay into fermion-antifermion pairs, and the heavier the fermion is the higher the branching ratio is. This makes the bottom-antibottom channel the most probable, as the Higgs is not heavy enough to decay into $t\bar{t}$ pairs. SM Higgs boson undergoes the $H \rightarrow b\bar{b}$ process 57% of the time. Events with this signature are the subject of the BSM search in this work.

1.4 Alternative theories

In this section, a few alternative ideas for solving the quite puzzling hierarchy problem and the dark matter problem are mentioned. These theories are still a field of active theoretical and experimental research. For both SUSY and the theories expanded on in this section many variants exist, and large portions of their parameter spaces are excluded based on the Higgs mass and other high and low energy observations. That being said, currently, there is no strongly favored theory for the explanation of these anomalies, and it is the responsibility of experimentalists to look under every stone they are strong enough to turn.

1.4.1 The anthropic approach

The anthropic principle is a philosophical approach to tackling the hierarchy problem as well as other similar questions. It relies on the consideration, that for a universe to form with suitable conditions to harbor intelligent life, a large number of conditions need to be met as far as the values of the fundamental constants go. Were Newton's constant or the fine structure constant a little different, we might not be here to observe the Universe and ask questions about why things are the way they are.

This approach is often framed in the context of multiverse theories, where our existence is supported by the infinite chances the different worlds had to randomly pick the parameters that are just right for us to be around and ask questions about them.

It is easy to see how the hierarchy problem fits well with this line of reasoning. From a scientific point of view, however, this approach is as good as giving up hope of finding an explanation as the claims of this theory cannot be experimentally verified. For the reason that search for deeper meaning is thus discouraged, this line of thinking is very divisive among physicists.

1.4.2 Compact extra dimensions

This idea is partially inspired by string theory, which predicts at least nine spatial dimensions. The only problem is, that our world seems to have a 3+1 dimensional spacetime. How could there be more spatial dimensions without us seeing things

embedded in this bigger structure penetrating our four-dimensional spacetime or being able to sense its presence by other means?

A possible answer is the scale difference between our four dimensions and the extra dimensions which could be subatomic in size. Their scale is usually taken to be of the order of the Planck scale ($l_P = \sqrt{G_N} = \frac{1}{m_p} \approx 1.6 \cdot 10^{-35}$ m) but they could be larger even down to the TeV scale ($\frac{1}{1\text{TeV}} \approx 20 \cdot 10^{-19}$ m) without conflicting any of our experimental findings. Direct measurements, like the Eöt-Wash experiment, managed to test Newtonian gravity only down to 44 μm distance [42].

A great analogy compares the experience of a rope-walker and an ant [43]. Imagine a rope of a few centimeters thickness tightened between two trees. A rope-walker trying to cross from one tree to the other can move only forwards or backward and hence sees the rope as a one-dimensional path. In comparison an ant on this rope experiences a two-dimensional surface being able to walk around it.

A large variety of theories exist based on which particles constitute an ant and which ones are the rope-walkers, or in other words which particles are allowed to propagate through the extra dimensions. Now suppose that gravitons are the only particles that can do so.

Such extra dimensions could account for the weakness of the gravitational force by introducing a characteristic distance scale l of their size [44]. For distances $r \gg l$, we must get the well-known Newton's law

$$F(r) = -G_{3+1} \frac{m_1 m_2}{r^2} \mathbf{e}_r \quad (1.43)$$

($G_{3+1} = G_N$) but the strength of the force is decreased according to Gauss' flux theorem as the gravitational flux fills the δ extra dimensions which bear a volume of $V_\delta = l^\delta$. This means that gravity can be a strong force with a coupling constant $\sqrt{G_{3+1+\delta}}$ only weakened by the volume of the extra dimensions to what we experience and describe with $\sqrt{G_{3+1}}$. For $r \ll l$, we must take the extra dimensions into account so – instead of Newton's – law we get the

$$F(r) = -G_{3+1+\delta} \frac{m_1 m_2}{r^{2+\delta}} \mathbf{e}_r \quad (1.44)$$

distance dependence. Connecting the two equations while taking the finiteness of the extra dimensions into consideration we have

$$-G_{3+1} \frac{m_1 m_2}{r^2} \mathbf{e}_r = -G_{3+1+\delta} \frac{m_1 m_2}{r^{2+\delta}} \mathbf{e}_r, \quad (1.45)$$

$$G_{3+1} = G_{3+1+\delta} \frac{1}{V_\delta}. \quad (1.46)$$

By translating the problem to the mass scale we have

$$M_{P_{1_{3+1}}}^2 = V_\delta M_{P_{1_{3+1+\delta}}}^{2+\delta}. \quad (1.47)$$

From where we can see that for large enough extra volume the real Plank mass $M_{P_{1_{3+1+\delta}}}$ can be brought down to the EW scale.

If a particle is allowed to propagate in the extra dimensions its mass is computed as

$$m_0^2 = E^2 - p_0^2 - p_1^2 - p_2^2 - p_3^2 - p_4^2 \dots - p_{3+\delta}^2, \quad (1.48)$$

which results in its mass appearing in our 3+1 dimensions to be

$$m^2 = E^2 - p_0^2 - p_1^2 - p_2^2 - p_3^2 = m_0^2 + p_4^2 \dots + p_{3+\delta}^2. \quad (1.49)$$

As our extra dimensions are compact we get a periodic boundary condition which requires the momenta to be quantized: $p_i = \frac{2\pi n_i}{l_i}$. Here $i \in \{4, \dots, 3 + \delta\}$, $n_i \in \mathbb{Z}$ and l_i is the size of the i^{th} extra dimension. The observed mass for these particles is hence

$$m^2(\{n_i\}) = m_0^2 + 2\pi \sum_{i=4}^{3+\delta} \left(\frac{n_i}{l_i} \right)^2 \quad (1.50)$$

giving rise to the Kaluza-Klein tower. This manifests itself as if there were multiple excitation modes of the elementary particles which is a characteristic property of theories with compact extra dimensions. Depending on the exact model this might consist of discrete levels or make up a continuous spectrum.

A similar but unrelated family of theories are the braneworld models which use normal-sized extra dimensions to explain the weakness of gravity [45]. In these theories, the SM particles are constrained to a 3+1 dimensional brane, while the gravitational force is allowed to propagate through the whole (3+N)+1 dimensional space called the bulk and is strong on a brane other than ours. This structure is illustrated in Figure 1.8. Similarly to the previous model much of the gravitational force "leaks" into the extra dimensions where our matter particles cannot enter. In this case, on the SM brane, we can only experience the weakened gravity at any distance scale.




Figure 1.8: An illustration of the braneworld model. The two gray rectangles represent the embedded 3+1 dimensional branes, while the direction perpendicular to them is the 5th (extra) dimension [45].

1.4.3 Conformal solution

A conceptually different idea [46] is that the corrections to the Higgs mass arise from the existence of a quadratic term in its potential. Without an explicit Higgs mass term in the Lagrangian, the hierarchy problem would not manifest. On the other hand, without a mass term, the Higgs potential does not guarantee a nonzero vacuum expectation value. The electroweak symmetry breaking is recovered by a mass term generated through radiative corrections according to the Weinberg-Coleman mechanism [47]. This effect is too weak however to account for the experimental results. A possible way to compensate for this is by postulating the existence of more than one Higgs fields.

1.4.4 Technicolor

Technicolor [48] is an extension of the standard model except that it does not have a scalar Higgs field. This theory is an alternative, dynamical, description of the electroweak symmetry breaking which is modeled after QCD where its name originates from.

In this framework, the Higgs is not an elementary particle but it is made up of a set of new fermions, the techniquarks, which interact through a new, strong-like, interaction. This gives an asymptotically free and technicolor confining theory. Modern technicolor theory takes some departures from QCD however. "Walking technicolor" theories (WTC) are a popular direction for this research, where the

coupling constants are attracted to fixed values in the low energy limit such that the breaking of the EW symmetry is realized.

The techniquark condensate Higgs breaks the EW symmetry but to give a mass to the SM fermions there has to be a new interaction communicating this effect. This is possible in extended technicolor theories.

This strategy makes the Higgs mass less sensitive to contributions due to other massive particles and solves the hierarchy problem.

1.4.5 Modified Newtonian dynamics

The modified Newtonian dynamics (MOND) [49] approach aims to address the astrophysical observations that are otherwise explained by the existence of dark matter.

The most significant such problem concerns the speed at which stars circle the galactic core as a function of their radial position. Based on the brightness of different parts of the galaxy a mass distribution can be estimated. However, the galactic rotation curve computed from this does not align with experimental data [27]. The latter suggests that there is significantly more mass in a galaxy than meets the eye.

MOND aims to resolve the discrepancy by assuming that the effects of gravity may deviate from those predicted by Newton's laws for extremely small accelerations [50]. While the different modifications successfully describe the angular speed profiles, many MOND models predict that the speed of gravitational waves will decrease below the speed of light. The multimessenger observation of binary neutron star mergers [51] has restricted the available models heavily.

Rotation curves far beyond the visible disc of galaxies were studied using weak gravitational lensing [52]. While the dark matter halo is expected to fall off in these regions, the extended rotation curves support a maintained effect, which favors MOND.

Finally, a fresh study of wide binary systems observed with the help of the Gaia satellite has been carried out in [53]. It is found that the data overwhelmingly prefers Newtonian gravity over MOND in its current form.

Chapter 2

Where the magic happens

CERN, the European Organisation for Nuclear Research was founded in 1954 just nine years after the end of the Second World War to help the unification of Europe and take a leading role in non-military nuclear research. The scientific achievements of the laboratory were fundamental in the evolution of modern particle physics, and the technologies developed for the experiments contribute to our everyday life in numerous ways. The laboratory, celebrating its 70th birthday this year, continues to act as one of the cornerstones in the world of high-energy experimental physics in the 21st century. Today CERN is a service provider for international collaborations by building and operating particle accelerators and hosting a diverse set of experiments.

2.1 The Large Hadron Collider at CERN

In the early '80s a decision was made to build the Large Electron Positron collider (LEP), a synchrotron with a circumference of 27 km at the main CERN site just outside of Geneva. For this, a tunnel 100 m underground had to be excavated which was a very ambitious plan and a huge investment ¹.

LHC was proposed as LEP's successor even before the LEP collider came online. LEP operated from 1989 till 2000 soon after which the installation of the LHC began. The new accelerator was first started up in 2008 but due to a serious

¹During the building of this tunnel the workers even hit an underground river which they had to freeze with liquid nitrogen to create a medium which they could then bore through [54].

magnet quenching² incident [54], damaging over 20 superconducting magnets, the first physics run was delayed until 2010³.

The lifetime of the Large Hadron Collider is divided into runs (Run1, Run2, etc.) and so-called long shutdowns when technical upgrades and maintenance work can be carried out. See Figure 2.1 for the planned milestones of LHC's near future. The runs all start with beam recommissioning and are interrupted by short technical stops for fine-tuning the equipment.

During Run1 (2010-2013) the LHC operated at a limited centre-of-mass energy which by 2012 reached 8 TeV for proton-proton collisions. In this period the Higgs boson was announced to be found independently by the CMS and ATLAS collaborations in the same event to a crowded auditorium. The Higgs boson was the last missing puzzle piece of the Standard Model and hence the discovery was an enormous advancement in particle physics.

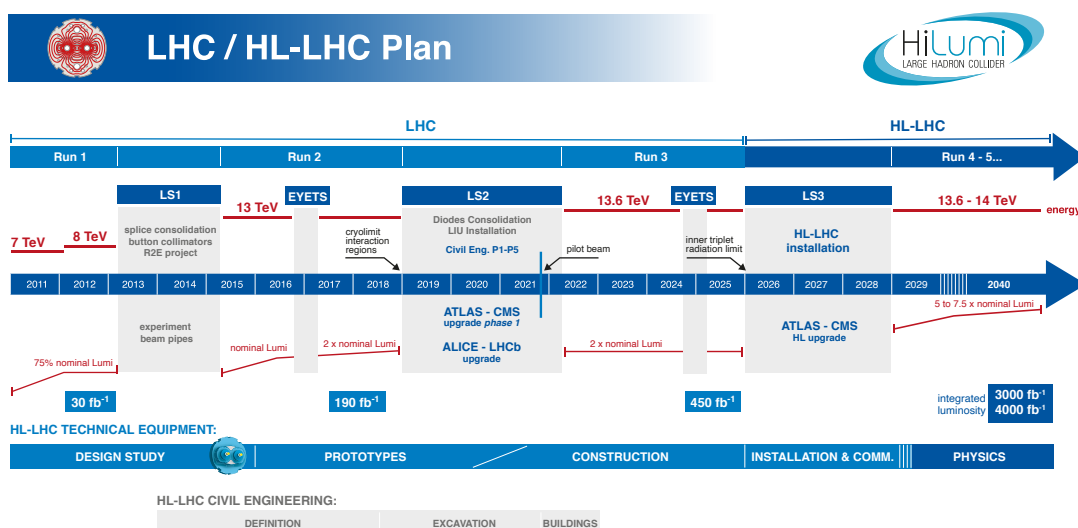


Figure 2.1: The LHC/HL-LHC (High Luminosity LHC) schedule with the planned shutdowns and data collection periods. The anticipated collision energies, instantaneous and integrated luminosities are also shown. Source⁴.

Run2 (2015-2018) started proton-proton collisions at a center-of-mass energy of $\sqrt{s} = 13$ TeV, approaching the design target of 14 TeV. This was maintained throughout the entire run due to the lack of time to retrain the dipole magnets

²A sudden termination of operation due to loss of superconductivity in a subsection of the coil.

³At the end of 2009 there was a commissioning run at 2.36 TeV collision energy.

⁴<https://hilumilhc.web.cern.ch/article/ls3-schedule-change>

for the higher currents required for 14 TeV. Run2 saw efforts to reach the design instantaneous luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ which was achieved in 2016 and was more than doubled in 2017 to $2.06 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$.

Currently, we are in the phase of Run3 which started in 2022 at a center-of-mass energy of $\sqrt{s} = 13.6 \text{ TeV}$ and is expected to last until the end of 2026, when the three-year-long long shutdown 3 (LS3) commences.

2.1.1 The CERN accelerator complex

Of course, the LHC is not a standalone accelerator and, like all synchrotrons, needs pre-acceleration. At the CERN site, this is provided by a cascade of four accelerators.

The physics analysis in this thesis uses data from proton-proton collisions which take up most of LHC's online time, so in case of ambiguity the facts in this subsection refer to this mode of operation. That being said the luminosity-related work presented in Chapter 3 also directly impacts the lead-lead physics results of the collaboration as well.

In the long shutdown between Run2 and Run3 the first stage of the accelerator chain was upgraded as part of the High Luminosity LHC (HL-LHC) project. The Linac2 (linear accelerator 2) proton accelerator was decommissioned and Linac4, a negative hydrogen ion (H^-) accelerator took its place, which now provides particles to all of CERN's proton accelerators.

The upgrade brings improvements to the maximal beam intensity achievable due to two factors. Firstly, the new linac accelerates particles to 160 MeV instead of only 50 MeV which in turn decreases the negative effect the beam's self-field (space charge) has on the stability of synchrotron orbits.

Secondly, previously the beam intensity in the second-in-chain accelerator, the Proton Synchrotron Booster (PSB) was enhanced by adding protons from Linac2 during up to 13 revolutions of the beam in the PSB. However, once an area in the position-momentum ($x-x'$) phase space is filled, it is not possible to add further protons into the same area, since the position of the particles in this space uniquely determines the effect the steering fields have on them. This leads to an increased emittance of the beams, hindering the maximum obtainable luminosity.

By adopting H^- ions in the linear accelerator a charge-exchange injection system can be used, which allows for the addition of new particles into the same phase-space area, overcoming the previously laid out limitation. A sketch of the new injector is shown in Figure 2.2.




Figure 2.2: A sketch of the charge-exchange injector system for the Proton Synchrotron Booster. Image from⁵.

The particles to be fed to the LHC ring start their journey bound in hydrogen molecules injected into a radio frequency H^- ion source, attached to Linac4, which is an 86-meter-long multi-stage radio frequency linear accelerator that also performs initial bunching. At the outlet, the ions have an energy of 160 MeV ($v \approx 0.5c$).

Next, the ions are transferred to the PSB, which is a stack of four synchrotrons with a radius of 25 meters. This architecture is chosen to produce the most amount of accelerated particles in the shortest time available at a reasonable cost. The incoming beam is deflected vertically to fill each of the rings using pulsed magnets.

The hydrogen anions traverse a stripping foil at the injection point designed to remove both electrons from 99% of the ions. The resulting protons are channeled into the PSB ring using dipole magnets, while hydrogen anions and atoms are directed to a beam dump, see Fig. 2.2.

Each ring accelerates its beam to 2 GeV ($v \approx 0.94c$) which are then combined for injection into the 36-meter-radius Proton Synchrotron (PS) with an output

⁵<https://cerncourier.com/a/the-ps-booster-hits-40/>

energy of 25 GeV ($v \approx 0.99926c$). The last element of the pre-acceleration chain is the "Nobel Prize Winner"⁶, 1100-meter-radius Super Proton Synchrotron (SPS) which pushes the beam to 450 GeV ($v \approx 0.9999975c$).

Finally, the bunches are transferred to the beam tubes of the LHC. Filling the 27-kilometer-circumference ($r = 4.3$ km) LHC takes 4 minutes and 20 seconds for each direction after which the protons reach the maximal energy of 6.8 TeV ($v \approx 0.999999989c$) in about 20 minutes. Figure 2.3 shows the complete CERN accelerator infrastructure.

Since the data used for the physics analysis in this thesis was recorded in Run2, we note that at that time Linac2 accelerated protons to 50 MeV, the PSB operated with an output energy of 1.4 GeV. The output energies of PS and SPS have not changed since Run2, and LHC operated at a proton beam energy of 6.5 TeV, which is equivalent to a center-of-mass energy of 13 TeV.

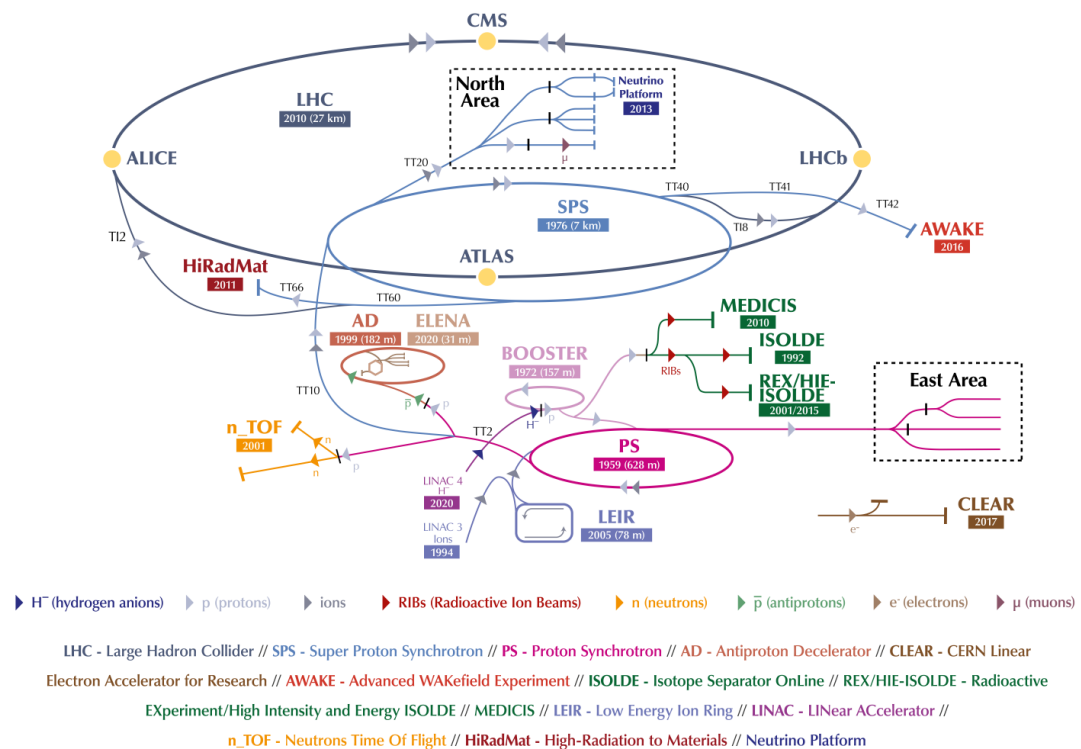


Figure 2.3: A not to scale schema of the accelerators and related experiments of CERN in 2022. Image adapted from. Source⁷.

⁶In 1983 the W and Z bosons were discovered using SPS which back then operated as a proton-antiproton collider ($Sp\bar{p}S$). In 1984, Carlo Rubbia and Simon van der Meer were awarded the Nobel Prize in Physics for their contribution to the project.

⁷<https://cds.cern.ch/record/2800984/files/>

2.1.2 The LHC collider

The LHC is a particle accelerator of the synchrotron type designed to primarily perform proton-proton collisions, but lead-lead and proton-lead are also carried out. Alternative options are also considered. For example, in Run2 xenon-xenon collisions were recorded, lead ions with a single electron were also injected into the machine, while Run3 foresees oxygen-oxygen collisions.

The LHC, unlike particle-antiparticle accelerators, cannot use the same magnetic field for guiding the beams traveling in opposite directions and hence it needs two beam tubes. Due to dimensional constraints, and to simplify cryogenic design a closely-spaced twin-bore design was adopted. The beam tubes intersect at four interaction points where detectors of the big experiments are located, see Figure 2.4.

The ring of LHC is actually a rounded octagon consisting of eight 2.45 km long arcs, each containing 154 bending (dipole) magnets (and of course the much smaller multipole focusing magnets) and eight straight "inserts". The role of these straight subsections depends on their location on the octagon. Such roles include the intersection (or collision) points, the locations of the accelerating radio frequency (RF) cavities, the injectors or the beam dump, for example. For a more comprehensive list see Figure 2.4.

The working units of the LHC are the sectors that stretch from the center of one octant to the next. They spread from one LHC site to the next one, thereby getting the names 12, 23, 34, and so on. These units are powered and cooled independently of one another.

For guiding and manipulating the beams, niobium-titanium superconducting magnets are used. These are kept at a 1.9 K temperature⁸ by cryogenic cooling. Of these, the most prominent ones are the 1232 15 m long dipole magnets weighing in at 35000 kg each. It is interesting to note, that the units containing the dipoles are also equipped with small multipoles for beam quality enhancement.

The dipole magnets can create a field of 8.3 T which requires an electric current in the coils of almost 12 kA. If we add up the energy of the magnetic field stored in the superconducting magnets we arrive at 9 GJ or 560 MJ per powering unit, of which there are two in each sector. This energy equals to roughly 130 kg

⁸It is often said that this is colder than outer space since the temperature of cosmic microwave background radiation was found to be approximately 2.7 K.

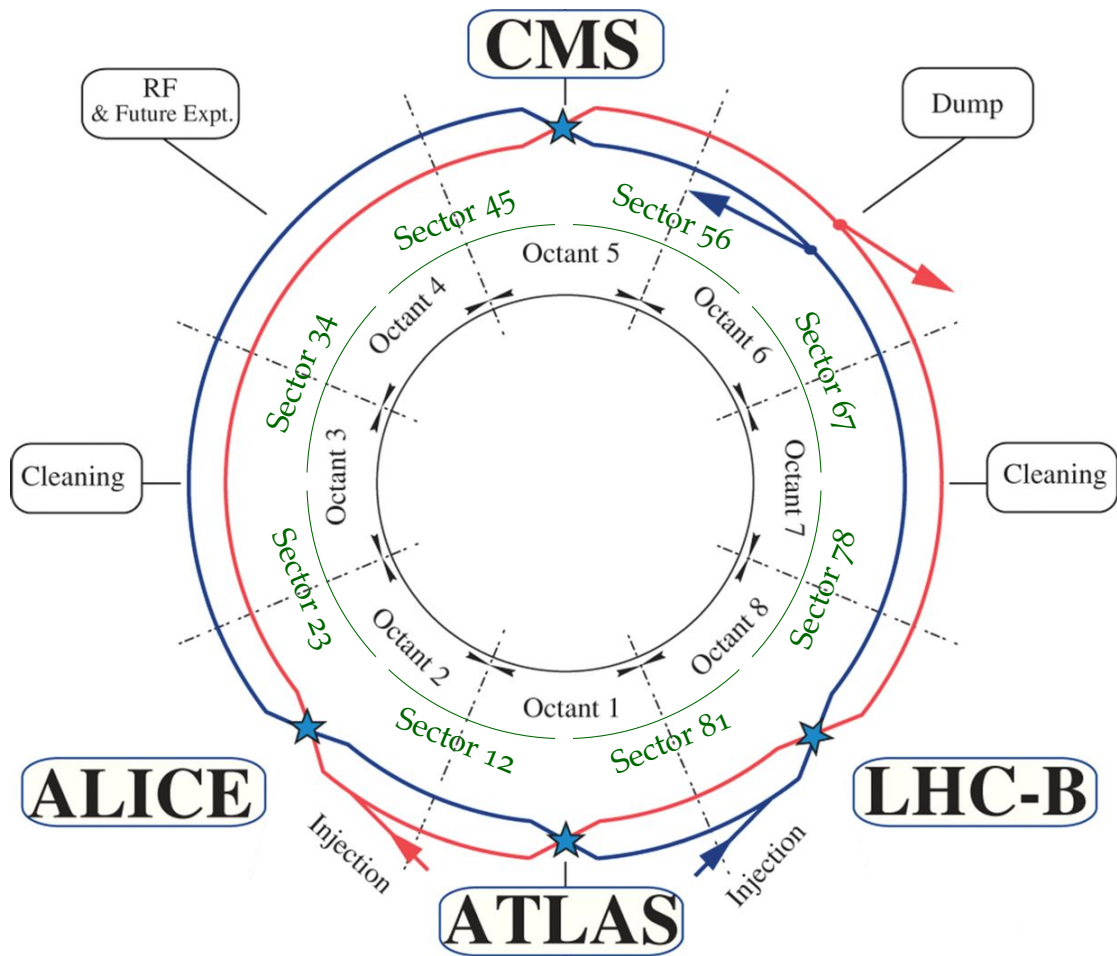


Figure 2.4: The layout of the LHC accelerator ring with the function of inserts in each octant, the sectors, and the big experiments (ATLAS, ALICE, CMS, LHC-B) listed. Figure adapted from [55].

of TNT. This makes an energy extraction system crucial for a safe operation.

In case of a quench occurring in a chain of series-powered dipole magnets, the quench protection system bypasses the affected magnet and rapidly de-excites the unaffected magnets by connecting them to an energy dump resistor that dissipates the energy as heat. The quench area of the affected magnet is spread out within the dipole using heaters to forego pinpoint damage.

Like in all synchrotrons, particles are organized into bunches to travel in the beam tube. A bunch contains around 10^{11} protons when the beam is new and each beam consists of a maximum of 2808 bunches with a minimum spacing of

25 ns.

From here it is easy to calculate that a single proton beam of LHC contains around 315 MJ of energy (at 7 TeV) which is comparable to the kinetic energy of a 40 m long full-speed TGV train (320 km/h). The same amount of energy can melt around 500 kg of copper, hence dumping the beams is not a trivial task either. Although, doing so quickly becomes unavoidable in case a dipole quench occurs, otherwise the beam could damage the machine.

To allow for a safe abortion of the operations there is a significant 3 μ s gap in the beam. This bigger gap enables the kicker magnets to turn on completely by the time the next bunch reaches the extraction point. Under usual data-taking conditions, the beam has a multi-level structure, with some bunch slots not filled, called the train structure. Due to the way the high-intensity beams are prepared, the bunches form 72-bunch long consecutive trains that follow one another with gaps of different but well-defined lengths. In special cases, however, the filling scheme can be adapted to the mode of operation and the purpose of the data taking, and various configurations (for example, with isolated bunches) can be loaded into the machine.

In CMS and ATLAS always the same bunch pairs collide because of their symmetric placement on the ring. The colliding bunch pairs are named according to their position on the LHC orbit by the beam crossing identifier (BCID or BX). On closer inspection, the bunches may have different properties depending on which interaction points they witness collisions at.

After around ten hours of operation, the beams get depleted and are dumped. This is followed by a pre-cycle phase, to reset the state of the magnets, required due to hysteresis as the magnetic field has to be significantly decreased for the new beam. In the next step, the injection system fills up the LHC in multiple rounds. This is followed by ramping up (accelerating) and a phase called "flat top" during which the beam is focused and collimated for the best luminosity. Only once this is done may the collisions restart. Each LHC fill is labelled by a unique number sequentially. We have reached Fill 10000 in August 2024.

2.2 The CMS detector

As previously discussed, the LHC is of the shape of an octagon and each straight section serves a different purpose, see Figure 2.4. The beam tubes cross paths on

four points along the ring. At these points collisions can be provided for LHC's large experiments. These experiments are:

ALICE A Large Ion Collider Experiment

ATLAS A Toroidal LHC Apparatus

LHCb Large Hadron Collider beauty experiment

CMS Compact Muon Solenoid

Of these, ATLAS and CMS are general-purpose detectors. They collect data for performing precision Standard Model measurements and in pursuit of the underlying theory beyond the Standard Model (BSM) searching for hints of SUSY and other exotic particles, for example, predicted by SM extensions with extra dimensions.

LHCb is an asymmetric construction optimized for flavor physics, while ALICE concentrates on heavy-ion physics. That being said, all four detectors have a diverse physics program and are constructed in a way to complement each other in many aspects. The big experiments act as independent projects with a carefully controlled flow of information between them, to guarantee independent results and the possibility to perform cross-checks.

CMS is located at LHC interaction point 5 (IP5). Its acronym stands for Compact Muon Solenoid. This name honors the magnet-system design, enabling the apparatus to be much smaller than ATLAS, as well as the large emphasis given to muon detection. CMS is only 25 m long and 15 m in diameter while ATLAS is 46 m long and has a diameter of 25 m.

As every general-purpose particle detector, CMS consists of multiple layers each of which has an important role in identifying particles and measuring their energy and momentum. The structure of the detector is illustrated in Figures 2.5, 2.6, 2.7 and 2.8.

CMS DETECTOR

Total weight : 14,000 tonnes
 Overall diameter : 15.0 m
 Overall length : 28.7 m
 Magnetic field : 3.8 T




Figure 2.5: A computer-generated cutaway image of the CMS detector with the main units marked. [56].




Figure 2.6: A slice of the detector's cross section with the typical trajectories of the final state particles projected on the image. A distance scale is shown on the top and the direction of the magnetic field is marked in two locations. Image sourced from [57].

Outgoing particles, created at the collision point, are produced with their momentum pointing in all possible directions, and hence CMS aims to provide a maximal coverage of the complete 4π solid angle. CMS uses a Cartesian coordinate system, shown in Fig. 2.7. The transverse plane (transverse to the beam direction, also called XY plane) is of special significance in the construction as here the vectorial sum of the momenta of the outgoing particles is zero. Studying this quantity enables the inference of the transverse direction of particles that escaped detection but carried a considerable amount of momentum. The invisible momentum component is called the missing transverse momentum (p_T^{miss}) or its absolute value is also referred to as the missing transverse energy (\cancel{E}_T or MET).




Figure 2.7: The CMS coordinate system with the angular coordinates of a vector originating from the interaction point (IP) shown. A not-to-scale beam pipe is also shown in blue.

In the absence of electromagnetic fields, particles would take a linear trajectory. This motivates the use of a spherical coordinate system for describing their initial momenta. Such a choice of system is also indicated in Figure 2.7. In practice, we do not use the polar angle θ , but instead, the pseudorapidity which is defined as

$$\eta = -\ln\left(\text{tg}\frac{\theta}{2}\right). \quad (2.1)$$

Pseudorapidity, in the ultrarelativistic limit, converges to the relativistic invariant rapidity

$$y = \frac{1}{2} \ln\left(\frac{E + p_{\parallel}}{E - p_{\parallel}}\right) \quad (2.2)$$

where p_{\parallel} is the longitudinal momentum component. This invariance makes η more suited to define an angular proximity between trajectories than θ . Since $\frac{d\eta}{d\theta}\bigg|_{\theta=90^\circ} = 1$, pseudorapidity and the azimuthal angle ϕ can be treated on equal

footing, giving rise to the angular distance metric $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$. Figure 2.8 illustrates the correspondence between θ and η .




Figure 2.8: A schematic view of the CMS detector sliced with an R - z plane. The figure shows only one quadrant. The hadron forward calorimeter and the tail catchers of the HCAL are not shown. The main subsystems are labeled unambiguously except for MB and ME which stand for the barrel and endcap muon trackers respectively. The red lines correspond to the pseudorapidity noted at the end of each line. The equivalent angles are noted in round brackets. [58].

The particles directly observed by the detector include photons, electrons, muons, and hadrons. Figure 2.6 shows a slice of the cross section of CMS and the typical trajectories particles follow inside the detector.

Most particles are not observed directly due to their short lifetime. The presence of these particles can be inferred by observing the properties of their decay products and reconstructing their invariant mass, for example. Other particles interact too weakly and hence they leave the CMS detector undetected. Such are neutrinos which regularly escape the detector without leaving any trace.

The next sections briefly introduce the main subsystems of the CMS detector which are organized into layers. Starting with the innermost layer, these are the

tracker, the electromagnetic calorimeter (ECAL), the hadron calorimeter (HCAL), the magnet, and the muon system.

2.2.1 The solenoid magnet

The CMS detector was built around its magnet-system [59] using it as its defining feature. This system consists of the world's largest operational superconducting solenoid and the iron yoke that guides the magnetic flux outside the solenoid. The latter stops the field from reaching dangerous levels in the experimental cavern and is reused to house the muon system.

The hollow cylinder-shaped magnet is 6 m in diameter which is large enough to accommodate the tracker, the ECAL, and the HCAL in the center, while the muon detectors, installed outside of the coil, fit in between the layers of the yoke. With this design, the muon detectors see a flipped magnetic field with respect to inner subsystems. A major advantage of this arrangement is the reduced amount of material traversed by particles arriving from the interaction point before reaching the calorimeters, and thus the minimization of the energy loss upstream.

As the superconducting magnet is the skeleton of CMS, the individual subdetectors have also accommodated cylindrical shapes as seen in Figures 2.6 and 2.8. This results in a natural division of the detector into a barrel and two endcaps.

The role of the magnetic field in these detectors is to bend the paths of the charged particles. Due to the Lorentz force, charged particles follow a helix trajectory in a homogeneous magnetic field, which is a superposition of a drift parallel to the field and a circular motion perpendicular to it. Examining the projection of the paths onto a plane perpendicular to the field, arcs are seen. An example of this is shown in Figure 2.6. As the radius of these arcs in a field of B strength can be expressed as

$$r = \frac{p_{\perp}}{qB}, \quad (2.3)$$

given r we can calculate the transverse momentum p_{\perp} of a particle with a charge q .

The magnet of CMS was designed to deliver a field of 4 T, twice the strength of the field in the inner detector of ATLAS. This means that the cyclotron radius (Eq. (2.3)) in CMS is halved compared to that in ATLAS and hence the detector's

dimensions can be scaled down proportionally preserving the same potential accuracy. This size difference is captured by the word "compact" in the detector's name.

When the detector was brought online, a decision had been made to use a decreased field strength, 3.8 T, in order to maximize its lifetime.

2.2.2 The beryllium beam pipe

Strictly speaking, the beryllium chamber belongs to the LHC. This is the intersection point of the two beam pipes of the collider and the collisions happen in these chambers.

Beryllium was chosen for this application for its low atomic number which guarantees that for high-energy particles, it is almost entirely transparent. This property is often referred to as having a small material budget, and in detector design, it is one of the leading considerations when choosing technologies.

Another key feature is its mechanical stiffness which allows for a powerful vacuum inside the tube while the outside is at atmospheric pressure. Finally, it is a non-magnetic material, which is crucial for handling its operating conditions.

2.2.3 The silicon tracker

The tracker [60] is the innermost component of CMS. Being closest to the collision region it receives the biggest radiation flux of all subsystems. Detectors and electronics in this region are required to have an extreme radiation hardness, high spatial resolution and, to keep up with the LHC, a fast readout.

The role of any tracker is to record points along the trajectories of charged particles when they pass through the sensitive material of the detector. Once the data on the individual crossings are gathered, computers reconstruct the path each particle took. Based on the geometric parameters of the track, the corresponding transverse momenta can be calculated, and the path can be extrapolated toward the beam line to see where the particles originated from. The source can either be a high-energy particle collision, beam particles colliding with residual gas in the beam pipe, decay products from activated materials, or even cosmic radiation.

Neutral particles are usually unaffected by the tracker but photons may convert to electron-positron pairs in its material or hadrons can induce nuclear reactions which are both stochastic processes and could even start early EM or hadronic showers. These phenomena, just as multiple scattering or Bremsstrahlung, are not desirable in a tracker.

To minimize these effects trackers are built of low-density materials and using the least possible amount even of those. CMS uses silicon for this purpose. The tracker is made up of two qualitatively different kinds of units, the pixel and the strip detectors, which both work very similarly to digital camera sensors. These are organized in a way that particles travel through 4 pixel and 10 strip layers before reaching the ECAL in the barrel, and 3 pixel and 12 strip layers in the endcaps. Before the pixel upgrade in 2017, one layer less was present. This subdetector is the largest silicon tracker ever built with a length of 5.8 m and a diameter of 2.5 m which gives a coverage of $|\eta| < 2.5$. The whole area of the tracker adds up to 210 m^2 .

The area of a single pixel is $100 \times 150 \text{ }\mu\text{m}^2$ while the strip pitch ranges from $80 \text{ }\mu\text{m}$ to $205 \text{ }\mu\text{m}$ and the total number of read-out channels is 124 and 10 million, respectively. The combined resolution of the silicon tracker at the interaction point, called the vertex resolution, is $13 \text{ }\mu\text{m}$ in the x and y directions, and $19 \text{ }\mu\text{m}$ in the z direction [61]. The location where several tracks converge when traced back to their origins is called a vertex (VTX) and is usually a sign of an interaction. By localizing the origin of particles, several coincidental collisions per bunch crossing can be analyzed. The number of average collisions per bunch crossing is called pileup.

This subdetector has a significant heat dissipation but due to the nature of silicon electronics, an optimal operation requires the tracker to be kept at low temperatures. For the pixel and the strip detectors, this means about -20 C° , which is achieved with the help of CO_2 evaporative cooling and fluorocarbon cooling systems.

For high-momentum (100 GeV) tracks in the central region of the detector, the transverse momentum resolution approaches the percent level, and the impact parameter resolution is about $10 \text{ }\mu\text{m}$.

2.2.4 The electromagnetic calorimeter

Calorimeters use a destructive method to measure energy. Incoming electrons and photons induce electromagnetic showers in the high-density material of the calorimeter while hadrons initiate a hadronic shower. The shower particles are then detected as they excite or ionize an active area of the detector which may or may not be the same as the high-density material in which the shower is created. These designs are distinguished accordingly as homogeneous or sampling calorimeters. While the former solution provides better accuracy in energy reconstruction, the latter allows for better segmentation for path-finding and separating particles.

The ECAL [62], which stands for electromagnetic calorimeter, is a homogeneous calorimeter containing 78,000 PbWO_4 scintillating crystals covering up to $|\eta| < 3$. The length of the crystals has been chosen in a way that showers created by photons and electrons are contained in the crystal while hadrons pass through it with minimal interactions to ultimately reach the HCAL. The crystals are typically $25X_0$ long, where $X_0 = 0.89$ cm is the radiation length of PbWO_4 .

The radiation length is the length in which an electron's energy is decreased e fold due to Bremsstrahlung, i.e. radiating away photons. The free path length of an energetic photon before it undergoes a pair production is comparable, approximately $9/7X_0$, hence this scale provides the characteristic size for electromagnetic showers.

The shower particles lose their energy exciting the crystal once their energy is small enough. The resulting scintillation light is then collected by avalanche photodiodes (APD) in the barrel region and vacuum phototriodes⁹ (VPT) in the endcaps. PbWO_4 has relatively low light yield ($\sim 1/100$ th compared to NaI). The amount of light produced is proportional to the energy of the particle but is highly temperature dependent ($-2\%/^\circ\text{C}$). The temperature of the ECAL, containing 100 tonnes of crystal, is maintained with an accuracy of 0.1°C at 18°C .

This construction allows for a designed energy resolution of

$$\left(\frac{\sigma_E}{E}\right)^2 \approx \left(\frac{0.12}{E}\right)^2 + \left(\frac{3\%}{\sqrt{E}}\right)^2 + (0.5\%)^2 \quad (2.4)$$

where E is measured in GeV [63]. The first term represents the electric noise of the detector, the second stands for stochastic processes like the shift of the first point

⁹These are used in the endcap region due to their higher radiation hardness and better magnetic field tolerance.

of interaction, while the constant term captures the lack of information on the irregularities of the detector. A major challenge is the change of the calorimeter response with radiation damage which necessitates continuous monitoring and careful calibration.

Although the CMS ECAL is a homogeneous calorimeter the Molière radius¹⁰ of the crystal is only 2.2 cm which provides a fine granularity in the transverse plane. The crystals themselves cover a solid angle of $\Delta\Phi\Delta\eta = 0.0175 \times 0.0175$ in the barrel region ($|\eta| < 1.48$) while the endcap region varies from 0.0175×0.0175 to 0.05×0.05 ($1.48 < |\eta| < 3$).

The PbWO_4 crystals provide a relatively quick signal which makes them suitable for trigger input.

The endcap region of the ECAL is also equipped with a so-called preshower subdetector (only in $1.65 < |\eta| < 2.6$) which is there to aid the distinction of high energy single photons and closely parallel photon pairs usually originating from decays of strongly boosted π^0 particles. The preshower unit consists of two lead absorbers ($2X_0$ and $1X_0$) and two silicon detector layers with 1.9 mm strip pitch which gives a much finer resolution than the 3 cm by 3 cm PbWO_4 crystals in this part of CMS. The silicon area of this subdetector is 16.5 m^2 built from 4300 units.

2.2.5 The hadron calorimeter

Since the electromagnetic radiation length (X_0) is much less than the similarly defined hadronic interaction length $X_0 \ll \lambda_I$, the energy of the hadrons is measured in a subsequent layer.

The hadronic calorimeter (HCAL) [64] is a brass-scintillator sampling calorimeter covering up to $|\eta| < 3$. A sufficient amount of absorbing material is crucial for the operation of HCAL to contain hadronic showers. This is however not easy to guarantee. In the barrel region ($|\eta| < 1.39$), just outside the magnet, a tail-catcher ($|\eta| < 1.3$) detects and measures the energy of high-energy showers that are not fully contained and thus escape the HCAL.

Hadron calorimeters suffer from many problems that degrade their angular and energy resolution.

- Hadronic showers are much wider than EM showers due to nuclear

¹⁰The radius of a cone that encompasses 90% of the energy of a fully contained electromagnetic shower.

reactions between the calorimeter and the incoming hadrons (on top of multiple scattering due to electromagnetic processes).

- Energy used for breaking up nuclei in the detector material is not measured.
- Many particles can be absorbed in the brass layer before reaching the active region.
- Stable neutral particles may escape the detector unnoticed.
- Muons created in hadronic showers escape the HCAL.

Due to these difficulties, the CMS HCAL has an energy resolution of

$$\left(\frac{\sigma}{E}\right)^2 \approx \left(\frac{100\%}{\sqrt{E}}\right)^2 + (5\%)^2 \quad (2.5)$$

where again, E is measured in GeV.

The scintillation light of the active layer is collected by wavelength-shifting (WLS) light guides and is transported by clear fibers to hybrid photodiodes (HPDs). This signal is summed up through several layers, called the tower, giving the energy estimate for the particles.

The $3 < |\eta| < 5$ region is covered by the hadron forward calorimeter (HF, see Figure 2.5). This part needs to be more radiation hard than the barrel or the endcaps and therefore it uses iron as the absorbing material and quartz fibers as the active region. The Cherenkov light emitted by the quartz fibers is read out by photomultipliers.

2.2.6 The muon system

The muon tracker's [65] purpose is similar to that of the silicon tracker but it only detects muons as, due to their unique physical properties¹¹, muons are the only charged particles that reach this layer. The designed coverage of the CMS muon system is $|\eta| < 2.4$.

Outside of the solenoid, the spatial resolution of the components is much less of a concern due to the longer lever arm. Therefore larger and cheaper but reliable and accurate gaseous detectors are used there.

¹¹Ultrarelativistic muons do not suffer from Bremsstrahlung as much as electrons due to their high mass. In the 1-100 GeV range, where most LHC muons fall they primarily lose energy as minimum ionizing particles in the detector's material. Also, they have a half-life of 2.2 μ s which allows for an almost 5 km flight before their decay (at 0.99c).

The system consists of drift tubes (DT) and resistive plate chambers (RPC) in the barrel region placed in between the layers of the flux-return yoke ($|\eta| < 1.2$), and of cathode strip chambers (CSC) and RPCs in the endcaps, as seen in Figure 2.8. Due to the high magnetic field and the good spatial resolution the muon system can measure momenta with great accuracy over a large kinematic range, the resolution for muons with transverse momenta up to about 100 GeV is 1% in the barrel and 3% in the endcap regions.

The RPCs are extremely fast (ns) detectors suitable for providing a well-segmented signal for the trigger independently from the other two types of detectors which also take a role in the trigger system.

2.2.7 The trigger system

When searching for new phenomena we are mainly interested in collisions with large momentum transfer or one might have extra requirements for the final states such as the presence of an energetic photon. Unfortunately, interesting events satisfying these requirements may be quite rare and necessitate so many collisions that not all of them can be recorded. This makes some automated real-time (online) selection necessary. This task is what the trigger system was created to accomplish. The frequency of bunch collisions at the LHC is 40 MHz and CMS generates approximately 1 MB of data for each such collision. The trigger aims to decrease the number of events to be written to tape to about a thousand per second.

The list of requirements set for an event for it to get saved are called trigger conditions, the process that verifies whether these hold is called the trigger. A good trigger system fulfills four criteria:

- Fast. The electronics need to keep up with the 40 MHz collision rate and dead time¹² needs to be minimal.
- Flexible. Setting new filters for a new kind of event topology should not be a problem.
- Redundant. The events of interest comply with several conditions and fulfilling any of these conditions results in saving the data.
- Efficient. Has a very low false negative rate while it considerably reduces the number of events to be recorded.

¹²The time for which DAQ (data acquisition) is down after a positive decision of the Level 1 trigger.

- Well understood. Designed in a way that allows careful study of the trigger performance, as an uncertainty in the trigger efficiency affects the precision of the physics measurements, potentially for a large set of analyses.

Of course, all triggers introduce some bias by saving only those events that match our requirements. To build an unbiased or minimum biased data set, triggers often record some events based on random selection or with very loose "zero bias" criteria only requiring the presence of colliding bunches. For lowering rates but with still a "minimum bias", some activity with very low energy thresholds might be required in the detector. This then can be used for commissioning and performance studies.

In some cases, more common processes are also of interest. For this purpose prescaled triggers are introduced, where only a certain fraction of the collisions satisfying the conditions are saved.

The CMS trigger [66, 67] is divided into two different levels. After each collision the data registered by the detector is locally stored in the read-out electronics until the extremely fast Level 1 trigger, running on custom-built target-oriented but largely programmable hardware (e.g. an ASIC or an FPGA), makes a decision of either dumping the data or forwarding it to the Level 2 (or High-Level) Trigger (HLT). The Level 1 trigger uses coarse information from the ECAL, HCAL, muon subsystems – as well as some smaller systems, like the Beam Pickup Timing eXperiment (BPTX) that monitors the bunches arriving to CMS – to look for signs of interesting physics. This step reduces the 40 MHz collision rate to about 100 kHz in 3 μ s.

In case of a positive Level 1 decision, the data arrives to the HLT, a farm of commercial computers where the full event is reconstructed, in around 0.2 s on average, using somewhat simplified algorithms to check whether the requirements imposed on them hold. Provided they do, the data is recorded on tape and is later (offline) reconstructed with the most sophisticated algorithms. The output rate of the HLT is 1 kHz on average.

For the measurement reported in Chapter 4, several trigger criteria were applied based on the presence of either an energetic photon (to select the signal sample), an energetic hadronic jet (to measure the photon trigger efficiency), or an isolated electron (to collect data in a background control region). On the other hand, the precision luminosity determination using the pixel detector relies on random and zero bias triggers, while the luminosity stability investigation on

muon barrel trigger objects.

2.2.8 Object reconstruction

The information recorded by CMS is initially a list of voltages, positions, and deposited charges that require interpretation and abstraction to make it suitable for humans to explore the secrets of the collisions. The process by which this information is transformed into particles with measured energies, momenta, charges, and other properties is called object reconstruction.

In CMS all final states are reconstructed using the particle flow method (PF) [68, 69]. PF event reconstruction aims to give more accurate results by combining the data acquired by the different subdetectors thereby enhancing the angular and energy resolution.

Electrons may lose energy via Bremsstrahlung or an interaction in the tracker material, before reaching the ECAL. With the PF method, a dedicated tracking algorithm called the Gaussian sum filter is used that allows significant energy loss at each detector layer. Moreover, the energy of the resulting photons is added to the measured electron cluster energy to recover the initial values. The energy resolution for electrons with transverse momentum of 45 GeV from Z boson decays ranges from 1.6 to 5%. It is generally better in the barrel region and depends on the momentum lost by bremsstrahlung before the electron reaches the ECAL.

Muons act as minimum ionizing particles in the energy range of interest. Using the PF method the energy that the muon deposited on its track before reaching the muon system can be taken into account.

A significant fraction of photons convert to electron-positron pairs in the detector material before reaching the ECAL and hence their reconstruction is not straightforward. Using the PF method these pairs originating from a secondary vertex at a detector layer are possible to find and can be used for reconstructing the photon energy.

This procedure is of particular significance for improving the jet energy reconstruction as sampling hadron calorimeters typically have a poor energy resolution. Jets are particle showers created in the process of the hadronization of a high-energy quark, antiquark, or gluon and in the decays of the initial hadrons. The momentum of charged hadrons can be measured very precisely in the tracker

while electrons, muons, and photons are reconstructed as discussed before. At this point, the neutral hadrons, which usually carry 10% of the jet's energy, are the only ones whose energy needs to be measured by the less accurate HCAL. Doing this and then building the jets in the reconstruction phase drastically improves the energy and the angular resolution of jet measurements.

The source of jets is most often unidentifiable. This is not true for the so-called b-jets (and with less purity, for c-jets) which are induced by a high-energy b or anti-b quark. Right after the hadronization, the shower will contain a b-hadron. These have a sufficient lifetime (typically $\tau \sim 10^{-12}$ s) to travel a few 100 μm before decaying (with the average decay length depending on their Lorentz-boost: $\beta\gamma c\tau$). Note that top quarks decay instantly to Wb even before forming a hadron.

The lifetime of hadrons that contain only light quarks depends on the force which governs their decay. Electromagnetic and strong decays give a half-life between 10^{-20} s and 10^{-16} s and so these particles decay in the interaction point.

On the other hand, weak decays, like those of the charged pions or kaons (not counting the K_S^0), give lifetimes of the order 10^{-8} s which is sufficiently long to seem stable from the detector's point of view.

In general, b-jets are recognizable by their characteristic properties, in particular since some of the jet particles originate from a secondary vertex (where the b-hadron decayed). With the particle flow method, such secondary vertices can be located within a jet utilizing the high precision of the inner tracker. B-hadrons, besides increasing the invariant mass of the jet, also frequently decay, via the weak interaction, to soft leptons. They also have different kinematic properties due to the harder fragmentation function resulting in smaller charged particle multiplicity and a narrower jet. In practice several different algorithms are developed to assign a discriminator value to all jets, classifying them as b-jet, c-jet, or light jet. The higher the b-jet discriminator value is, the more likely that the jet originated from a b quark.

The reconstruction of b-jets is important for SUSY searches to identify the decays of top quarks where more than 99% goes to Wb , and Higgs bosons which have a 57% $b\bar{b}$ decay branching ratio.

The missing transverse energy is calculated as the negative vectorial sum of the transverse momenta of the reconstructed objects that belong to that particular event. Due to momentum conservation, this sum should be zero. If it is significantly different from zero, then likely some particles left the detector

without leaving a trace. This could be either due to the poor instrumentation of a detector region or the fact that the particle does not interact with the detector material.

Of course, jet mismeasurements can also cause a high p_T^{miss} signal but these can be partially suppressed by rejecting events where the angle between the transverse momentum of a jet and the missing transverse momentum vector is small. Similar problems may arise in case of object misidentification, for example, when due to a jet punch-through some jet particles escape the HCAL and are registered as muons.

High-energy leptons are also used in searches for supersymmetry. Electrons and muons are directly observed by the subdetectors of CMS as mentioned before, however, tau leptons decay much quicker ($\tau_\tau = 2.9 \cdot 10^{-13}$ s) which still allows them to travel a sufficient distance to be identified by a secondary vertex or a large transverse impact parameter track. Tau leptons, having a mass of 1.7 GeV, are the only leptons heavy enough to decay hadronically. This occurs in 2/3 of the cases (typically to one or three charged mesons, which are predominantly pions, often accompanied by up to two neutral pions and a ν_τ) while about 35% of the tau leptons decay into an electron or a muon and the appropriate neutrinos. Taus decaying in one of the hadronic channels can be found by searching for low-multiplicity, collimated jets, and decay products originating from a secondary vertex or having a large impact parameter [70].

Chapter 3

Precision luminosity at the CMS experiment

Collider experiments are foundational for modern-day high-energy physics, especially when it comes to precision measurements of the SM. Numerous institutes operate globally, conducting measurements in close collaboration with one another either through direct correspondence, via networks of the scientists working in the field, or by cross-checking their independent results.

One of the most important grounds on which crosschecks between experiments or between an experiment and theory stand is the concept of luminosity. At CMS, the Beam Radiation, Instrumentation and Luminosity (BRIL) project is responsible for real-time monitoring of this quantity, as well as for the accurate calibration online and offline of detectors contributing to the measurements. They work in close collaboration with the Luminosity Physics Object Group, which is under Physics Coordination, to provide the most accurate luminosity determination for physics studies. As part of my Experimental Physics Responsibilities (EPR), compulsory for all signing members of the collaboration, I was lucky enough to work with these overlapping teams of excellent scientists.

During the time of my association with CMS, I had a chance to do BRIL shifts in the CMS Control Room at CERN, dip my toes into data quality monitoring of a BRIL detector, gather a lot of experience with one of the more untamed and dominant source of systematic effects, the non-factorization bias, and work on most things under the Sun related to luminosity after becoming the contact person for the Run2 combined calibration paper, under the code name LUM-20-001 [71], which is currently under collaboration approval. The supporting

material of the publication under preparation is almost 200 pages, and my work included the coordination of the work of about 40 scientists who contributed in one way or another.

3.1 Cross sections and luminosity

In particle physics the result of a collision is fundamentally stochastic, one may not end up with the same result twice even if the initial conditions of the experiment were the same down to the quantum states. It is never a boring day with collisions! The probability of these processes is quantified with the help of the cross section of the process at interest denoted by σ . This quantity has a unit of area and is often measured in barns (b, $1 \text{ b} = 10^{-28} \text{ m}^2$).

A classical example considers two marbles, both with a radius of R . If our event of interest is the two marbles colliding, then their impact parameter must be $d < 2R$ for this to happen. From one marble's perspective, the area it needs to hit is $4\pi R^2$. If we consider that the outcome can not only be an elastic scattering, but the marbles may shatter with a 25% probability, we may say, that a πR^2 effective area belongs to the destructive outcome, while $3\pi R^2$ belongs to a benign bounce off. Of course, this split may depend on the center-of-mass kinetic energy of the marbles: $\sigma(E_K)$. One may rightfully argue, that the outcome must also depend on the impact parameter, but for now, this dependence is averaged over. It is also possible to examine the outcome of elastic marble scattering as a function of reflection angle (θ) and with that, the concept of the differential cross section can be introduced: $\frac{d\sigma}{d\theta}$.

In reality, elementary particles and even larger composite structures such as hadrons, nuclei, or even molecules lack the well-defined edge marbles bear, however, an effective cross section can still be associated with them for every outcome of interest.

In practice collisions are not achieved on a one-on-one basis, but by colliding beams containing a large number of particles with one another or with a fixed target. Now our focus will be on the former, as is the case in LHC.

Taking our classical example one step further, we examine the collision of two cylindrical clouds of marbles. By the cloud notion we aim to signify, that the density of the groups is small enough, that the marbles do not cover one another to a significant degree when looked at from the other group's perspective. If

the two clouds have n marbles each and a diameter of ρ , then looking from one beam's perspective the other has the effective area of $n\sigma$, the expected number of collisions is $N = n\sigma \cdot n/(\pi\rho^2)$ as the incoming n marbles are evenly distributed over an area of $(\pi\rho^2)$.

Notice that the expression for N is separable into the cross section, which only depends on the properties of our marbles (or particles), and $n^2/(\pi\rho^2)$ which only depends on the properties of the clouds (beams). This latter quantity is the toy analogy of luminosity. The larger this number is, the more of a certain interaction can be observed.

Luminosity (L), interchangeably called instantaneous luminosity, describes the rate at which outcomes with a given cross section are realized at a collider experiment,

$$\frac{dN}{dt} = \sigma L. \quad (3.1)$$

A common unit used to describe L is $\text{cm}^{-2}\text{s}^{-1}$ but $\text{Hz}/\mu\text{b}$ is also often used ($1 \text{ Hz}/\mu\text{b} = 10^{30} \text{ cm}^{-2}\text{s}^{-1}$). Similarly to our toy example, it can be computed from the beam parameters. For bunches with a Gaussian profile in the XY plane, colliding head-on, it is given by

$$L = \frac{N_1 N_2}{4\pi\sigma_x \sigma_y} n_b \quad (3.2)$$

where N_1 and N_2 is the number of particles in the colliding bunches, n_b is the number of bunch crossings per second, and σ_x and σ_y are the width of the beam in the transverse directions.

To describe the number of interactions expected to happen over a given data collection period, integrated luminosity is introduced.

$$L_{\text{int}} = \int_{t_1}^{t_2} L(t) dt \quad (3.3)$$

which is conventionally expressed in the inverse unit of cross sections. For reference, the total integrated luminosity in Run2 is approximately 137 fb^{-1} , the total proton-proton cross section is about 80 mb (as predicted by the Pythia8 event generator [72]), while the cross section of Higgs boson creation in proton collisions is about 55 pb at 13 TeV [73].

Integrated luminosity is one of the cornerstones for cross section measurements and, hence, comparisons with theory and other experiments.

To have better statistics, it is in general desirable to have a higher luminosity. The LHC was designed to provide a luminosity of $L = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with bunch

proton multiplicities of $1.1 \cdot 10^{11}$. This was reached in mid-2016 and even doubled in late 2017 by increasing the bunch intensities to $2.3 \cdot 10^{11}$ protons. This of course increases the number of interactions per bunch crossing (pileup) which in turn makes event reconstruction and data handling more challenging.

A proposed upgrade for the Large Hadron Collider, with first collisions after 2029, is the High Luminosity LHC which would ultimately deliver $L = 7.5 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (see Figure 2.1). This amounts to 15 million Higgs particles a year with an ordinary schedule of operation. The increase in luminosity would be achieved mainly through the upgrade of the beam optics to achieve higher particle density in the collision regions.

However, not only the total data volume counts for the physics program, but also the precision with which we determine its integrated luminosity. CMS has set the ambitious goal to reach a luminosity precision of 1% per year during the HL-LHC period [74] which implies an envelope of 0.6% for the absolute calibration of the luminometers. As I will show in this chapter, I played a prominent role in getting close to this very ambitious target. From the about 1.5-2% calibration uncertainty in 2017-2018 when I joined the effort, we delivered more than a factor 2 improvement, reaching 0.64% for the 2018 data sample.

This is essential for the physics success of CMS [75]. Luminosity affects all LHC measurements where the cross section of a given process is targeted, including precision tests of the standard model and searches for new phenomena, as the normalization factor for the measured event rates. It has become the dominant source of experimental systematic uncertainty in the most accurate SM physics results which include studies of W [76], Z [77], and Higgs boson [78], as well as top quark pair [79] production. Its role will only become more significant at the HL-LHC, as shown in Figure 3.1.

While with lower precision, stable luminosity information in real-time is also essential for LHC operations to optimize the performance of the accelerator, especially for luminosity leveling which in Run3 became the standard procedure to keep the peak instantaneous luminosity of the LHC at the highest level the experiments can process. Real-time luminosity is also necessary to ensure efficient data taking (e.g., by properly setting the trigger configuration of the experiment).

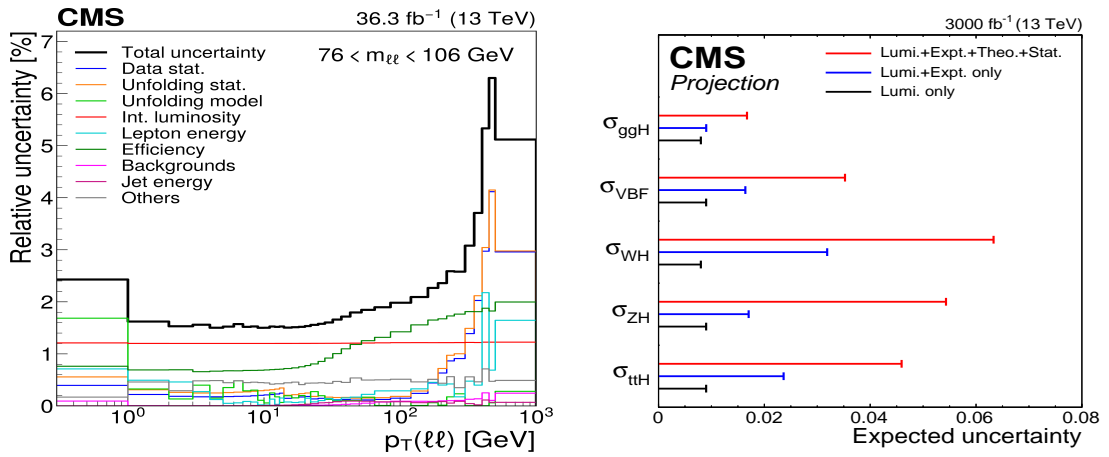


Figure 3.1: The left-hand side panel shows the relative systematic uncertainties in the differential Drell–Yan cross section measurement as a function of the dilepton transverse momentum $p_T(\ell\ell)$ at 13 TeV energy [77]. The black line shows the total uncertainty, and the colored lines correspond to the various contributions. The luminosity uncertainty (red line) dominates below 40 GeV. On the right, the expected uncertainties in the Higgs boson production cross sections per production mode with 3000 fb^{-1} of data expected at HL-LHC are plotted: contributions from the target 1% luminosity uncertainty (black bars), from the total experimental systematic uncertainty (blue), and the total uncertainty (red) including experimental statistical (0.8% for the gluon-gluon fusion process, ggH) and theoretical sources [74].

3.2 Luminometers in CMS

When the cross section for an easily identifiable process is very accurately known, the luminosity measurement simplifies to counting these, so-called standard candle events with the best possible accuracy. The luminosity is given by $L = N/\sigma$. This is the case in electron-positron colliders, where due to the extreme precision of quantum electrodynamics (QED) the cross section of the Bhabha scattering is accurately known.

Unfortunately, for hadron colliders using standard candle processes is not a viable option for absolute luminosity measurement as the uncertainties related to the internal structure of the proton bear a heavy influence on the accuracy of all theoretical cross sections. Evidently, ion colliders also suffer from this curse.

In the CMS experiment, specialized detectors, so-called luminometers are used for monitoring L that make use of the fact that the frequency of any well-defined final state is proportional to the luminosity. These detectors either count the number of hits or other objects via some indirect means, or directly measure the hits or the amount of energy deposited in them. This makes a relative measurement possible. The locations of the CMS luminometers are shown in Figure 3.2. They include the pixel luminosity telescope (PLT), the beam condition monitor with fast readout (BCM1F), the pixel detector using cluster counting (PCC) or reconstructed vertex counting (VTX), the hadron forward calorimeter measuring the occupancy (HFOC) or the transverse energy sum (HFET), the muon drift tubes (DT) in the barrel region, and the ionization chambers of the radiation monitoring system for the environment safety (RAMSES).

The aim is to obtain rate measurements for each colliding bunch pair separately, since at an accuracy of the percent level many corrections rely on the precise properties of the bunches; hence, this is essential information for delivering accurate measurements. That being said, not all luminometers are capable of fulfilling this task, namely DT and RAMSES provide only measurements integrated over all colliding bunches of the orbit. These measurements are nevertheless very valuable for their low occupancy and excellent linearity.

The detectors can be classified as real-time (online) or offline depending on how readily their data is available. In Run2, the only offline luminometers were PCC and RAMSES (and VTX, which is not used in general), while for Run3 RAMSES became available online.




Figure 3.2: A computer-generated cutaway image of the CMS detector with the luminometers marked. [56].

A short description of each luminosity detector follows, introducing their Run2 status. Some of the most significant qualities of the detectors are mentioned in these descriptions, but the list is by far not exhaustive nor does it reflect the difficulties of obtaining a data quality sufficient for a percent level accuracy. As a reference papers [80–82] are recommended.

The pixel luminosity telescope

The pixel luminosity telescope (PLT) [83] is one of the few dedicated detectors engineered by BRIL purely for luminosity measurement. The detector is situated in the cooled volume of the tracker just outside of the pixel endcaps. It is made up of 48 silicon pixel front-end modules each having a resolution of 80 by 52 pixels and a size of 8 by 8 millimeters. The modules are organized into so-called telescopes. Each telescope contains three parallel sensor planes that are positioned close to perpendicularly to the beam pipe pointing towards the interaction point (IP). The 16 telescopes are situated in two groups on either side of the interaction point distributed evenly along a circle.

To decrease the false rate due to beam halo, activation, and other backgrounds, the PLT only counts triple coincidences where hits are registered in all three planes of a telescope. These typically correspond to a particle track originating from the interaction point. Since PLT needs to be read out at the collision frequency, 40 MHz, it is not feasible to transmit the hits per pixel in all cases, much less to perform tracking for each hit. Instead, a so-called "fast-or" signal is generated that registers a hit for the telescope if there is at least one hit in each plane. To make the selection more exclusive, only the central part of the sensors are used as active area, and the areas are shifted slightly with respect to each other to point the telescope at the IP more accurately.

Since the number of hits in a detector follows a Poisson distribution

$$p_n = \frac{\mu^n e^{-\mu}}{n!}$$

the rate measured by the detector can be deduced from the frequency of no hits (p_0) as

$$\mu = -\ln p_0. \quad (3.4)$$

To compute p_0 typically 2^{14} revolutions of the LHC are integrated

$$p_0 = 1 - \frac{N_{\text{hit}}}{2^{14}}.$$

This procedure is called zero counting, and detectors operating based on this principle are called zero counting detectors. Zero counting is less sensitive to cases where two or more separate signals overlap in the detector and produce only one reconstructed observable and thus is expected to be more linear. A drawback of this approach is that for very high luminosities the probability p_0 may get close to zero, where the logarithm becomes very steep and the method becomes increasingly more prone to statistical fluctuations and backgrounds. This scenario is called zero starvation. The rate is computed independently for all telescopes, and finally, the values are averaged.

At a high luminosity, accidental coincidences become significant and PLT starts to deviate from the linear response [83] which must be corrected for an accurate measurement.

PLT, being only 1.75 m away from the IP, is subjected to a high dose of radiation over its lifetime. This leads both to degradation of the sensor material but also presents challenges for the readout and data transfer circuitry. These issues are addressed in Section 3.3.

The beam condition monitor with fast readout

The beam condition monitor with fast readout (BCM1F) is the second purpose-built detector system using a handful of different semiconductor sensor technologies. It aims to measure both luminosity and machine-induced background (MIB) using zero counting. BCM1F is located on the same support carriage as PLT. In the 2016-17 extended year-end technical stop, the sensors were replaced, and for the rest of Run2 it operated with 10 silicon, 10 polycrystalline diamond (pCVD), and 4 single-crystal diamond (sCVD) sensors. For Run3, a full silicon-pad version was prepared and active cooling was installed to limit the radiation damage.

The BCM1F counts hits on the semi-conductor pad sensors with an excellent 6.25 ns time resolution that allows for a temporal separation between the MIB arriving with the incoming beam and the outgoing collision products. The time difference between the two signals at the detector's position, 1.8 m away from the IP, is approximately 12.5 ns.

The hadron forward calorimeter

The hadron forward calorimeter (HF) is not a BRIL detector, but a dedicated readout system was installed before Run2 that provides an independent data path to save the information necessary for luminosity measurement for each collision at 40 MHz, minimizing the dependence on the central CMS timing and control distribution and data acquisition systems. For luminosity measurement only two of the total twelve rings of the detector covering $|\eta| = 3.15 - 3.5$ is used to keep the occupancy consistent in the sub-modules and ensure better response linearity.

The signal of HF can be interpreted in two different ways. Either using zero counting for each sub-unit of the detector to calculate the average μ (called HFOC method), or measuring the sum of transverse energy (E_T) deposited in the two rings (HFET method), which is also proportional to the luminosity.

Both methods suffer from contributions to the signal from out-of-time processes (OOT), of which there are two kinds. Type 1 is due to spillover of the electronic signal into the next 25 ns bunch slot, while Type 2 is due to the activation of the material which results in an exponentially decaying residual signal after each collision. The effect is corrected by decorrelating with the single bunch response function.

HFET is more sensitive to aging due to radiation, while HFOC has a more pronounced non-linearity at high luminosity as it starts to suffer from zero starvation in this regime.

The PCC method

Pixel cluster counting (PCC) uses the fact that the number of individual hits (~clusters) in the pixel detector is proportional to the luminosity. The large area of the detector provides suitable statistical accuracy, and the low density of hits in the outer layers guarantees sound linearity. The measurement is limited by the readout which happens through the central data acquisition system and is subject to the finite rate and bandwidth of the trigger. The final measurement has therefore less statistical precision than featured by the other per-bunch capable luminometers operating at 40 MHz for the smallest time windows, however, integrated over a larger period of time the difference becomes negligible with respect to other sources of uncertainty.

For the vdM calibration (Sec. 3.4), special triggers are implemented to record high-rate unbiased data for five colliding bunch pairs only (in Run2, nine in Run3) to match the per bunch statistical precision of other detectors. This is to be compared to typically 32-144 colliding bunches in pp vdM fills, with many different bunch families defined by the history of the bunches in the accelerator chain and their collision pattern in LHC.

The PCC luminosity, unlike the so-far discussed detectors, is not available online. The data needs to undergo the reconstruction step first.

For luminosity measurement, only the most stable modules are used that operate without glitches throughout the entire year. Towards the end of 2017 many modules (installed the year before during the pixel detector upgrade) were affected by issues related to the radiation aging of the elements of the power supply chain, which were then replaced for the 2018 operation.

Like HF, the PCC method is also subject to significant contributions from OOT processes.

The vertex counting method

Counting the individual proton-proton collisions is a reliable source of luminosity in a low pileup environment, while at high luminosity, the collisions may happen too close to one another for the tracker to resolve them, thus a non-linear response due to vertex merging is observed. The vertex counting (VTX) method is therefore only a viable option at low pileup and especially in the luminosity calibration fill (see Sec. 3.4) as a cross-check of the calibration of other detectors. The primary benefit of this method is that it has a negligible background.

The muon drift tubes

DT uses the rate of short muon tracks (stubs) in the barrel muon drift tubes for luminosity measurement. This detector does not provide bunch-by-bunch values, only an orbit-integrated luminosity is available. Because of this, and since the statistical precision in the calibration fill is too low for this detector, it is later cross-calibrated to other detectors. The stability and linearity of DT are in general excellent and as such it is used as a reference in these aspects.

RAMSES

The Radiation Monitoring System for the Environment and Safety is part of a supervisory system installed at the start of the LHC era. The CMS cavern has 10 ionization chambers installed, each filled with 3 l of atmospheric pressure air, which measure the level of radiation near the CMS. The detectors are mainly sensitive to 50 keV to 7 MeV photons and report an ambient dose equivalent rate between 5 μ Sv to 500 mSv which allows for tracking the conditions during stops and operation alike.

Although RAMSES was not designed to measure luminosity, because of the low occupancy and low radiation environment, it is less affected by linearity and stability issues than most of the other detectors. Similarly to DT, this detector measures only orbit-integrated luminosity and is cross-calibrated to a reference luminometer. It is used as a reference in some of the systematic studies investigating biases in the luminosity integration step.

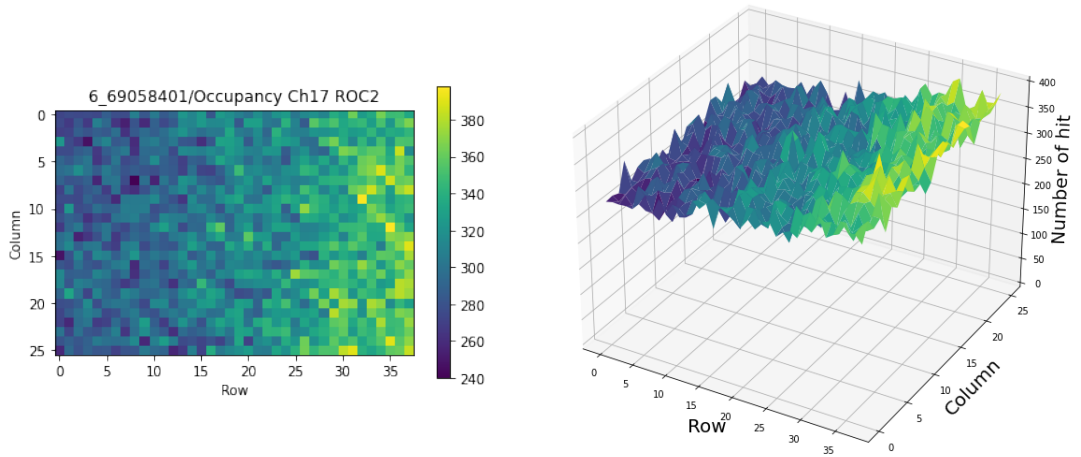
3.3 Data quality monitoring for PLT

For an accurate luminosity measurement data quality monitoring of the detectors is essential. Some of these instruments, such as the PLT operate in close proximity of the interaction point in an intense radiation area. Under these conditions, the active detector material undergoes aging due to the radiation displacing nuclei in the semiconductor bulk material leading to an increased depletion voltage, leakage current, and a decreased charge collection efficiency. At the same time, the readout electronics are also affected. Ionizing particles can flip the value of a register of microelectronics components potentially leading to undefined behaviour. Such events are dubbed single-event upsets (SEU) and are in general not permanent, a reset of the unit can fix them. SEUs are much more abrupt than the slow aging of the detector material and pose the potential danger of polluting our data if not handled correctly.

My first excursion to the field of data quality monitoring targeted the use of the PLT signal amplitude and cluster size as part of the monitoring system. I aimed to identify observables that can indicate in an automated way if aging-related performance loss requires the adjustment of the global operating conditions, such as the high voltage.

Then, in 2020 we set out to create an early warning system for such cases with Gabriella Pásztor, and our shared supervisee Milan Fejes BSc student. The report resulting from this work [84, 85] became part of the PLT detector performance paper [83].

The study was based on pixel-level diagnostic data recorded by PLT. As mentioned in Section 3.2, PLT reports fast-or data for every bunch crossing, but the full pixel-level information is also read out at 3.3 kHz frequency. I created occupancy maps for all 2016, '17, and '18 LHC fills with a five-minute integration window when collisions were available. This resulted in over three million maps, a large enough dataset for machine learning applications. Since the labeling of this dataset was not viable for us, the most appropriate unsupervised anomaly detection approach was searched for.



(a) As a 2D heat map

(b) As a 3D surface

Figure 3.3: A healthy occupancy map in different visual representations. Source [84].

Healthy occupancy maps are free of features (Figure 3.3). The bins are Poisson-distributed random numbers with an expected value that resembles a tilted flat surface. We attempted to develop a set of indicators that are sensitive to deviations from this behavior. The baseline was removed either using a surface fit or by removing the average of the neighboring pixels computed by convolution with the matrix

$$\frac{1}{8} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

The two methods can be thought of as a global or a local trend subtraction. From the resulting matrix, a set of 31 variables was computed, that were rotation, and translation invariant, and were corrected for the average hit count to remove the effect of the changing instantaneous luminosity. Among the studied quantities were the normalized standard deviation $\sigma(M)/\sqrt{\langle M \rangle}$, the normalized difference between the minimum and maximum $(\max(M) - \min(M))/\sqrt{\langle M \rangle}$, the normalized standard deviation of the column and row averages, and more. A full table of the features is found in [84].

According to the best practices, in machine learning the descriptors are standardized using their distributions in data to negate the possible scale differences between the variables. Figure 3.4 shows a pair of occupancy maps and the corresponding normalized feature vectors. Note that more extreme imperfections lead to larger normalized feature amplitudes. For training only maps with at least 50 hits were used.

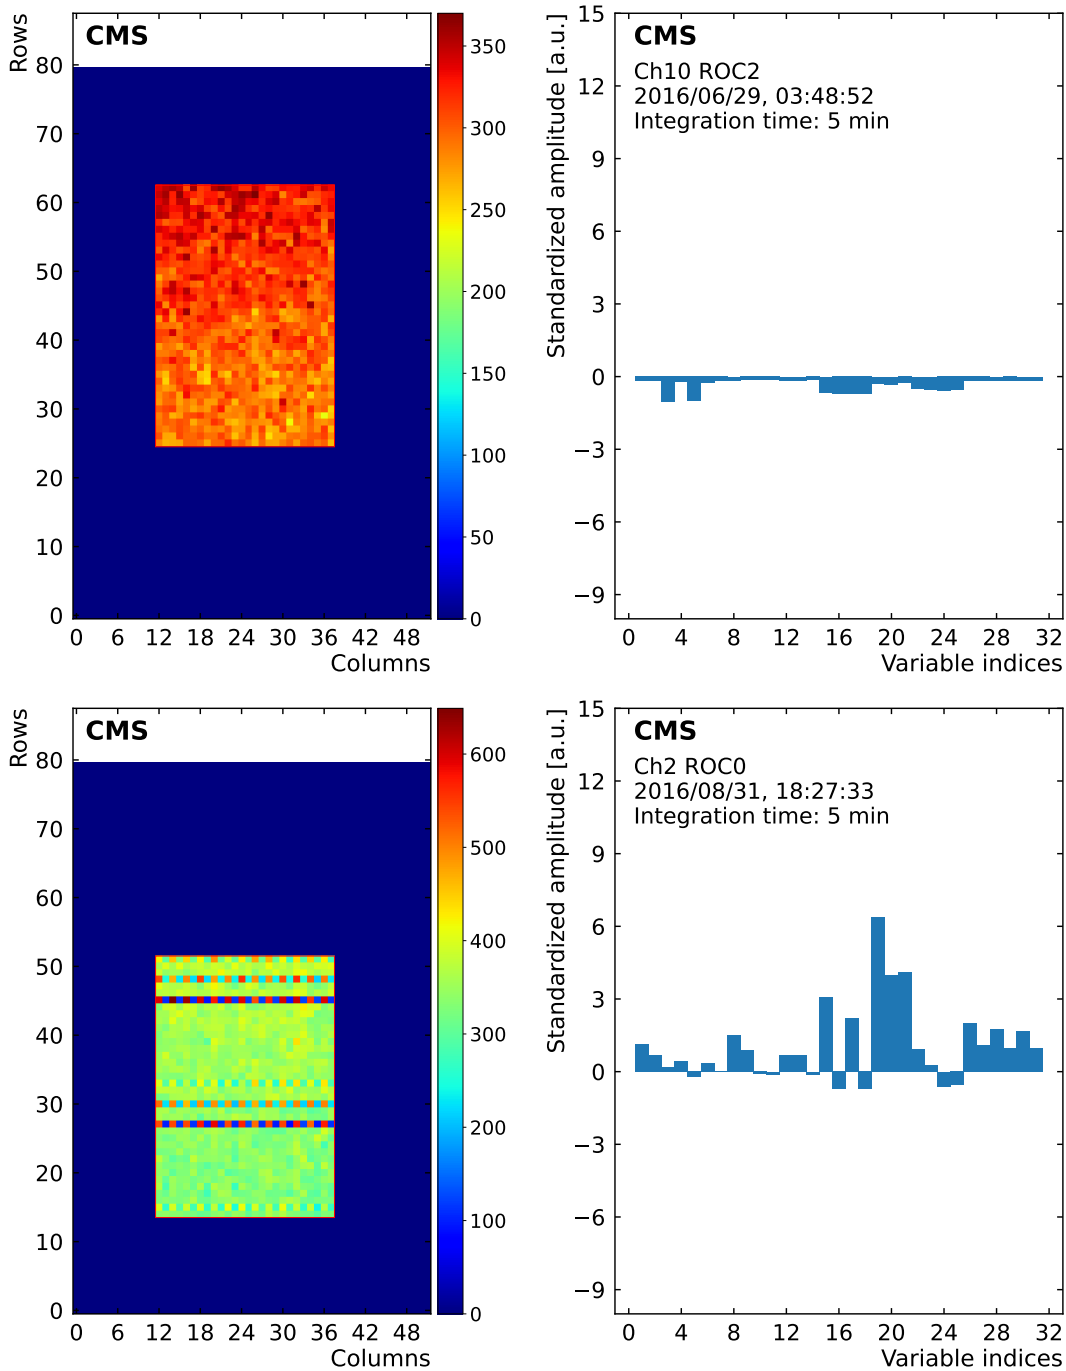


Figure 3.4: The top plots show an occupancy map for a single module during a period of good operation and the corresponding values of the 31 variables used as input to the K-means clustering. The bottom plots show the same quantities for a period when the pixel data was not correctly decoded, resulting in line errors in the occupancy plot. [83]

Finally, the K-means clustering algorithm [86] was used to sort the dataset into qualitatively different segments, with some sets belonging to a correct mode of operation while others visually demonstrating signs of erroneous behavior. The K-means operates on m -dimensional vectors, the coordinates of which are called features in machine learning terminology, but we referred to them so far as “variables”. The K-means method attempts to group the vectors into K clusters such that the combined spread of the clusters,

$$\mathcal{L}_{\text{WSS}} = \sum_c \sum_{x \in c} |x - \langle c \rangle|^2, \quad (3.5)$$

is minimized. Here, $x \in \mathbb{R}^m$ is a vector in the space of the $m = 31$ normalized discriminating variables representing a single occupancy map, the c sets, called clusters, realize a partitioning of the maps, and $\langle c \rangle = (\sum_{x \in c} x)/|c|$ is the center of the cluster.

In this method, K , the number of clusters is a hyperparameter and is to be determined by the user. Higher values of K produce more granularity in the variable space, thus increasing the chance that each cluster is homogeneous. On the other hand, the chance of splitting genuinely similar groups into multiple clusters also increases.

The most common way of data-driven determination of K is the elbow method. In this procedure, one plots the training loss \mathcal{L}_{WSS} as a function of K and picks the K value where there is a sharp flattening of the graph. This procedure makes sure that one only adds new clusters while they are most beneficial.

Our goal is a bit different from what the elbow method is meant to achieve. We wish to separate the healthy and the abnormal data as best as possible, but we are not terribly worried if similar error modes are split into several clusters. This means that we wish to increase K as much as possible without splitting the cluster of healthy maps.

In this study, therefore, a different K determination procedure is applied. It is assumed that the data is proper most of the time, hence, at least while K is relatively small, the biggest cluster is made up of healthy maps. Plotting the size of the largest cluster as a function of K , we choose the value of K , where a sudden decrease occurs.

Figure 3.5 shows the metrics of the training procedure using all three years separately and the full Run 2 data set. The resulting cluster numbers are $K_{2016} = 11$, $K_{2017} = 10$, $K_{2018} = 12$, and $K_{\text{Run 2}} = 13$, only slightly dependent on

the period.

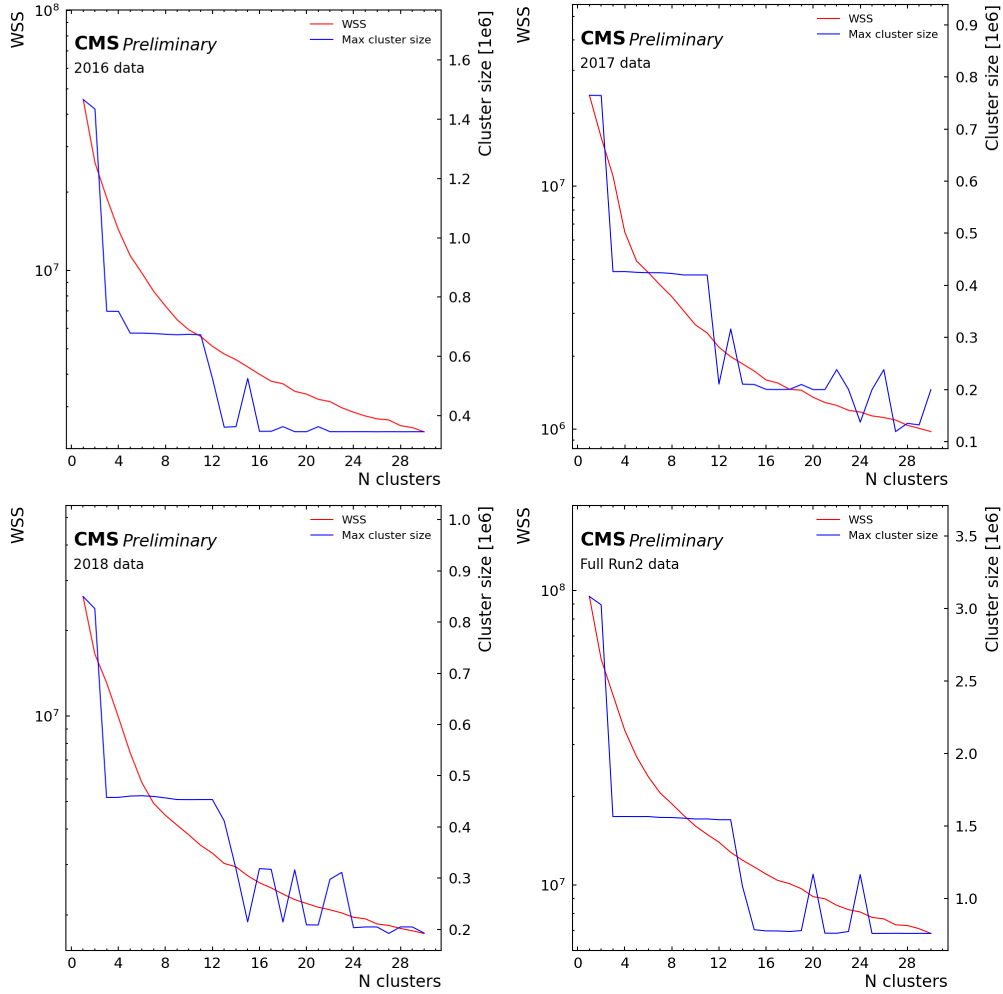


Figure 3.5: The K-means training metric \mathcal{L}_{WSS} and the size of the largest cluster as a function of K for the three years (2016-2018) separately and for the full Run 2 dataset. [84]

We evaluated the performance of the clustering by randomly plotting 100 occupancy maps from each of the resulting groups and examining the plots by eye for abnormal features. Good occupancy maps were accumulated in the largest cluster with a purity of approximately 95%. The remaining clusters contain occupancy maps with different imperfections, but no healthy map was seen in these clusters in our sampling. In the future, this performance evaluation could be made more systematically using occupancy maps artificially generated in a toy Monte Carlo simulation that I have set up.

The results of the occupancy map analysis for the global Run 2 training are summarized in Table 3.1. The description provides for each cluster a rough classification of the error mode, such as different number of pixels with very low or two high occupancy, row or column errors, crater-like defects, and other types

Cluster index	Cluster size	Fraction [%]	Description
1	1543097	50.07	95% healthy maps
2	794726	25.79	One or two point defects
3	521634	16.93	Two to approximately five point defects, a few line errors
4	103630	3.36	A large number of point defects
5	34532	1.12	Line- and crater-like defects
6	32848	1.07	Address decoding errors and gain calibration patterns
7	17969	0.58	A large number of point defects, extended inactive pixel groups
8	15755	0.51	Address decoding errors
9	10391	0.34	Address decoding errors
10	4500	0.15	Extended inactive pixel groups
11	1190	0.04	Maps where only a very few pixels have hits
12	828	0.03	Address decoding errors
13	594	0.02	Address decoding errors, maps where only a very few pixels have hits

Table 3.1: Description of the clusters resulting from using the K-means segmentation algorithm on the full Run 2 data with $K = 13$.

of issues. This method thus allows to develop an automated recovery algorithm, quickly fixing issues that could significantly affect the data quality. Other issues with only a small effect on the global performance (such as a single temporarily dead pixel) could be safely ignored.

In future iterations of the study, attempts could be made to reduce the number of features used for the training, as according to a principal component analysis (PCA) 95% of the variance of the data is contained in an 11-dimensional subspace. This could improve the resilience of the method against statistical fluctuations. It is currently not clear how much influence these errors bear on the fast-or-based luminosity, which could also be a direction for future explorations.

I implemented an online version of this analysis into the C++-based CMS data acquisition framework, XDAQ. With the classification tool deployed online, the most offending electronics problems, such as the address decoding error

visualized in Fig. 3.4 could be flagged for an operator to intervene, and after some testing, even automatic actions, such as a restart could be scheduled at the closest possible time when operations allow.

3.4 The van der Meer method

The state-of-the-art method used for calibrating luminometers and establishing the absolute luminosity scale is the van der Meer (vdM) method [87], originally proposed by the Nobel prize winning physicist for the use at the Intersecting Storage Rings (ISR) at CERN. This procedure is carried out in a dedicated fill using special beam optics to decrease the beam-related systematic uncertainties. The special settings result in smaller bunch intensities, wider and well-separated bunches, and hence decreased luminosity. While in proton-proton physics conditions, the typical pileup is above 50, it is around 0.5 in the vdM fill. The bunches undergo special preparation to shape them as close to a Gaussian profile as possible, and in most cases, a zero crossing angle is set at the interaction point between the beamlines.

Separating the beams by Δx and Δy in the transverse plane, the luminosity for a single bunch is given as

$$L(\Delta x, \Delta y) = n_1 n_2 f_{\text{LHC}} \int_{\mathbb{R}^2} dx dy b_1\left(x - \frac{\Delta x}{2}, y - \frac{\Delta y}{2}\right) b_2\left(x + \frac{\Delta x}{2}, y + \frac{\Delta y}{2}\right), \quad (3.6)$$

where n_1 and n_2 are the number of particles in the colliding bunches, b_1 and b_2 are the normalized particle density functions, and f_{LHC} is the accelerator's revolution frequency ($f_{\text{LHC}} = 11245.6$ Hz).

Integrating over $\Delta x, \Delta y$ one obtains

$$\int_{\mathbb{R}^2} d\Delta x d\Delta y L(\Delta x, \Delta y) = n_1 n_2 f_{\text{LHC}}, \quad (3.7)$$

because the integration according to x and Δx commutes and the convolution of the b_1 and b_2 probability densities is a probability density, which integrates to unity.

Since the R rate measured by the luminometers is proportional to L we introduce σ_{vis} so that $R = \sigma_{\text{vis}} L$. Using Eq. 3.7, we get

$$\int_{\mathbb{R}^2} d\Delta x d\Delta y L(\Delta x, \Delta y) = \frac{1}{\sigma_{\text{vis}}} \int_{\mathbb{R}^2} d\Delta x d\Delta y R(\Delta x, \Delta y) = n_1 n_2 f_{\text{LHC}}, \quad (3.8)$$

and hence

$$\sigma_{\text{vis}} = \frac{1}{n_1 n_2 f_{\text{LHC}}} \int_{\mathbb{R}^2} d\Delta x d\Delta y R(\Delta x, \Delta y). \quad (3.9)$$

The number of protons, n_1 and n_2 are measured using beam current transformers (DCCT [88] and FBCT [89, 90]) and are cross-checked with the longitudinal density monitor (LDM [91]), while f_{LHC} is a precisely known parameter. With this, the calibration constant can be obtained if the integral in Eq. 3.9 is known.

In case R is factorizable, i.e.

$$R(\Delta x, \Delta y) = f(\Delta x)g(\Delta y), \quad (3.10)$$

then

$$\begin{aligned} R(\Delta x, 0) &= f(\Delta x)g(0) \quad \text{and} \quad R(0, \Delta y) = f(0)g(\Delta y), \\ R(\Delta x, \Delta y) &= \frac{R(\Delta x, 0)}{g(0)} \frac{R(0, \Delta y)}{f(0)} = \frac{R(\Delta x, 0)R(0, \Delta y)}{R(0, 0)}. \end{aligned} \quad (3.11)$$

With this the integral can be written as

$$\int_{\mathbb{R}^2} d\Delta x d\Delta y R(\Delta x, \Delta y) = \frac{1}{R(0, 0)} \int_{\mathbb{R}} d\Delta x R(\Delta x, 0) \int_{\mathbb{R}} d\Delta y R(0, \Delta y). \quad (3.12)$$

In the van der Meer procedure the shapes of $R(\Delta x, 0)$ and $R(0, \Delta y)$ are measured using the so-called vdM scans. A standard vdM measurement consists of two scans where the beams are first separated in the x direction while they are kept aligned in the orthogonal direction ($\Delta y = 0$). The beams are moved symmetrically, as suggested by Eq. 3.6, from $\Delta x = -6\sigma_b$ to $+6\sigma_b$ separation in 25 steps, where σ_b signifies the beam width. The integration time at each step lasts in general around 30 seconds. The scan is then repeated with the x and y axes swapped. The results are in general normalized by the bunch current, to simplify the calculations: $R_n = R/(n_1 n_2)$.

This procedure is repeated several times during a calibration fill in order to study the reproducibility of the results. In some cases, the order of the scans is swapped, the number of steps, the range of separation, or the integration time of the steps are altered to study different aspects of the method. The calibration fill also has several other scan types that aim to constrain some of the systematics of the measurement, such as the background rate in the absence of collisions (non-collision rate), which is measured simply by separating the two beams to $6\sigma_b$ along both axes at the same time, hence setting the probability of collisions to zero. This is called the Super Separation. The structure of a typical vdM fill is shown in Figure 3.6.

The vdM calibration is normally performed once every year for each beam type and collision energy (i.e. separately for pp, pPb, PbPb, as applicable). This happens in the same special LHC fill that ATLAS uses for their calibration. While one experiment performs its scans, the other interaction points must keep their beams head-on, to avoid introducing orbit instabilities and beam-beam deflection-related systematic effects, hence the large gaps in Figure 3.6.

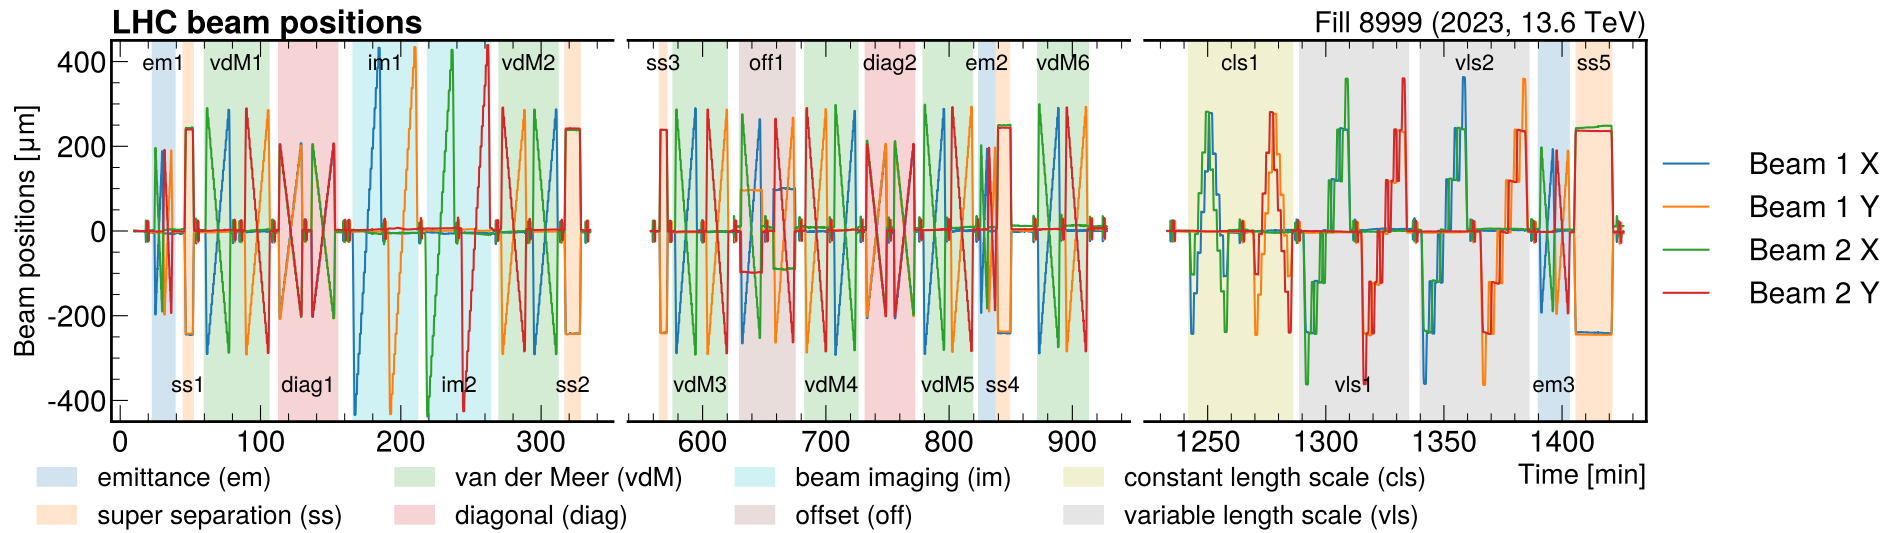


Figure 3.6: The vertical (Y) and horizontal (X) beam positions during the 2023 pp vdM program at 13.6 TeV. Several different types of scans were performed which are color-coded at the bottom of the figure. [92].

The integrals in Eq. 3.12 are calculated using analytical fit functions. After using a simple Gaussian function with a peak of P and a width of Σ to perform a fit to the normalized rates, one gets

$$\sigma_{\text{vis}} = \frac{1}{f_{\text{LHC}}} 2\pi \frac{P_x + P_y}{2} \Sigma_x \Sigma_y, \quad (3.13)$$

as in case of a perfect fit we get $P_x = R_n(0, 0)$ using the x scan.

Since this formula is very convenient, an effective width is obtained for all other fit functions $f(\Delta x, \theta)$ as

$$\Sigma_x = \frac{1}{\sqrt{2\pi}f(0, \theta)} \int_{\mathbb{R}} d\Delta x f(\Delta x, \theta), \quad (3.14)$$

where θ symbolizes the fit parameters of f . This is called the CapSigma parameterization. At the time I joined BRIL the single Gaussian (SG) and double Gaussian (DG, the sum of two SG) fit functions were implemented and widely used. Since then I introduced a number of one-dimensional fit functions to better describe the observed non-Gaussian shapes. The parameterization of these functions is discussed in Appendix A. For my work in [71], I performed a systematic survey of the implemented fit functions (see Fig. 3.7) and the fourth-order polynomial-modulated Gaussian function (poly4G, Eq. A.27) was found to give the best description of the data in 2017 and 2018. An illustration of the vdM fits is shown in Figure 3.8.

3.5 Systematics affecting the vdM calibration

While the focus of this chapter is the non-factorization bias discussed in the next section, it is essential to give a short description of some of the systematic effects impacting the vdM calibration that also affect one or more of the factorization measurements. The beam current determination, while essential for the final luminosity precision does not directly affect these and thus are not discussed here. In CMS publications all the relevant effects are expanded on in detail [80–82].

The position of the beams is an input to all vdM fits in the form of beam separation. The nominal position of the beams is specified by setting the current in the calibrated deflector dipole magnets on either side of CMS. The calibration of these magnets is not perfect, and there might be a scale difference between the true beam movement and the nominal expectation. This effect is quantified

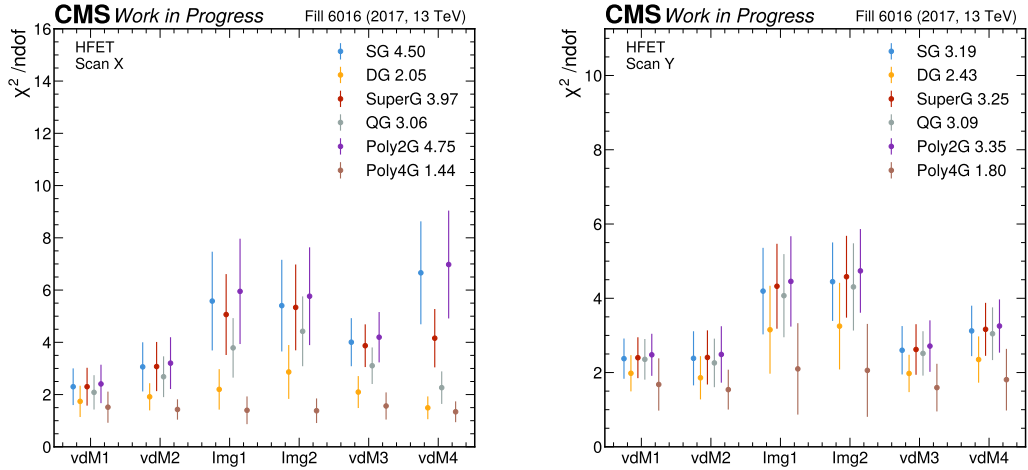


Figure 3.7: The reduced χ^2 of vdM fits for each of the scans in 2017 for the X (left) and the Y (right) scans. The fit functions used are the Single Gaussian (SG), Double Gaussian (DG), Super Gaussian (SuperG), q -Gaussian (QG), second- and fourth-degree polynomial-modulated Gaussians (Poly2G and Poly4G). The results are shown for the HFET detector. The reduced χ^2 is averaged over all BCIDs, the error bars indicate the standard deviation of the individual values. The numbers in the legend show the average over all scans. The label "lmg" stands for imaging scans, detailed in Sec. 3.7

by the **length scale calibration** that relates the nominal beam positions to the vertex position information in the CMS tracker coordinate system. In 2015 and 2016 the average length scale (LS) of the two beams was measured in the X and Y directions, while in 2017 and 2018 a new type of scan (called variable length scale in Fig. 3.6) allowed for a determination of the scale factors separately for each beam.

The beam positions are also impacted by **orbit drifts** (OD), which are uncommanded systematic or stochastic movements of the beams. In LHC, two beam position monitor systems (BPM) help determine this effect, which are however also subject to their own length scale effects and are calibrated similarly to the nominal positions. The DOROS (diode orbit and oscillations system) BPMs [93, 94] are located close to the IP on the first quadrupole magnet, while the arc BPMs [95] as the name suggests are located in the LHC arcs closest to CMS.

The orbit drift is the dominant source of uncertainty in the traditional nominal length scale calibration according to Monte Carlo closure studies I performed for [71], where a lot of effort was put into decreasing this effect.

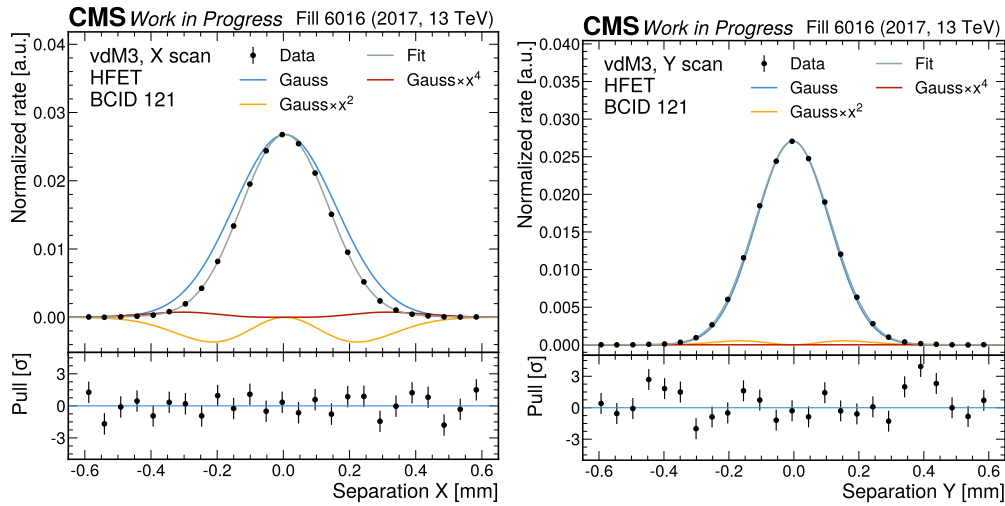


Figure 3.8: An example of a one-dimensional vdM fit using a fourth-order polynomial-modulated Gaussian function. The normalized rate is the transverse energy sum in the HF detector divided by the product of the beam currents ($n_1 n_2$). It is shown as a function of the beam separation, on the left figure for a scan in X, while on the right for a scan in Y. The bottom panels show the residual of the fit in units of the statistical uncertainty of the data. This quantity is also known as the pull. The data was recorded in the 2017 pp vdM fill at 13 TeV collision energy [71].

Finally, the beams are deflected by one another's electromagnetic fields depending on their proximity and charge distributions, which is called the **beam-beam deflection** (BB). Part of this effect is measured by the BPMs, but since they provide an orbit-integrated position, and the non-colliding bunches do not suffer from BB, the effect is diluted. The exact deflection can be computed from the beam parameters [96].

The mutual electromagnetic influence of the beams does not only modify their separations but also their charge distributions and consequently the luminosity. This is called the optical or **dynamic beta** (DB) effect and is computed using simulation [96]. The impact of the beam-beam deflection and dynamic beta effect on the vdM measurement have an opposite sign, hence a level of compensation takes place and their uncertainties are anticorrelated.

The shape of the beams also changes in isolation with time, most notably due to the slow **change of the emittance**¹ of the beams, but more complex changes

¹The transverse emittance is a measure of the average spread of particle coordinates within a bunch along a coordinate, (e.g. x), in the position - momentum, or more precisely in the $(x, dx/ds)$ plane with s being the position along the orbit. It has a dimension of length. The magnet configuration and the emittance determine the bunch width.

of the colliding particle distributions are also observed (Section 3.7). The change of emittance can be indirectly tracked with vdM scans using the measured Σ parameters. I computed the impact of this effect on the vdM calibration, and for the first time in CMS, a related correction was applied.

3.6 Emittance change

Figure 3.9 shows the evolution of the peak and CapSigma vdM parameters as a function of time. In the vdM method, the x and y scans are completed sequentially, which means that quickly changing beams open up the possibility for a biased measurement, since the parameters obtained from the x and y scans no longer describe the same $R(\Delta x, \Delta y)$. By fitting trend lines on the parameters one can interpolate the parameters to the middle of the scan and examine the relationship between

$$\sigma_{\text{vis1}} = 2\pi \frac{P_x(t_x) + P_y(t_y)}{2} \Sigma_x(t_x) \Sigma_y(t_y)$$

and

$$\sigma_{\text{vis2}} = 2\pi \frac{P_x(t_0) + P_y(t_0)}{2} \Sigma_x(t_0) \Sigma_y(t_0),$$

where t_0 is the average of the head-on times of the x and y scans (t_x and t_y , respectively). This analysis is performed for each BCID separately. The correction associated to this effect $\sigma_{\text{vis2}}/\sigma_{\text{vis1}} - 1$ is shown in Figure 3.10. In 2017, this suggests a 0.4% differential correction between vdM2, which had an inverted scan order, and the rest of the scans. This is shown in Figure 3.11. The spread of the central values is decreased after the correction. The effect of the correction is 0.15 and 0.22% in 2017 and 2018 respectively, no additional uncertainty is assigned.

Of course, the shape of the beam does not only change between the two halves of an $x - y$ scan pair, but also during a single scan. Based on simulations carried out using a single-Gaussian model the impact of this is $\ll 0.1\%$ and we disregard it.

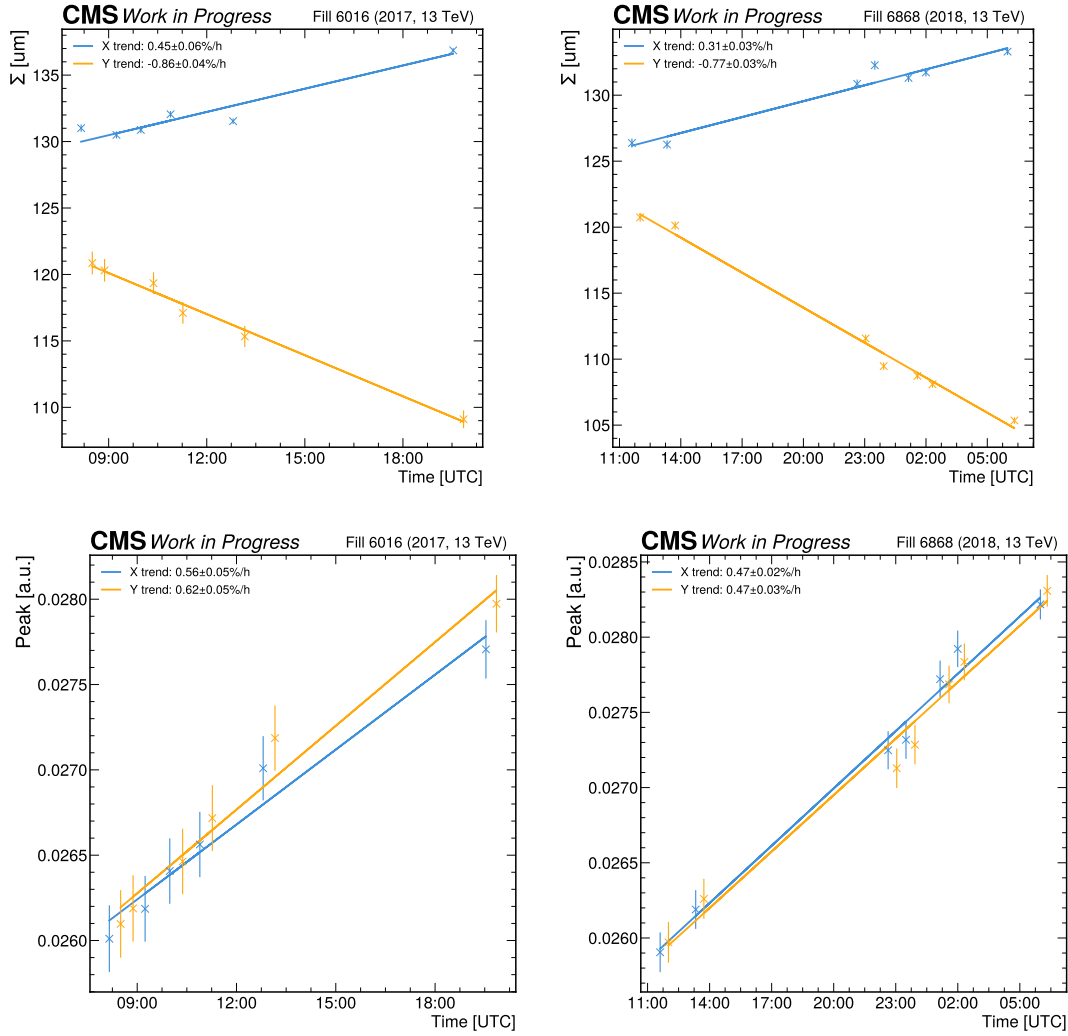


Figure 3.9: The evolution of the vdM parameters as a function of time for 2017 (left) and 2018 (right) as measured by the HFET luminometer. The markers signify the average value over the BCIDs. The top row shows the beam overlap width (Σ) while the bottom row shows the peak measured using HFET. A trend line is fit to the data to extract the rate of change.

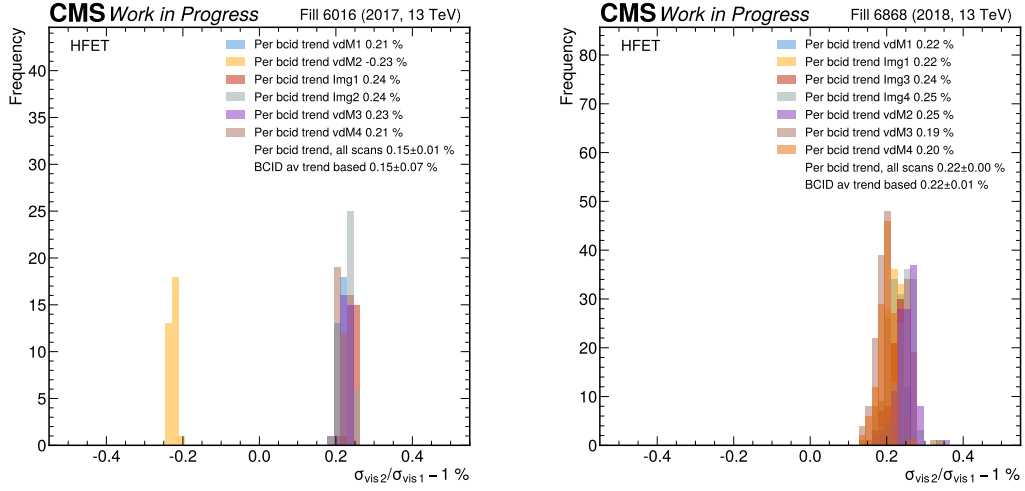


Figure 3.10: The impact of the changing vdM parameters on the calibration constant histogrammed per BCID. In 2017, $vdM2$ scan pair had a reverse $Y-X$ order, which flips the sign of the correction.

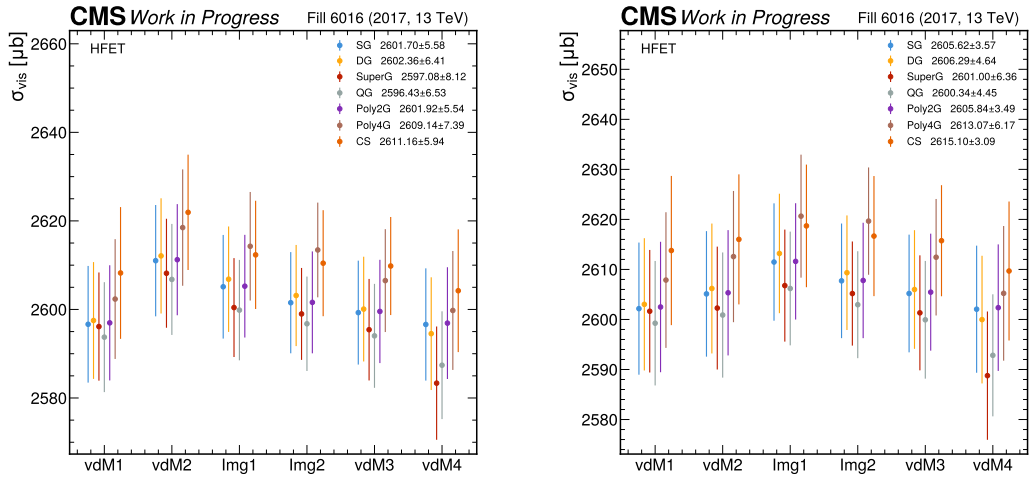


Figure 3.11: The σ_{vis} value extracted using the different available fit functions for each of the scans in 2017 before (left) and after (right) applying the emittance change correction for the HFET detector. CS stands for a cubic spline interpolator. σ_{vis} is averaged over all BCIDs, the error bars indicate the standard deviation of the individual values. The numbers in the legend show the average, as well as the standard deviation over all scans.

3.7 Non-factorization bias

In practice, the factorization assumption (Eq. 3.10) does not hold at the sub-percent precision CMS aims to achieve. The calibration constant obtained from Eq. 3.13 does not match the still exact result in Eq. 3.12 to a sufficient accuracy. Concerns in [97] were already raised in this regard. In general, but especially in the 2018 preliminary luminosity measurements [82], non-factorization was one of the most dominant contributors to the final uncertainty.

At CMS three different methods have been implemented to tame this effect. The earliest method used in [80–82, 98, 99], the imaging scan analysis [100, 101], uses a set of four vdM-like scans, the imaging scans (see Fig. 3.6). In each scan one of the beams is stationary and the other performs the scan from a typically $-4.5\sigma_b$ separation to $+4.5\sigma_b$ in 19 steps. During the scan steps, for five BCIDs (nine in Run3), high-rate zero bias data is recorded to reconstruct the primary interaction vertices. The vertex information collected in the four scans is used to reconstruct the b_1 and b_2 bunch density functions in a combined fit. Using these, the bias of the vdM measurement, expressed by the ratio from Eq. 3.12

$$\frac{\int_{\mathbb{R}^2} d\Delta x d\Delta y R(\Delta x, \Delta y)}{\frac{1}{R(0,0)} \int_{\mathbb{R}} d\Delta x R(\Delta x, 0) \int_{\mathbb{R}} d\Delta y R(0, \Delta y)} \quad (3.15)$$

can be computed exactly. Both the numerator and the denominator can be acquired from the model parameters analytically, but in practice, the denominator is obtained using the same 1D fit functions that were used in the actual vdM procedure to also account for potential mismodeling in the fits. This is carried out in simulated vdM scans.

The beam imaging method models the b_i density functions as a sum of 2D Gaussians, integrating out the Z coordinate, and supports up to three components currently. With this, highly complex $L(\Delta x, \Delta y)$ shapes can be described. However, obtaining a good quality fit is not trivial and the method achieves this by performing a large number of fits from different initial parameters. At present, efforts concentrate on improving this aspect of the procedure.

Another consideration is that this method only provides a correction for the bunch pairs that have zero bias vertex data, although in principle bunches could have a different factorization behavior depending on their history in the accelerator chain and whether they are subject to collisions at other experiments along the LHC. This is demonstrated in [102].

Moreover, since a specialized scan is required, the correction reflects the

status quo at a point in time and not during the entirety of the fill, although we have shown that the factorization evolves with time [80, 92]. Theoretically, imaging scans could be used in place of vdM scans, as they can be analyzed as two vdM pairs, but this approach has not been adopted, in part because there are uncertainties which have a more pronounced effect on imaging scans.

The other two methods, the 2D rate fit method, or offset scan method, and the luminous region method are introduced in more detail in the following two sections as these are both methods that I introduced and developed in CMS. By the time of writing this thesis, several young colleagues have picked up these methods to apply and further develop them for newer publications [71, 92, 102] and achieved spectacular results.

3.8 The 2D rate fit method

In order to avoid the factorization assumption, one may think of using the two-dimensional integral in Equation 3.12 to obtain the calibration constant, σ_{vis} . Unfortunately, sampling the entire $R(\Delta x, \Delta y)$ function is not feasible in general. This is because such scans take a long amount of time, on the scale of hours, and since the emittance and shape of the beam changes during the fill the resulting dataset may not end up being self-consistent. The most important constraint is, however, the time budget, as vdM fills are in general very busy and utilized up to the last second. That being said, a grid-like scan of $R(\Delta x, \Delta y)$ was recorded in 2024 and awaits a careful analysis.

A possible solution to this problem is a partial sampling of the off-axis separations by so-called offset scans and diagonal scans. The extra information obtained can be used to constrain two-dimensional fit functions that describe non-factorizable features.

It is important to note that this method is the only one of the three described here that does not rely on vertex information. This has two advantages. As no special trigger is required, measurements can be performed for all colliding bunch pairs and thus there is no (intended or unintended) bias from the bunch selection. Furthermore, vdM calibration data for ion collisions is collected in physics-like conditions with the beam width not inflated, as the instantaneous luminosity for ion collisions is already very low even with physics optics (well below a pileup of 0.1). This means that the characteristic beam size is close to the

vertex resolution, both being around $20\mu\text{m}$. Therefore, using vertices to perform beam "tomography" is severely limited in precision.

Offset scans were already performed in Run1 [97], originally intended for the study of non-linear effects, but in Run2 only the 2017 and 2018 proton-proton and 2015 and 2018 lead-lead vdM programs featured them. Offset scans are identical to the vdM scan, only the beams are separated to a fixed distance in the non-scanning plane as well. Diagonal scans sample the $(\Delta x, \Delta y)$ diagonally. The first diagonal scans were performed in the 2018 PbPb calibration fill based on my suggestion. Since then it has been part of all proton-proton calibration fills in Run3 as well. The beam positions during diagonal and offset scans are shown in Figure 3.6 at around 110 and 620 minutes into the program respectively. Figure 3.12 demonstrates the measured rates in the separation space.




Figure 3.12: An illustration of the offset (left) and diagonal (right) scans combined with their neighboring vdM scans using rate data recorded in the 2018 PbPb vdM fill [103]. For better legibility, a one-dimensional fit function is displayed for each scan, in green for the off-axis scans and red for the vdM scans.

3.8.1 Data consistency and rate matching

In order to perform a two-dimensional fit, the off-axis scans are paired up with the closest-in-time vdM scan. In general, this can be both before, or after the off-axis scan. Comparing results between multiple options allows for a reproducibility check while taking the average removes possible biases due to the emittance change.

In the left-hand side panel of Fig. 3.12 it is seen, that the offset scan crosses the perpendicular vdM scan. At these intersections, the VdM curve has a steep

slope while the offset scan is near its maximum which results in an increased sensitivity to a relative shift in separation parallel with the VdM scan. For this reason, orbit drifts are of primary concern for this measurement.

The separations of each scan are corrected using the standard procedure utilizing the beam position monitors, but not all the discrepancy is resolved. Performing a χ^2 minimization for the difference between the interpolated rate of the two scans in the intersection point using all BCIDs as a function of the offset distance, we are given a residual drift value. For vdM X and offset Y the χ^2 function is

$$\chi^2 = \sum_i \frac{1}{\sigma_i^2} \left(f_{i,\text{vdM}_X}(\Delta x) - f_{i,\text{off}_Y}(0) \right)^2, \quad (3.16)$$

where i is the BCID index, f is a one-dimensional fit function used to fit the vdM and offset scan rate curves for interpolation, and Δx stands for the true offset value in X. σ_i is the uncertainty that takes into account how well the VdM and offset f functions reproduce the rates on both sides of the intersection.

The procedure is performed separately for each luminometer. It is found, that these results are always highly consistent. For the 2018 heavy ion data the value of the χ^2 function and the resulting offsets is shown in Figure 3.13.

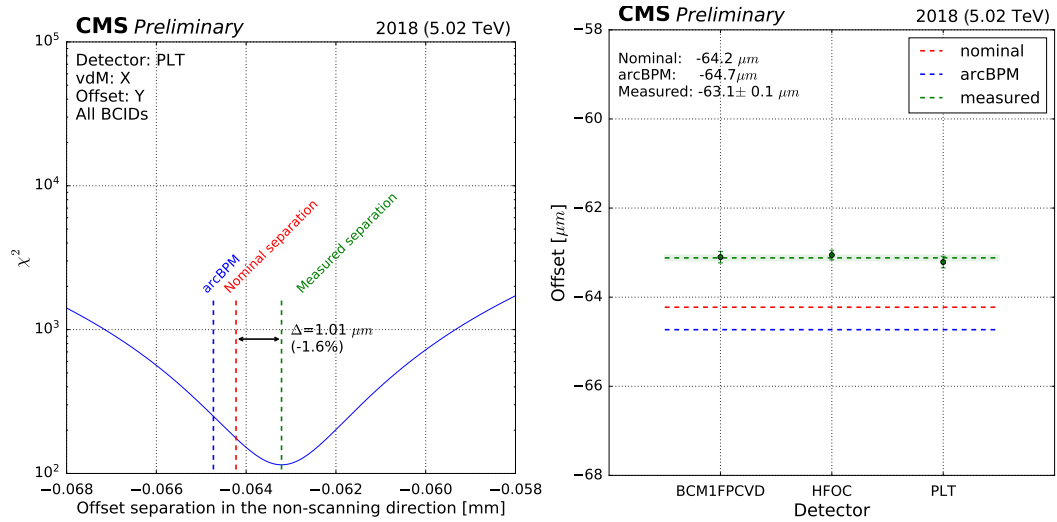


Figure 3.13: The figures show intermediate steps for obtaining the rate matching offset value. First, a χ^2 function is constructed using all BCIDs as seen in the left panel. The minimum gives the measured value for the offset. The consistency across detectors, as seen in the right panel, shows that this is not a detector or a statistical effect.

With this method the misalignment along the vdM scans can be estimated, however, no intersection between the scans is sensitive to the perpendicular

direction. For this reason, the average of the two differences derived in this procedure is taken as the characteristic size of the residual misalignment of the offset scans in all directions, and this is handled as a systematic uncertainty. In case the difference is found to be s , then it is taken into account by shifting each offset scan with respect to its corrected (or nominal) position by $\pm s$ in Δx and Δy , creating eight systematic versions in total, and repeating the subsequent 2D fits for each.

Diagonal scans only meet near the peak, where they are not sensitive to such shifts at all. The misalignment systematic derived from the offset scans is applied to the diagonal scans as well.

3.8.2 Fit functions

The simplest two-dimensional function applicable to beam shapes, that incorporates a level of non-factorizability is the 2D Single Gaussian (SG) function with a correlation term, conventionally parametrized as

$$SG_{2D}(x, y) = \frac{V}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}} \exp\left(-\frac{r^2}{2}\right), \quad (3.17)$$

where V stands for the integral of this function and $\rho \in (-1, 1)$ is the correlation parameter. The exponent is defined using

$$r^2 = \frac{\left(\frac{x-\mu_x}{\sigma_x}\right)^2 + \left(\frac{y-\mu_y}{\sigma_y}\right)^2 - 2\rho\frac{x-\mu_x}{\sigma_x}\frac{y-\mu_y}{\sigma_y}}{1-\rho^2}, \quad (3.18)$$

with σ_i and μ_i being the width and the center for $i \in \{x, y\}$. The contour lines of an SG function are ellipses. A non-zero value of ρ corresponds to a stretch and the rotation of the contour lines in a way that their main axes no longer align with the coordinate axes. This effect is shown in Figure 3.14.




Figure 3.14: The effect of the ρ parameter on the contours of a Single Gaussian function or any radially parametrized function (left) and of the p parameter on the rate in a Super Gaussian function (right). All curves on the right are normalized to have the same amplitude and integral.

For this function, the non-factorization bias can be determined analytically,

$$\sigma_{\text{vis}}^{\text{true}} / \sigma_{\text{vis}}^{\text{vdM}} = 1 / \sqrt{1 - \rho^2}. \quad (3.19)$$

Notice that the correction in this case is always a positive number $\sigma_{\text{vis}}^{\text{true}} > \sigma_{\text{vis}}^{\text{vdM}}$. For small values of ρ the correction can be simplified by Taylor expansion to $\sigma_{\text{vis}}^{\text{true}} / \sigma_{\text{vis}}^{\text{vdM}} \approx 1 + \rho^2 / 2$.

In the 2018 proton-proton calibration, the ρ parameter was found to be only around 0.05 with no discernible pattern across the BCIDs. At the same time, since this fill had two offset scans approximately 12 hours apart, a change in the ρ value could be established as shown in Figure 3.15. This measurement directly observed a beam shape change beyond the emittance evolution for the first time at CMS.

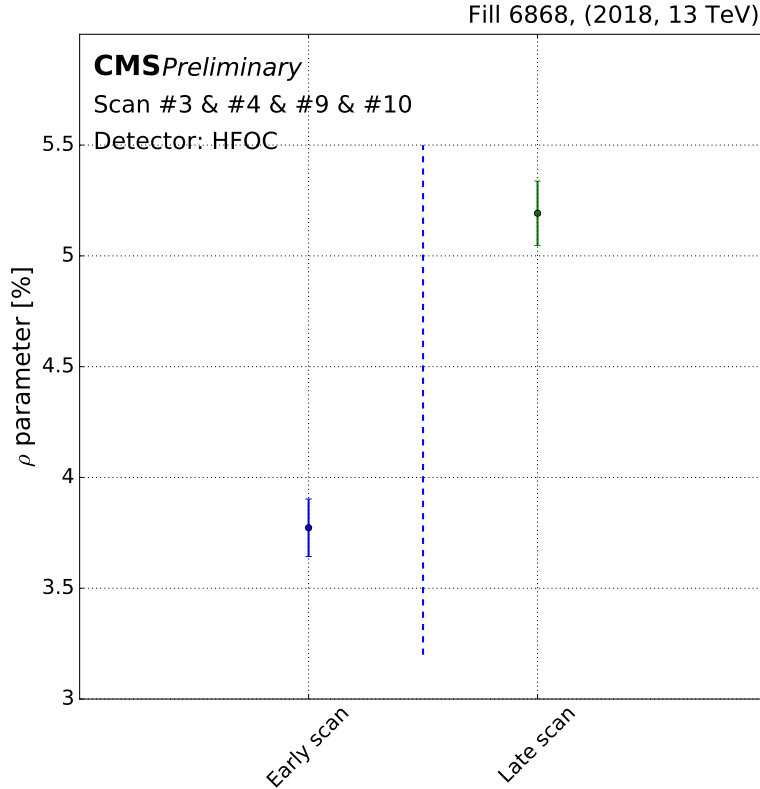


Figure 3.15: The first direct observation of beam shape evolution at CMS. The plot shows a definite change in the linear correlation parameter of the single Gaussian fit function with data recorded in the 2018 vdM calibration fill. The time difference between the two scans is approximately 12 hours. [104]

A big drawback of the SG model is, that it does not account for a difference in the width of the offset and vdM scans, which has been observed consistently. For a more accurate modeling of the true non-factorisation more complex fit functions are necessary. To this end, I have devised and created suitable parametrizations for several two-dimensional functions that are all feasible candidates to describe $R(\Delta x, \Delta y)$. The detailed definition and normalization of these functions are found in Appendix A.

A natural progression from the SG function is the Double Gaussian function, where two SG components are summed up either with a shared center or separate ones. This model is however poorly constrained by the off-axis scans and its parameters often have large uncertainties. A more stable non-Gaussian fit function is achieved by setting the radial profile of a two-dimensional shape

to follow a Super Gaussian function (see App. A.2):

$$\text{SupG}_{2D}(x, y) = V \frac{p}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}\Gamma(1/p)2^{1/p}} \exp\left(-\frac{(r^2)^p}{2}\right), \quad (3.20)$$

where V represents the integral of the function, $p > 0$ controls how flat the peak of the bell curve is, and r^2 is our familiar exponent (Eq. 3.18). Figure 3.14 demonstrates the effect of the p parameter.

As an alternative way of generalizing one-dimensional functions to higher dimensions Appendix A.2 also presents a product construction for this function,

$$\text{SupG}_{1D \times 1D}(x, y) \sim \text{SupG}_{1D}(x'; p_1) \text{SupG}_{1D}(y'; p_2), \quad (3.21)$$

where x' and y' are rotated coordinates to introduce the non-factorization. This way the two main axes are allowed to differ in terms of peak flatness.

One peculiar property of product construction is that, with the exception of the Gaussian function, the resulting contours are not elliptical. Instead, they develop four distinct more or less rounded corners. Because of this and since no different flatness in vdM scans in X and Y was observed, these functions have been neglected in Run2 studies but have been later introduced for Run3 to possibly reduce the dependence of the result on the models used.

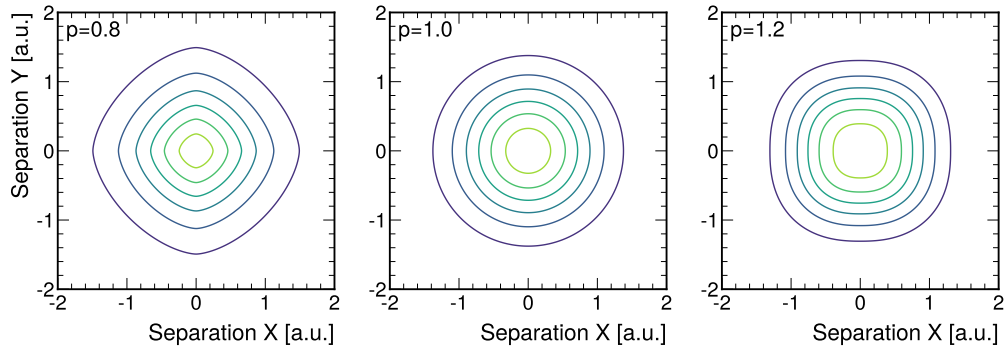


Figure 3.16: The contour lines of the two-dimensional Super Gaussian function using the product construction for different p values. The $p = 1$ scenario is equivalent to a simple Gaussian function.

In Appendix A.2 the two-dimensional generalization of the q -Gaussian and the polynomial-modulated Gaussian is also discussed. Beyond these functions, the possibility of a 2D Gaussian modulated by a 2D polynomial is explored. While this approach breaks away from the radial parametrization and gives up the elliptical contour lines, it does not necessarily suffer from the angles

product constructions exhibit either. With that, these functions have too many free parameters to be well constrained by a single set of off-axis scans but may prove useful when analyzing the grid scan recorded in 2024.

The convergence of functions is not only checked via the χ^2/dof value but also by examining the correlation of selected fit parameters across detectors. Such parameters are the widths and, for the Super and q-Gaussian functions, p and q , respectively.

3.8.3 Deriving the correction

The final corrections are derived using simulated vdM scans using the fitted function in place of $R(\Delta x, \Delta y)$. The procedure is carried out for each detector, bunch pair, and fit model independently.

All 2D fit functions are parametrized in a way that one of the fit parameters (V) is the integral of the function. The fitted shapes are sampled in the regular vdM pattern and an uncertainty is assigned to each point based on uncertainties measured in real data. This is only to establish their relative importance in the fit, the points are not randomized within these uncertainties. The curves are fit using the same one-dimensional fit function that is used in the vdM procedure. Finally, the resulting σ_{vis} is compared to the V parameter to determine the bias of the vdM measurement on the input 2D shape.

To propagate the uncertainty of the fit parameters to the correction result the fit parameters are randomized with their full covariance matrix to take the correlations into account, thus morphing the fitted shape slightly before deriving a correction in the simulated vdM scan. Repeating this procedure several times, the spread of the obtained correction gives an insight into the stability of this result.

On introspection, it was found that, in the case of using the 2D Super Gaussian fit function and the 1D SG vdM fit curve, the correction value has a maximum as a function of the p parameter of the 2D fit function at $p \approx 0.9$. Hence randomizing p around this value would result in a very small spread in the corrections, much smaller than for example a $p = 1.1$ central value with the same uncertainty would give. This introduces a potential bias favoring corrections resulting from shapes with p close to 0.9 if the spreads (as uncertainties) are used for weighted averaging. Figure 3.17 shows the correction as a function of p for several 1D

models describing the vdM fit curves. It appears that SG, the second-degree polynomial-modulated Gaussian (poly2G) and, to some extent, DG are affected. The lack of this behavior from the 1D Super Gaussian vdM profile indicates that this is a mismodeling artifact.




Figure 3.17: The multiplicative correction to σ_{vis} for a 2D Super Gaussian model of $R(\Delta x, \Delta y)$ with $\rho = 0$ as a function of p . Results are shown for the following vdM fit curves: Single Gaussian, Double Gaussian, Super Gaussian (1D), and a second- and fourth-degree polynomial-modulated Gaussian. Note, that although the 1D Super Gaussian always gives a perfect fit, the correction resulting from the intrinsic non-factorizability of the 2D shape is still similar to the poly4G corrections. Near $p = 1$ a small fit instability is seen.

3.8.4 Uncertainties

In general, this study considers the following uncertainties:

Residual misalignment: The systematic position shifts derived using the rate matching procedure are applied to the off-axis scans before the 2D fits. The spread of the resulting corrections quantifies the sensitivity of the method to this amount of shift.

Model uncertainty: Selecting a model always introduces some assumptions about how the never-before-seen regions of the $R(\Delta x, \Delta y)$ behave. Models with sufficiently large degrees of freedom to forego this problem are however not sufficiently well constrained by our sparse data. To establish a result that is more model-agnostic, the models are averaged over and their spread provides a related uncertainty. Models with clear issues, such as SG which lacks the appropriate degrees of freedom or functions with too many parameters to give stable fit results, are not used in this average.

Disagreement between detectors: As the geometry of the overlap shape is an intrinsic property of the colliding bunches, these differences are due to detector-specific effects, like background handling. A standard deviation between the results is taken.

Bunch-by-bunch disagreement: Theoretically the bunches may have different non-factorization properties and in fact, they were shown to have differences in [102]. Since the correction is applied to the bunch-averaged luminosity, the error on the mean is computed and counted with the final uncertainty as a conservative choice.

Finally, in case there are multiple scans, one must decide on a method to extrapolate the results to the scans and calculate the uncertainty of the extrapolation accordingly taking correlations between the measurements with shared scans into account. A trend line was utilized for this purpose in the 2018 PbPb calibration using two diagonal scans and an offset scan paired with their neighboring vdM scans as shown in Figure 3.18.

Departing from the previous general description, in this work the Super Gaussian result was used as a central value for its stability, and the model uncertainty was derived as the average absolute distance of the trend lines of the alternative models from the central one.

With this result, the first time-dependent non-factorization corrections were published in CMS [103].

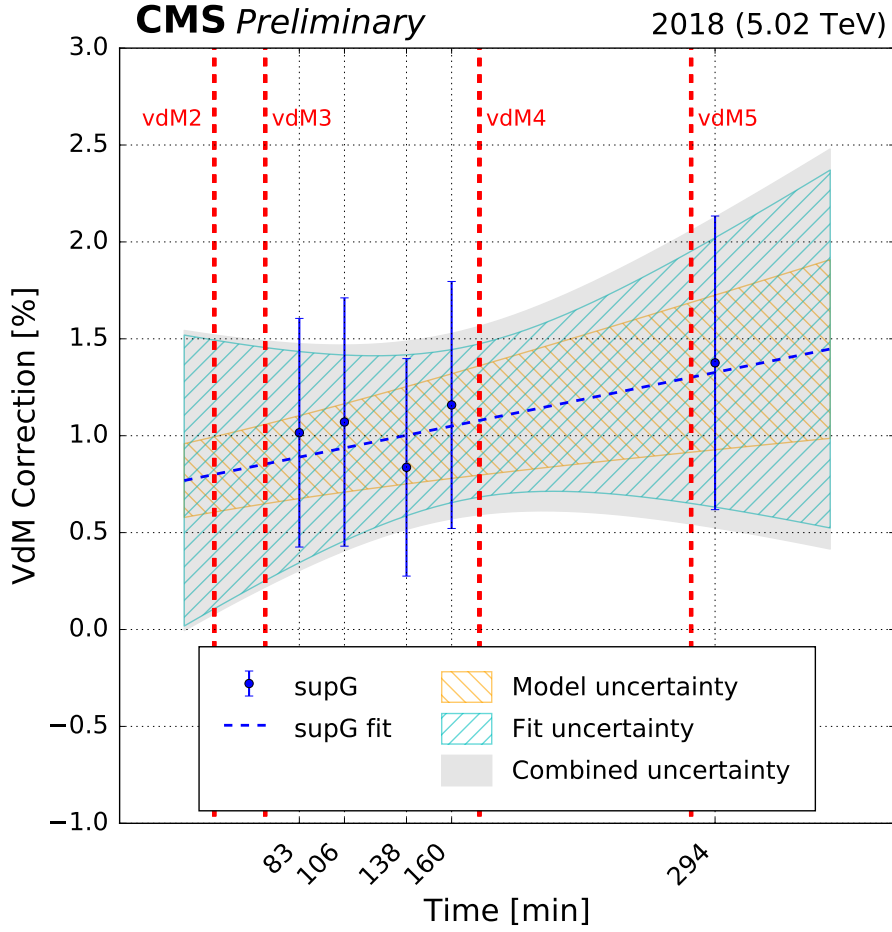


Figure 3.18: Non-factorization correction for all studied vdM +offset and vdM +diagonal scan combinations in PbPb fill 7443. The blue markers represent the Super Gaussian fit results, the error bars correspond to the orbit drift, detector and BCID-related uncertainties added in quadrature. In this plot, the BCID standard deviation is used. The dashed line shows the straight line fit to all scan combinations, with the solid light grey area denoting the total uncertainty, correlated among the scan combinations. The hatched areas denote subcomponents of the total uncertainty due to the fit (blue) and model choice (orange). The vertical lines mark the time of the vdM scans. [103]

3.9 The luminous region method

The luminous region (LR) method, first used in ATLAS and ALICE [105, 106], aims to reconstruct the bunch proton densities in three dimensions to infer the characteristics of the beam overlap shape and estimate the bias of the vdM measurement for such a shape using simulated vdM fits.

To make this reconstruction possible, the technique makes use of the three-

dimensional distribution of the primary vertices emerging at each scan step of the scan. This distribution is called the luminous region or beamspot. The luminous region data is utilized to augment the luminometer and beam position data traditionally used in vdM fits.

This approach does not depend on special "scan configurations", such as imaging or offset scans, and therefore it is capable of sampling the bias as a function of time with each scan appropriate for a vdM fit. With that being said, the need for vertex information restricts the number of BCIDs this method can obtain a correction for.

In this work, an independent derivation of the luminous region parameters is presented for general three-dimensional Gaussian proton density functions in Appendix B and for a simplified, faster-to-evaluate model in Appendix C. To handle complex fit functions and test different parametrizations, a dynamic χ^2 building and fitting framework is created with vectorized calculations. A closure test of the LR method is performed using Monte Carlo simulations to exclude possible biases and determine the stability of the delivered results. Finally, the method is applied to the Run2 proton-proton calibration fills. The results are published for 2015 and 2016 in [80] and are part of the in-approval combined 2017-2018 paper [71]. The method was applied to the 2023 proton-proton calibration by N. Rab, whose MSc research [107] on this subject I co-supervised. The results were compared to those of the 2D fit method derived by A. Fehérkuti [92]. A striking agreement is found over a large range of correction values between the two independent methods as shown in Figure 3.19. An encouraging result.

3.9.1 Input data

Beamspot parameters

In the vdM fill the CMS tracker takes zero bias data gated on five BCIDs (nine in Run3). Using this data the primary vertices (VTX) are reconstructed separately for each BCID. The vertices can be accumulated for any time window. The three-dimensional spatial distribution of the primary vertices is called the luminous region or beamspot. In our case, the vertices are collected for the duration of each scan step.

A maximum-likelihood (ML) fit with a three-dimensional Gaussian function is used to extract the nine beamspot parameters. Three of these parameters describe




Figure 3.19: The measured non-factorization correction using the luminous region method in 2023 for all scans plotted for selected fit functions (see Table 3.2 for the naming convention). The models with the best and second best reduced χ^2 value are shown in purple and gray. The green markers represent the results of the 2D rate fit method based solely on luminometer rate data. The X axis is time-ordered.

the beamspot center (μ_x, μ_y, μ_z) while the rest parametrize the symmetric covariance matrix ($\Sigma \in \mathbb{R}^{3 \times 3}$): $\sigma_x, \sigma_y, \sigma_z, \rho_{xy}$ and two tilt parameters t_{xz} and t_{yz} . The full (nine-by-nine) covariance matrix describing the fit uncertainty is also saved.

The final model is:

$$p(\mathbf{r}) = \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma + \Sigma^{\text{res}})}} \exp\left(-\frac{1}{2}(\mathbf{r} - \mu)(\Sigma + \Sigma^{\text{res}})^{-1}(\mathbf{r} - \mu)\right) \quad (3.22)$$

where $\mathbf{r}, \mu \in \mathbb{R}^3$, $p(\mathbf{r})$ is the probability density of vertices, μ represents the center, and

$$\begin{aligned} \Sigma_{xy} &= \rho_{xz} \sigma_x \sigma_y, \\ \Sigma_{xz} &= t_{xz}(\sigma_x^2 - \sigma_z^2) - t_{yz} \Sigma_{xy}, \\ \Sigma_{yz} &= t_{yz}(\sigma_z^2 - \sigma_y^2) + t_{xz} \Sigma_{xy}, \end{aligned} \quad (3.23)$$

the diagonal terms are the squares of the widths, and Σ^{res} stands for the vertex resolution, which is determined by the tracker group. This may depend on the position of the vertex or the number of tracks emerging from it for example.

By including Σ^{res} into the fit the resulting beamspot parameters are automatically unfolded with the VTX resolution, which is the primary benefit of a maximum likelihood fit for this task.

In practice, tunes used for the vertex reconstruction in the vdM fill are not perfectly applicable to the conditions, and the vertex resolution is not modeled to sufficient accuracy for the LR method. This was uncovered in a comparison of data processed in the 2016 prompt and legacy reconstruction campaigns. For this reason, the LR method uses a free-floating term to compensate for this effect.

First, note that a sensibly small deviation from the nominal VTX resolution only impacts σ_x and σ_y of all of our observables, since σ_{vtx} is only 5-7 times smaller than the beamspot width, while $\sigma_{\text{vtx}} \ll \sigma_z$. The beamspot centers are not changed and neither was there a change to be found in the tilts on comparing 2016 vertex data produced using different tunes.

This motivates the transformations

$$\begin{aligned}\sigma_x &\rightarrow \sigma'_x = \sqrt{\sigma_x^2 + \Delta_{\text{res}}^x |\Delta_{\text{res}}^x|}, \\ \sigma_y &\rightarrow \sigma'_y = \sqrt{\sigma_y^2 + \Delta_{\text{res}}^y |\Delta_{\text{res}}^y|}.\end{aligned}\tag{3.24}$$

With this parametrization both an over and an underestimation of the resolution can be compensated. The central results in the study use $\Delta_{\text{res}} = \Delta_{\text{res}}^x = \Delta_{\text{res}}^y$.

Normalized rates

The LR method also makes use of the normalized rates as measured by one of the luminometers. The normalized rates need to be corrected for the measured and properly corrected beam current, the non-collision background, but nothing else. While the dynamic beta effect is usually handled as a rate correction in vdM scans, the LR uses a different approach (see Sec. 3.9.3). The background subtraction can be also carried out by using a constant in the fit. Any per-bunch luminometer is appropriate as a data source. Usually, PCC is used for the central results which are cross-checked against PLT.

Beam positions

Initial attempts used the beam-beam separation instead, however especially for the imaging scans using the beam position leads to significantly better fit quality.

Instead of the nominal positions, the method uses beam positions as measured by the beam position monitors, which hence provides a step-by-step orbit drift correction. This is followed by the correction for the beam-beam deflection (BB) and finally the length scale is corrected using different factors for the two beams when available.

The sensitivity of the result to these corrections is tested. This is particularly important for BB since the beam position monitor picks up some of the beam-beam deflection and this way a slight overcorrection of BB could happen. However, this is outweighed by the improved fit quality that using the beam position monitor directly provides. In Section 3.9.8, the effect of all corrections and model choices are summarized and illustrated by Figure 3.31.

Bunch lengths

The analysis uses beam length information measured by the Longitudinal Beam Synchrotron Radiation Light monitor (BSRL) [108] as a new input in the combined analysis of 2017-2018 [71] to better constrain the bunch densities in the longitudinal direction. This is important, because while the longitudinal width of the beamspot constrains the bunch length via the relationship

$$\sigma_z^{\text{beam spot}} = \frac{\sqrt{(\sigma_z^{\text{b1}})^2 + (\sigma_z^{\text{b2}})^2}}{2}, \quad (3.25)$$

which holds for SG densities, there remains a trade-off-type freedom between σ_z^{b1} and σ_z^{b2} which is only very weakly constrained by other inputs if at all.

Although these parameters bear no direct influence on the derived corrections in the case of a zero crossing angle, the stability of the fit is improved by the extra input.

Since the BSRL does not measure the Gaussian width of the bunches directly, rather it computes this value based on a corrected full width at half maximum (FWHM) measurement, the lengths are not used directly in our fits. As [108] mentions, in special beam conditions, the FWHM-based procedure can lead to inaccuracies. Instead, this information is used to constrain the ratio of the Gaussian length of the bunches in the fit:

$$\Delta\chi^2 = \sum_{i \in \text{steps}} \frac{1}{\sigma_{\text{unc}, i}^2} \left(\frac{\sigma_{z,i}^{\text{b1}}}{\sigma_{z,i}^{\text{b2}}} - \frac{L_i^{\text{b1}}}{L_i^{\text{b2}}} \right)^2,$$

where $\sigma_{z,i}$ is the standard deviation of the beam in the i^{th} step in the z direction computed by projecting out the transverse directions, and L_i is the BSLR measurement interpolated to the i^{th} step, where σ_{unc} is the standard deviation of the measured ratios around the quadratic interpolation line. Interpolation is necessary since the BSLR measurement has a five-minute integration window.




Figure 3.20: The bunch length values (left) and the ratio (right) as measured by the Beam Synchrotron Radiation Light monitor (BSLR) in Fill 6016 for BCID 41. The points are interpolated using a second-order polynomial. The std value in the legend signifies the square root of the mean quadratic deviation from the interpolation line.

A histogram with a 50 ps resolution of the synchrotron light intensity is also available, however, the non-trivial transfer function of the measurement chain does not make it possible to include this shape into the fit directly at the current stage.

3.9.2 Fit models

In this study, it is assumed that the proton density of the bunches can be well described as the sum of a few three-dimensional Gaussian functions.

$$b(\mathbf{r}) = \sum_{i=1}^n a_i g(\mathbf{r}, \mu_i, \Sigma_i) = \sum_{i=1}^n a_i g_i(\mathbf{r}) \quad (3.26)$$

where $\sum_{i=1}^n a_i = 1$ barycentric coordinates describe the amplitudes of each component and $g(\mathbf{r}, \mu, \Sigma_i)$ is a three-dimensional normalized Gaussian density function. This convention has the benefit that $\int b(\mathbf{r}) d^3\mathbf{r} = 1$. Nota bene, the coefficients a_i are not all necessarily positive, but physical solutions necessitate $b(\mathbf{r}) > 0$ for $\forall \mathbf{r} \in \mathbb{R}^3$.

The use of Gaussian components is also motivated by the general adequacy of the shape for a cloud of particles but most importantly the luminosity and beamspot parameters for proton densities described by these functions can be easily computed analytically (see Appendix B). This makes numerically stable and fast computations possible.

In [80] it was found, that it is sufficient to use a shared centre model, that is to take $\mu_i = \mu_j$ for the components of a given bunch.

While in 2015 and 2016 [80] the method used a spherical parametrization [109] for the covariance matrix of the Gaussian components and allowed the components to bear independent tilts, in the 2017-2018 study it was found that with an Euler parametrization, using shared Euler angles between the components, the fit quality is similarly good. The same approach is used for the 2023 analysis. This parametrization is achieved by sandwiching the diagonal covariance matrix between rotation matrices

$$R_{yz} R_{xz} R_{xy} \Sigma R_{xy}^{-1} R_{xz}^{-1} R_{yz}^{-1}.$$

The order of transformations is such that the alignment of the Gaussian components around the longitudinal axis can be easily released without ruining the alignments in the tilts.

Further on, it is more comfortable to reference these models with abbreviated names. Table 3.2 introduces these abbreviations for the most important models and gives an insight into the number of fit parameters for each in the 2015-2016 and 2017-2018 studies.

The “*sepA” class of models only releases the Euler parameter describing the orientation of the components in the transverse direction (in the frame of the bunch), but not in the $x - z$ or $y - z$ direction. This choice was made, as the variables relating to the longitudinal tilts are relatively free of structures and are well described, while the transverse misalignment of the components could improve the description of the transverse beamspot and luminosity-related features that contain the most structure. Fitting with all angles released makes convergence more difficult with no significant benefits in return.

While the 2015-2016 work studied models with $n = 4$ Gaussian components at most, the current study goes up to $n = 5$ and also introduces the Super Gaussian (supG) family of models. This latter feature is an attempt at fitting more Gaussian components while keeping the number of parameters moderate. The Super Gaussian function is a one-dimensional non-Gaussian bell-like curve

2015–2016					
Short name	Long name	n_c	Means	Tilts	n_p
SG	Single Gaussian	1	–	–	20
DG	Double Gaussian	2	shared	separate	34
TG	Triple Gaussian	3	shared	separate	48
QG	Quadruple Gaussian	4	shared	separate	62
TGsep	Triple Gaussian Separate Means	3	separate	separate	60
QGsep	Quadruple Gaussian Separate Means	4	separate	separate	80
2017–2018					
Short name	Long name	n_c	Means	Euler Angles	n_p
SG	Single Gaussian	1	–	–	20
supG	Super Gaussian	2	shared	shared	22
DG	Double Gaussian	2	shared	shared	28
DsupG	Double Super Gaussian	4	shared	shared	32
TG	Triple Gaussian	3	shared	shared	36
QG	Quadruple Gaussian	4	shared	shared	44
PG	Penta Gaussian	5	shared	shared	52
TGsepM	TG Separate Means	3	separate	shared	48
QGsepM	GQ Separate Means	4	separate	shared	62
TGsepA	TG Separate Angles	3	shared	separate	40
QGsepA	QG Separate Angles	4	shared	separate	50

Table 3.2: Abbreviated names for the most often used models in the Run2 LR studies. The column n_c stands for the number of Gaussian components per bunch. n_p shows the number of fit parameters taking the the bunch density parameters for both beams, their amplitudes, σ_{vis} , and the vertex resolution bias into account, to show the minimum number of parameters for a fit of a given complexity.

with only one more parameter (p) compared to a Gaussian function (see Eq. A.10). The Super Gaussian function reduces to a simple Gaussian function at $p = 1$.

The supG model in this study is constructed as $f(x', y')g(z')$, where the (x', y', z') are the coordinates rotated by the Euler angles describing the bunch orientation. The longitudinal shape g_z is modeled as a Gaussian, while the tangential directions (f) are modeled as a two-dimensional function swept out by rotating a one-dimensional Super Gaussian. Since the luminous region method can only use Gaussian components, this is approximated as the sum of two two-dimensional Gaussian functions. The latter approximation is carried out offline and the results are condensed into a look-up table. The approximation is only accurate near $p = 1$ and deteriorates much quicker for lower p values than for higher ones, but the aim of this approximation is not to describe an actual

Super Gaussian accurately but to describe a flatter peak, or a thicker tail with only one additional parameter, which is achieved. A few examples are shown in Figure 3.21.

Since the supG model is implemented as a sum of two Gaussian components, the double Super Gaussian (DsupG) can be used to seed TG and QG functions by releasing the appropriate parameters and for TG discarding the component with the smallest amplitude. These models are abbreviated as “DsupG>TG” and “DsupG>QG”.

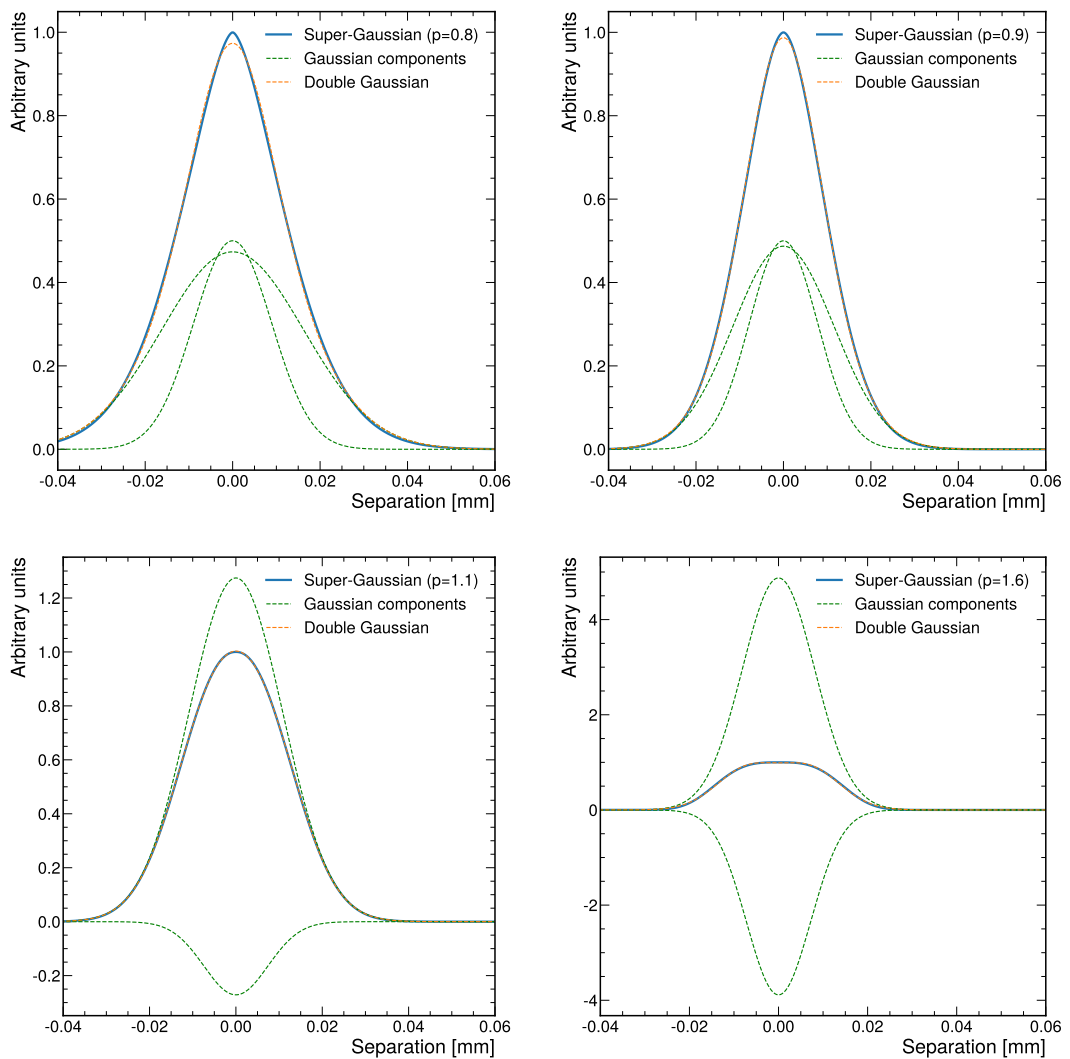


Figure 3.21: Double Gaussian approximation of a unit amplitude Super Gaussian function at different p values with the constituent Gaussian components plotted. The Super Gaussian function reduces to a simple Gaussian function at $p = 1$. The approximation is not perfect, but it allows for a description of a flat top or a thicker tail with only one extra parameter.

Besides the parameters of the bunch density shapes, other parameters include

σ_{vis} (the volume parameter) and the vertex resolution bias term Δ_{res} which can be fit separately in x and y . For the central result, this parameter is shared for x and y . Optionally a constant term can be incorporated into the analysis to model the background of the luminometers (C).

3.9.3 From observables to bunch proton densities

Appendix B details how the beam density functions determine the luminometer rates and luminous region parameters. The calculation also incorporates a correction for the dynamic beta effect using the m_β factors obtained from simulation based on [96]. The gradual change of the beam emittances can be compensated for using the same mechanism. These calculations are easily vectorizable which guarantees optimal performance. The implementation is cross-checked against a Monte Carlo numerical integration-based calculation and a good agreement is found.

Such a connection between the parameters and the observables allows us to use a minimizer, in this case Minuit [110], to find the optimal values for our model parameters.

$$\left. \begin{array}{l} \text{Fit parameters} \left\{ \begin{array}{l} \sigma_{\text{vis}} \\ \mathbf{a}_i^{b1}, \mathbf{a}_j^{b2} \\ \mu_i^{b1}, \mu_j^{b2} \\ \Sigma_i^{b1}, \Sigma_j^{b2} \\ \Delta_{\text{res}} \\ C \\ \text{Beam positions} \\ m_\beta^x, m_\beta^y \end{array} \right. \right\} \rightarrow \left\{ \begin{array}{l} \text{Normalized rates} \\ \text{Beamspot parameters} \\ L^{b1}/L^{b2} \end{array} \right.$$

The χ^2 function makes use of the full covariance matrix of the beamspot parameters obtained in the maximum likelihood fit. In general, it is written as

$$\begin{aligned} \chi^2 = & \sum_i \left(\frac{R_{\text{measured}}^{(i)} - R_{\text{model}}^{(i)}}{R_{\text{unc.}}^{(i)}} \right)^2 + \\ & + \sum_i (\mathbf{P}_{\text{measured}}^{(i)} - \mathbf{P}_{\text{model}}^{(i)}) (\mathbf{P}_{\text{cov. mx}}^{(i)})^{-1} (\mathbf{P}_{\text{measured}}^{(i)} - \mathbf{P}_{\text{model}}^{(i)}) + \quad (3.27) \\ & + \text{Penalty terms,} \end{aligned}$$

with the extra constraint on the bunch length ratio (Eq. 3.9.1) added in the new results. Here $R^{(i)} \in \mathbb{R}$ stands for the normalized rate, $\mathbf{P}^{(i)} \in \mathbb{R}^9$ is a

nine-dimensional vector of the beamspot parameters, and $P_{\text{cov. mx}}^{(i)} \in \mathbb{R}^{9 \times 9}$ is the covariance matrix of these parameters as calculated in the ML fit, all at the i -th step of the scan.

Penalty terms are introduced to help the fit converge to physical solutions. These terms aim to penalize non-strictly positive and non-bell-like proton bunch density functions. Unfortunately, the positivity of the sum of a handful of Gaussian functions is not easily guaranteed. To compute the penalty term both bunch densities are sampled across their centre-of-mass in a line along all non-redundant directions of the $\{(a, b, c) \mid a, b, c \in \{-1, 0, 1\}\}$ vector set up to $5\sigma_b$. All negative values are added up, squared, and added to the χ^2 function with a w_n penalty multiplier. The linear scans are also used to test for double peaks. The difference between the average of the peaks, and the value of the minimum between peaks is squared and added to the χ^2 function with a w_v penalty multiplier. The penalty is in general zero near the result, but it protects the fit from slipping into non-physical local minima during the optimization.

For fitting the density functions, the computed beamspot covariance matrix, which is in the standard $\mathbb{R}^{3 \times 3}$ representation, can be converted to the parametrization used by the input (see Sec. 3.9.1). This approach was used in 2015-2016. In this case, the square root of the diagonal matrix elements needs to be computed, which is one of the few cases in which invalid values can be encountered in the computation. Although near the optimum these are all positive, the fitting may attempt to explore regions where they are not.

Even though Minuit handles these cases, in 2017-2018 the reverse approach is adopted and the ML fit results are converted to the standard matrix representation. In this case, the uncertainties and correlations of the input beamspot parameters must also be recomputed. This is carried out before the fitting using a toy Monte Carlo simulation.

It is quite common, that the ML fit does not converge properly if there are too few vertices. This is often the case for the outermost points of the vdM curve. In this case, the χ^2 function is constructed without beamspot information for these separations, but the luminometer data is still used to constrain the fit.

In case the beamspot fits are missing for the outermost three points, for a 25-step vdM scan pair we get $2 \cdot (25 + (25 - 6) \cdot 9) = 396$ observables to constrain our fits. This is well above the number of fit parameters indicated in Table 3.2.

3.9.4 Sequential fitting

Due to the large number of model parameters the convergence of the fit is not guaranteed out of the box. To mitigate this the initial parameters need to be chosen carefully for each of the models hence the sequential fitting method is introduced. The exact procedure has changed in many ways between the 2015-2016 and the 2017-2018 studies, the current chapter reflects the latter conditions.

To obtain the initial parameters for the Single Gaussian model the beamspot parameters measured at the nominal head-on position are used

$$\mu_i^{b1} = \mu_i^{b2} = \mu_i^{lr} \quad i \in \{x, y, z\}, \quad (3.28)$$

$$\sigma_i^{b1} = \sigma_i^{b2} = \sqrt{2}\sigma_i^{lr} \quad i \in \{x, y, z\}, \quad (3.29)$$

and the correlation and tilt parameters are initialized at a zero value. The initial parameter for σ_{vis} is computed as

$$\sigma_{vis} = 8\pi\sigma_x^{lr}\sigma_y^{lr}, \quad (3.30)$$

and Δ_{res} and C are initialized at zero as well.

The procedure fits models in order of increasing complexity continuously building on the results of the previously performed fit. For bunch density models containing $N > 1$ Gaussian components, the initial parameters are derived from the previous result by adding a new Gaussian, based on the dominant component of the sum, with a small positive or negative amplitude and an increased or decreased width. Several combinations of the amplitude and width modifications are tried in a grid-search-like manner by using a truncated optimization round before selecting the best initial parameters for this model based on the χ^2 . Figure 3.22 illustrates the concept.

After a full fit is performed with the selected initial values, the uncertainty of the parameters is estimated using the Hesse routine of Minuit. In case the computation fails, the method attempts to find a more suitable minimum by reoptimizing a few times from a slightly changed position.

In general, it can be said, that even if the parameter uncertainties are not of adequate quality according to Minuit, the χ^2 is usually not worse than for the previous model, and hence the description is not bad, even if overfitting in this case is questionable. Biases related to overfitting are covered by the Monte Carlo closure test (see Sec. 3.9.7).

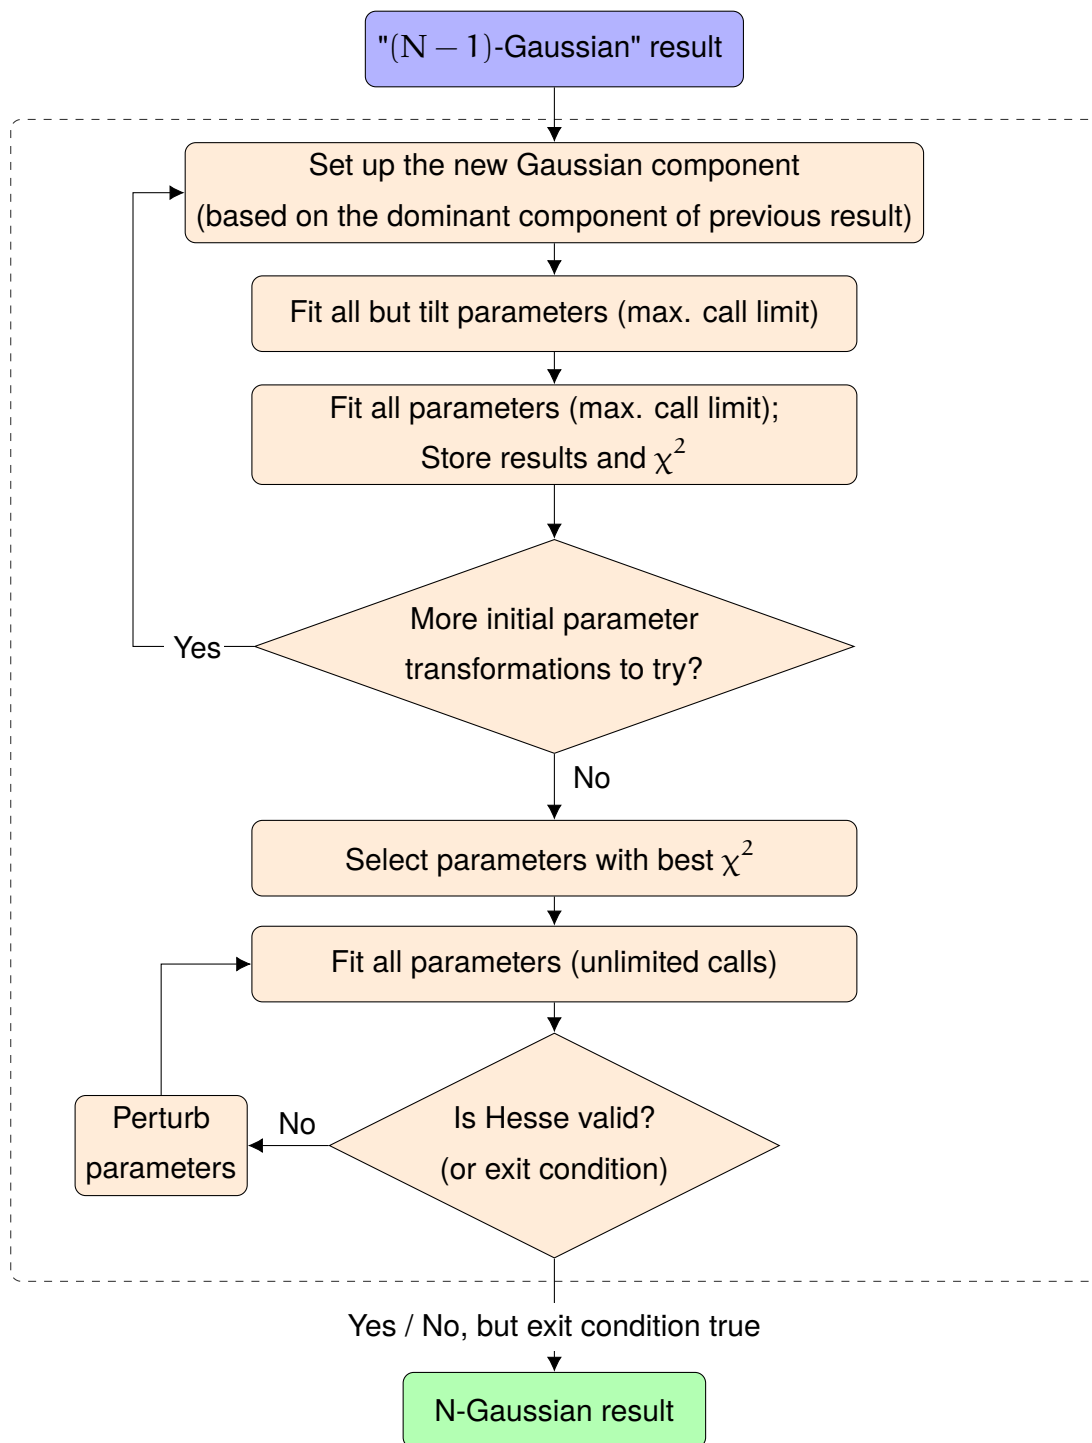


Figure 3.22: This flowchart summarizes an atomic step of the sequential fitting procedure. The exit condition at the Hessian computation stops the iteration if the number of attempts reaches a threshold.

Quick fitting

The sequential fitting procedure is clearly time-consuming, hence a new “quick fit” procedure has also been devised starting with the 2017-2018 study. This relies on the quick convergence of simplified models and testing a large number

of reasonable initial parameters for the optimization. The quick fit method models all Gaussian components as a product: $g_x(x)g_y(y)g_z(z)$. The beamspot parameter calculations and a simple implementation for this case are detailed in Appendix C.

In this model, the main axes of the bunch are always aligned to the tracker coordinate system. Even with these simplifications, the sum of several such components can result in a non-factorizable shape, which can describe the most apparent features seen in the beamspot parameters, such as the W shape in the width aligned with the scanning direction (see Fig. 3.26). On the other hand, these assumptions simplify the beamspot parameter and luminosity calculations greatly.

The optimum of such a simplified model can be used as the initial parameters to a full model, which can obtain results as good or better than the sequential fitting method. This is of course conditioned on the quality of the simplified model, which in turn depends on the number and quality of the tested initial parameters.

3.9.5 Fit results

Figures 3.23 and 3.24 show the fit result for BCID 281 in the first vdM scan pair of the 2017 calibration fill. Of the luminous region parameters only the panel belonging to the x scan is shown.

3.9.6 Deriving the vdM correction factors

By fitting our model to the observables, we obtain the bunch shape parameters and their uncertainties as well. One of the fit parameters is σ_{vis} . Eq. B.39 allows for the calculation of the normalized rates at any beam positions, hence the rate values for the vdM separations can be computed to perform simulated vdM scans as described in Section 3.8.3.

3.9.7 Closure test

While a closure test in [80] resulting in a systematic of 0.6% was already performed, with the reworked fitting procedure, new models, and a better understanding of

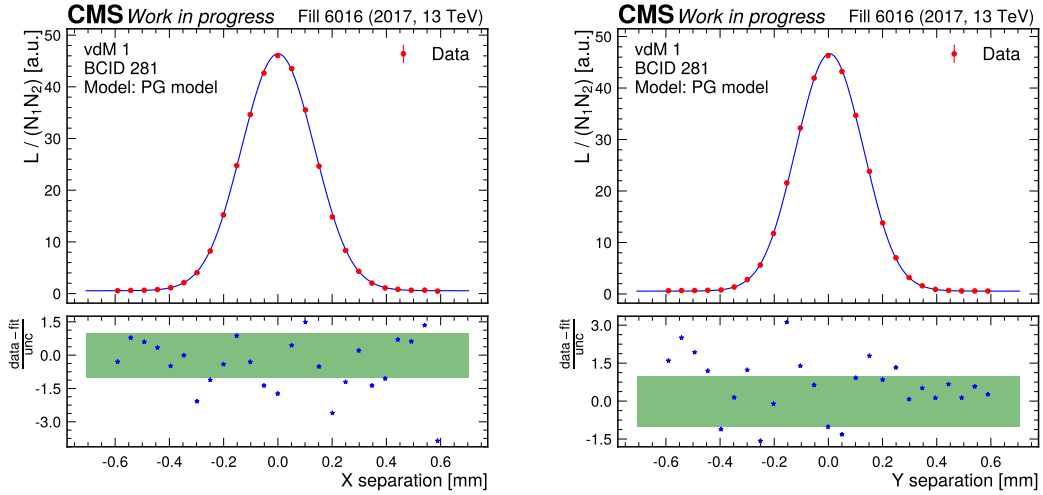


Figure 3.23: The normalized rate of the PCC luminometer (red) for the first x and y vdM scans in 2017 as a function of the beam-beam and length scale corrected beam separation measured by the DOROS beam position monitor and a Penta Gaussian (PG) model fitted on the data (blue). The normalized residuals of the model can be seen in the bottom panel of the subfigures, the green band signifies one σ in the deviations.

the possible proton density shapes as an input, a new closure test is carried out for the 2017-2018 study.

The aim of this test is to quantify the performance of the luminous region method on a broad range of possible proton density shapes using Monte Carlo simulation. The truth (proton density) models (TM) tested in this section are the SG, the TG, and the QG models. While the SG does not describe data well, evaluating the capacity of the resulting observables to constrain higher-order fit models does probe the LR method for overfitting. The DG function is still not a good fit in general, while TG and QG conform to the data successfully.

To challenge the LR method, over a thousand sets of random parameters are selected for each TM from a broad distribution governed by a "tune" which is to be described later. The random parameters define a truth bunch shape, for which the "truth correction" can be calculated using a simulated vdM scan. On the other hand, the exact luminosity and luminous region parameters can also be computed for each point in a vdM scan pattern. The resulting values are randomized based on uncertainties measured in actual data to obtain a toy dataset. This toy data is then treated identically to actual data, using the luminous region method as a black box. Finally, the fit correction is compared to the truth correction to establish the bias. The bias and the reduced χ^2 value are

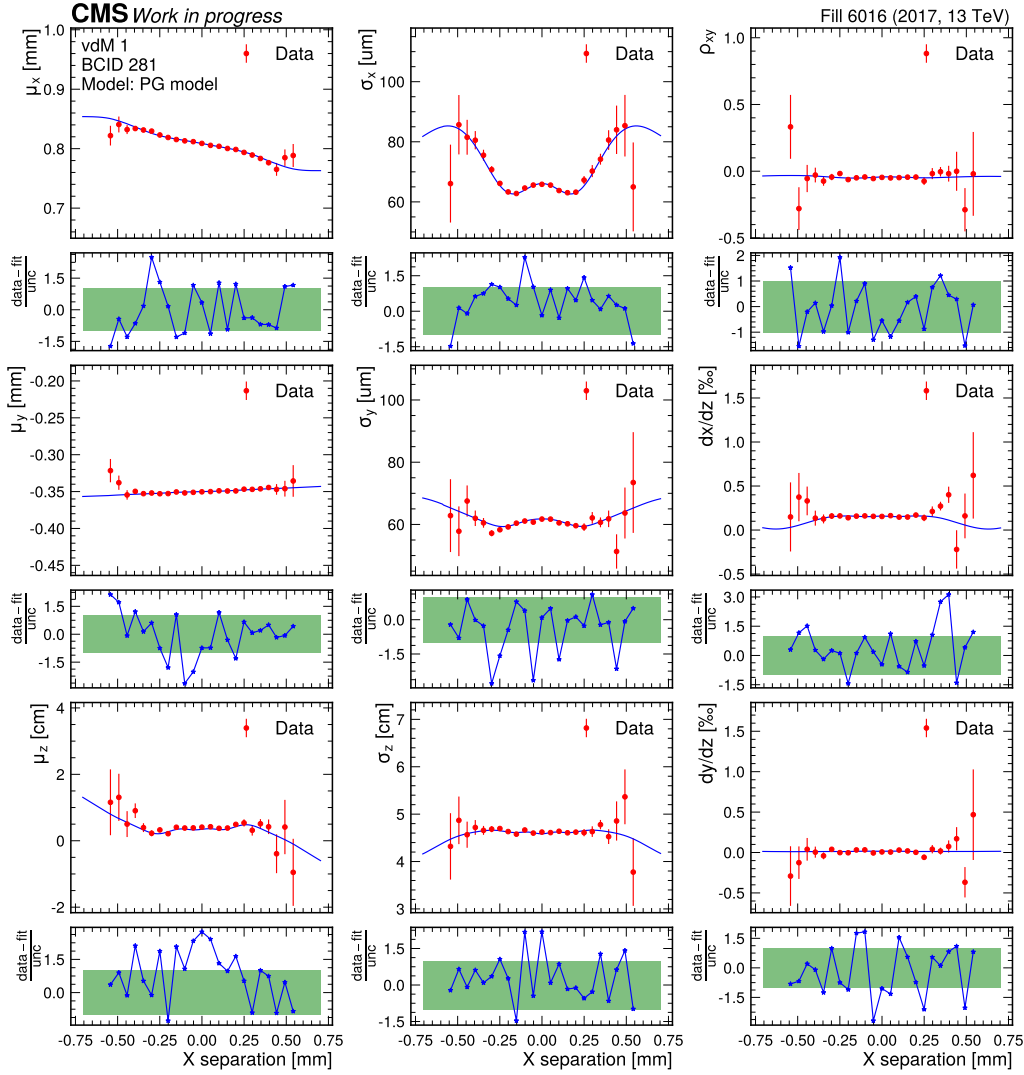


Figure 3.24: The parameters of the 3-dimensional beamspot ellipsoid as a function of the beam-beam and length scale corrected beam separation measured by the DOROS beam position monitor and a Penta Gaussian (PG) model fitted on the data (blue) for the first vdM scan in the X direction in 2017. The normalized residuals of the model can be seen in the bottom panels of the subfigures, the green band signifies one sigma range. Left column: Beamspot center in x , y , and z . Middle column: Beamspot width in x , y , and z . Right column: Tilt parameters describing the orientation of the ellipsoid axes in $x - y$, $x - z$ and $y - z$.

saved. The procedure is illustrated by the flowchart in Figure 3.25. A fit of the toy data using the truth parameters as initial values is also performed.

In [80] the analytical mode was always a TG function with random parameters drawn from reasonably wide Gaussian distributions located in a way that the distributions covered the physical parameter space. The resulting shapes were tested for their correctness and non-physical instances were rejected.

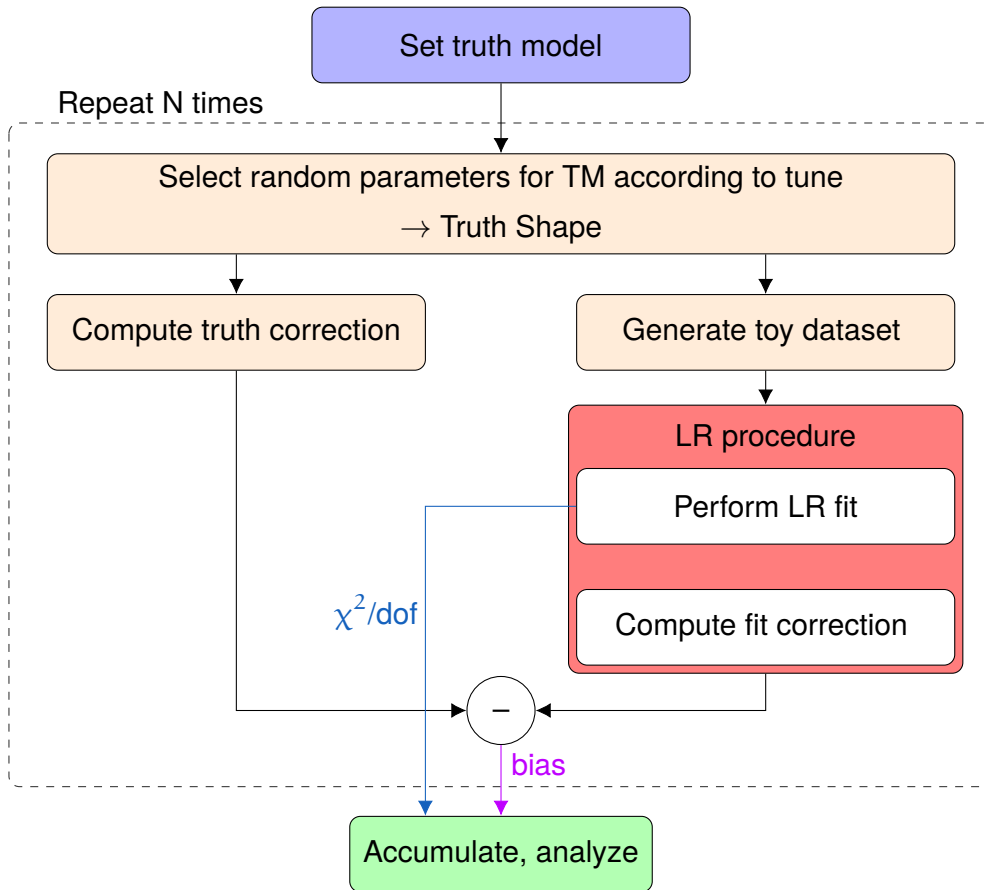


Figure 3.25: Flowchart of the luminous region closure test procedure.

This time around a more nuanced approach uses six different tunes for the models:

- SIMPESHAPE_1_SHM_SHA
- SIMPESHAPE_3_SHM_SHA
- SIMPESHAPE_4_SHM_SHA
- WSHAPE_3_0.99_SHM_SHA
- WSHAPE_3.5_0.99_SHM_SHA
- WSHAPE_4_0.99_SHM_SHA

In all cases “SHM” stands for shared means and “SHA” stands for shared Euler angles. The “SIMPLE” class of models uses the old approach described in the previous paragraph. The digit in the name signifies the number of Gaussian components in a single bunch proton density function.

The “WSHAPE” class of models was inspired by the realization, that, while we have little information about the actual proton density distributions, the simple shapes rarely exhibit one of the most characteristic features of our observables, the W shape in the width of the luminous region in the scanning direction (see

Figure 3.26). This means that the simple models while allowing for a wide variety of possible densities, might not be specific enough for our purposes.

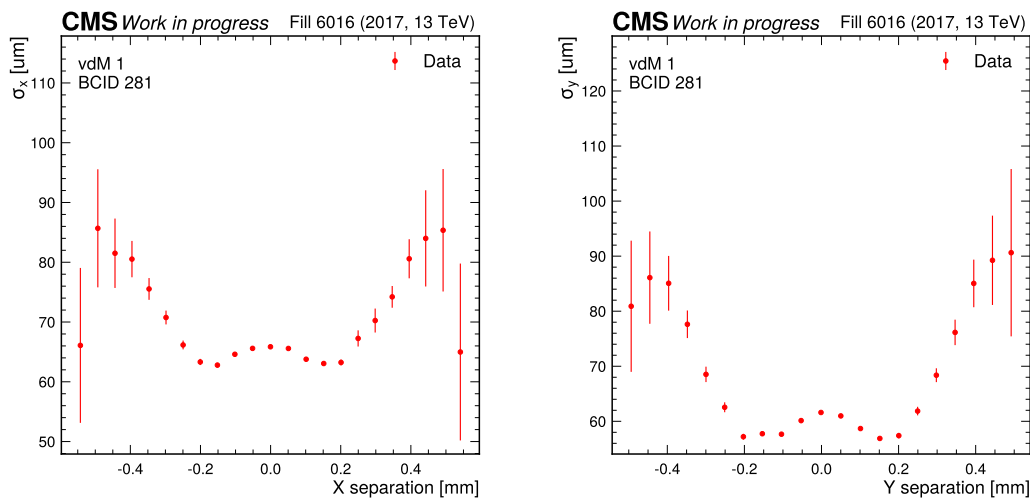


Figure 3.26: The width of the three-dimensional beamspot ellipsoid in the scanning direction in Fill 6016 recorded in the vdM1 scan pair for BCID 281. The W-like structure is a frequent characteristic of this observable at CMS. The same feature is not always well-defined in the non-scanning direction.

The first number in the “WSHAPE” name signifies the number of Gaussian components in a single bunch. It is observed, that many of the TG WSHAPE models, although are non-negative, and exhibit the W feature properly, also show a small decrease in density near the center of the distribution, hence suffering from a double peak structure, which is deemed non-physical. However often these shapes need not be disregarded as adding a fourth, very small Gaussian component can compensate for this deficiency, while barely changing the derived observables. Models that have undergone this fix have 3.5 in their names instead of 3.

The second number in the name of the model, currently 0.99 in all cases, signifies the strength of the W feature as the maximum value for the ratio between the minimum of the W and the local maximum at the head-on position.

All six tunes were run with between 1000 and 1600 random parameter combinations, totaling up to a few years of CPU time on a regular PC. A set of examples of the test results for the SG model are shown in Figure 3.27. While the SG fit model accurately describes Single Gaussian-like data with very good convergence properties, it fails spectacularly for more complex truth models, which is expected.

In the actual measurement, the correction is derived from the model with the best χ^2/dof value. The closure of the best-fit models with the SIMPLE-SHAPE_4_SHM_SHA tune is shown in Figure 3.28. The bias of the method is $< 0.1\%$ over the entire range of the truth correction value, however, the local standard deviation of the biases reaches up to 0.3% .

The closure between all truth and fit models is summarized in Figure 3.29. These tables only include fits with $\chi^2/\text{dof} < 20$, which is a very loose cut, since in practice the reduced χ^2 does not even reach 4. From the tables, we conclude, that the more complex models converge properly giving a good average reduced χ^2 and a negligible expected bias, especially when the complexity of the model matches the complexity of the truth shape. The best-fit model closely matches the truth parameter-based fits in both the expected bias and the standard deviation, which seems to indicate that the convergence is close to ideal, and the method may be statistically limited, either requiring more precision for the individual measurement points or more complex scanning patterns. The only exception to these observations is presented by observing SIMPLESHAPE_1_SHM_SHA carefully. Here the more complex models show a larger variance due to some level of overfitting. Still, the standard deviation of the bias provided by the best-fit model is comparable to that of the other tunes.

The final closure result is derived as the quadratic sum of the mean bias and the standard deviation, which is the same as the expected quadratic deviation from zero. Figure 3.30 presents these results. Based on the largest number in the last column, corresponding to the best-fit model, a closure uncertainty of 0.25% is assigned as a conservative choice.

3.9.8 Results and uncertainties

Several different options to the flow of this analysis have been presented in this section. The impact of these choices on the goodness of fit and the correction with respect to the central result is evaluated. The results are shown in Figure 3.31.

An explanation of the rows of the figure follows.

Independent $\alpha\gamma$ tilts: The central result uses shared Euler angles in all three axes. This entry shows the effect of fitting Euler angles independently in the transverse plane with respect to the bunch's reference frame. The PG model was not fit with this option for lack of convergence. While for TG

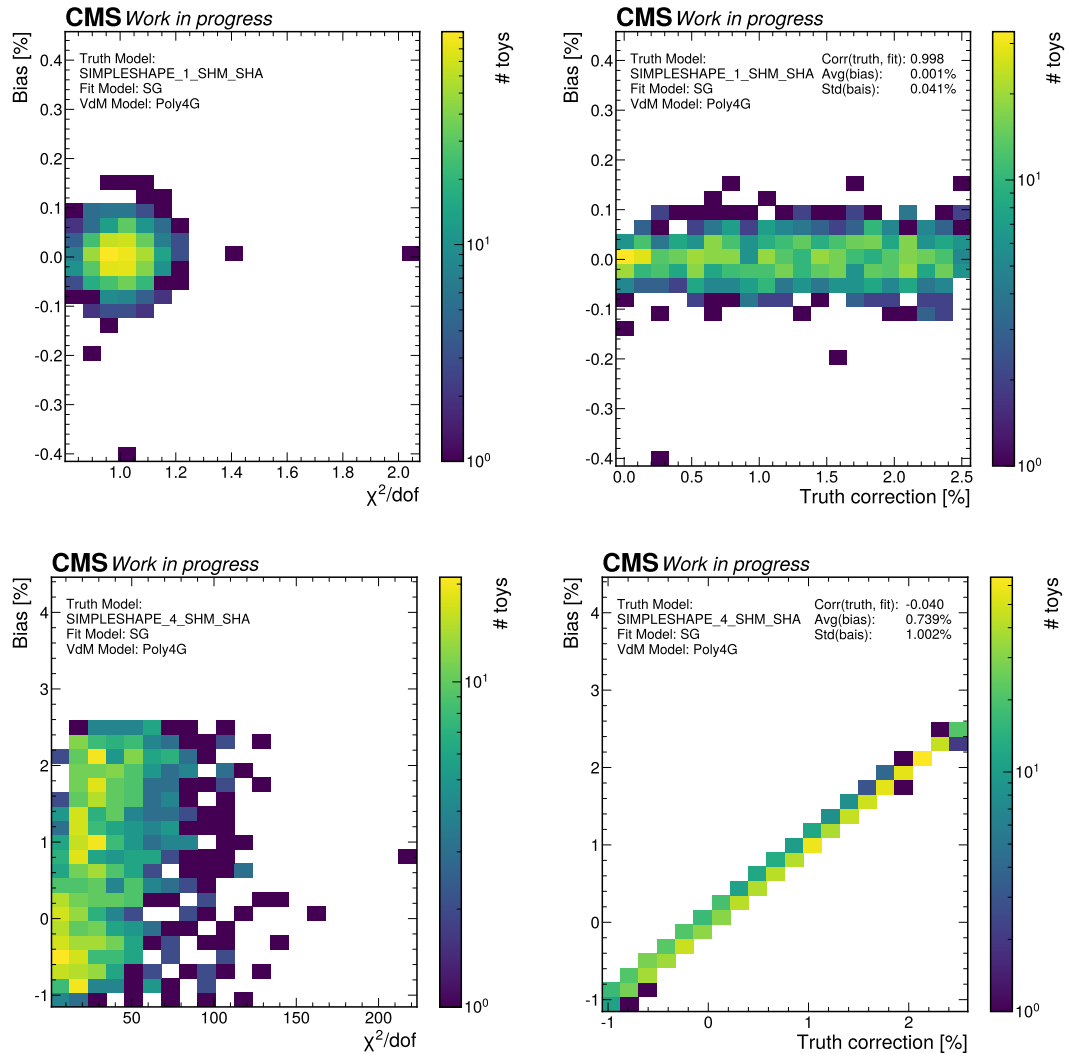


Figure 3.27: Closures for the SG fit model. Results with the Single Gaussian-like truth model (top row) are very convergent and the bias is negligible ($< 0.1\%$). Using a more complex truth model, such as a Quarto Gaussian-like shape (bottom row), it is clear that the SG fit model does not adequately describe the data, and both the reduced χ^2 and bias are large accordingly.

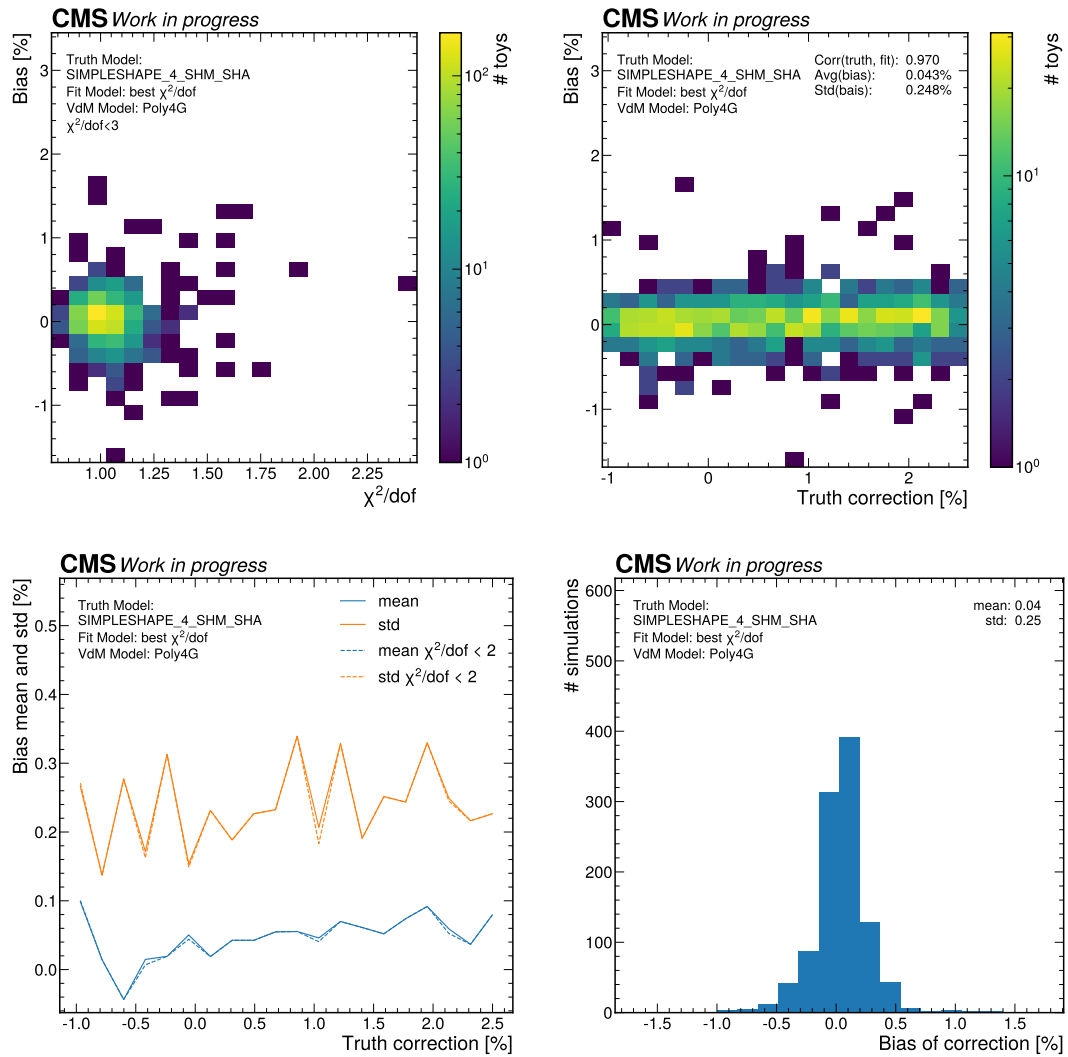


Figure 3.28: Closures of the best-fit model to the SIMPleshape_4_SHM_SHA truth model. For the top left subfigure outliers with $\chi^2/\text{dof} > 3$ were discarded to better show the distribution.

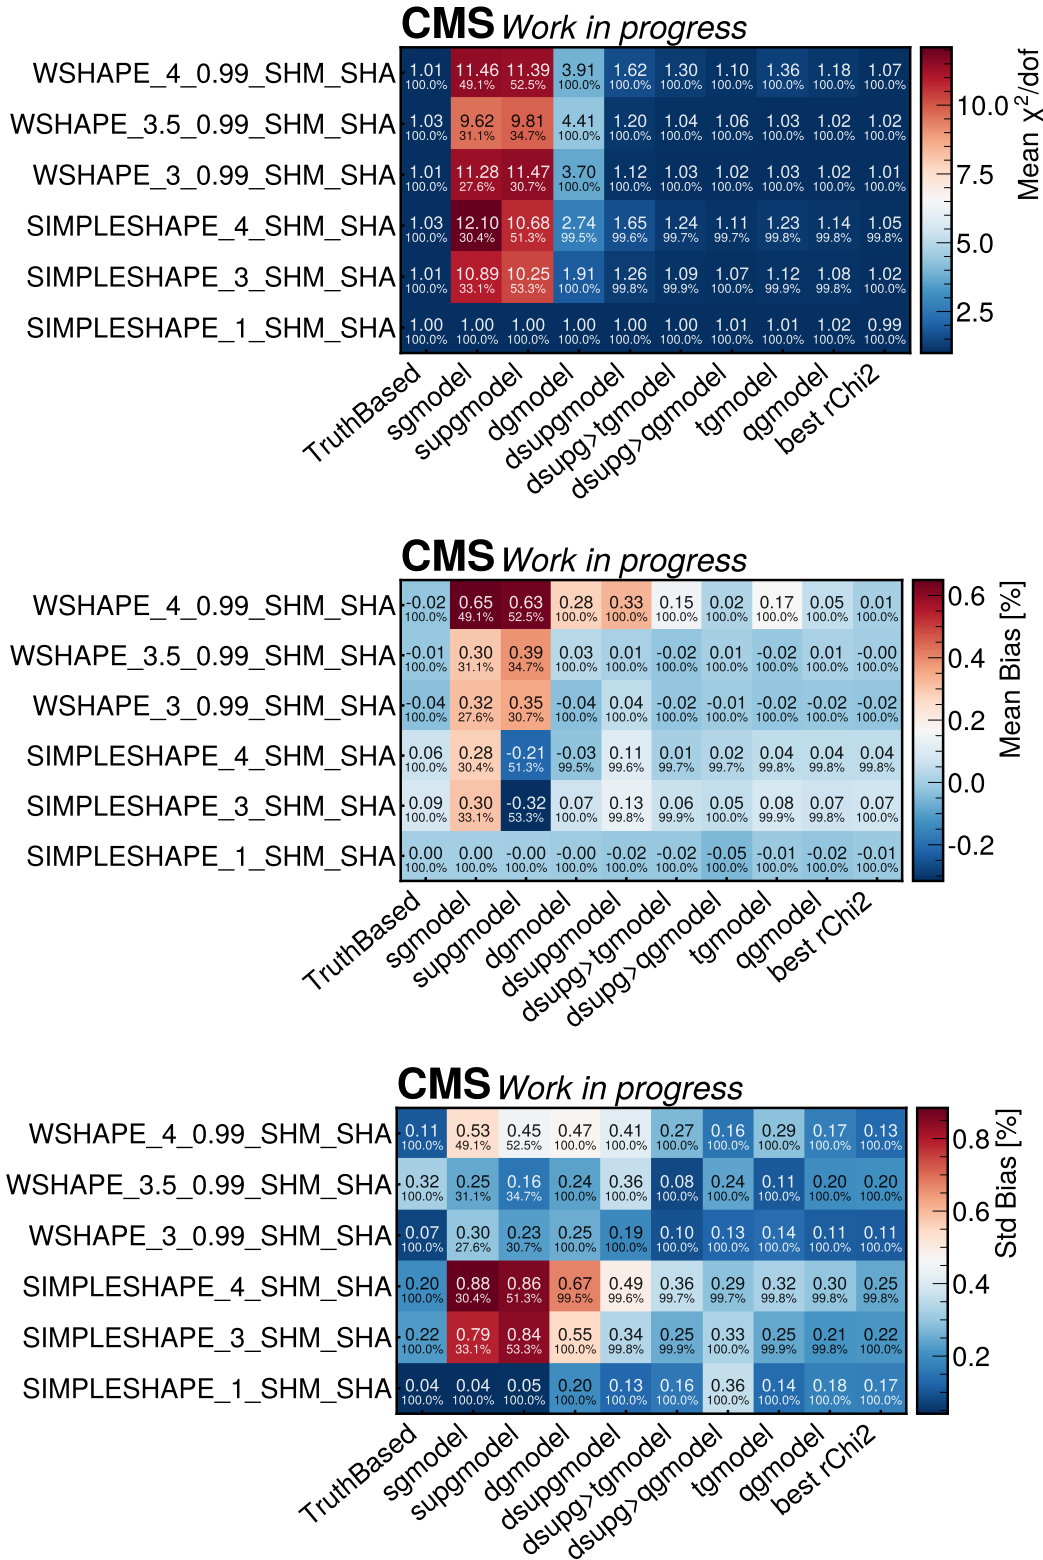


Figure 3.29: The results of the closure test for the luminous region method. Rows represent the different truth models, while the columns stand for the fit functions. The top plot shows the average χ^2/dof , the middle plot shows the mean bias as a percentage of σ_{vis} , while the bottom plot shows the standard deviation of the bias in the same units. Only fits with $\chi^2/\text{dof} < 20$ are used to fill the tables. The fraction of fits passing this criterion is displayed as a percentage at the bottom of each cell.




Figure 3.30: The summary of the closure test for the luminous region method showing the full closure uncertainty for all truth to fit model pairings. Rows represent the different truth models, while the columns stand for the fit functions. Only fits with $\chi^2/dof < 20$ are used to fill the table. The fraction of fits passing this criterion is displayed as a percentage at the bottom of each cell.

models the impact is insignificant, for QG models an average shift of up to 0.2% is observed.

Independent centers: The central result uses shared means. This entry shows the effect of fitting the centroids of all Gaussian components independently. The PG model was not fit with this option for lack of convergence. The effect on the correction is negligible.

Use PLT instead of PCC: The central result uses PCC as the main source of luminosity measurement. This entry shows the effect of swapping this input to a different luminometer. The effect on the correction is negligible.

No σ_b modulation: The central result uses beam width modulation to both model the dynamic beta effect as well as the emittance change. This entry shows the effect of turning off this modulation. The effect on the correction is negligible.

$\Delta_{res} = 0$: The central result uses a free-floating term for the vertex resolution bias in the transverse direction. This entry shows the effect of constraining this parameter to zero. This is the largest among the examined effects, which demonstrates the need for this parameter.

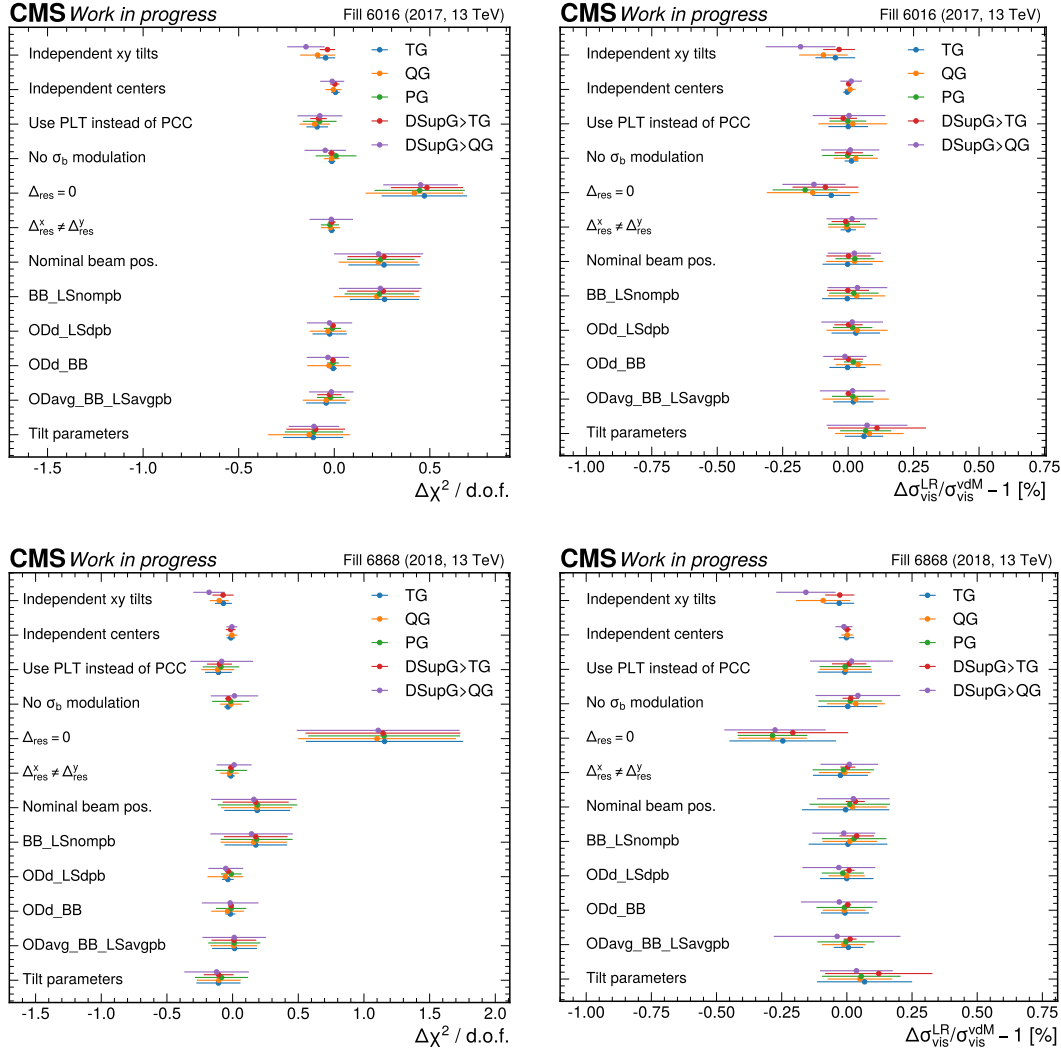


Figure 3.31: The impacts of systematic changes on the mean of the reduced χ^2 (left) and the mean correction (right) for 2017 (top) and 2018 (bottom). See the text for explanation of the labels.

$\Delta_{\text{res}}^x \neq \Delta_{\text{res}}^y$: The central result uses the same vertex resolution bias term in the x and y directions. This entry shows the effect of fitting the two directions separately. The effect on the correction is negligible, demonstrating the adequacy of the $\Delta_{\text{res}}^x = \Delta_{\text{res}}^y$ choice.

Nominal beam positions: The central result uses beam position corrected with Doros BPM for orbit drift, BB correction, and per beam Doros length scale (ODd_BB_LSdpb). This entry shows the effect of turning off all corrections on the beam position measurement. The effect on the correction is negligible, but the fit quality suffers due to this change.

BB_LSnombp: This entry shows the effect of not using OD correction. The effect on the correction is negligible, but the fit quality suffers due to this change.

ODd_LSdpb: This entry shows the effect of turning off the beam-beam correction. The effect on the correction and fit quality is negligible.

ODd_BB: This entry shows the effect of turning off the length scale correction. The effect on the correction and fit quality is negligible.

ODavg_BB_LSavgpb: This entry shows the effect of changing the OD input from Doros BPM to the average of arc and Doros BPM and changing the length scale correction accordingly. The effect on the correction and fit quality is negligible.

Tilt parametrization: The central result transforms the tilt parametrization of the provided beamspot data to standard covariance matrix parametrization and uses this for fitting. This entry shows the effect of using the original parametrization instead. The effect on the correction and fit quality is small but consistent across the years. Until a better understanding is achieved a 0.1% uncertainty is assigned.

Based on these plots, uncertainties are assigned to the tilt parametrization (0.1%) and the Euler angle sharing (0.2%) choices noting that these effects are anticorrelated.

The final corrections derived using the best-fit model are shown in Figure 3.32. The internal consistency of the result was checked by incorporating off-axis scans into the LR fit when available and comparing the correction to the result where only the on-axis scan is used. As in 2017 only a single point of comparison is available with large uncertainties, the differences are averaged over the two years and 0.1% uncertainty is assigned.

The combined fits are also compared to the corresponding 2D rate fit results performed by K. Farkas and are found to be within uncertainties in all cases. No additional uncertainty is assigned.

Comparing the 2D rate fit method to itself between using all colliding bunches or only using those five for which luminous region corrections are available, the only difference is seen in the second half of fill 6868. Since in 2018 applying the luminous region correction decreased the scan-by-scan variance, it is believed, that this might be just statistical fluctuation instead of a genuine difference between the behavior of the bunches. Nonetheless, in 2018 a conservative 0.16% uncertainty is added.

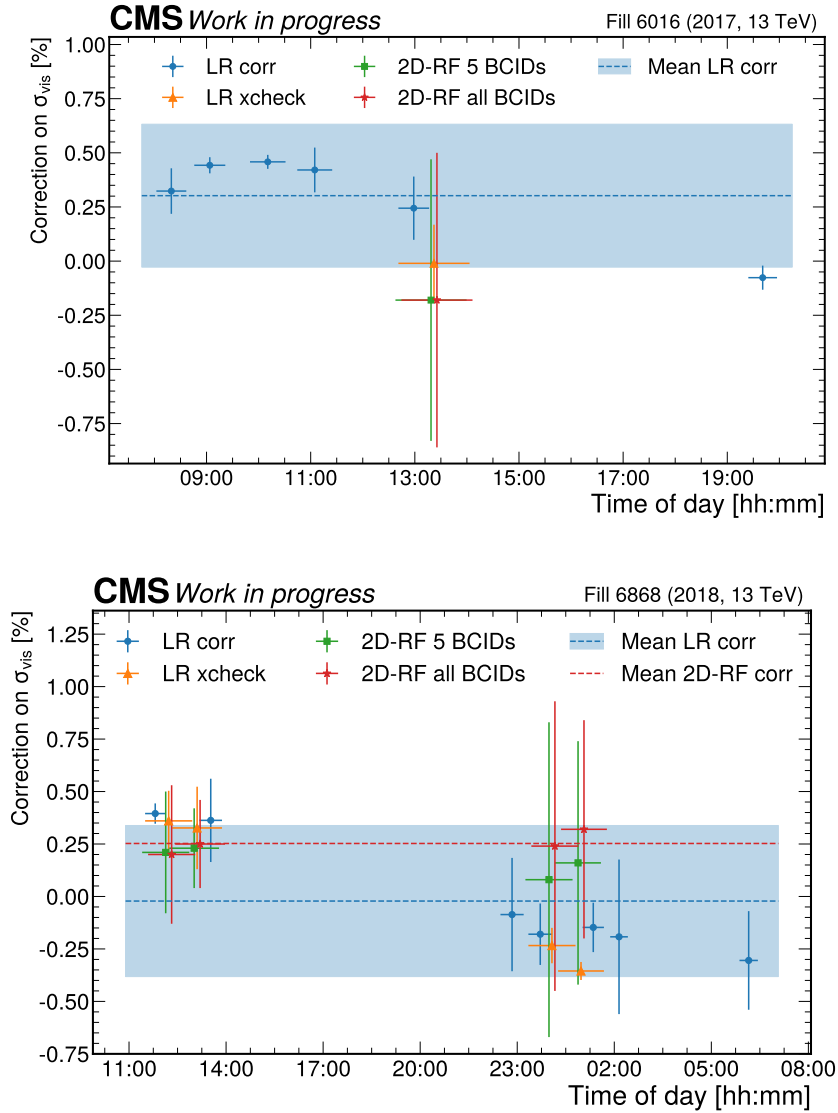


Figure 3.32: The non-factorization correction using the luminous region (LR) method for 2017 (top) and 2018 (bottom). The blue markers indicate the values used for the scan-by-scan correction, the vertical error bars correspond to the standard deviation of the per-bunch results. In all cases, the horizontal error bars indicate the time intervals span by the scans used for the measurement. The blue region shows the total uncertainty of the average correction. The orange marker is for LR fits incorporating off-axis scans. The green markers show the 2D rate fit method (RF) results computed for the five BCIDs that the LR uses, while the red markers show the RF results for all colliding bunches. The vertical RF error bars show the full uncertainty.

Taking the closure uncertainty (0.25%) into account the final uncertainty of the method in 2017 and 2018 is 0.33% and 0.36% respectively. The preliminary factorization uncertainties for these years were 0.8% and 2.0% [81, 82].

Figure 3.33 shows the final result for 2015-'2016, where no off-axis scan data was available. The cross-check is performed against the beam imaging method, the results are found compatible. The LR uncertainty is dominated by the closure test (0.6%) [80]. The trends recovered by the luminous region method are shifted to match the average correction of the two methods at the time of the beam imaging scans. The smaller of the two uncertainties is kept, instead of using the error on the mean, as a conservative choice. The final uncertainty is 0.5% in both years, an improvement with respect to the 2.0% and 0.9% preliminary results [98, 99].

3.10 Normalization results

In luminosity publications, normalization refers to the derivation of the σ_{vis} calibration constant. In the work revisiting 2017-2018, as a result of the new vdM fit function and the scan-by-scan corrections of the non-factorization, the emittance change correction and the more nuanced orbit drift corrections [71],




Figure 3.33: The non-factorization correction using the luminous region and the beam imaging method method for 2015 (left) and 2016 (right). The error bars on the orange points signify the standard deviation of the per-BCID results, the shaded areas show the total uncertainty for both methods and the combination. In 2017 a linear fit, while in 2018 a constant was used to match the trend in the luminous region results.

the scan-by-scan variation was decreased from the preliminary 0.8% to 0.24% for the 2017 calibration, while for the 2018 calibration, a similar level was maintained going from 0.32% to 0.25%.

The biggest improvement in all years comes from a large decrease in the non-factorization systematic uncertainty as seen in Table 3.3, but for results concerning 2017 and before, the combination of the beam-beam deflection and dynamic beta uncertainties also contributes significantly.

Year	Preliminary		Final	
	Non-fact. [%]	Total norm. [%]	Non-fact. [%]	Total norm. [%]
2015	1.5	2.3	0.5	1.3
2016	0.9	1.5	0.5	1
2017	0.8	1.6	0.33	0.73
2018	2	2.1	0.36	0.64

Table 3.3: A summary of the preliminary and revisited non-factorization and total normalization uncertainties for the proton-proton luminosity calibrations of Run2.

3.11 Linearity and stability

The accuracy of the measurement is not guaranteed by a set of well-established calibration constants in the vdM fill. By using the same calibration in the high-luminosity data collection regime (physics conditions) we rely on the linear response of our detectors.

Additionally, since the calibration is performed once a year, we must monitor and gradually correct every time- or dose-dependent inefficiency of the detectors.

3.11.1 Emittance scans

A large amount of effort is put towards developing techniques that allow for an intrinsic correction of the response of the detectors since this allows for a cross-check between independent quantities as the final step. One of the techniques used to this end in [80–83] is the emittance scan analysis.

It is customary to perform a pair of emittance scans at the beginning of each fill. These are vdM-like scans performed in physics conditions. Although the

reduced number of scan points, the lack of control of many systematic sources, and missing reproducibility checks do not allow for these measurements to be used for calibration, by comparing the measurements to each other the relative efficiency of the detector can be tracked. That is, for a linear detector. In case of a finite non-linearity, the obtained σ_{vis} also depends on the instantaneous luminosity, which may bias the measurement. In Appendix D I show in great detail, that for a non-linearity of

$$l^m = l^0(1 + \alpha l^0), \quad (3.31)$$

where l^m and l^0 stand for the measured and true value of the single bunch instantaneous luminosity (SBIL) and α is the measure of non-linearity per unit of SBIL, the dependence is approximately

$$\sigma_{\text{vis}}^m \approx \sigma_{\text{vis}}^0(1 + 0.41\alpha l_{\text{peak}}^0),$$

where l_{peak}^0 stands for the SBIL at the head-on position. The previous assumption was that there is no such factor multiplying α .

Since the luminosity of the colliding bunch pairs has some spread, it is also possible to extract the non-linearity by plotting the obtained σ_{vis} values as a function of the peak SBIL. However, the formula above also highlights, that the slope of the line fitted to the aforementioned points is only about half of the true non-linearity despite the initial expectations. Missing this, one could end up correcting only half of the non-linearity of their detector.

The implications of this formula have been taken into account for the efficiency correction of PLT in the 2017 dataset in [71]. Based on my work presented in Appendix D, strictly behind closed doors, this factor is jokingly referred to as the Major factor in BRIL. A similar relationship for the beam overlap width (Σ) is also derived.

3.11.2 The relative non-linearity of PLT

Due to accidental coincidences, the PLT detector is susceptible to overcounting at high luminosities which leads to a non-linear response. The behavior can be modeled using Equation 3.31. On the other hand, HFET has a sufficiently linear response to provide a reference for PLT to be corrected against. DT and RAMSES are deliberately not considered to keep them independent for the final uncertainty measurement.

The gradual decrease of luminosity in a fill allows for a comparison of the relative response of the detectors. In each fill the ratio of luminosity as measured by PLT and HFET is plotted as the function of the per bunch luminosity as measured by HFET. To compute the ratio, the data is averaged over 15 lumi sections (LS) first. A lumi section, representing 2^{18} LHC orbits or about 23.31 seconds, is the standard granularity (and the basic time unit) for long-term luminosity information.

The resulting data is fit with a straight line $\alpha x + b$. As the two luminometers are calibrated $b \approx 1$, but it is taken into account for the determination of the relative non-linearity, $\alpha = a/b^2$, to account for possible efficiency changes over time. The results are smoothed using a quadratic interpolation function as not all fills had a large enough luminosity range to perform a reliable fit.

This quantity is used to revert the quadratic formula in Eq. 3.31 and obtain a correction. Formerly $l^m(1 - \alpha l^m)$ was used for such applications which deviates from the proper solution by up to 0.5% in certain regions of 2017. The relative linearity of the two detectors is shown in Figure 3.34 before and after the correction. The average single bunch instantaneous luminosity in 2017 is 5.3 Hz/ μb , hence the effect of the correction reaches up to 6%!

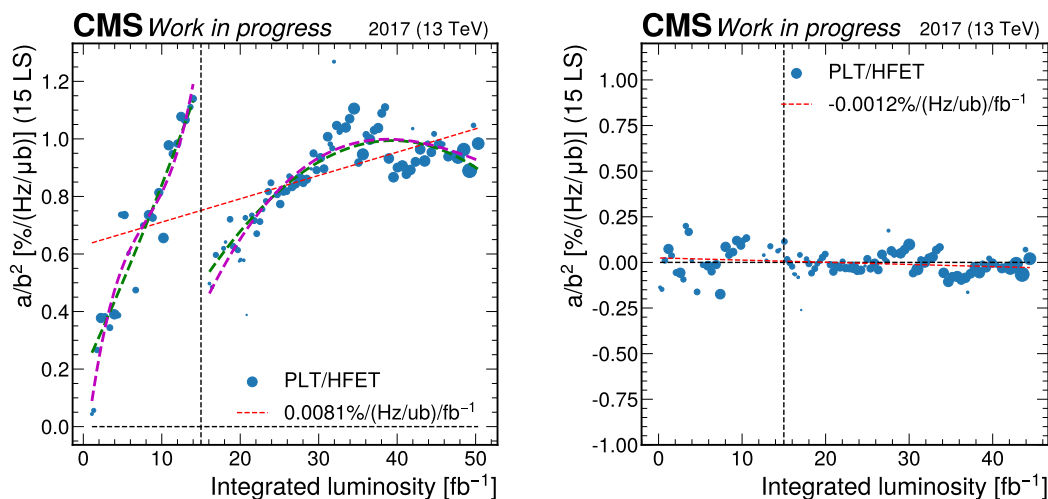


Figure 3.34: The relative residual non-linearity between PLT and HFET in 2017 before (left) and after (right) the correction of the PLT response. The size of the dots indicates the integrated luminosity in each fill. On the left plot, the values are interpolated using a second-order polynomial (green) and a smoothing spline (purple) which give consistent results. A general trend line is also shown, which however only describes the data properly after the correction.

3.11.3 The mean luminosity approach

In previous luminosity publications [80–82, 102] it was customary to obtain the final luminosity using the single most trusted luminometer available at the given point in time, and the consistency uncertainty was derived in a comparison to the second-best luminometer.

With the improved trust in the PLT, PCC, HFET, and HFOC luminometers, the average luminosity of these detectors is introduced. In practice, a cleaned mean is used where detectors deviating more than 10% from the local median value are not taken into account for the averaging.

This procedure rejects outliers effectively and provides a stable mean luminosity. The composition of the result depends on both the availability of the luminometers as well as the momentary stability of the individual measurements. Details of the composition are shown in Table 3.4.

2017	
source	contribution [%]
HFET+PCC+PLT+HFOC	89.62
HFET+PCC+HFOC	4.82
HFET+PLT+HFOC	3.94
PCC+PLT+HFOC	1.20
HFET+HFOC	0.22
PLT+HFOC	0.11
2018	
source	contribution [%]
HFET+PCC+PLT+HFOC	76.74
HFET+PCC+HFOC	15.81
HFET+PLT+HFOC	5.55
HFET+PCC+PLT	1.30
HFET+HFOC	0.46

Table 3.4: The contribution of different detector combinations to the total mean luminosity in 2017-'18 data. Sources with a contribution of less than 0.1% are not added to the table.

To gauge the consistency of the independently calibrated detectors used in the averaging, their ratio to the mean is computed per 50 LS, see Figure 3.35.

Only data that is included in the cleaned average is used for each detector. A good agreement is found over time, especially in 2018. The mean and standard deviation of the ratios are shown in the bottom plot. The spread of the means is assigned as a conservative consistency uncertainty, 0.43% and 0.39%.

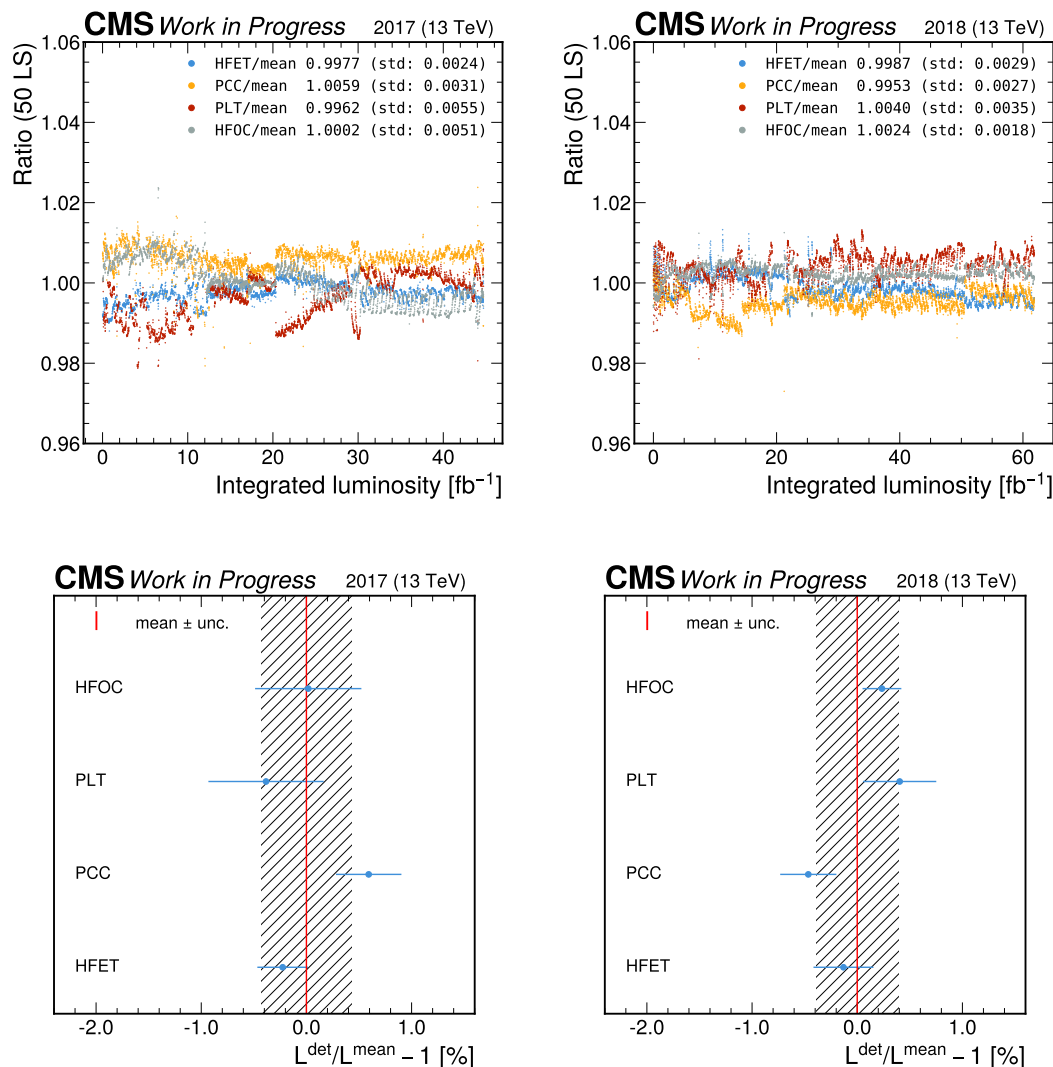


Figure 3.35: The ratio of HFET, HFOC, PLT and PCC to their cleaned average per 50 LS for 2017 (left) and 2018 (right). The upper row shows the ratio as a function of integrated luminosity, while the lower row displays the mean value of the ratios. The error bars represent the spread of the ratio distribution, and the hatched area shows the standard deviation of the points.

The linearity of the average luminosity is compared to DT and RAMSES which are cross-calibrated detectors that serve as linearity standards. The measurement is carried out as described in Sec. 3.11.2, and the results are shown in Figure 3.36 for RAMSES, which showed the larger deviations from zero in both years. To give a conservative uncertainty, the larger of the mean and the standard deviation

is multiplied by the average single bunch instantaneous luminosity. In 2017 and 2018 this was 5.33 and 5.11 Hz/ μb leading to a 0.48% and a 0.45% uncertainty respectively.

The mean luminosity approach played an important role in decreasing the total integration uncertainty from 1.7% to 0.64% and from 1.3% to 0.60% for the two years respectively, but this result could not have been achieved without the new linearity and out-of-time corrections, and the new veto list for PCC in 2018, that brought the detectors to the same level of reliability.

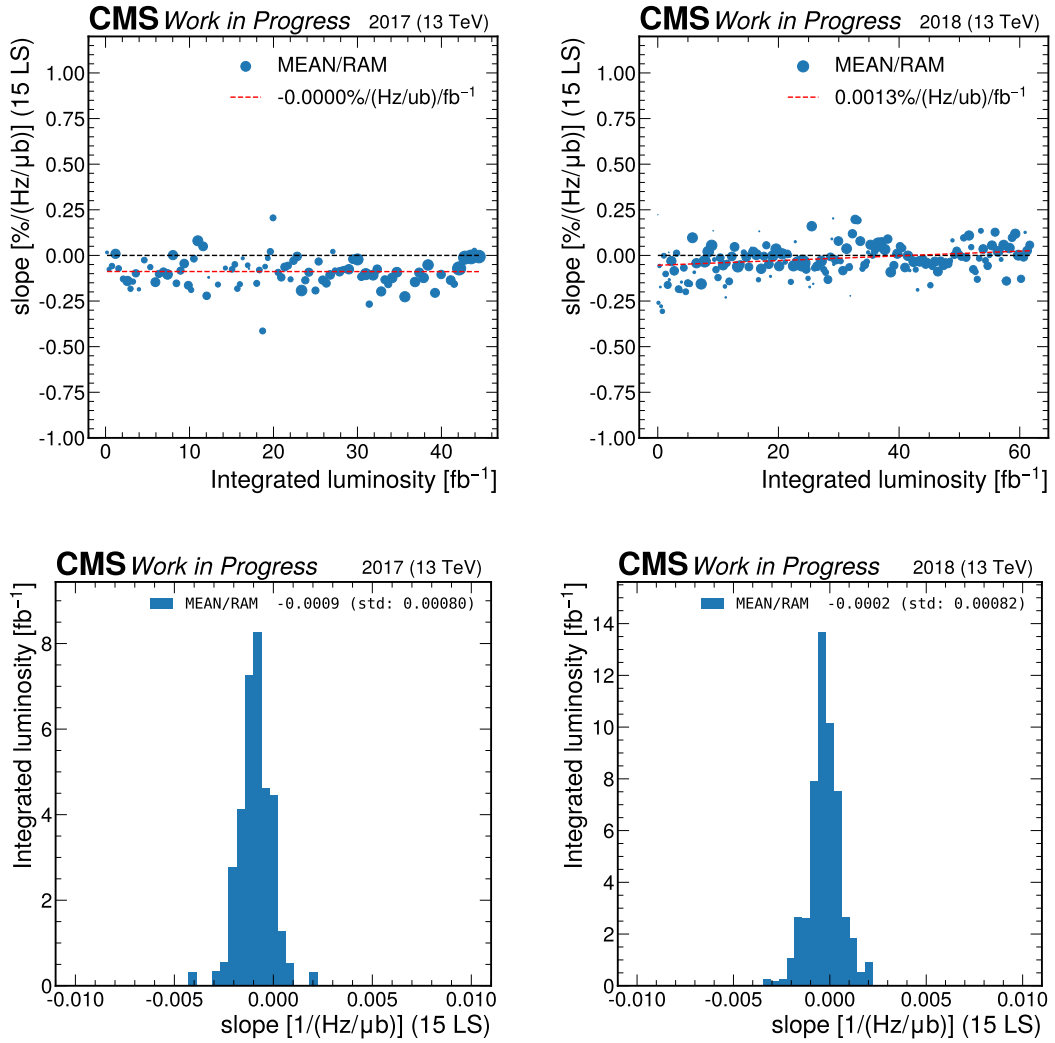


Figure 3.36: The trends (top) and histograms (bottom) of the relative non-linearity between the mean luminosity and RAMSES for 2017 (left) and 2018 (right). In the top panels, the size of the dots indicates the integrated luminosity in each fill. The histograms are weighted by luminosity. The number in the legend indicates the mean and standard deviation of the histogram.

3.12 Summary

Luminosity is a key ingredient to precision Standard Model measurements. As a demonstration two versions of a plot are shown in Figure 3.37. The preliminary version is found here² while the updated version is in [111]. The plots showcase the agreement between the measured and the theoretical weak gauge boson production cross sections. The top panel uses the preliminary luminosity calibration [81] published a bit more than half a year after the data taking for 2017 finished. In the new 2017-2018 study [71] the detectors are recalibrated, and both the central value and the uncertainty of the luminosity are updated. The new version of the plot in the bottom panel using the updated central value of the luminosity for the 13 TeV dataset showcases a drastic improvement in the consistency of the theoretical model and the measured values. This shows the importance of the immense progress that was made for the understanding of the absolute luminosity calibration with the vdM method and of the luminometer response from the quick preliminary results within a year of the data taking to the final precision study many years in the works.

During my work at BRIL I contributed to the paper revisiting the 2015-2016 luminosity measurement which decreased the total uncertainty of the years from 2.7% to 1.6% and from 2.5% to 1.2% respectively, and played a significant role in the new 2017-2018 calibration as the coordinator of the publication. The total uncertainties decreased from 2.3% to 1.0% and from 2.5% to 0.9% respectively, which brings the combined Run2 uncertainty to 0.7%.

A large contribution to these improvements is the direct result of my work presented in this chapter, especially, but far from uniquely, related to the estimation of the non-factorization bias arising from the non-Gaussian bunch proton densities.

I also participated in the luminosity calibration for lead-ion collisions that resulted in a 1.7% accuracy for the 2018 dataset, by developing the first method that could achieve a percent-level precision on the non-factorization bias in Pb+Pb data [103, 112]. No similar result was available previously, but it could be compared to the 2016 proton - nucleus calibration that reached a 3.5% total uncertainty of which 2.5% came from the non-factorization bias [113].

²<https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined>





Figure 3.37: Summary of the production cross section of weak gauge bosons, measured by CMS using preliminary 2017 luminosity (top) and updated luminosity values (bottom) for the 13 TeV dataset, plotted against the pp centre-of-mass energy ranging from 2.76 to 13 TeV. The error bars around the experimental data points represent the total uncertainty of the measurement. The measurements are compared with theoretical predictions (black lines) obtained at N³LO in QCD using the MSHT20an³LO PDF set. The grey band shows the envelope from normalization and factorization scale variations.

Chapter 4

In pursuit of new physics

The beneficial properties of R-parity conserving GMSB SUSY theories are discussed in Section 1.3. The experimental signature of a large class of these theories are available for probing at the LHC using statistical methods. In this work, a search for signs of supersymmetry in final states with an energetic photon, a Higgs boson decaying to $b\bar{b}$ pair and large transverse missing momentum is presented using the Run2 data set. Such a final state is abbreviated as GbbMET.

The study is carried out in a collaboration with several colleagues. Their work is also briefly described. My contribution lies in the study of the lost lepton background, formulating the statistical model used for interpreting the results and optimizing the analysis cuts based on the expected reach, hence these aspects are highlighted.

Previous work probing the GbbMET channel were published by ATLAS using 20.3 fb^{-1} 8 TeV Run1 data [114] and 36.1 fb^{-1} of 13 TeV Run2 data [115], and by CMS using the full Run2 dataset with a different analysis strategy in [116]. The current analysis is expected to extend the reach of the so-far presented exclusions for the neutralino masses set in gaugino pair production and shows the strongest expected limits along the $m_{\tilde{q}} = m_{\tilde{\chi}_1^0}$ mass diagonal in the gluino pair production case.

4.1 Analysis strategy

To infer the presence of a non-SM process an excess in the number of events in a particular final state compared to the SM predictions can be looked for.

Determining the SM prognosis is often called background (BG) estimation or background modeling. To gain a better signal-to-background ratio a signal region is defined based on the properties of the reconstructed objects (Section 4.4) in the final state. The signals of interest are introduced in Section 4.3 and the signal region is defined in 4.5.

Monte Carlo (MC) simulations are frequently used for background determination. These simulations undergo post production tuning to ensure their conformity to the data. Still, in general, control regions are used to cross-check or anchor their yields to data near the region of interest. The second approach is often called semi-data-driven and is utilized in Section 4.7, but fully data-driven methods also exist, such as the ABCD, applied in Section 4.6.3, or the tag-and-probe methods that are primarily used to measure selection efficiencies or fake rates as outlined in 4.6.2. The background estimation of this analysis is discussed in Sections 4.6-4.7.

This analysis uses the "cut-and-count" method to probe for excesses in the recorded data, where the number of events in the signal region is compared to the SM expectation. The significance of the excess is then determined using statistical hypothesis testing. The interpretation of the result is discussed in Section 4.8.

4.2 Data samples

This search exploits the 137 fb^{-1} collision data recorded at a 13 TeV centre-of-mass energy by the CMS detector at LHC in Run2. The 2015 dataset constitutes an exception, as it is not used due to its small size. The Run2 period is not homogenous as far as the conditions and performance of the detector go, four mean eras are distinguished for the period utilized: 2016preVFP, 2016postVFP, 2017 and 2018 with the corresponding luminosities of 19.5, 16.8, 41.5, and 59.8 fb^{-1} .

2016 is divided into two regions due to the strip tracker readout chips experiencing saturation at high luminosity. This was mitigated by lowering the pre-amplification gain for the readout chip mid-year. The pixel detector was upgraded during the end-of-year technical stop between the 2016 and 2017 data taking periods.

Besides the detector conditions, the change in pileup in the three years

is also significant, as shown in Figure 4.1, which affects the trigger and the reconstruction efficiency.

The SUSY signal and SM background processes are studied using Monte Carlo (MC) simulated samples, generated centrally by the experiment, separately for each of the four eras to accommodate for these differences. The simulations use state-of-the-art MC event generators [117–119] and a full model of the CMS detector implemented in Geant4 [120]. For the SUSY signal grid, partially parametrized "fast" simulations were employed. Differences with the detailed "full" simulation were studied and a related uncertainty was added to the statistical model.




Figure 4.1: The calculated distribution of the average number of interactions per crossing (pileup) for pp collisions for all years in Run2. For the calculation, a minimum bias proton-proton cross section of 80 mb is used, which is the "LHC standard" computed with the Pythia Monte Carlo event generator.

4.3 Signals

The targeted signal is described by simplified GGM models where the SUSY particles are created either by gluino pair production or a gaugino pair production. The two scenarios are handled independently. The Feynman diagrams of the related processes are seen in Figure 4.2. These models are indexed by the mass of the sparticles appearing in the diagrams, $(m_{\tilde{q}}, m_{\tilde{\chi}^0})$ and $m_{\tilde{\chi}^0}$ respectively. The properties of the models are studied using Monte Carlo samples. The MC

⁰<https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>

samples are generated from $m_{\tilde{\chi}_1^0} = 300$ GeV to 1600 GeV in 53 equidistant points for the electroweak production, and in a two-dimensional grid with a varying step size that covers the $127 < m_{\tilde{q}} < 2790$ GeV and $800 < m_{\tilde{\chi}_1^0} < 2800$ GeV region for $m_{\tilde{\chi}_1^0} < m_{\tilde{q}}$ with 951 points in total. The grid is illustrated in Figure 4.3.




Figure 4.2: The Feynman diagrams of the signals considered. The left side shows the strong production channel ($T5qqqqHg$) while the right side panel shows the simplified electroweak production channel ($TChiNg$).

As mentioned in Sec. 1.3.2, in the GMSB framework the lightest supersymmetric particle (LSP) is always the gravitino (\tilde{G}) while the next-to-lightest sparticle (NLSP) is the lightest neutralino ($\tilde{\chi}_1^0$). The gravitino is light. With its mass in the eV - GeV range, it can serve as warm dark matter (DM), either alone or as one of the components of DM. In simulations at the LHC, its mass is typically set to 1 GeV but the collider phenomenology is largely insensitive to the chosen value.

In the models with the strong production channel, a pair of gluinos are created which decay to colored SM particles and the NLSP. Depending on the mass difference $m_{\tilde{q}} - m_{\tilde{\chi}_1^0}$, the events may have different levels of hadronic activity. The NLSP decays to an LSP and either a photon or a Higgs boson with $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow \tilde{G}\gamma) = \mathcal{B}(\tilde{\chi}_1^0 \rightarrow \tilde{G}H) = 50\%$ branching ratios. This analysis concentrates on the case where the final state contains one of each, which is one-half of the total cases. All Higgs boson decays are included in the simulations but we concentrate on the dominant decay to b quark pairs ($\mathcal{B}(H \rightarrow b\bar{b}) = 57\%$). Depending on the momentum of the Higgs boson, the resulting b-jets may be picked up as separate objects, or they may be so close that they merge. Finally, the stable LSP escapes detection while carrying a large momentum, which leads to a significant p_T^{miss} completing our signature.

In the weak production scenario, the cases where $\tilde{\chi}_1^0$ is more bino- or higgsino-like are handled using the same sample, only the cross-section of the process is




Figure 4.3: The mass grid of the Monte Carlo samples used for the direct search for gluino pair production.

changed accordingly. Depending on the nature of the NLSP, different gaugino pairs can be produced in the hard interaction, which are almost mass-degenerate with the lightest neutralino in this model. The predicted cross sections [121] are shown in Fig. 4.4 for the various processes. Only soft radiation is emitted in the chargino and second lightest neutralino decay to the NLSP. The exact production is hence hidden in the bulge in the Feynman diagram. In our sample, the lightest neutralino always decays by emitting a photon on one arm of the process, while on the other we have $\mathcal{B}(\tilde{\chi}_1^0 \rightarrow H\tilde{G}) = \mathcal{B}(\tilde{\chi}_1^0 \rightarrow Z\tilde{G}) = 50\%$. The events are filtered using truth information for the Higgs final states, the cross sections do not need modification. The kinematic behavior of these events is similar to the strong production case with a low $m_{\tilde{q}} - m_{\tilde{\chi}_1^0}$ difference, and thus low jet activity.




Figure 4.4: Theoretical cross sections for sparticle pair production at NLO+NLL approximation for electroweak gaugino and sfermion production and at NNLO+NNLL approximation for strong production processes. Source¹.

4.4 Object definitions

The CMS object reconstruction procedure aims to serve analyzers with a wide range of requirements for efficiency and purity. Objects reconstructed are labeled with a set of quality indicators based on which the users can refine their own selections. Often the specific values picked are motivated by the availability of centrally produced correction factors that match the object selection efficiencies in the Monte Carlo samples to that measured in data. These are the so-called scale factors (SF) that are typically used as event weights with values close to 1. Whenever SFs are applied, their measurement uncertainties are also propagated to the results. A frequently used quality indicator is the identification (ID) working point. This indicates the purity of the sample going from low to high (and correspondingly the efficiency from high to low) usually labeled as "loose", "medium", and "tight". Many analyses, such as this one, also use custom kinematic or particle objects. The selections made for this analysis follow.

¹<https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections>

Leptons

Since events with leptons constitute a possible SM background in this search, leptons are vetoed in the search bins. For this reason, the imposed selection cuts are all very loose, on top of requiring that the objects stay within the recommended limits defined by the detector acceptance. Events with leptons are collected separately to be used for background estimation (Section 4.7).

Let's first define a few variables that are particularly useful to select leptons.

- "Mini-isolation" sums the transverse momentum of all particle flow objects within a cone around the lepton, with the cone size varying as a function of the lepton p_T : $R = 0.2$ when $p_T < 50\text{GeV}$; $R = 10\text{GeV}/p_T$ when $50 < p_T < 200\text{ GeV}$; and $R = 0.05$ otherwise. The sum is corrected for the average energy density in the event scaled by the cone's "effective area", and then normalized by the lepton p_T . This isolation variable is particularly well suited to catch isolated leptons in events with significant hadronic activity but also performs well in clean environments.
- Impact parameter measures the distance of closest approach of the lepton track to the primary vertex. It can be calculated in three dimensions (d_{3D}) or its projections can be considered in the transverse plane (d_0) and in the longitudinal direction (d_z).
- The significance of the measured value of the impact parameter is calculated by normalizing it with the measurement uncertainty.

Muon objects must have $p_T > 5\text{ GeV}$ and $|\eta| < 2.4$, satisfy the loose ID criteria, and have a "mini-isolation" below 0.2. The 3D IP significance is required to be lower than 4, and the impact parameters must satisfy $d_z < 0.1\text{ cm}$, $d_0 < 0.05\text{ cm}$, following the central recommendation. A good description of muon selection in simulations requires the use of scale factors both for identification and for isolation efficiency.

Electrons must have $p_T > 10\text{ GeV}$, $|\eta| < 2.5$, and mini isolation < 0.2 . The central cut-based loose ID criteria are applied. Electrons are rejected if they are close in angular distance to selected muon candidates $\Delta R(e, \mu) < 0.3$. Electrons necessitate the use of reconstruction and identification efficiency scale factors.

Tau leptons are required to have $p_T > 20\text{ GeV}$, $|\eta| < 2.3$, and the central deep learning based ID algorithm is used with a loose cut to select them. Taus are

rejected if they are close to a selected muon or electron object: $\Delta R(\tau, \mu) < 0.3$ or $\Delta R(\tau, e) < 0.3$. On top of the identification efficiency scale factors, they also need an energy scale correction applied in the analysis. As for all corrections necessary to match the simulations to data, the uncertainty on this is propagated to the final results.

Photons

As the photon is one of the central objects in this work, the events are selected using a photon trigger (see Sec. 4.5). The minimum photon transverse momentum is established slightly above the trigger threshold to ensure a stable and high trigger efficiency: $p_T^\gamma > 210$ GeV. We use an identification algorithm based on multivariate analysis (MVA) that guarantees a 90% efficiency to increase the statistics. Photons from the region $1.4442 < |\eta| < 1.566$ where the barrel and the endcap subdetectors meet are excluded. If a photon leaves a hit in the pixel detector, it is labeled as a **fake photon** and is set aside to be used for background estimation (Section 4.6.2). Photons are rejected if they are close to selected leptons ($\Delta R < 0.3$). The centrally recommended energy scaling and smearing corrections are applied. Identification and pixel seed efficiency scale factors are also used.

Jets

The central jet clustering algorithm, anti-kt [122], is used for this study with radii $R = 0.4$ and $R = 0.8$ to obtain the so-called AK4 and AK8 (or fat) jet categories. In both cases, $|\eta| < 2.4$ and a tight ID is required which protects against noise in the detector and badly measured jets.

Fat jets must have $p_T > 200$ GeV, while for AK4 jets the cut is $p_T > 30$ GeV. For soft jets with $p_T < 50$ GeV a tight (loose) pileup ID is used in 2016 (2017-18) to remove instances coming from pileup. Jets are removed from the selection if they overlap with leptons or photons within an angular distance of $\Delta R < 0.4$ and $\Delta R < 0.8$ for AK4 and AK8 jets, respectively. The centrally recommended jet energy scaling and smearing corrections are applied.

All AK4 jets are assigned a b-tag discriminator using the DeepJet algorithm [123]. The analysis uses the loose and medium working points for event selection. A b-tagging efficiency correction is used based on the central recommendation.

Fat jets are scored by the ParticleNet double b-tagger algorithm [124] that

identifies merged b-tagged jets originating from the decay of a heavy boson using also jet substructure discriminants. Later, the loose working point is used for selection.

The centrally recommended energy scaling and smearing corrections are applied.

Transverse mass

The transverse mass of the most energetic photon and the missing transverse energy is defined as

$$M_T^2 = 2p_T^{\text{miss}} E_T^\gamma (1 - \cos |\Delta\phi(\vec{p}_T^{\text{miss}}, \vec{p}_T^\gamma)|).$$

In the case when the leading photon is a misidentified electron coming from a $W \rightarrow e\nu$ decay, and the missing transverse momentum originates from the undetected neutrino ($p_T^{\text{miss}} = p_T^\nu$), this quantity has a cut off at the W boson mass, i.e. $M_T < M_W$. This can be used for $W \rightarrow e\nu$ background rejection, both in the case of direct W production and top quark decays ($\mathcal{B}(t \rightarrow Wb) \approx 99.8\%$).

4.5 Event selection and binning

Trigger

The events are selected using an un-prescaled single photon trigger with the lowest possible energy threshold requirement in each era of Run2. For a higher efficiency, triggers with higher energy threshold but looser ID requirements were included in a logical OR. For the majority of the dataset, the trigger threshold was 200 GeV (except in 2016, where it was 165 GeV). The efficiency of the trigger is measured using data samples selected by an independent jet trigger as a function of the reconstructed and identified photon kinematic quantities. It was found to be above 99% (and fairly independent of the photon properties) above the 210 GeV photon p_T requirement. The uncertainty on the trigger efficiency is 1%.

Base cuts

The following conditions must hold for all events in the signal region:

- At least one well reconstructed interaction vertex is present.
- The event passes a set of very loose event cleaning cuts which ensure that the missing transverse momentum is well measured.
- At least one photon with $E_T > 210$ GeV is found. Fake photon events are used for background estimation.
- Leptons are vetoed. Leptonic events are used for background estimation.
- At least two AK4 or one AK8 jet is present.
- $p_T^{\text{miss}} > 300$ GeV. Lower p_T^{miss} events are used for background estimation.
- For all b-tagged AK4 jets, the distance in azimuth angle between the jet and the missing transverse momentum vector $\Delta\phi(\vec{p}_T^{\text{miss}}, \vec{p}_T^{\text{bjet}}) > 0.3$ to veto events where significant contribution to p_T^{miss} comes from neutrinos emerging from semi-leptonic B hadron decays.

The efficiency of the base cuts to select the SUSY signal grows with the sparticles masses. At the highest mass values, it is about 45% both for gluino and for gaugino pair production. At lower masses the efficiency drops (for example for EW gaugino production with a 400 GeV neutralino, it is only about 5%), however, this is compensated by the higher production cross section values (see Figure 4.4).

Higgs tagging

This procedure aims to tag events with a $H \rightarrow b\bar{b}$ decay based on the invariant mass of the final state b-jets. Two cases are distinguished. In the **boosted** case (B), due to the large momentum of the Higgs boson, the direction of the two b-jets are so close that they are reconstructed as a single AK8 jet, while in the **resolved** case (R) they can be separated into two AK4 jets.

In the boosted case an AK8 jet is considered a Higgs candidate, if it has at least a loose double b-tag score and an invariant mass in the 80 to 160 GeV range, while it is considered a fake Higgs candidate if the invariant jet mass is in the right range, but the b-tag discriminator is not sufficiently large. For this analysis, the fat jets are sorted in decreasing order according to their b-tag score, and the first that falls into the 80 to 160 GeV range is picked.

In case no AK8 jet in the appropriate mass range is found, the AK4 jets are examined in a similar manner. In the resolved case, the AK4 jets are sorted in decreasing order according to their b-jet discriminators. The first AK4 jet is picked as one of the possible decay products and the rest are tested against

it. The pair must have at least $\sum \vec{p}_T > 100$ GeV. The first couple for which $80 < m_{jj} < 160$ GeV is considered a proper Higgs candidate if one of the pair is at least a medium and the other is at least a loose b-tagged jet. Otherwise, it is considered a fake Higgs candidate.

If no Higgs or fake Higgs candidate is found, the event is rejected. Fake Higgs events are used for background estimation.

The Higgs tagging efficiency when both AK8 and AK4 selections are considered can be as high as 60% at high sparticle masses. For EW production, it is always higher than 45% for the studied neutralino masses (>300 GeV). For gluino production, the lowest tagging efficiency of about 30% is observed at the lowest gluino mass studied (700 GeV) for low values of the neutralino mass (200 GeV).

4.5.1 The signal and validation regions

In the signal region, the events are subdivided into several categories based on the missing transverse energy, the Higgs tag category, and the number of extra jets (N_j). In the boosted case, N_j is defined as the number of AK4 jets not overlapping with the AK8 Higgs candidate ($\Delta R > 0.8$). In the resolved case, the two AK4 jets forming the Higgs candidate are subtracted and $N_j = N_{AK4} - 2$. In the $N_j < 4$ region, $M_T > 100$ GeV is required to reduce the W background where the daughter electron is misidentified as a photon.

Table 4.1 shows the resulting binning. The basic structure of the table has a four-bin periodicity, each one bracket higher in p_T^{miss} , except for the very last group, where all bins are merged for lack of statistics. The "QCD validation region" is mainly dominated by the multijet events, while the global validation region has a similar composition to the signal region.

	QCD VR				Global VR				Signal Region				
Bin ID	1	2	3	4	1	2	3	4	1	2	3	4	5
Higgs tag	R	B	R	B	R	B	R	B	R	B	R	B	R&B
Extra jets	$N_j < 4$		$N_j \geq 4$		$N_j < 4$		$N_j \geq 4$		$N_j < 4$		$N_j \geq 4$		$N_j \geq 0$
p_T^{miss} [GeV]	$100 < p_T^{\text{miss}} < 200$				$200 < p_T^{\text{miss}} < 300$				$300 < p_T^{\text{miss}} < 500$				$p_T^{\text{miss}} > 500$

Table 4.1: The binning of the signal region and the two validation regions.

4.6 Background processes

The SM background contributions after the event selection – considering that a significant amount of missing transverse energy and an energetic photon must be present – can be grouped into four categories:

"Z($\nu\nu$)" events In the case of an invisible Z boson decay, $Z \rightarrow \nu\bar{\nu}$, the missing energy originates from the undetected neutrino pair and the photon can arise from initial state radiation. In MC simulation, this group includes all single and diboson processes with a $Z \rightarrow \nu\bar{\nu}$ decay (e.g., Z, ZZ, ZW).

Misidentified electron (" $e \rightarrow \gamma$ ") events In this category, p_T^{miss} is either due to a neutrino, for example when a $W \rightarrow e\nu$ decay happens, or it might originate due to mis-reconstructed final state objects, e.g. jets. In MC simulation, this group includes all processes where a prompt electron is produced (e.g., $W \rightarrow e\nu$, $Z/\gamma^* \rightarrow ee$) and then misidentified during the analysis.

"Lost lepton" (LL) events In the leptonic decay of weak gauge bosons p_T^{miss} can arise if a lepton is not reconstructed due to either the acceptance or some imperfections of the instrumentation, or the unavoidable inefficiencies of the reconstruction algorithm. If this is accompanied by an energetic photon radiation either in the initial or the final state, a signal-like event is produced. Alternatively, a jet can also be misidentified as a photon. In MC simulation, this group includes all processes where a prompt lepton (electron, muon, tau) is produced, including single and pair production of W, Z/ γ^* , and top quark.

Multijet "QCD" events In these events, the missing momentum arises primarily from poorly reconstructed jets, with either a real direct photon or a jet misidentified as a photon. In MC simulation, this group includes all QCD processes with no direct lepton, including direct multijet production or fully hadronic decay of top quark pairs, as well as single or pair production of weak gauge bosons followed by hadronic decay.

The overlap between the first three categories is systematically removed, giving preference to the categories according to their listing, i.e. Z($\nu\nu$) highest.

The composition of the signal region and the validation regions are demonstrated in Figure 4.5 with the help of Monte Carlo simulation, as described above. It is seen, that the signal region inhabits an extreme corner of the kinematic

phase space, and very few SM events are left as background. On one hand, this is beneficial, since with a good signal efficiency this guarantees a good signal-to-background ratio. On the other hand, it makes it difficult to find sidebands with large statistics, and in turn, background estimations are very cumbersome. In fact, the two leading backgrounds are estimated by incorporating their control regions into the statistical model that describes our observed data, using the signal and background predictions. This ensures that all background uncertainties are adequately handled. A short explanation of each contribution follows with a focus on the lost lepton estimation.

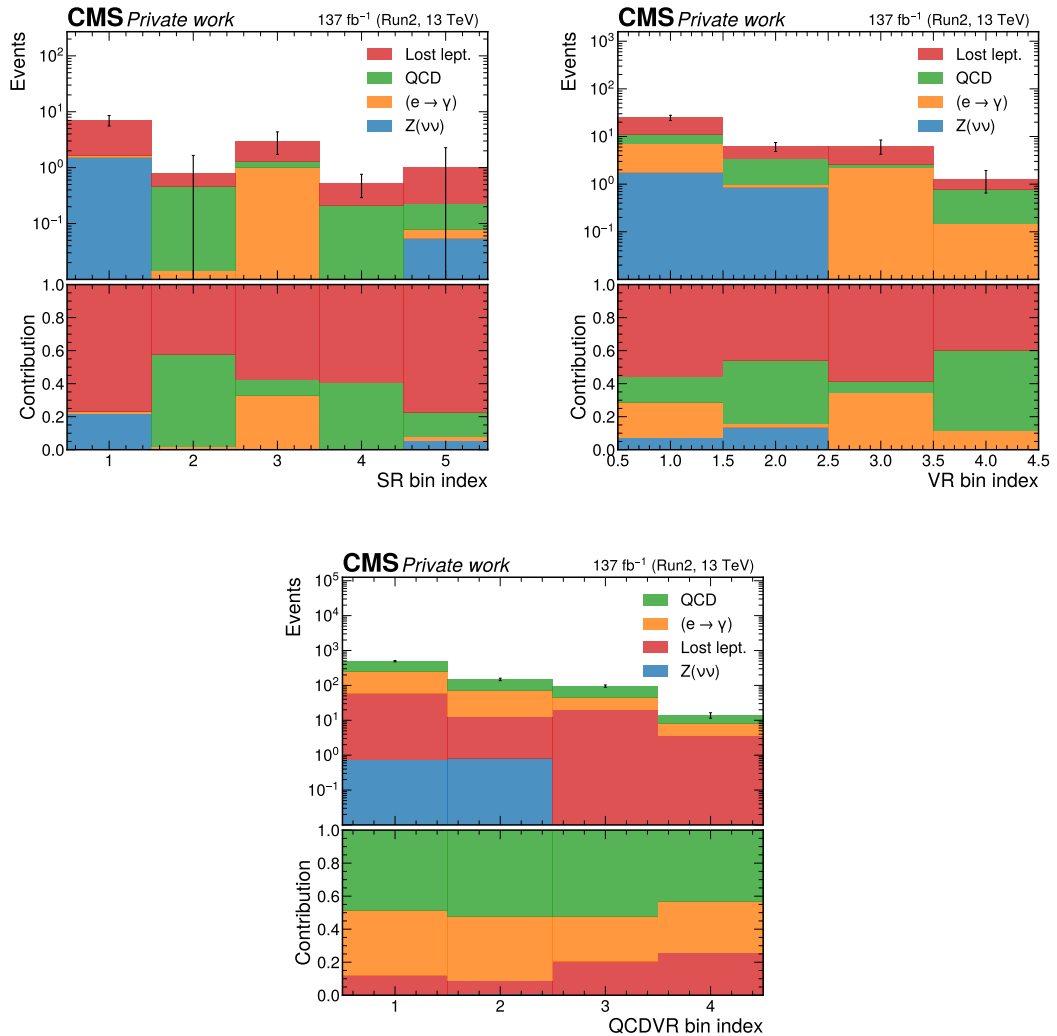


Figure 4.5: The composition of the signal and validation regions based on Monte Carlo simulation. The four background categories are explained in the text. The black error bars indicate the statistical uncertainty on the total prediction.

4.6.1 The $Z(\nu\nu)$ background

The Z boson decays to neutrinos in about 20% of the time, leading to a p_T^{miss} equivalent to the transverse momentum of the Z boson. Such events can contribute to our signal bins when the quarks giving rise to the Z boson in the hard interaction also radiate away a few gluons and a high-energy photon. A Feynman diagram shows an example of such a tree topology in Figure 4.6. Since this is a very rare event, it is estimated using Monte Carlo simulation which takes into account all corrections described above related to the modelling of the physics objects. It is important to note that special care is taken to match the distribution of pileup in all simulated samples to that observed in the data, as well as the uncertainty on this procedure. Additionally, the theoretical uncertainty on the cross section (5%) is also considered.




Figure 4.6: A tree level $Z(\nu\nu)$ topology possibly contributing to the signal region.

The prediction is cross-checked with a semi-data-driven estimation which extrapolates data from two different control regions to the signal region. In these control regions, besides all base cuts, an e^+e^- or a $\mu^+\mu^-$ pair is required with $80 < m_{\ell\ell} < 100$ GeV. The $\sum \vec{p}_T$ of the leptons is counted in with p_T^{miss} for the event selection. Since the statistics is poor, the Higgs tag requirement is inverted in one of the control regions. The extrapolation is performed using a Monte Carlo transfer factor. A good agreement is found within the statistical uncertainties. Nonetheless, the largest deviation between the MC prediction and the semi-data-driven estimates (15%) is assigned as an additional systematic uncertainty to the modeling of this background.

4.6.2 The misidentified electron background

Events where electrons are erroneously identified as photons can also constitute a background. In a $W \rightarrow e\nu$ decay, if the electron "fakes" a photon and the

neutrino leads to significant p_T^{miss} , the GbbMET signal topology can arise. This can happen similarly to the $Z(\nu\nu)$ case, if an initial quark radiates a gluon which splits into a b-quark pair.

This background is estimated using a control region with identical cuts and binning to the SR and VRs, only the pixel seed veto for the photon is inverted (fake photon). The events are extrapolated to the signal region using a transfer (or fake) factor measured in data in a "tag-and-probe" setup using $Z \rightarrow ee$ events. For this measurement, a single electron trigger is used with $p_T^e > 27$ GeV. An electron with tight ID requirement and matched in dR to the trigger electron acts as the "tag". Another electromagnetic object plays the role of the "probe", choosing it such that the tag-probe reconstructed invariant mass is consistent with the Z boson mass. To form the numerator (denominator) of the transfer factor, we count pairs where the "probe" is the leading (fake) photon. The transfer factor is measured in bins of (η, Φ) , E_T and the number of reconstructed vertices in the event. The transfer factor is assigned a 50% systematic uncertainty to account for the possible bias due to the factorized $2D \times 1D \times 1D$ parametrization and to cover for the statistical and systematic uncertainties of the per bin values.

4.6.3 The multijet background

The multijet background is one of the dominant backgrounds in this analysis. The b-jets could be genuine or a result of misidentification, and the photon might be a real photon created in a radiative process or a misidentified jet. In this scenario, the p_T^{miss} is due to jet energy mismeasurement which biases the p_T balance. These conditions are all rather infrequent, but the sizable cross section of the QCD processes still makes the yield significant. Multiple cuts were tested on the azimuth angle between the \vec{p}_T^{miss} vector and different jet objects to reduce the counts. Finally $\Delta\phi(\vec{p}_T^{\text{miss}}, \vec{p}_T^{\text{b-jet}}) > 0.3$ was selected.

This contribution is estimated in a data-driven way by the so-called ABCD method. Here two uncorrelated variables are used to define four regions with one of them, say D, being the signal region. The rate in a neighboring region to D can be extrapolated into D using the remaining two control regions to compute a transfer factor:

$$N_D = N_B \frac{N_C}{N_A}$$

In this case, p_T^{miss} and the "fakeness" of the Higgs tag are used as the two variables. The layout of the related regions is shown in Figure 4.7. The A' and C' regions

are used for systematic uncertainty estimation, and the validation regions have their dedicated bin-by-bin "A" regions (A_{VRs}) as well.

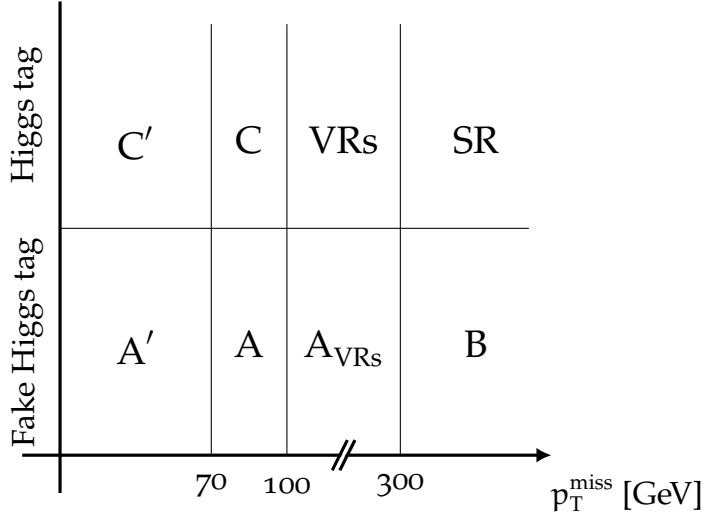


Figure 4.7: The control regions of the multijet ABCD background estimation method. Each region has a bin-by-bin sub-structure conforming to their corresponding "D regions" which can be the signal region (SR) or any of the two validation regions (VRs), as shown in Table 4.1. The detailed binning convention for the ABCD control regions is shown in Table 4.2.

Region	A'		C'		A		C		$A_{VR(QCD)}$				$A_{VR(Global)}$				B				B								
Bin ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
Higgs tag	FR	FB	R	B	FR	FB	R	B	FR	FB	R	B	FR	FB	R	B	FR	FB	FR	FB	FR	FB	FR	FB	FR	FB	FR	FB	FR & FB
N_j	< 4		≥ 4		< 4		≥ 4		< 4		≥ 4		< 4		≥ 4		< 4		≥ 4		< 4		≥ 4		≥ 0				
p_T^{miss} [GeV]	< 70				70 – 100				100 – 200				200 – 300				300 – 500				> 500								

Table 4.2: Binning structure of ABCD control region histogram. R and B indicate the resolved and boosted Higgs tag categories, respectively. FR indicates the fake resolved, and FB the fake boosted Higgs tag categories.

The correlations between the two variables is not expected to be exactly zero as b-jet on average have a larger number of neutrinos due to semi-leptonic B hadron decays. The effect is estimated using fully hadronic Monte Carlo samples (including QCD multijet, jet + photon, hadronically decaying top quark pair and weak gauge boson events):

$$K = \frac{N_{SR}^{MC} N_A^{MC}}{N_B^{MC} N_C^{MC}}$$

the value of K is measured in the QCD validation region with $100 < p_T^{miss} < 200$ GeV, as at higher p_T^{miss} the MC statistics is too poor to do so. This value is

measured using both the A-C and the A'-C' regions. The correlation is found to be smaller when the control regions are closer to the signal region (i.e. A and C). A systematic uncertainty of 16% is assigned accordingly. The statistical error of K is negligible with respect to this number.

Finally, the estimate in the signal region is

$$N_D = KN_B \frac{N_C}{N_A}. \quad (4.1)$$

As mentioned in the introduction, the control regions of this method are incorporated into the likelihood function of the statistical model, and the estimation is performed using free-floating parameters to ensure a correct modeling of the uncertainties.

4.7 The lost lepton background

SM events featuring a prompt lepton, for example from a $W \rightarrow \ell \nu_\ell$ decay, and accompanied by high hadronic activity may mimic the signature of our signal. The p_T^{miss} can be a result of the non-reconstructed lepton escaping the detector, a ν , or jet energy mismeasurement, while b-jets can arise in top quark decays or be a result of QCD radiation followed by gluon splitting, or false identification due to the large number of jets. Figure 4.8 shows a simple example.




Figure 4.8: A tree-level diagram of a W boson decay with the charged lepton lost (i.e. not observed in the detector) and a photon and a gluon radiated by the initial state quarks taking part in the hard interaction. Such an event can contribute to our signal region.

The lost lepton background (LL) is defined as the events where the lepton is not reconstructed for any other reason than being identified as a photon. This latter case is handled separately in Section 4.6.2. The yield of this process is estimated using a control region where the lepton veto is inverted. The counts are extrapolated using purely leptonic Monte Carlo events.

In the decay of the W boson, all three flavors of leptons can be created. Electrons and muons are stable at distances that concern the CMS experiment. On the other hand, tau particles are quick to decay further and are never seen directly by the detector. Approximately one-third of the time the decay is leptonic, producing an electron or a muon in almost equal (17.8% and 17.2%, respectively) proportions. These are mostly inseparable from genuine electron/muon events individually. The rest of the decays are hadronic, which are difficult to distinguish from narrow QCD jets.

Since our classification of the events can only rely on final state observations, we distinguish the following cases:

- No reconstructed (reco) lepton found (signal region, SR)
- Highest p_T reco lepton is an electron (control region: electrons, CR_e)
- Highest p_T reco lepton is a muon (control region: muons, CR_μ)
- Highest p_T reco lepton is a hadronic tau (control region: taus, CR_τ)

The combination of these disjoint control regions is $CR_{LL} = CR_e \cup CR_\mu \cup CR_\tau$. Events in this phase space rarely contain more than one lepton, so the highest p_T requirement is only present to resolve rare ambiguities.

While the definition of the electron and hadron categories are straightforward, we considered two different definitions for the hadronic tau region. Either a tau MVA tagger with a loose cut or an isolated track (isoTrack) reconstructed object is used, both with their auxiliary cuts further discussed in Section 4.7.1. The effectiveness of this change in definition is tested in MC using truth information. It is found that using the tau MVA tagger leads to a better discriminating power over the isoTrack method.

For the background estimation, two alternative versions of the same sideband extrapolation method are considered. The main difference lies in the choice of whether lepton flavors are considered individually or grouped together in the calculation. The first method addresses the fact, that hadronic tau leptons are much harder to identify than muons for example, while the second approach has statistical benefits at the cost of potential systematic effects. The difference is found to be small in the region of interest with respect to the statistical uncertainties. As a safety margin the average normalization difference, 20%, is taken as a systematic uncertainty. This is discussed in 4.7.2.

4.7.1 Selecting a hadronic tau definition

To inspect the efficiency and purity of the selected events using either of the strategies, "truth-reco response matrices" were created using the two hadronic tau definitions:

Tau candidate identified by an MVA tagger with a loose working point

- $|\eta| < 2.3$
- Not overlapping with an electron or muon within $\Delta R < 0.3$
- $p_T > 20$ GeV

Charged isolated track (isoTrack) fulfilling the constraints:

- $|\eta| < 2.5$
- Not overlapping with a muon or electron within $\Delta R < 0.3$
- Originates from the primary vertex with $d_z < 0.1$ and $d_{xy} < 0.1$ impact parameters
- Consistent with a W or Z decay: $m_T(p_T^{\text{miss}}, p_T^{\text{iso}}) < 100$ GeV
- Fulfills one of the two requirements for the relative charged hadron isolation within a cone of $R = 0.3$ ($I_{0.3}^{\text{ch}}$):
 - Tagged as an electron, muon or tau with $p_T > 20$ GeV and $I_{0.3}^{\text{ch}} < 0.2$
 - Not tagged as an electron, muon or tau with $p_T > 25$ GeV and $I_{0.3}^{\text{ch}} < 0.1$

The "truth" labels are assigned to a given event as follows: All truth events are searched for a charged prompt lepton. In the rare cases where multiple such particles are present, the one with the highest p_T is chosen. If this lepton is an electron or muon we assign them this label. In case it is a tau particle the truth particles are searched for tau decay products. (The direct tracing of the decay chain is not possible due to incomplete truth information.) The decay products are sorted into three groups: electrons, muons, hadrons. The category with the highest average p_T gives the label for the event. In most cases two of the three categories will be empty, this condition is to resolve ambiguous cases.

The "reco" labels are assigned based on the highest p_T lepton object from the electron, muon, tau or the electron, muon, isoTrack triplets.

Figure 4.9 shows a set of truth-reco response matrices filled based on the previously described labels. In general, the left panels use a tau MVA tagger for identifying hadronic tau decays, while the right-hand side plots use isolated

tracks for the same purpose. In order to assess the purity of the reco objects, all MC samples are used which are considered in this analysis as a potential source of background. Events, where the leading reco photon is a misidentified electron, are filtered out using truth information.

All events filled into the histograms have passed the base cuts of the analysis. Cuts used for binning in the search region were not applied. This means, that these response matrices are integrated over the number of extra jets, the Higgs candidate category, and p_T^{miss} . For the bottom row of Figure 4.9, a $p_T^{\text{miss}} > 200$ GeV cut was also included as the high p_T^{miss} requirement improves the purity of the reco objects drastically. Note, that the global validation region also has a lower p_T^{miss} limit of 200 GeV, while the signal region starts above 300 GeV, so it is reasonable to concentrate on this MET range. The non-leptonic contamination for low values of p_T^{miss} can also be seen on the panels of Figure 4.10.

From the response matrices, we can conclude, that the tau MVA tagger is superior with respect to the isolated track method for selecting hadronic tau decay events in terms of efficiency, capturing approximately 38% of the truth events instead of only 25% in the isolated track case, while its purity is also significantly better, 88% instead of 49%.

The contamination of the reco objects with non-leptonic events in the $p_T^{\text{miss}} > 200$ GeV region is about 1.5% for electrons, 2% for muons, and 2% for tau leptons. This is negligible with respect to the statistical uncertainties in the signal region.

4.7.2 Estimating the LL background

The lost lepton background is extrapolated from a control region ($\text{CR}_{\text{LL}}^{\text{SR}}$ and $\text{CR}_{\text{LL}}^{\text{VR}}$) that has the same cuts and binning as the SR (and the Global VR) only the lepton veto is inverted. The method is not applied below $p_T^{\text{miss}} < 200$ GeV, as the purity of the control region is not sufficient. For simplicity we expand the estimation for the SR only, the LL estimation in the Global VR is handled identically.

The following formula encapsulates the method for a single search bin:

$$N_{\text{LL}}^{\text{SR}} = (N_{\text{data}}^{*\rightarrow\ell} - N^{0\rightarrow\ell}) \cdot \frac{N_{\text{MC}}^{\ell\rightarrow 0}}{N_{\text{MC}}^{\ell'\rightarrow\ell}} \quad (4.2)$$

where the superscripts in the formula indicate the truth \rightarrow reco requirements.

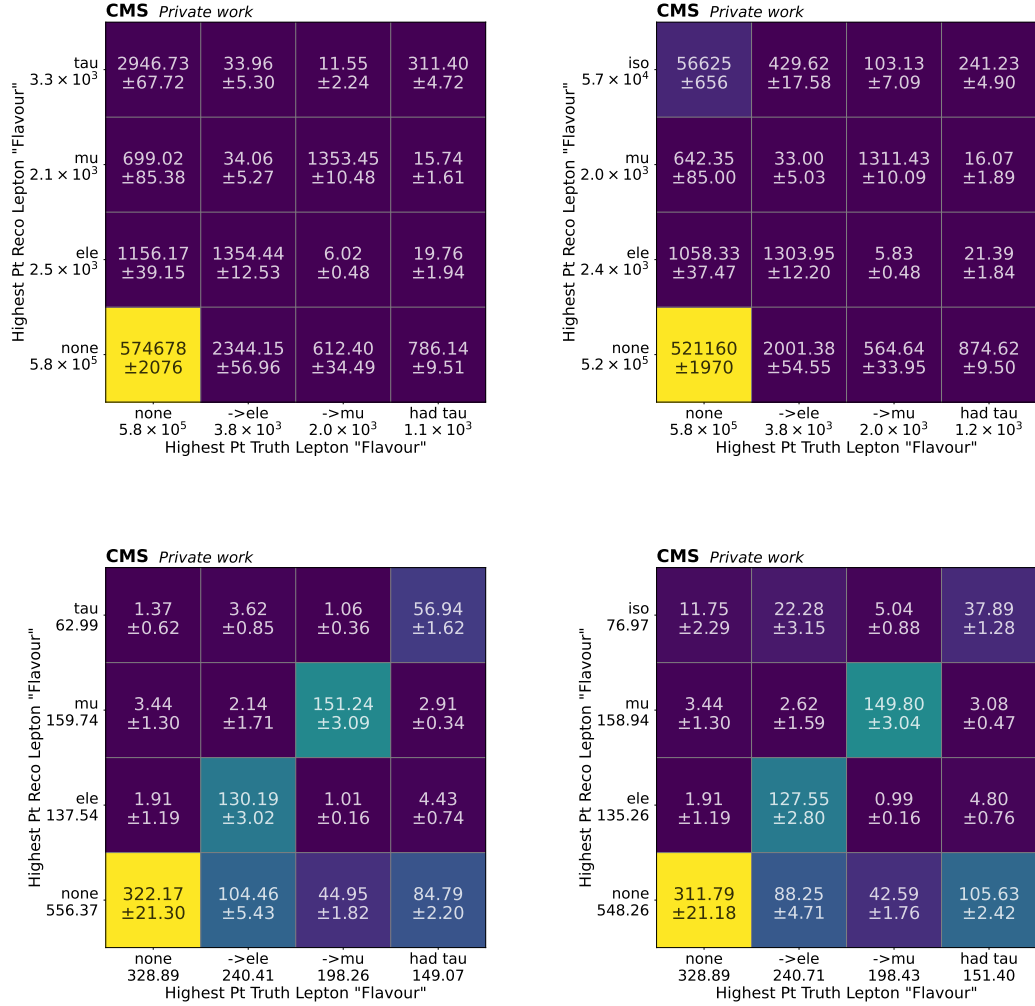


Figure 4.9: The truth-reco response matrix containing the expected number of events in the Run 2 data set after all analysis cuts except for the cuts on the binning variables for all Monte Carlo samples, including both leptonic and hadronic ones. For the bottom row $p_T^{miss} > 200$ GeV is applied. The left-hand side panels use a tau MVA tagger for identifying hadronic tau decays, while the right-hand side plots use isolated tracks for the same purpose. The number of truth objects fluctuates slightly from left to right, due to the scale factors that depend on the list of reco objects in the event.

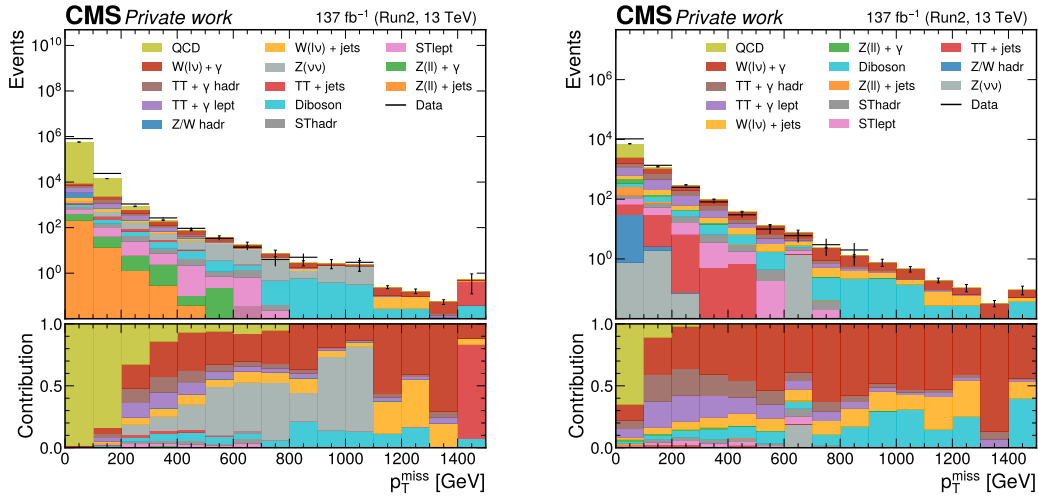


Figure 4.10: The composition of the events passing all analysis cuts except for the binning cuts without a constraint on leptons (left) and with at least one reconstructed lepton required (right) as a function of p_T^{miss} . In the legend, QCD labels the multijet and the jet+photon events. For events with leptonic decays of the W and Z the decay products are given in paranthesis. TT stands for top quark pair production. "hadr" and "lept" indicates whether the event is fully hadronic or has a final state propt lepton. Diboson covers electroweak diboson events such as WW, WZ, ZZ. ST stands for electroweak single top production. The data (black dots with error bars) is compared to the sum of all MC simulated samples.




Figure 4.11: The composition of the lost lepton control regions for the Global VR (left) and the SR (right). For the legend, see Fig. 4.10.

For example, $\ell' \rightarrow \ell$ means, that both a truth and a reco lepton are required in the event, but the two objects do not need to be matched in flavor. The other possible contributions to the CR_{LL} can be removed (e.g. $Z(\nu\nu)$), but the control region bins have a high purity in truly leptonic events as shown in Figure 4.11 and other contributions are considered to be zero. In some bins $N_{\text{data}}^{*\rightarrow\ell} = 0$. In this case, the prediction in this bin is zero, and uncertainty is taken to be the MC ratio times the one sigma Poisson confidence limit, $-\ln(0.32) \approx 1.1$.

One may wonder if computing the transfer factor (TF) simultaneously for all lepton flavors is justified, since, by their nature, detecting muons is simpler (more efficient and more pure) than detecting tau particles. Indeed Figure 4.12 demonstrates the difference.

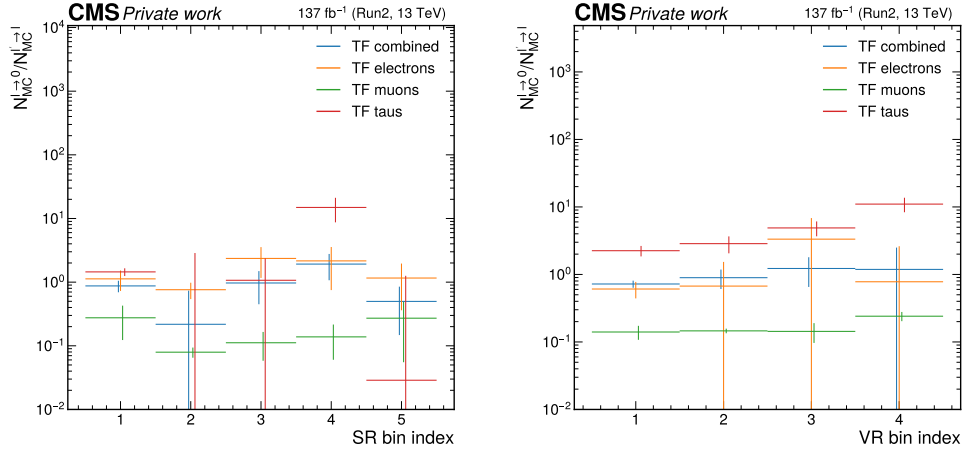


Figure 4.12: The per flavor transfer factors as well as the combined transfer factor in the SR (left) and the VR (right). For better legibility, the horizontal position of the error bars has been staggered.

Since the truth-reco response matrix is diagonally dominant, the lost lepton contribution from the different flavors could be separately estimated the following way (here τ stands for the hadronic tau category):

$$\begin{aligned}
 N_{LL}^{SR} = & (N_{\text{data}}^{*\rightarrow e} - N^{0\rightarrow e}) \frac{N_{MC}^{e\rightarrow 0}}{N_{\ell\rightarrow e, MC}^{CR}} \\
 & + (N_{\text{data}}^{*\rightarrow \mu} - N^{0\rightarrow \mu}) \frac{N_{MC}^{\mu\rightarrow 0}}{N_{\ell\rightarrow \mu, MC}^{CR}} \\
 & + (N_{\text{data}}^{*\rightarrow \tau} - N^{0\rightarrow \tau}) \frac{N_{MC}^{\tau\rightarrow 0}}{N_{\ell\rightarrow \tau, MC}^{CR}}
 \end{aligned} \tag{4.3}$$

A comparison of the two background estimation variants to pure MC is shown in Figure 4.13. All three methods give compatible results within the uncertainties.

Comparing the integral of the two estimations, in the VR a difference of $15 \pm 70\%$ is found, while in the signal region, the difference is $45 \pm 200\%$. Similar agreement is observed in the ABCD control region for bins with $p_T^{\text{miss}} > 200$ GeV. For the ABCD binning convention, see Table 4.2.

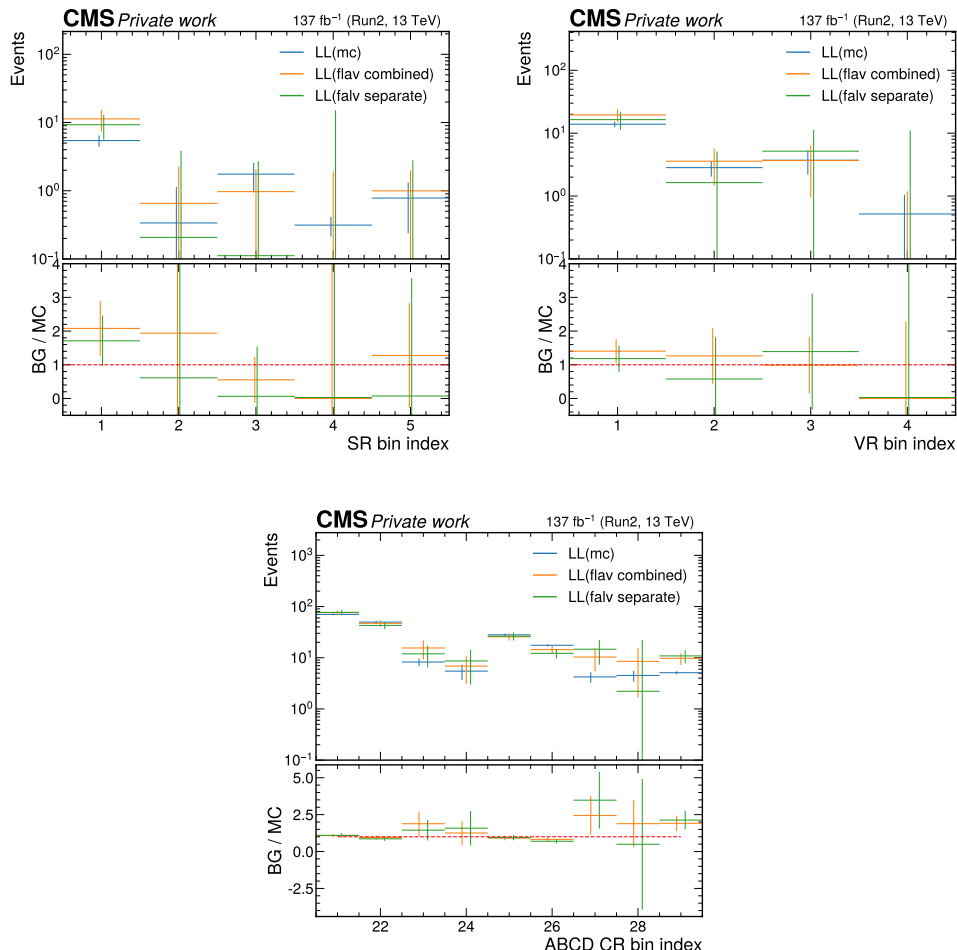


Figure 4.13: Plots of the lost lepton background prediction using the combined or the per-flavor transfer factors compared to pure MC in the SR (left), the Global VR (right), and the ABCD control region for the bins with $p_T^{\text{miss}} > 200$ GeV (bottom).

Since lost leptons are the most significant SM contributors to the signal regions, the importance of properly modeling the uncertainties of this background is paramount. Because of this, similarly to the ABCD procedure, the control region is incorporated into the statistical model of the interpretation and a combined maximum likelihood fit determines the final background. In the fit a 20% uncertainty is assumed on the theoretical cross-section of the leptonic processes.

In this procedure, the combined flavor estimation is used, with a 25% uncertainty to account for the mean deviation between it and the per flavor estimation. This uncertainty is insignificant with respect to the statistical error

bars.

The final value of the background in the signal region integrated over the five bins is 14 ± 6 events using the flavor combined estimate, which corresponds to a $0.10 \pm 0.04 \text{ fb}^{-1}$ fiducial acceptance.

4.7.3 Validation of the background estimates

Using the background estimates for the four categories of contributions, we can compare the observed data in the validation regions with our estimates. In Fig. 4.14 a good agreement is shown both in the QCD-dominated VR and in the global VR.

4.8 Statistical interpretation

As the analysis is still under collaboration review, it is currently blinded, which means that data in the signal region is not available for the analyzers to work with (and thus can not be shown here). This is a standard procedure in the LHC collaborations to eliminate the possibility of unintentional bias in the measurement design.

The analysis cuts were initially set up based on the final state signature studied in simulation and subsequently optimized based on the expected exclusion limit, which is a good indicator of the sensitivity of the analysis.

For exclusions that depend only on one parameter, like the NLSP mass in the case of gaugino pair production, this is straightforward. For the strong production scenario, the "excluded area" quantity was developed, which quantifies the power of the method by a single number as the area of the parameter space which is probed and excluded in units of TeV^2 .

The state-of-the-art way of evaluating the validity of models is statistical hypothesis testing. In this setting, two hypotheses are pitted against one another, the null hypothesis (H_0), which corresponds to the model to be disproven, and the alternative hypothesis (H_1). The procedure either rejects H_0 or fails to do so, which is however not a proof of the validity of H_0 .

The rate of type I error (α) is the probability of rejecting H_0 whereas it is true while the rate of type II error (β) is the probability of failing to reject H_0 even

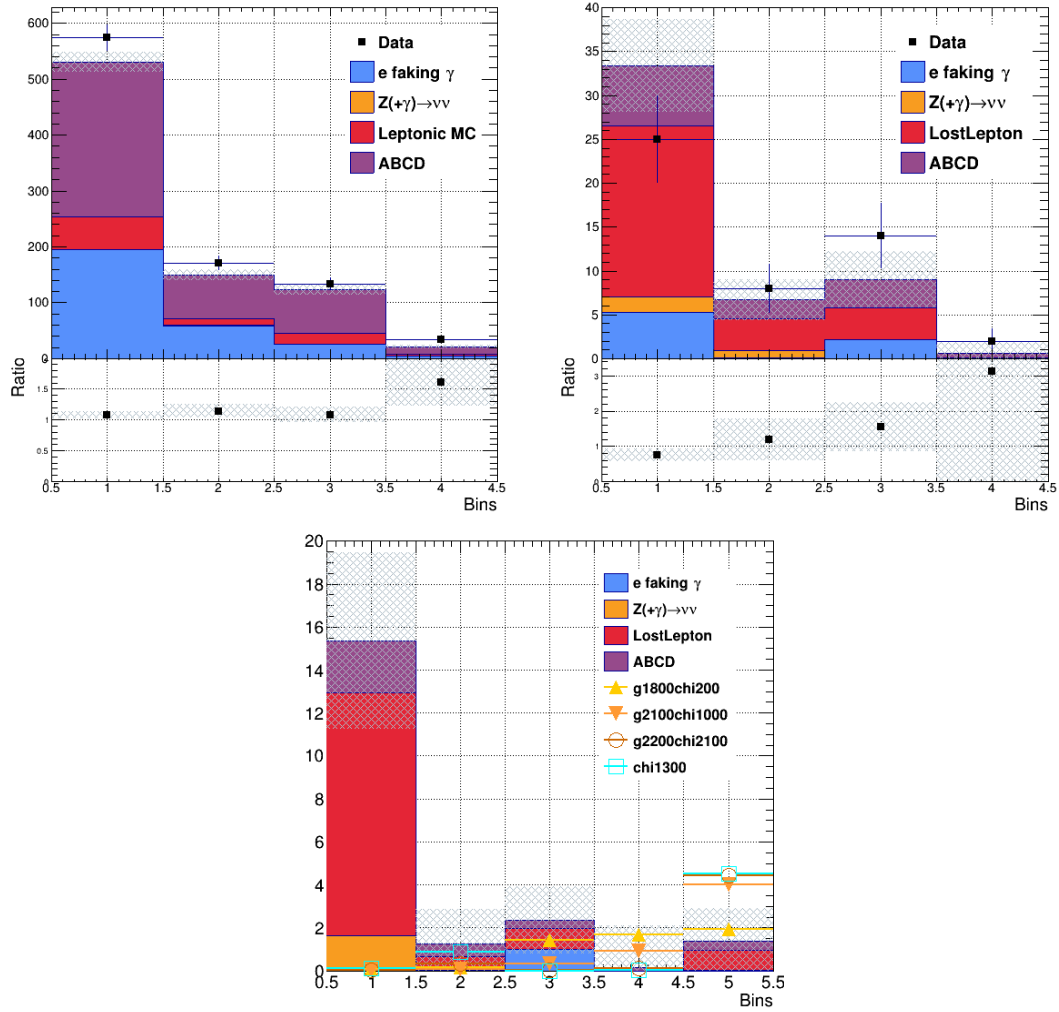


Figure 4.14: Data driven background estimation in the validation and signal regions, with overlaid data in the VRs. In the QCD VR, the leptonic background is estimated using MC since in the low p_T^{miss} region the lost lepton method is unreliable. In the SR, the data-driven background estimations and the MC expectation for a few signal points are shown. The signal samples are labeled according to the mass point: for example, g1800chi200 stands for gluino pair production with a gluino mass of 1800 GeV and an NLSP neutralino with a mass of 200 GeV. Similarly, chi1300 stands for gaugino pair production with an NLSP neutralino mass of 1300 GeV. The uncertainties are purely statistical. Figure by M. Bartok.

though it is not true. $1 - \alpha$ is called the confidence level and $1 - \beta$ is the power of the test.

In a discovery setting the null hypothesis is that there is no significant deviation from the background ($H_0 = H_b$), while the alternative hypothesis is that both the signal and background are present ($H_1 = H_{s+b}$). As announcing false findings in particle physics is highly undesirable, the standard threshold for discoveries is 5σ which corresponds to $\alpha \approx 2.9 \cdot 10^{-7}$.

In the current exclusion setting the roles are reversed. A signal strength parameter r is introduced and the null hypothesis becomes $H_0 = H_{r+s+b}$. The alternative hypothesis is the background-only model $H_1 = H_b$. In fact, with r as a parameter, this becomes a family of hypotheses tests that allows for the computation of an upper limit on the signal strength parameter. In the literature, usually, exclusion limits with a 95% confidence level are reported, which means that only in 5% of the cases, can the model at the limit be falsely rejected.

In all cases, the decision is based on the observed value of the test statistics. This is a one-dimensional function computed from the observed data. A test statistics could be as simple as the sum of all yields, but there are significantly more powerful choices available.

In the interpretation of CMS data, the test statistics are based on Wilk's likelihood-ratio test statistics.

$$t_r = \frac{p(x|r)}{p(x|\hat{r})}, \quad (4.4)$$

where r is the signal strength parameter and p is the probability distribution of the x observables. \hat{r} represents the r value that maximizes the probability for a given x . When p is regarded not as the distribution of x , but as a function of the r parameter for a fixed x , it is called a likelihood: $\mathcal{L}(r) = \mathcal{L}(r|x) = p(x|r)$. In this case, \hat{r} is called the maximum likelihood estimator for r .

If there are more unknown parameters influencing the behavior of the statistical model beyond the signal strength, these are called nuisance parameters and are noted as θ , hence $\mathcal{L}(r, \theta)$. These parameters are profiled out of the likelihood for computing the test statistics by maximalizing the likelihood according to θ for fixed values of r

$$\sup_{\theta} \mathcal{L}(r, \theta) = \mathcal{L}(r, \hat{\theta}(r)).$$

Using t_r one can only set up two-sided limits, hence it is not directly applicable

for setting upper exclusion limits on the signal strength. The possibility of negative \hat{r} values is also a problem, as in this case, rejecting the $r = 0$ hypothesis is a possibility while still not being compatible with the signal. Instead

$$\tilde{q}_r = \begin{cases} -2 \ln \frac{\mathcal{L}(r)}{\mathcal{L}(r=0)}, & \text{if } \hat{r} < 0 \\ -2 \ln \frac{\mathcal{L}(r)}{\mathcal{L}(\hat{r})}, & \text{if } 0 < \hat{r} < r \\ 0 & \text{if } r < \hat{r} \end{cases} \quad (4.5)$$

is used, which corrects these behaviors. Finally, CMS results use the CL_s modified frequentist approach [125]. The CL_s method solves the problem, that large downward fluctuations in the data may result in the exclusion of models that the study is not even sensitive to. Picking an r value and computing the \tilde{q}_r^{obs} observed value of \tilde{q}_r , one may define

$$CL_{s+b} = p(\tilde{q}_r \geq \tilde{q}_r^{\text{obs}} | r), \quad \text{and} \quad CL_b = p(\tilde{q}_r \geq \tilde{q}_r^{\text{obs}} | r = 0),$$

and lastly $CL_s = CL_{s+b}/CL_b$. The basic idea is, that when the observation is similarly unlikely under the $r = 0$ and $r \neq 0$ hypotheses, then CL_s converges to one. The distribution of the \tilde{q}_r in the large sample size limit is known [125], while for analyses with low statistics toy Monte Carlo events can help uncover the distribution. The limit on the signal strength is established iteratively. The test statistics is computed with an initial r value, and based on whether CL_s is larger or smaller than α , a new r is tested, finally, interpolation is applied to find the limit where $CL_s = \alpha$.

To obtain expected limits, toy data is generated with $r = 0$ based on the statistical model, and the limit is computed several times. It is customary to present the median expected exclusion value as well as the ones belonging to the 2.5, 16, 86, and 97.5% percentiles.

To build the test statistics the $p(x|r, \theta)$ probability distribution must be known. Since the data counts in a bin follow a Poisson distribution, for n data counts we get

$$p(n|r, s, b) = \text{Poisson}(n|rs + b) \text{Gauss}(s|\mu_s, \sigma_s) \text{Gauss}(b|\mu_b, \sigma_b) \quad (4.6)$$

in case the s signal and b background yields have an expected value of μ_s and μ_b and a Gaussian uncertainty of σ_s and σ_b .

The model can incorporate control bins as well

$$p(n, n_{\text{CR}}|r, c) = \text{Poisson}(n|rs + cb) \text{Poisson}(n_{\text{CR}}|cb_{\text{CR}}) \quad (4.7)$$

which is the exact mechanism by which the lost lepton background estimation is incorporated into the model. In this simple example, the signal and background

contributions have a negligible uncertainty and hence the Gaussian uncertainty terms are not shown and c establishes a scaling connection between the background in the signal region and the background in a control bin where no signal contribution was added (for the sake of simplicity).

The simplified implementation of the ABCD method for a single signal bin is

$$\begin{aligned}
 p(n, n_A, n_B, n_C | r, a, b, c, k) = & \text{Poisson}(n | rs + k \frac{bc}{a}) \\
 & \cdot \text{Poisson}(n_A | a) \\
 & \cdot \text{Poisson}(n_B | b) \\
 & \cdot \text{Poisson}(n_C | c) \\
 & \cdot \text{Gauss}(k | K, \sigma_K),
 \end{aligned} \tag{4.8}$$

which performs Equation 4.1 in the fit in the process of marginalizing the a, b, c , and k parameters. Of course, in practice, both the signal region and the control bins have all four backgrounds and the signal contamination considered.

The GbbMET analysis has five search bins, but, with the 5 lost lepton control regions and the 13 ABCD control regions included a total of 23 bins are incorporated into the statistical model with over 120 nuisance parameters taking the statistical and systematic uncertainties into account.

The statistical computation outlined here was implemented using the Combine [126] likelihood description and fitting framework. Blind upper limits were computed using the AsymptoticLimits subroutine. In this case, in the signal region, the expected background is used instead of the observation. The procedure uses the asymptotic formulae for the distribution of \tilde{q}_r . The results are shown on the left-hand side panels of Figure 4.15. Three versions of the statistical model are demonstrated. Turning our attention to the strong production channel (upper left), the first model only contains statistical uncertainties and the multijet and lost lepton backgrounds are calculated externally. The second version already contains the control regions but still no systematic uncertainties. The improvement is clear, the expected sensitivity reach is shifted by about 50 GeV towards higher masses. The third version introduces the systematic uncertainties. The median expected exclusion and the one-sigma quantiles barely change. The differences are even less pronounced in the weak production channel (lower left).

The asymptotic approach is checked against an alternative method (Hybrid-New) deriving the distributions using Monte Carlo simulation in a few mass

points using the largest statistical model. While this procedure is more accurate, especially in a low statistics regime, it is much more time-consuming, up to 6 hours instead of about a minute for the asymptotic limits per signal mass point of which there are hundreds to cover the full simulated signal grid. The simulations were carried out using cloud computing. The left-hand side plots in Figure 4.15 show the results. In general, the median expected values are within the uncertainties, and while the MC-based method gives tighter uncertainties, the use of the asymptotic approach is justified.

The final median expected exclusion lower limits on the mass of the NLSP for the weak production of gaugino pairs are $m_{\tilde{\chi}_1^0} = 1350$ GeV for wino-like neutralinos and $m_{\tilde{\chi}_1^0} = 1200$ GeV in the hino-like case. The expected production cross section upper limit at high neutralino mass is about 0.3 fb. For the strong production of gluino pairs, the result with the expected excluded cross section indicated is shown in Figure 4.16. In a large parameter region, the sensitivity reaches 0.3 to 1 fb, with the most stringent limits achieved close to the gluino - neutralino mass diagonal, where the AK8 tagging is the most sensitive. The expected exclusion for the gluino mass reaches up to $m_{\tilde{g}} = m_{\tilde{\chi}_1^0} = 2250$ GeV.

The impact of each of the nuisance parameters was evaluated. First, a maximum likelihood fit is performed using all nuisance parameters. Instead of the observed event counts, toy experiments with input from the control region data are used to keep the signal region blind. Once the post-fit uncertainty of all nuisances is obtained, they are fixed to their $\pm 1\sigma$ locations one-by-one and the rest of the parameters are fit again each time. The amount by which the signal strength parameter shifts in this process is indicative of the study's sensitivity to the nuisance parameter in question. The result is examined for representative points near the exclusion limit and no qualitative difference is found. The average for the 30 most significant nuisance parameters is reported in Figure 4.17 with no signal injected and with $r = 1$ signal injected in the toy data generation. It is observed that the impact tables are dominated by statistical uncertainties (stat.*) and the uncertainties of the free-floating coefficients implementing the lost lepton (lambda*) and ABCD (alpha*, beta*, gamma*) methods.

Nuisance parameters representing statistical uncertainties are always named stat.(process name).(bin name). The processes can be "Signal", "Egamma", "QCD", "LL" and "Znu", while the bin name contains a region name and a bin index. "sb" for the SR, "LL" for the lost lepton CR, "A", "B" and "C" for the ABCD control region. The bins in the lost lepton control region and the signal region

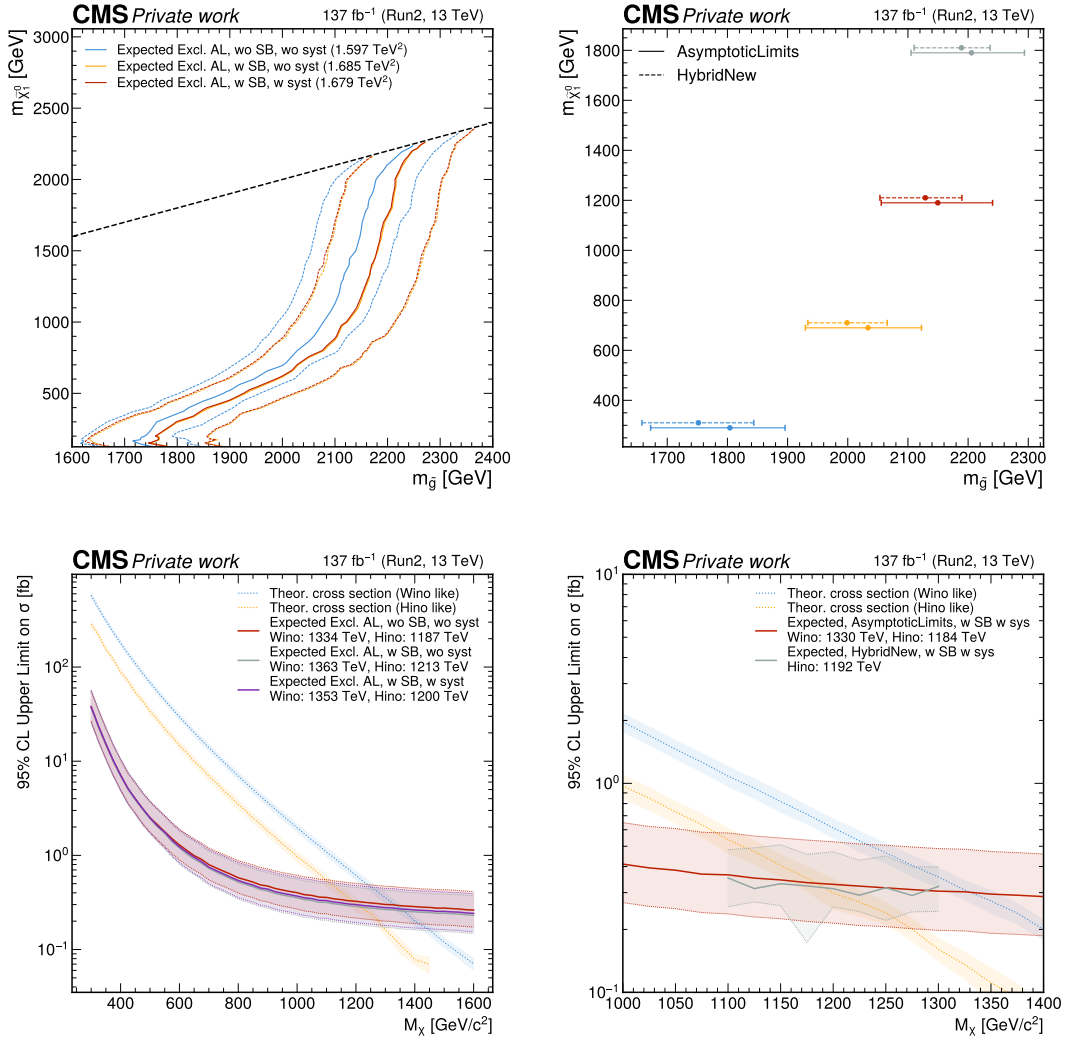


Figure 4.15: The median expected exclusion limits with the one sigma quantiles illustrating the effect of the CR and the systematics uncertainty inclusion to the interpretation (left) and the comparison between the statistical tests relying on asymptotic formula and Monte Carlo methods (right). The upper plots display the results for the strong gluino pair production channel in the gluino - NLSP mass plane, while the lower plots are for the weak gaugino pair production channel as a function of the NLSP mass. For the weak production the theoretical cross section predictions with their uncertainties and the expected upper limit on the cross section are compared to derive the mass limit.

are indexed from zero, while the ABCD control region is indexed according to Table 4.2.

The λ_i parameter multiplies the LL contribution in the i^{th} bin of the SR and the LL CR, while α^* , β^* , γ^* stand for the QCD contribution in bins of A, B, or C kind of the ABCD CR, finally a formula written similarly to Eq. 4.8 stands for the SR QCD contribution.

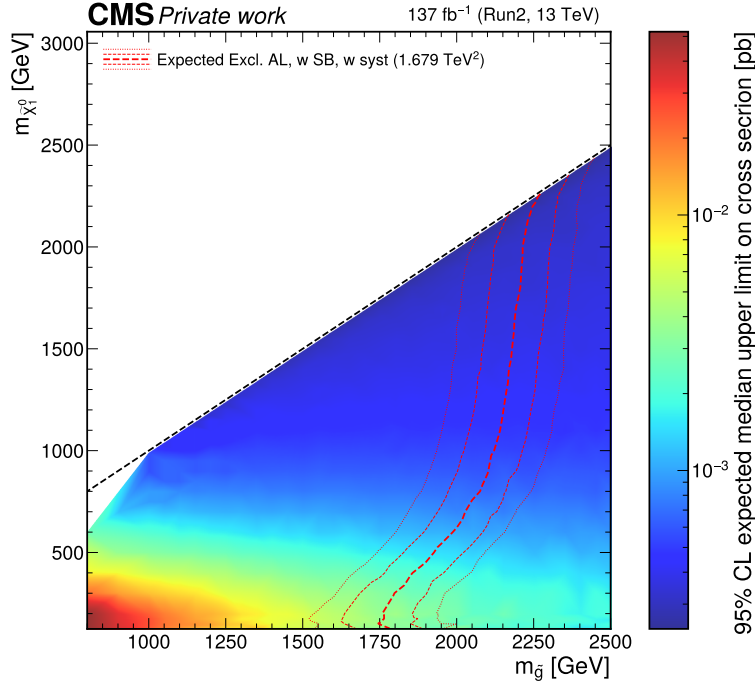



Figure 4.16: The median expected exclusion for the strong production channel with the one and two sigma percentiles indicated and the map of the excluded cross section values as a function of the model parameters.

Looking at the top plot, the three parameters with the most significant impact all belong to the last signal region bin. This is not surprising, as near the exclusion limit this bin has the most sensitivity. The nuisance parameters standing for systematic uncertainties follow the CMS standard conventions. There are one or more of these related to all scale factors, energy and efficiency corrections, depending on whether the uncertainty is taken as correlated across years. The ones in the figures are "CMS_LUM24006_LLFlat", which describes the 25% systematic uncertainty of the lost lepton background correlated across all signal bins. The "CMS_LUM24006_EgammaFlat" which describes the 50% correlated uncertainty of the $e \rightarrow \gamma$ background, the "CMS_pileup_13TeV" which represents the pileup reweighting uncertainty, the "CMS_LUM24006_ZnunuFlat" which stands for the 15% correlated uncertainty of the $Z(\nu\nu)$ background, "CMS_btag_DeepFlav_light" which correspond to one of the systematic uncertainties related to the DeepJet b-tagging efficiency. The AK4 jet b-tagging uncertainty is divided into a light jet (u,d,s) efficiency and a heavy jet (b,c) efficiency uncertainty part. To account for mis-tagging rates, both need to be taken into account. Finally, "CMS_m_eff_id" stands for the muon identification efficiency-related systematic uncertainty.



30




Figure 4.17: The expected pulls and impacts of the 30 most significant nuisance parameters evaluated and averaged over six representative points along the gluino pair production expected exclusion curve and three near the gaugino pair production limits. The top plot shows the result with zero signal injected into the toy data generation, while the bottom plot shows the result with $r = 1$ signal injected.

Figure 4.17 also reports the shift of all nuisance parameters with respect to their initial value relative to both the input uncertainty σ_I and relative to $\sqrt{\sigma_I^2 - \sigma^2}$, where σ is the post fit uncertainty of the parameter. The second ratio is called the pull and in a constrained likelihood fit to data it is expected to follow a standard normal distribution. In the current case, since the toy data is generated based on the expected yields, all pulls are close to zero. For free-floating parameters, only the central value and the fit uncertainty are shown.

All results look healthy, no unexplained feature can be seen in the statistical interpretation, neither when it is performed with no signal injected ($r = 0$), nor with a SUSY signal close to the sensitivity limit ($r = 1$).

4.9 Prospects

We have completed the analysis preparation for the search of pair produced gluinos and gauginos predicted by GMSB SUSY models. Higgs boson final states provide in these models a window to new physics.

Our work (labeled as CMS-SUS-24-006) is under collaboration approval and is expected to progress toward publication in the coming months. It is expected to extend the CMS sensitivity to the studied models, especially for gaugino pair production, where we foresee – in the case when no long-awaited excess appears – raising the limit from $m_{\tilde{\chi}_1^0} > 950$ GeV [116] to $m_{\tilde{\chi}_1^0} > 1200$ GeV in the higgsino-like case, and for gluino pair production when the mass difference to the NLSP is small an approximately 50 GeV improvement is anticipated. In terms of gluino production cross section sensitivity in this highly boosted region, we reach 0.24 fb which is almost a factor two improvement on the previous published result.

Conclusion

In this dissertation, my contributions to the search for BSM physics on both ends of the precision scale are presented. Experiments probing for deviations from the standard model via cross section measurements require accurate knowledge of the data volume they processed, which is described by luminosity. I worked on many aspects of improving the luminosity measurement at CMS. I contributed significantly to a large number of CMS results which I presented at international conferences [127, 128] and which earned me a "CMS Award 2022". I was selected as the coordinator and corresponding author for the final Run2 luminosity effort which has about 40 participating scientists and reanalyzed the 2017-2018 data to achieve an over twofold increase in accuracy [71].

Accurate luminosity starts at the source, hence my investigations into machine learning-based data quality monitoring for the Pixel Luminosity Telescope [83] as a co-supervisor for M. Fejes are also reported.

I highlighted the powerful tools I developed and characterized to study and constrain one of the leading sources of luminosity uncertainty, the non-factorization bias. The 2D rate fit method [129] and the luminous region method demonstrated their compatibility in Run2 and Run3, and they either decreased the luminosity uncertainty in all Run2 eras [71, 80] or established their own measurement for the first time [102, 103, 112, 130].

I detailed some of the work I carried out for the 2017-2018 measurements, such as the introduction and systematic testing of new fit curves, the introduction of a correction for the gradual emittance change, the derivation of the emittance scan modification factor, the linearity correction of PLT or the first use and characterization of the dynamic mean of the individual luminometer results as the best measure of instantaneous luminosity. These efforts played a major role in achieving a per year uncertainty below 1% and the record-breaking combined Run2 uncertainty of 0.7%. The new calibrations provided for 2017 pp data at 13

TeV resolved the outstanding conflict between theory predictions and preliminary CMS observations for the inclusive W and Z boson cross sections [131, 132].

In an attempt to probe the validity of the SM with a direct search, I utilized the CMS Run2 dataset at 13 TeV energy to investigate the production of SUSY particles in the context of two simplified models employing gauge-mediated symmetry breaking (GMSB) and motivated by the prospect of providing a gravitino dark matter candidate [133]. I investigated two tau-identification methods to improve the rejection of a component of the most significant SM background in the signal region, the lost lepton background. I developed a semi-data-driven measurement of this background and considered the systematic effects of the different properties of the lepton flavors. I created a statistical model for the interpretation of the results which was also employed to optimize the analysis procedures. To ensure the proper treatment of uncertainties, I incorporated the two most dominant backgrounds into the statistical model which I used for hypothesis testing to gauge the sensitivity of the analysis in terms of upper limits on the production cross sections and obtain constraints on the SUSY particle masses. The obtained cross section limits reach below 0.3 fb which is about a factor of two better than previously published results from CMS for mass-degenerate gaugino pair production, as well as for gluino pair production when the mass difference between the gluino the lightest neutralino diminishes.

Appendix A

Normalization and CapSigma parametrization

A.1 Double Gaussian with shared means

The normalization of the sum of several Gaussian functions is trivial. In the one-dimensional case we can derive the CapSigma parametrization as follows:

$$f(x) = A_1 \exp\left(-\frac{(x-\mu)^2}{2\sigma_1^2}\right) + A_2 \exp\left(-\frac{(x-\mu)^2}{2\sigma_2^2}\right) \quad (\text{A.1})$$

$$\int_{\mathbb{R}} f(x) dx = \sqrt{2\pi}(A_1\sigma_1 + A_2\sigma_2) \stackrel{!}{=} \sqrt{2\pi}P\Sigma \quad (\text{A.2})$$

We define the peak (P), amplitude fraction (F) and width ratio (R) parameters

$$P = A_1 + A_2, \quad (\text{A.3})$$

$$F = \frac{A_1}{A_1 + A_2}, \quad (\text{A.4})$$

$$R = \frac{\sigma_1}{\sigma_2}, \quad (\text{A.5})$$

which are the quantities the double Gaussian fit function is historically parametrized by in the BRIL group. Expressing Σ from Eq. A.2 we get

$$\Sigma = \frac{A_1\sigma_1 + A_2\sigma_2}{A_1 + A_2} = F\sigma_1 + (1-F)\sigma_2 = (FR + (1-F))\sigma_2, \quad (\text{A.6})$$

and consequently

$$\sigma_2 = \frac{\Sigma}{FR + (1-F)}, \quad (\text{A.7})$$

$$\sigma_1 = \frac{\Sigma R}{FR + (1-F)}. \quad (\text{A.8})$$

The final form of the CapSigma parametrization is

$$f(x) = P \left\{ F \exp \left(- \frac{(x - \mu)^2}{2 \left(\frac{\Sigma R}{FR + 1 - F} \right)^2} \right) + (1 - F) \exp \left(- \frac{(x - \mu)^2}{2 \left(\frac{\Sigma}{FR + (1 - F)} \right)^2} \right) \right\}. \quad (\text{A.9})$$

A.2 Super Gaussian

The super Gaussian function takes the same form as a regular Gaussian, but allows for an exponent different from two, in the current parametrization $2p$:

$$f(x) = A \exp \left(- \frac{1}{2} \left(\frac{(x - \mu)^2}{\sigma^2} \right)^p \right), \quad (\text{A.10})$$

$$\int_{\mathbb{R}} f(x) dx = A \sigma \int_{\mathbb{R}} \exp \left(- \frac{x^{2p}}{2} \right) dx = 2A \sigma \int_0^{\infty} \exp \left(- \frac{x^{2p}}{2} \right) dx, \quad (\text{A.11})$$

using the $u = x^{2p}/2$ substitution, we get

$$\int_{\mathbb{R}} f(x) dx = A \sigma 2^{\frac{1}{2p}} \frac{1}{p} \int_0^{\infty} u^{\frac{1}{2p}-1} \exp(-u) du = A \sigma \frac{2^{\frac{1}{2p}}}{p} \Gamma \left(\frac{1}{2p} \right), \quad (\text{A.12})$$

where Γ stands for Euler's Gamma function. With this we get the following expression for CapSigma:

$$\Sigma = \sigma \frac{2^{\frac{1}{2p}}}{\sqrt{2\pi p}} \Gamma \left(\frac{1}{2p} \right). \quad (\text{A.13})$$

Expressing σ , it is possible to parametrize the shape using A , μ , Σ and p .

A.2.1 Normalization in N dimensions

In the N dimensional case, with radial parametrization $x, \mu \in \mathbb{R}^N$, $\Sigma \in \mathbb{R}^{N \times N}$,

$$r^2 = \sum_{i,j} (x - \mu)_i \Sigma_{ij}^{-1} (x - \mu)_j, \quad (\text{A.14})$$

$$f(x) = A \exp \left(- \frac{(r^2)^p}{2} \right). \quad (\text{A.15})$$

Transforming the variables with $z = \Sigma^{-\frac{1}{2}}(x - \mu)$ and then changing to spherical coordinates, we get

$$\int_{\mathbb{R}^N} f(x) d^N x = \sqrt{|\Sigma|} \int_{\mathbb{R}^N} f_z(z) d^N z = \sqrt{|\Sigma|} \int_0^{\infty} r^{N-1} dr \int_{S^{N-1}} d^{N-1} \Theta f_r(r). \quad (\text{A.16})$$

Performing the spherical integral gives the surface of an N dimensional sphere:

$$\int_{S^{N-1}} d^{N-1} \Theta = \frac{2\pi^{\frac{N}{2}}}{\Gamma(\frac{N}{2})}. \quad (\text{A.17})$$

The radial integral is computed using a change of variables, $u = r^{2p}/2$,

$$\int_0^\infty r^{N-1} \exp\left(-\frac{r^{2p}}{2}\right) dr = \frac{2^{\frac{N}{2p}}}{2p} \int_0^\infty u^{\frac{N}{2p}-1} \exp(-u) du = \frac{2^{\frac{N}{2p}}}{2p} \Gamma\left(\frac{N}{2p}\right). \quad (\text{A.18})$$

The full integral equates to

$$\int_{\mathbb{R}^N} f(x) d^N x = A \sqrt{|\Sigma|} \frac{2\pi^{\frac{N}{2}}}{\Gamma(\frac{N}{2})} \frac{2^{\frac{N}{2p}}}{2p} \Gamma\left(\frac{N}{2p}\right). \quad (\text{A.19})$$

For $N=2$, the fit function with this shape takes the form

$$f(x, y) = V \frac{p}{2\pi\sigma_x\sigma_y\sqrt{1-\rho^2}\Gamma(1/p)2^{1/p}} \exp\left(-\frac{(r^2)^p}{2}\right), \quad (\text{A.20})$$

where V represents the integral of the function.

A.2.2 Normalization for product constructions

In some cases it might be desirable to have a different p in the different directions. One way of achieving this goal is a product construction

$$f(x, y; p_x, p_y) = f(x'; p_x) f(y'; p_y), \quad (\text{A.21})$$

where x' and y' are a rotated version of x and y to introduce some level of non-factorisation.

In this case, the integral of $f(x, y; p_x, p_y)$ is written as

$$\int_{\mathbb{R}^2} dx dy f(x'; p_x) f(y'; p_y) = \int_{\mathbb{R}^2} dx dy f(x; p_x) f(y; p_y) = \int_{\mathbb{R}} dx f(x; p_x) \int_{\mathbb{R}} dy f(y; p_y), \quad (\text{A.22})$$

where the one dimensional integrals were discussed in Eq. A.12. This factorization of the integral holds for all product constructions regardless of the one dimensional shapes used, therefore it is not discussed in other cases.

A.3 Polynomial-modulated Gaussian

The one-dimensional generalized poly-Gaussian formula can be written as

$$f(x) = A \left(1 + \sum_{i=1}^{\infty} a_i \left(\frac{x-\mu}{\sigma}\right)^i\right) \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad (\text{A.23})$$

where a_i are the coefficients characterizing the polynomial. The central moments of the Gaussian distribution are computed as

$$\frac{1}{\sqrt{2\pi}\sigma} \int_{\mathbb{R}} (x - \mu)^n \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right) dx = \begin{cases} 0 & \text{if } n \text{ is odd,} \\ (n - 1)!! \sigma^n & \text{else,} \end{cases} \quad (\text{A.24})$$

using integration by parts, where the notation $!!$ signifies the double factorial

$$n!! = n(n - 2)(n - 4) \dots 5 \cdot 3, \text{ or } n!! = n(n - 2)(n - 4) \dots 4 \cdot 2,$$

depending on parity. Using this result, we get

$$\int_{\mathbb{R}} f(x) dx = A\sqrt{2\pi}\sigma \left(1 + \sum_{i=1}^{\infty} a_{2i}(2i - 1)!!\right). \quad (\text{A.25})$$

With this CapSigma can be written as

$$\Sigma = \sigma / \left(1 + \sum_{i=1}^{\infty} a_{2i}(2i - 1)!!\right). \quad (\text{A.26})$$

For Poly4G functions ($n=4$)

$$f(x) = A \left(1 + a_2 \left(\frac{x - \mu}{\sigma}\right)^2 + a_4 \left(\frac{x - \mu}{\sigma}\right)^4\right) \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right). \quad (\text{A.27})$$

This gives $\Sigma = \sigma / (1 + a_2 + 3a_4)$.

A.3.1 Radial polynomial-modulated Gaussian normalization

In the N dimensional case, with radial parametrization $x, \mu \in \mathbb{R}^N, \Sigma \in \mathbb{R}^{N \times N}$,

$$r^2 = \sum_{i,j} (x - \mu)_i \Sigma_{ij}^{-1} (x - \mu)_j, \quad (\text{A.28})$$

$$f(x) = A \left(1 + \sum_{i=1}^{\infty} a_i (r^2)^i\right) \exp\left(-\frac{r^2}{2}\right). \quad (\text{A.29})$$

Transforming the variables with $z = \Sigma^{-\frac{1}{2}}(x - \mu)$ and then changing to spherical coordinates, we get

$$\int_{\mathbb{R}^N} f(x) d^N x = \sqrt{|\Sigma|} \int_{\mathbb{R}^N} f_z(z) d^N z = \sqrt{|\Sigma|} \int_0^{\infty} r^{N-1} dr \int_{S^{N-1}} d^{N-1} \Theta f_r(r). \quad (\text{A.30})$$

Performing the spherical integral gives the surface of an N dimensional sphere

$$\int_{S^{N-1}} d^{N-1} \Theta = \frac{2\pi^{\frac{N}{2}}}{\Gamma(\frac{N}{2})}. \quad (\text{A.31})$$

The radial integral is computed separately for each degree of the polynomial using the substitution $u = \frac{r^2}{2}$

$$\int_0^\infty r^{i+N-1} \exp\left(-\frac{r^2}{2}\right) dr = \int_0^\infty \exp(-u)(2u)^{\frac{i+N-1}{2}} \frac{1}{\sqrt{2u}} du = \quad (\text{A.32})$$

$$2^{\frac{i+N-2}{2}} \int_0^\infty \exp(-u) u^{\frac{i+N-2}{2}} du = 2^{\frac{i+N-2}{2}} \Gamma\left(\frac{i+N}{2}\right). \quad (\text{A.33})$$

Which leaves us with

$$\int_{\mathbb{R}^N} f(x) d^N x = A \sqrt{|\Sigma|} \frac{2^{\frac{N}{2}} \pi^{\frac{N}{2}}}{\Gamma\left(\frac{N}{2}\right)} \left(\Gamma\left(\frac{N}{2}\right) + \sum_{i=1}^{\infty} a_i 2^{\frac{i}{2}} \Gamma\left(\frac{i+N}{2}\right) \right). \quad (\text{A.34})$$

For the two-dimensional poly4G function,

$$f(x) = A \left(1 + a_2 r^2 + a_4 (r^2)^2 \right) \exp\left(-\frac{r^2}{2}\right). \quad (\text{A.35})$$

This gives $V = 2\pi A \sqrt{|\Sigma|} (1 + 2a_2 + 8a_4)$.

A.3.2 2D polynomial-modulated Gaussian normalization

For simplicity, a case with $\mu_x = \mu_y = 0$ is considered, since this choice has no bearing on the integral of our function. With this r^2 takes the form

$$r^2 = \frac{1}{1-\rho^2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} - 2\rho \frac{xy}{\sigma_x \sigma_y} \right), \quad (\text{A.36})$$

and a general two-dimensional polynomial-modulated Gaussian is written as

$$f(x, y) = A \left(\sum_{i=0, j=0} a_{i,j} \left(\frac{x}{\sigma_x} \right)^i \left(\frac{y}{\sigma_y} \right)^j \right) \exp\left(-\frac{r^2}{2}\right). \quad (\text{A.37})$$

The integral can be performed piecewise for the summands

$$\int_{\mathbb{R}^2} f(x, y) dx dy = A \left(\sum_{i=0, j=0} a_{i,j} \int_{\mathbb{R}^2} \left(\frac{x}{\sigma_x} \right)^i \left(\frac{y}{\sigma_y} \right)^j \exp\left(-\frac{r^2}{2}\right) dx dy \right). \quad (\text{A.38})$$

Introducing $u = x/\sigma_x$, $v = y/\sigma_y$ and

$$s^2 = \frac{u^2 + v^2 - 2\rho uv}{1-\rho^2}, \quad (\text{A.39})$$

and performing the associated change of variables, we get

$$\int_{\mathbb{R}^2} f(x, y) dx dy = A \sigma_x \sigma_y \left(\sum a_{i,j} \int_{\mathbb{R}^2} u^i v^j \exp\left(-\frac{s^2}{2}\right) du dv \right). \quad (\text{A.40})$$

The integrand can be written as a derivative

$$u^i v^j \exp\left(-\frac{s^2}{2}\right) = \partial_a^i \partial_b^j \exp\left(-\frac{s^2}{2} + au + bv\right) \Big|_{\substack{a=0 \\ b=0}}. \quad (\text{A.41})$$

$$\int_{\mathbb{R}^2} u^i v^j \exp\left(-\frac{s^2}{2}\right) dudv = \left[\partial_a^i \partial_b^j \int_{\mathbb{R}^2} \exp\left(-\frac{s^2}{2} + au + bv\right) dudv \right] \Big|_{\substack{a=0 \\ b=0}}, \quad (\text{A.42})$$

where the integral can be performed by transforming the exponent into a whole square in u and v plus a constant:

$$-\frac{s^2}{2} + au + bv \stackrel{!}{=} -\frac{s'^2}{2} + c, \quad (\text{A.43})$$

where s'^2 is written using the same structure as s^2 (see Eq. A.39) but with u and v shifted $u' = u + \Delta u$ and $v' = v + \Delta v$.

$$\begin{aligned} (1 - \rho^2)s'^2 &= (u + \Delta u)^2 + (v + \Delta v)^2 - 2\rho(u + \Delta u)(v + \Delta v) = \\ &= (1 - \rho^2)s^2 + 2(\Delta u + \rho\Delta v)u + 2(\Delta v + \rho\Delta u)v + \Delta u^2 + \Delta v^2 + 2\rho\Delta u\Delta v. \end{aligned} \quad (\text{A.44})$$

Reading off the coefficients of the linear terms, we get

$$a = -\frac{\Delta u + \rho\Delta v}{1 - \rho^2}, \quad (\text{A.45})$$

$$b = -\frac{\Delta v + \rho\Delta u}{1 - \rho^2}, \quad (\text{A.46})$$

and, therefore,

$$\Delta u = \rho b - a, \quad (\text{A.47})$$

$$\Delta v = \rho a - b. \quad (\text{A.48})$$

Combining with Eq. A.43, we get

$$\begin{aligned} c &= -\frac{\Delta u^2 + \Delta v^2 + 2\rho\Delta u\Delta v}{2(1 - \rho^2)} = \\ &= -\frac{(\rho a - b)^2 + (\rho b - a)^2 + 2\rho(\rho a - b)(\rho b - a)}{2(1 - \rho^2)}. \end{aligned} \quad (\text{A.49})$$

And hence, the intermediate integral is

$$\int_{\mathbb{R}^2} \exp\left(-\frac{s^2}{2} + au + bv\right) dudv = e^c \int_{\mathbb{R}^2} \exp\left(-\frac{s'^2}{2}\right) dudv = 2\pi\sqrt{1 - \rho^2}e^c. \quad (\text{A.50})$$

This leaves us with

$$\int_{\mathbb{R}^2} u^i v^j \exp\left(-\frac{s^2}{2}\right) dudv = 2\pi\sqrt{1 - \rho^2} \left(\partial_a^i \partial_b^j e^c \right) \Big|_{\substack{a=0 \\ b=0}} = 2\pi\sqrt{1 - \rho^2} n_{ij}, \quad (\text{A.51})$$

which is an elementary, yet still a tedious calculation. For $0 \leq i \leq j \leq 8$, the polynomials n_{ij} are brought to a closed form with the help of a symbolic computation tool, sympy, see Table A.1. With this, the final result is

$$\int_{\mathbb{R}^2} f(x, y) dx dy = A \left(\sum_{i=0, j=0} a_{i,j} \int_{\mathbb{R}^2} \left(\frac{x}{\sigma_x} \right)^i \left(\frac{y}{\sigma_y} \right)^j \exp \left(-\frac{r^2}{2} \right) dx dy \right) \quad (\text{A.52})$$

$$= A 2\pi \sqrt{1 - \rho^2} \sigma_x \sigma_y \left(\sum_{i=0, j=0} a_{i,j} n_{ij} \right). \quad (\text{A.53})$$

A.4 q-Gaussian

The q-Gaussian function is created by using a q-exponential

$$e(x, q) = [1 + (1 - q)x]_+^{\frac{1}{1-q}}, \quad (\text{A.54})$$

where

$$[x]_+ = \begin{cases} x & \text{if } x > 0, \\ 0 & \text{else,} \end{cases} \quad (\text{A.55})$$

instead of the natural exponential function, which is returned in the $q = 1$ limit. A normalized q-Gaussian function takes the form

$$q(x) = \frac{\sqrt{\beta}}{C_q} e(-\beta(x - \mu)^2, q), \quad (\text{A.56})$$

where

$$C(q) = \begin{cases} \frac{2\sqrt{\pi}\Gamma\left(\frac{1}{1-q}\right)}{(3-q)\sqrt{1-q}\Gamma\left(\frac{3-q}{2(1-q)}\right)} & \text{if } -\infty < q < 1, \\ \frac{\sqrt{\pi}\Gamma\left(\frac{3-q}{2(1-q)}\right)}{\sqrt{q-1}\Gamma\left(\frac{1}{q-1}\right)} & \text{if } 1 < q < 3, \\ \sqrt{\pi} & \text{if } q = 1. \end{cases} \quad (\text{A.57})$$

For $q \geq 3$, the expression is not normalizable.

With the CapSigma parametrization we require $\sqrt{2\pi}A\Sigma = V = 1$. Using $e(0, q) = 1$, we get

$$\sqrt{2\pi} \frac{\sqrt{\beta}}{C_q} \Sigma = 1, \quad (\text{A.58})$$

and, therefore,

$$\frac{C_q}{\sqrt{2\pi}\Sigma} = \sqrt{\beta}. \quad (\text{A.59})$$

The one-dimensional q-Gaussian function with the CapSigma parametrization is then written as

$$f(x) = A e \left(-\frac{C(q)^2}{2\pi\Sigma^2} (x - \mu)^2, q \right). \quad (\text{A.60})$$

$i \backslash j$	0	1	2	3	4	5	6	7	8
0	1	0	1	0	3	0	15	0	105
1	ρ	0	0	3ρ	0	15ρ	0	105ρ	0
2		$1+2\rho^2$	0	0	$3+12\rho^2$	0	$15+90\rho^2$	0	$105+840\rho^2$
3			$9\rho+6\rho^3$	0	0	$45\rho+60\rho^3$	0	$315\rho+630\rho^3$	0
4				$9+72\rho^2+24\rho^4$	0	$45+540\rho^2+360\rho^4$	0	$315+5040\rho^2+5040\rho^4$	0
5					$225\rho+600\rho^3+120\rho^5$	0	$1575\rho+6300\rho^3+2520\rho^5$	0	$1575+37800\rho^2+75600\rho^4+20160\rho^6$
6						$225+4050\rho^2+5400\rho^4+720\rho^6$	0	$11025\rho+66150\rho^3+52920\rho^5+5040\rho^7$	$1575+37800\rho^2+75600\rho^4+20160\rho^6$
7							$11025\rho+66150\rho^3+52920\rho^5+5040\rho^7$	0	0
8								$11025+352800\rho^2+1058400\rho^4+564480\rho^6+40320\rho^8$	0

Table A.1: The r_{ij} polynomials used for the normalization of two-dimensional polynomial-modulated Gaussian function for $0 \leq i \leq j \leq 8$. Note, that r is symmetric in its indices, $r_{ij} = r_{ji}$.

```

1 import numpy as np
2 from scipy.special import gammaln
3
4 def e_q(x, q):
5     thr = 0.001
6     if abs(q-1)<thr:
7         a = e_q(x, 1.0+thr*(1+1e-6))
8         b = e_q(x, 1.0-thr*(1+1e-6))
9         return ((a-b)/(2*thr*(1+2e-6))*(q-(1.0-thr*(1+1e-6))))+b)
10    else:
11        v = (1+(1-q)*x)
12        v = v * (v > 0)
13        return v**(1/(1-q))
14
15 def c_q(q):
16     thr = 0.001
17     if abs(q-1)<thr:
18         # return np.exp(x)
19         a = c_q(1.0+thr*(1+1e-6))
20         b = c_q(1.0-thr*(1+1e-6))
21         return ((a-b)/(2*thr*(1+2e-6))*(q-(1.0-thr*(1+1e-6))))+b)
22    elif q<1:
23        return 2 * np.pi**0.5 / ( (3-q) * (1-q)**0.5 ) * np.exp(gammaln(1/(1-q)) - gammaln((3-q)/(1-q)/2) )
24    else:
25        return np.pi**0.5 / (q-1)**0.5 * np.exp( gammaln((3-q)/(q-1)/2) - gammaln(1/(q-1)) )
26
27 def qGauss1D(x, A, mu, capsigma, q):
28     beta = c_q(q)**2/capsigma**2/(2*np.pi)
29     return A * e_q( - beta * (x-mu)**2, q)

```

Listing 1: The implementation of the CapSigma parametrized q -Gaussian function.

In practice, the implementation of $C(q)$ and $e(x, q)$ is not trivial, as although the full expression is convergent, some constituent functions diverge near $q = 1$. This is mitigated by linear interpolation in a small neighbourhood of $q = 1$. The domain where the interpolation is required can be decreased by using the log gamma function instead of the gamma function in $C(q)$. A code example is shown in Listing 1.

A.4.1 2D q -Gaussian normalization

The two-dimensional q -Gaussian fit function is based on the N -dimensional generalized formula in [134]. In this case, for $q < 2$,

$$f(\mathbf{r}) = \frac{V}{C_2(q)\sqrt{|\Sigma|}} e\left(-\frac{1}{4-2q}(\mathbf{r}-\boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{r}-\boldsymbol{\mu}), q\right), \quad (\text{A.61})$$

where $\mathbf{r}, \boldsymbol{\mu} \in \mathbb{R}^2, \Sigma \in \mathbb{R}^{2 \times 2}$, the integral of $f(\mathbf{r})$ is V , and

$$C_2(q) = \begin{cases} \frac{4-2q}{1-q} \pi \frac{\Gamma\left(\frac{2-q}{1-q}\right)}{\Gamma\left(\frac{2-q}{1-q}+1\right)} & \text{if } -\infty < q < 1, \\ \frac{4-2q}{q-1} \pi \frac{\Gamma\left(\frac{1}{q-1}-1\right)}{\Gamma\left(\frac{1}{q-1}\right)} & \text{if } 1 < q < 2. \end{cases} \quad (\text{A.62})$$

Implementation difficulties are handled by interpolation and the use of the log gamma function similarly to the one-dimensional case.

Appendix B

Full derivation of the luminous region parameters

B.1 The zero crossing angle case

In this section, the colliding bunches circulating in the two beams are differentiated by the index $p \in \{+, -\}$ displayed in round braces. Using this notation, by $a + (p)b$ we mean $a + b$ for the parameters of one bunch, while $a - b$ for the other.

For the sake of simplicity, the calculation takes the orbit frequency and the bunch proton multiplicities to be one.

For indexed quantities, we introduce the colon notation. If $v \in \mathbb{R}^n$ then for $1 \leq k < l \leq n$ the vector $v_{k:l} \in \mathbb{R}^{l-k+1}$ is a projection of v into the subspace span by the base vectors of index k to l . For example, for $v \in \mathbb{R}^4$ the symbol $v_{1:3}$ denotes the vector created from the first three coordinates of v . The notation generalizes to matrices analogously.

Denote the proton density function of our colliding bunches with $b^{(+)}(r)$ and $b^{(-)}(r)$ ($r \in \mathbb{R}^3$) at the bunch positions $(0, 0, 0)$ (the origin in the tracker coordinate system).

During a given vdM scan step the bunches are separated in the x and y directions. This is usually done in a symmetric fashion with respect to the head-on position but there are exceptions such as the imaging scan. Furthermore, some effects, such as the orbit drift can introduce further deviations. For this reason separate displacement vectors are used for the two beams:

$$\Delta^{(+)}, \Delta^{(-)} \in \mathbb{R}^2 \times \{0\}.$$

The unit vectors $e_i \in \mathbb{R}^3$ are also introduced:

$$\Delta^{(\pm)} = \begin{pmatrix} \Delta_1^{(\pm)} \\ \Delta_2^{(\pm)} \\ 0 \end{pmatrix} \quad \text{and} \quad e_i = \begin{pmatrix} \delta_{1i} \\ \delta_{2i} \\ \delta_{3i} \end{pmatrix},$$

where δ_{ij} stands for the Kronecker Delta.

If no bunch distortion takes place due to electromagnetic effects and no bunch shape evolution happens due to accelerator dynamics, the resulting distribution of the primary vertices, i.e. the beamspot shape or luminous region shape can be written as

$$B(\mathbf{r}) = 2 \int_{\mathbb{R}} \mathbf{b}^{(+)}(\mathbf{r} - \Delta^{(+)} - \mathbf{e}_3 t) \mathbf{b}^{(-)}(\mathbf{r} - \Delta^{(-)} + \mathbf{e}_3 t) dt. \quad (\text{B.1})$$

While B is also a function of the Δ offsets and hence $B(\mathbf{r}; \Delta^{(+)}, \Delta^{(-)})$ is the correct way to address it, for brevity this is sometimes omitted. The multiplier of 2 in this definition is just convention. It's purpose is revealed when we discuss how the rates are computed. One last thing to notice is that B is not a proper probability density function i.e. $N = \int B(\mathbf{r}) d^3r \neq 1$, rather N is proportional to the luminosity. However, since N is finite and $B > 0$ the normalized version, $B(\mathbf{r})/N$, is a proper probability density function of the primary vertices.

Having obtained the beamspot density function the beamspot parameters ($\hat{\mu}$ and $\hat{\Sigma}$) and the rate of our luminometer (\hat{R}) are simple to calculate:

$$\hat{\mu}(\Delta^{(+)}, \Delta^{(-)}) = \frac{\int_{\mathbb{R}^3} \mathbf{r} B(\mathbf{r}) d^3r}{\int_{\mathbb{R}^3} B(\mathbf{r}) d^3r}, \quad (\text{B.2})$$

$$\hat{\Sigma}(\Delta^{(+)}, \Delta^{(-)}) = \frac{\int_{\mathbb{R}^3} (\mathbf{r} - \hat{\mu}) \otimes (\mathbf{r} - \hat{\mu}) B(\mathbf{r}) d^3r}{\int_{\mathbb{R}^3} B(\mathbf{r}) d^3r}, \quad (\text{B.3})$$

$$\hat{R}(\Delta^{(+)}, \Delta^{(-)}) = \varepsilon \int_{\mathbb{R}^3} B(\mathbf{r}) d^3r, \quad (\text{B.4})$$

where, of course, the rate depends on the efficiency and acceptance of the luminometer we are using, hence the constant ε . With this knowledge under our belt, the volume of the beam-beam convolution shape is calculated.

Using Eq. 3.9 for the definition of σ_{vis} with

$$\Delta = \begin{pmatrix} \Delta_1 \\ \Delta_2 \\ 0 \end{pmatrix}$$

as the separation vector of the two beams in the x-y plane, we get

$$\begin{aligned}
\sigma_{\text{vis}} &= \int_{\mathbb{R}^2} d\Delta_1 d\Delta_2 \hat{R}(\Delta/2, -\Delta/2) \\
&= \varepsilon \int_{\mathbb{R}^2} d\Delta_1 d\Delta_2 \int_{\mathbb{R}^3} B(r, \Delta/2, -\Delta/2) d^3r \\
&= \varepsilon \int_{\mathbb{R}^3} d^3r \int_{\mathbb{R}^2} d\Delta_1 d\Delta_2 B(r, \Delta/2, -\Delta/2) \\
&= \varepsilon \int_{\mathbb{R}^3} d^3r \int_{\mathbb{R}^2} d\Delta_1 d\Delta_2 \int_{\mathbb{R}} dt b^{(+)}(r - \Delta/2 - e_3 t) b^{(-)}(r + \Delta/2 + e_3 t) \\
&= \varepsilon \int_{\mathbb{R}^3} d^3r \underbrace{\int_{\mathbb{R}^3} d\Delta_1 d\Delta_2 \int_{\mathbb{R}} d\Delta_3 b^{(+)}(r - \Delta/2 - e_3 \Delta_3/2) b^{(-)}(r + \Delta/2 + e_3 \Delta_3/2)}_{\mathbb{R}^3 \text{ convolution of } b^{(+)} \text{ and } b^{(-)}} \\
&= \varepsilon,
\end{aligned}$$

where the last equation is a consequence of the fact that the full convolution of two density functions is also a density function. So the coefficient (ε) that connects the rates to the beamspot volume is exactly σ_{vis} , which is a nice touch and motivates the multiplier of 2 in Eq. B.1.

In case $b^{(\pm)}(r)$ is the sum of three dimensional Gaussian functions (see Eq. 3.26), this calculation is possible to carry out analytically. Firstly, notice that $B(r)$ is possible to compute by breaking it up to component wise convolutions:

$$B(r) = 2 \int_{\mathbb{R}} b^{(+)}(r - \Delta^{(+)} - e_3 t) b^{(-)}(r - \Delta^{(-)} + e_3 t) dt \quad (\text{B.5})$$

$$= 2 \sum_{i,j} a_i^{(+)} a_j^{(-)} \int_{\mathbb{R}} g_i^{(+)}(r - \Delta^{(+)} - e_3 t) g_j^{(-)}(r - \Delta^{(-)} + e_3 t) dt, \quad (\text{B.6})$$

where $g_i^{(p)}$ abbreviates the normalized Gaussian components with the different means and covariance matrices:

$$g_i^{(p)}(r) = g(r, \mu_i^{(p)}, \Sigma_i^{(p)}). \quad (\text{B.7})$$

Secondly, notice that the convolution of normal distributions is still a normal distribution even if it is performed only in one of the directions.

We now examine how such a convolution is calculated along the t axis between two

components:

$$\int_{\mathbb{R}} \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma^{(+)})}} \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma^{(-)})}} \exp \left(-\frac{1}{2} \underbrace{\sum_{p \in \{+;- \}} (r - \tilde{\Delta}^{(p)} - (p)e_3 t)(\Sigma^{(p)})^{-1}(r - \tilde{\Delta}^{(p)} - (p)e_3 t)}_A \right) dt, \quad (\text{B.8})$$

where

$$\tilde{\Delta}^{(p)} = \mu^{(p)} + \Delta^{(p)} \quad (\text{B.9})$$

was introduced for a more compact notation. To perform this integral, the terms in A need to be reordered. As a first step, the summands in A are rewritten as quadratic terms in r_i and t :

$$\begin{aligned} A &= \sum_{p \in \{+;- \}} (r - \mu^{(p)} - \Delta^{(p)} - (p)e_3 t)(\Sigma^{(p)})^{-1}(r - \mu^{(p)} - \Delta^{(p)} - (p)e_3 t) \\ &= \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(+)} \\ r_2 - \tilde{\Delta}_2^{(+)} \\ r_3 - \tilde{\Delta}_3^{(+)} \\ t \end{pmatrix}^T \begin{pmatrix} (\Sigma^{(+)})^{-1} & -(\Sigma^{(+)})_{1,3}^{-1} \\ & -(\Sigma^{(+)})_{2,3}^{-1} \\ \text{-----} & -(\Sigma^{(+)})_{3,3}^{-1} \\ -(\Sigma^{(+)})_{3,1:3}^{-1} & (\Sigma^{(+)})_{3,3}^{-1} \end{pmatrix} \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(+)} \\ r_2 - \tilde{\Delta}_2^{(+)} \\ r_3 - \tilde{\Delta}_3^{(+)} \\ t \end{pmatrix} + \\ &+ \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(-)} \\ r_2 - \tilde{\Delta}_2^{(-)} \\ r_3 - \tilde{\Delta}_3^{(-)} \\ t \end{pmatrix}^T \begin{pmatrix} (\Sigma^{(-)})^{-1} & (\Sigma^{(-)})_{1,3}^{-1} \\ & (\Sigma^{(-)})_{2,3}^{-1} \\ \text{-----} & (\Sigma^{(-)})_{3,3}^{-1} \\ (\Sigma^{(-)})_{3,1:3}^{-1} & (\Sigma^{(-)})_{3,3}^{-1} \end{pmatrix} \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(-)} \\ r_2 - \tilde{\Delta}_2^{(-)} \\ r_3 - \tilde{\Delta}_3^{(-)} \\ t \end{pmatrix}, \quad (\text{B.10}) \end{aligned}$$

where the matrices are constructed in $\mathbb{R}^{4 \times 4}$ and symmetric. Let us introduce shorthands for our matrices of interest

$$(\mathcal{M}^{(+)})^{-1} = \begin{pmatrix} & (\Sigma^{(+)})^{-1} & & -(\Sigma^{(+)})_{1,3}^{-1} \\ & & & -(\Sigma^{(+)})_{2,3}^{-1} \\ & & & -(\Sigma^{(+)})_{3,3}^{-1} \\ \text{-----} & -(\Sigma^{(+)})_{1,3}^{-1} & -(\Sigma^{(+)})_{2,3}^{-1} & -(\Sigma^{(+)})_{3,3}^{-1} \\ & & & (\Sigma^{(+)})_{3,3}^{-1} \end{pmatrix}, \quad (\text{B.11})$$

$$(\mathcal{M}^{(-)})^{-1} = \begin{pmatrix} & (\Sigma^{(-)})^{-1} & & (\Sigma^{(-)})_{1,3}^{-1} \\ & & & (\Sigma^{(-)})_{2,3}^{-1} \\ & & & (\Sigma^{(-)})_{3,3}^{-1} \\ \text{-----} & (\Sigma^{(-)})_{1,3}^{-1} & (\Sigma^{(-)})_{2,3}^{-1} & (\Sigma^{(-)})_{3,3}^{-1} \\ & & & (\Sigma^{(-)})_{3,3}^{-1} \end{pmatrix}. \quad (\text{B.12})$$

As a second, step A needs to be transformed into a sum of a whole square and a constant:

$$A = \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \\ t - \mu_4 \end{pmatrix}^T \mathbf{M}^{-1} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \\ t - \mu_4 \end{pmatrix} + C, \quad (\text{B.13})$$

where μ_i and C are constants to be determined.

After comparing the quadratic, linear and constant terms in Eq. B.10 and Eq. B.13, we arrive to

$$\begin{aligned} \mathbf{M}^{-1} &= (\mathbf{M}^{(+)})^{-1} + (\mathbf{M}^{(-)})^{-1}, \\ \begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix} &= (\mathbf{M}^{-1})^{-1} \left\{ (\mathbf{M}^{(+)})^{-1} \begin{pmatrix} \tilde{\Delta}_1^{(+)} \\ \tilde{\Delta}_2^{(+)} \\ \tilde{\Delta}_3^{(+)} \\ 0 \end{pmatrix} + (\mathbf{M}^{(-)})^{-1} \begin{pmatrix} \tilde{\Delta}_1^{(-)} \\ \tilde{\Delta}_2^{(-)} \\ \tilde{\Delta}_3^{(-)} \\ 0 \end{pmatrix} \right\}, \\ C &= \begin{pmatrix} \tilde{\Delta}_1^{(+)} \\ \tilde{\Delta}_2^{(+)} \\ \tilde{\Delta}_3^{(+)} \\ 0 \end{pmatrix}^T (\mathbf{M}^{(+)})^{-1} \begin{pmatrix} \tilde{\Delta}_1^{(+)} \\ \tilde{\Delta}_2^{(+)} \\ \tilde{\Delta}_3^{(+)} \\ 0 \end{pmatrix} + \begin{pmatrix} \tilde{\Delta}_1^{(-)} \\ \tilde{\Delta}_2^{(-)} \\ \tilde{\Delta}_3^{(-)} \\ 0 \end{pmatrix}^T (\mathbf{M}^{(-)})^{-1} \begin{pmatrix} \tilde{\Delta}_1^{(-)} \\ \tilde{\Delta}_2^{(-)} \\ \tilde{\Delta}_3^{(-)} \\ 0 \end{pmatrix} - \begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix}^T \mathbf{M}^{-1} \begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix}. \end{aligned}$$

Now we can perform the integration according to t . Notice that this is a one-dimensional Gaussian integral in t and, hence, we can use the formula

$$\int_{\mathbb{R}} dt \exp\left(-\frac{(t-m)^2}{2\sigma^2} + c\right) = \sigma e^c \sqrt{2\pi}. \quad (\text{B.14})$$

Separating the terms containing t in two steps, we get

$$A = (\mathbf{M}^{-1})_{4,4} (t - \mu_4)^2 \quad (\text{B.15})$$

$$+ 2(t - \mu_4) (\mathbf{M}^{-1})_{4,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.16})$$

$$+ \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix}^T (\mathbf{M}^{-1})_{1:3,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.17})$$

$$+ C \quad (\text{B.18})$$

and opening line B.16

$$A = (M^{-1})_{4,4}(t - \mu_4)^2 \quad (\text{B.19})$$

$$+ 2t(M^{-1})_{4,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.20})$$

$$- 2\mu_4(M^{-1})_{4,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.21})$$

$$+ \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix}^T (M^{-1})_{1:3,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.22})$$

$$+ C. \quad (\text{B.23})$$

$$(\text{B.24})$$

Forming the whole square plus constant structure for t , we get

$$A = (M^{-1})_{4,4} \left\{ t - \mu_4 + \frac{1}{(M^{-1})_{4,4}} (M^{-1})_{4,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \right\}^2 \quad (\text{B.25})$$

$$- \frac{1}{(M^{-1})_{4,4}} \left\{ (M^{-1})_{4,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \right\}^2 \quad (\text{B.26})$$

$$+ \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix}^T (M^{-1})_{1:3,1:3} \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} \quad (\text{B.27})$$

$$+ C. \quad (\text{B.28})$$

Mapping this to Eq. B.14 and reordering some terms gives

$$\sigma^{-2} = (M^{-1})_{4,4}, \quad (\text{B.29})$$

$$c = \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix}^T Q \begin{pmatrix} r_1 - \mu_1 \\ r_2 - \mu_2 \\ r_3 - \mu_3 \end{pmatrix} + C, \quad (\text{B.30})$$

where

$$Q = (M^{-1})_{1:3,1:3} - \frac{1}{(M^{-1})_{4,4}} (M^{-1})_{4,1:3} \otimes (M^{-1})_{4,1:3}. \quad (\text{B.31})$$

Notice that the most expensive operation, the inversion of M^{-1} , is independent of the offsets which can accelerate this computation greatly. Also notice that these operations are all vectorizable and, hence, highly optimized linear algebra libraries can be used.

At this point we are ready to write the result of the integration in Eq. B.8:

$$\int_{\mathbb{R}} g_i^{(+)}(\mathbf{r} - \Delta^{(+)} - \mathbf{e}_3 t) g_j^{(-)}(\mathbf{r} - \Delta^{(-)} + \mathbf{e}_3 t) dt = C_{(ij)} g(\mathbf{r}, \mathbf{m}_{(ij)}, S_{(ij)}), \quad (\text{B.32})$$

where

$$\mathbf{m}_{(ij)} = \boldsymbol{\mu}_{1:3}, \quad (\text{B.33})$$

$$S_{(ij)} = \left((M^{-1})_{1:3,1:3} - \frac{1}{(M^{-1})_{4,4}} (M^{-1})_{4,1:3} \otimes (M^{-1})_{4,1:3} \right)^{-1}, \quad (\text{B.34})$$

and

$$\begin{aligned} C_{(ij)} &= \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma_i^{(+)})}} \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma_j^{(-)})}} \\ &\quad \cdot \frac{1}{\sqrt{(M^{-1})_{4,4}}} e^C \sqrt{2\pi} \sqrt{(2\pi)^3 \det(S_{(ij)})} \\ &= \frac{1}{2\pi \sqrt{\det(\Sigma_i^{(+)}) \det(M^{-1}) \det(\Sigma_j^{(-)})}} e^C, \end{aligned} \quad (\text{B.35})$$

where M^{-1} , $\boldsymbol{\mu}$ and C are calculated using $\boldsymbol{\mu}_i^{(+)}$, $\Sigma_i^{(+)}$ and $\boldsymbol{\mu}_j^{(-)}$, $\Sigma_j^{(-)}$, i.e., the corresponding Gaussian parameters.

Also notice that since g is a Gaussian function, we can just read off the result of the linear and quadratic integrals

$$\int_{\mathbb{R}^3} \mathbf{r} C_{(ij)} g(\mathbf{r}, \mathbf{m}_{(ij)}, S_{(ij)}) d^3 \mathbf{r} = \mathbf{m}_{(ij)} C_{(ij)}, \quad (\text{B.36})$$

$$\int_{\mathbb{R}^3} (\mathbf{r} - \mathbf{m}_{(ij)}) \otimes (\mathbf{r} - \mathbf{m}_{(ij)}) C_{(ij)} g(\mathbf{r}, \mathbf{m}_{(ij)}, S_{(ij)}) d^3 \mathbf{r} = S_{(ij)} C_{(ij)}, \quad (\text{B.37})$$

and

$$\sigma_{\text{vis}} \int_{\mathbb{R}^3} C_{(ij)} g(\mathbf{r}, \mathbf{m}_{(ij)}, S_{(ij)}) d^3 \mathbf{r} = \sigma_{\text{vis}} C_{(ij)}. \quad (\text{B.38})$$

Applying the linear property of integration to Eq. B.2-B.4, we arrive to the simple forms

$$\hat{\mathbf{R}}(\Delta^{(+)}, \Delta^{(-)}) = \sigma_{\text{vis}} \int_{\mathbb{R}^3} \mathbf{B}(\mathbf{r}) d^3 \mathbf{r} = \sigma_{\text{vis}} \sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)}, \quad (\text{B.39})$$

$$\hat{\boldsymbol{\mu}}(\Delta^{(+)}, \Delta^{(-)}) = \frac{\int_{\mathbb{R}^3} \mathbf{r} \mathbf{B}(\mathbf{r}) d^3 \mathbf{r}}{\int_{\mathbb{R}^3} \mathbf{B}(\mathbf{r}) d^3 \mathbf{r}} = \frac{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)} \mathbf{m}_{(ij)}}{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)}}, \quad (\text{B.40})$$

and

$$\begin{aligned}
\hat{\Sigma}(\Delta^{(+)}, \Delta^{(-)}) &= \frac{\int_{\mathbb{R}^3} (\mathbf{r} - \hat{\boldsymbol{\mu}}) \otimes (\mathbf{r} - \hat{\boldsymbol{\mu}}) B(\mathbf{r}) d^3\mathbf{r}}{\int_{\mathbb{R}^3} B(\mathbf{r}) d^3\mathbf{r}} \\
&= \frac{\int_{\mathbb{R}^3} \mathbf{r} \otimes \mathbf{r} B(\mathbf{r}) d^3\mathbf{r}}{\int_{\mathbb{R}^3} B(\mathbf{r}) d^3\mathbf{r}} - \hat{\boldsymbol{\mu}} \otimes \hat{\boldsymbol{\mu}} \\
&= \frac{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)} (S_{(ij)} + \mathbf{m}_{(ij)} \otimes \mathbf{m}_{(ij)})}{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)}} - \hat{\boldsymbol{\mu}} \otimes \hat{\boldsymbol{\mu}} \\
&= \frac{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)} S_{(ij)}}{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)}} + \frac{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)} (\mathbf{m}_{(ij)} \otimes \mathbf{m}_{(ij)} - \hat{\boldsymbol{\mu}} \otimes \hat{\boldsymbol{\mu}})}{\sum_{i,j} \mathbf{a}_i^{(+)} \mathbf{a}_j^{(-)} C_{(ij)}}.
\end{aligned} \tag{B.41}$$

B.1.1 The dynamic β^* correction

In the current model, the focusing/defocusing effect of the electromagnetic interaction between the colliding bunches is handled by multiplying their width by a separation-dependent factor $m_\beta^x(\Delta^{(+)} - \Delta^{(-)})$ and $m_\beta^y(\Delta^{(+)} - \Delta^{(-)})$,

$$\Sigma_i^{(\pm)} \rightarrow \Sigma_i'^{(\pm)} = \begin{pmatrix} m_\beta^x & 0 & 0 \\ 0 & m_\beta^y & 0 \\ 0 & 0 & 1 \end{pmatrix} \Sigma_i^{(\pm)} \begin{pmatrix} m_\beta^x & 0 & 0 \\ 0 & m_\beta^y & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{B.42}$$

This, in the zero crossing angle case, leads to

$$\mathbf{M}^{-1} \rightarrow \mathbf{M}'^{-1} = \begin{pmatrix} (m_\beta^x)^{-1} & 0 & 0 & 0 \\ 0 & (m_\beta^y)^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \mathbf{M}^{-1} \begin{pmatrix} (m_\beta^x)^{-1} & 0 & 0 & 0 \\ 0 & (m_\beta^y)^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{B.43}$$

The same formula applies for $(\mathbf{M}^{(\pm)})^{-1}$, and, hence,

$$\Sigma_{(ij)} \rightarrow \Sigma'_{(ij)} = \begin{pmatrix} m_\beta^x & 0 & 0 \\ 0 & m_\beta^y & 0 \\ 0 & 0 & 1 \end{pmatrix} \Sigma_{(ij)} \begin{pmatrix} m_\beta^x & 0 & 0 \\ 0 & m_\beta^y & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{B.44}$$

and $\mu_{(ij)}$ and $C_{(ij)}$ should be calculated using \mathbf{M}'^{-1} and $(\mathbf{M}'^{(\pm)})^{-1}$.

B.2 The non-zero crossing angle case

Currently, the vdM procedure is performed with vanishing crossing angles for the main vdM programs at the highest pp energies. For future use cases, the non-zero crossing

angle scenario is formulated. The angle of the orbits with the z axis of the tracker system is denoted by $\Theta^{(p)}$.




Figure B.1: A visualization of the crossing angle at the interaction point. In the CMS experiment, the crossing plane coincides with the x - z plane. The red and blue lines signify the two orbits, and the ellipsoids indicate two colliding bunches. The arrows on the orbits indicate the direction of the bunch propagation. The crossing angle ($\Theta = \Theta^{(+)} + \Theta^{(-)}$) is enlarged for the purposes of visualization.

$$\int_{\mathbb{R}} \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma^{(+)})}} \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma^{(-)})}} \exp \left(-\frac{1}{2} \underbrace{\sum_{p \in \{+;- \}} (\mathbf{r} - \tilde{\Delta}^{(p)} - (p)\mathbf{e}_p t)(\Sigma^{(p)})^{-1}(\mathbf{r} - \tilde{\Delta}^{(p)} - (p)\mathbf{e}_p t)}_A \right) dt,$$

where

$$\mathbf{e}_+ = \begin{pmatrix} \sin(\Theta_+) \\ 0 \\ 0 \\ \cos(\Theta_+) \end{pmatrix}, \quad \mathbf{e}_- = \begin{pmatrix} \sin(\Theta_-) \\ 0 \\ 0 \\ \cos(\Theta_-) \end{pmatrix}.$$

To perform this integral, we need to reorder the terms in A . As a first step, we write the summands in A as quadratic terms in r_i and t . Just as previously, an M matrix is

introduced to aid with the calculation

$$(\mathbf{M}^{(p)})^{-1} = \left(\begin{array}{ccc|c} & & & (\mathbf{M}^{(p)})_{1,4}^{-1} \\ & (\boldsymbol{\Sigma}^{(p)})^{-1} & & (\mathbf{M}^{(p)})_{2,4}^{-1} \\ \hline & & & (\mathbf{M}^{(p)})_{3,4}^{-1} \\ \hline (\mathbf{M}^{(p)})_{4,1}^{-1} & (\mathbf{M}^{(p)})_{4,2}^{-1} & (\mathbf{M}^{(p)})_{4,3}^{-1} & (\mathbf{M}^{(p)})_{4,4}^{-1} \end{array} \right),$$

where

$$\begin{aligned} (\mathbf{M}^{(p)})_{4,1}^{-1} &= (\mathbf{M}^{(p)})_{1,4}^{-1} = \sin(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{1,1}^{-1} - (p) \cos(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{3,1}^{-1}, \\ (\mathbf{M}^{(p)})_{4,2}^{-1} &= (\mathbf{M}^{(p)})_{2,4}^{-1} = \sin(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{1,2}^{-1} - (p) \cos(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{3,2}^{-1}, \\ (\mathbf{M}^{(p)})_{4,3}^{-1} &= (\mathbf{M}^{(p)})_{3,4}^{-1} = \sin(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{1,3}^{-1} - (p) \cos(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{3,3}^{-1}, \\ (\mathbf{M}^{(p)})_{4,4}^{-1} &= \sin^2(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{1,1}^{-1} + \cos^2(\Theta^{(p)}) (\boldsymbol{\Sigma}^{(p)})_{3,3}^{-1}. \end{aligned}$$

With this,

$$\begin{aligned} \mathbf{A} &= \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(+)} \\ r_2 - \tilde{\Delta}_2^{(+)} \\ r_3 - \tilde{\Delta}_3^{(+)} \\ t \end{pmatrix}^T (\mathbf{M}^{(+)})^{-1} \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(+)} \\ r_2 - \tilde{\Delta}_2^{(+)} \\ r_3 - \tilde{\Delta}_3^{(+)} \\ t \end{pmatrix} + \\ &+ \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(-)} \\ r_2 - \tilde{\Delta}_2^{(-)} \\ r_3 - \tilde{\Delta}_3^{(-)} \\ t \end{pmatrix}^T (\mathbf{M}^{(-)})^{-1} \begin{pmatrix} r_1 - \tilde{\Delta}_1^{(-)} \\ r_2 - \tilde{\Delta}_2^{(-)} \\ r_3 - \tilde{\Delta}_3^{(-)} \\ t \end{pmatrix}, \end{aligned}$$

and, hence, the problem is brought to a shape identical to the previous case, the calculation proceeds in an identical manner.

Appendix C

Simplified derivation of the luminous region parameters

For this calculation, we model the bunch proton densities as the sum of factorizable Gaussian components

$$f(x, y, z) = \sum_i a_i g^{(i)}(x, y, z) = \sum_i a_i g_x^{(i)}(x) g_y^{(i)}(y) g_z^{(i)}(z), \quad (\text{C.1})$$

where g represents normalized Gaussian functions with different means and widths indexed by i and the direction and $\sum_i a_i = 1$. N.B. the sum of such factorizable components can lead to highly non-factorizable shapes. We also assume a crossing angle of zero degrees. With these assumptions, the expression Eq. B.8 factorizes in x , y and z .

First, we examine the transverse directions. Here x and y are handled identically, so only one of these directions is discussed. For simplicity, g_1 and g_2 signify a Gaussian component in beam one and beam two respectively.

$$g_1(x)g_2(x) = N_1 \exp\left(-\frac{(x - \mu_1)^2}{2\sigma_1^2}\right) N_2 \exp\left(-\frac{(x - \mu_2)^2}{2\sigma_2^2}\right) \quad (\text{C.2})$$

$$= N_1 N_2 \exp\left(-\frac{(x - \mu_1)^2}{2\sigma_1^2} - \frac{(x - \mu_2)^2}{2\sigma_2^2}\right), \quad (\text{C.3})$$

where $N_i = 1/\sqrt{2\pi\sigma_i^2}$ is the normalization term. Reordering the exponent:

$$\frac{(x - \mu_1)^2}{2\sigma_1^2} + \frac{(x - \mu_2)^2}{2\sigma_2^2} = \frac{\sigma_2^2(x - \mu_1)^2 + \sigma_1^2(x - \mu_2)^2}{2\sigma_1^2\sigma_2^2} = \quad (\text{C.4})$$

$$= \frac{\sigma_2^2 x^2 - 2x\sigma_2^2\mu_1 + \sigma_2^2\mu_1^2 + \sigma_1^2 x^2 - 2x\sigma_1^2\mu_2 + \sigma_1^2\mu_2^2}{2\sigma_1^2\sigma_2^2} = \quad (\text{C.5})$$

$$= \frac{(\sigma_2^2 + \sigma_1^2)x^2 - 2x(\sigma_1^2\mu_2 + \sigma_2^2\mu_1) + (\sigma_2^2\mu_1^2 + \sigma_1^2\mu_2^2)}{2\sigma_1^2\sigma_2^2} = \quad (C.6)$$

Transforming the expression to a complete square in x plus a constant:

$$= \frac{(\sigma_2^2 + \sigma_1^2)(x - (\sigma_1^2\mu_2 + \sigma_2^2\mu_1)/(\sigma_2^2 + \sigma_1^2))^2}{2\sigma_1^2\sigma_2^2} + \frac{(\sigma_2^2\mu_1^2 + \sigma_1^2\mu_2^2 - (\sigma_1^2\mu_2 + \sigma_2^2\mu_1)^2/(\sigma_2^2 + \sigma_1^2))}{2\sigma_1^2\sigma_2^2}. \quad (C.7)$$

Hence, $g_1(x)g_2(x)$ also describes a bell curve with the parameters:

$$\begin{aligned} A &= N_1 N_2 \exp\left(-\frac{(\sigma_2^2\mu_1^2 + \sigma_1^2\mu_2^2 - (\sigma_1^2\mu_2 + \sigma_2^2\mu_1)^2/(\sigma_2^2 + \sigma_1^2))}{2\sigma_1^2\sigma_2^2}\right) \\ \mu &= (\sigma_1^2\mu_2 + \sigma_2^2\mu_1)/(\sigma_2^2 + \sigma_1^2) \\ \sigma^2 &= (\sigma_1^2\sigma_2^2)/(\sigma_2^2 + \sigma_1^2) \end{aligned} \quad (C.8)$$

Performing the computation in the z direction:

$$\int_{\mathbb{R}} N_1 \exp\left(-\frac{(z - \mu_1 + ct)^2}{2\sigma_1^2}\right) N_2 \exp\left(-\frac{(z - \mu_2 - ct)^2}{2\sigma_2^2}\right) dt = \quad (C.9)$$

$$= N_1 N_2 \int_{\mathbb{R}} \exp\left(-\frac{\sigma_2^2(z - \mu_1 + ct)^2 + \sigma_1^2(z - \mu_2 - ct)^2}{2\sigma_1^2\sigma_2^2}\right) dt. \quad (C.10)$$

Expanding parts of the numerator of the exponent:

$$\begin{aligned} \sigma_2^2(z - \mu_1 + ct)^2 &= \sigma_2^2(z^2 + \mu_1^2 + (ct)^2 - 2\mu_1 z - 2\mu_1 ct + 2zct), \\ \sigma_1^2(z - \mu_2 - ct)^2 &= \sigma_1^2(z^2 + \mu_2^2 + (ct)^2 - 2\mu_2 z + 2\mu_2 ct - 2zct). \end{aligned} \quad (C.11)$$

Transforming to a complete square in t , we get

$$\begin{aligned} (\sigma_1^2 + \sigma_2^2) \left\{ (ct)^2 + 2 \frac{\sigma_1^2(\mu_2 - z) + \sigma_2^2(z - \mu_1)}{\sigma_1^2 + \sigma_2^2} ct + \left(\frac{\sigma_1^2(\mu_2 - z) + \sigma_2^2(z - \mu_1)}{\sigma_1^2 + \sigma_2^2} \right)^2 \right\} - \\ - \frac{(\sigma_1^2(\mu_2 - z) + \sigma_2^2(z - \mu_1))^2}{\sigma_1^2 + \sigma_2^2} + (\sigma_1^2 + \sigma_2^2)z^2 + \sigma_1^2\mu_2^2 + \sigma_2^2\mu_1^2 - 2(\sigma_1^2\mu_2 + \sigma_2^2\mu_1)z. \end{aligned} \quad (C.12)$$

The complete square integrates out to

$$\sqrt{2\pi} \sqrt{\frac{\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \frac{1}{c}}, \quad (C.13)$$

the remaining exponent is

$$(\sigma_1^2 + \sigma_2^2)z^2 - 2(\sigma_1^2\mu_2 + \sigma_2^2\mu_1)z + \sigma_1^2\mu_2^2 + \sigma_2^2\mu_1^2 - \frac{((\sigma_2^2 - \sigma_1^2)z + \sigma_1^2\mu_2 - \sigma_2^2\mu_1)^2}{\sigma_1^2 + \sigma_2^2} = \quad (C.14)$$

$$\begin{aligned}
&= \left(\sigma_1^2 + \sigma_2^2 - \frac{(\sigma_2^2 - \sigma_1^2)^2}{\sigma_1^2 + \sigma_2^2} \right) z^2 - 2 \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right) z + \\
&\quad + \sigma_1^2 \mu_2^2 + \sigma_2^2 \mu_1^2 - \frac{(\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1)^2}{\sigma_1^2 + \sigma_2^2}.
\end{aligned} \tag{C.15}$$

Reordering the coefficient of the quadratic term:

$$\left(\sigma_1^2 + \sigma_2^2 - \frac{(\sigma_2^2 - \sigma_1^2)^2}{\sigma_1^2 + \sigma_2^2} \right) = \frac{(\sigma_2^2 + \sigma_1^2)^2 - (\sigma_2^2 - \sigma_1^2)^2}{\sigma_1^2 + \sigma_2^2} = \frac{4\sigma_2^2 \sigma_1^2}{\sigma_1^2 + \sigma_2^2}. \tag{C.16}$$

Transforming the expression to a complete square in z plus a constant:

$$\begin{aligned}
&\frac{4\sigma_2^2 \sigma_1^2}{\sigma_1^2 + \sigma_2^2} \left\{ z^2 - 2 \frac{\sigma_1^2 + \sigma_2^2}{4\sigma_2^2 \sigma_1^2} \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right) z \right\} + \\
&\quad + \sigma_1^2 \mu_2^2 + \sigma_2^2 \mu_1^2 - \frac{(\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1)^2}{\sigma_1^2 + \sigma_2^2} =
\end{aligned} \tag{C.17}$$

$$\begin{aligned}
&= \frac{4\sigma_2^2 \sigma_1^2}{\sigma_1^2 + \sigma_2^2} \left\{ z - \frac{\sigma_1^2 + \sigma_2^2}{4\sigma_2^2 \sigma_1^2} \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right) \right\}^2 - \\
&\quad - \frac{\sigma_1^2 + \sigma_2^2}{4\sigma_2^2 \sigma_1^2} \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right)^2 + \\
&\quad + \sigma_1^2 \mu_2^2 + \sigma_2^2 \mu_1^2 - \frac{(\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1)^2}{\sigma_1^2 + \sigma_2^2}
\end{aligned} \tag{C.18}$$

Hence, the result of the integral in Eq. C.9 is a bell curve with the parameters

$$\begin{aligned}
&A = N_1 N_2 \sqrt{2\pi} \sqrt{\frac{\sigma_1^2 \sigma_2^2}{\sigma_1^2 + \sigma_2^2}} \frac{1}{c} e^E, \\
&E = \frac{\sigma_1^2 + \sigma_2^2}{8\sigma_2^4 \sigma_1^4} \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right)^2 - \\
&\quad - \frac{1}{2\sigma_2^2 \sigma_1^2} \left(\sigma_1^2 \mu_2^2 + \sigma_2^2 \mu_1^2 - \frac{(\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1)^2}{\sigma_1^2 + \sigma_2^2} \right), \\
&\mu = \frac{\sigma_1^2 + \sigma_2^2}{4\sigma_2^2 \sigma_1^2} \left(\frac{\sigma_2^2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2} (\sigma_1^2 \mu_2 - \sigma_2^2 \mu_1) + \sigma_1^2 \mu_2 + \sigma_2^2 \mu_1 \right), \\
&\sigma = \frac{1}{2} \sqrt{\sigma_1^2 + \sigma_2^2}.
\end{aligned} \tag{C.19}$$

For obtaining the result of Eq. B.8, the amplitudes in Eq. C.8 for x and y and in Eq. C.19 for z need to be multiplied, taking into account that the symbols in the formulas, such as N_i , differ in the different directions, so in fact N_i^x , N_i^y , N_i^z would be the most appropriate way of indexing them after this step. The means and widths are obtained directly. Finally, the beamspot parameters for the composite beams are computed using Equations B.39-B.41.

A vectorized implementation of this calculation is shown in Listing 2. The code accepts parameters for two bunch proton density functions as lists of their Gaussian components.

The individual components are represented using seven parameters: the volume, three centres and three widths. It is required that the volumes in a bunch add up to unity. The parameters are calculated at different beam positions. The positions are represented as a 2D array, where the first index represents the steps, while the second stands for the x , y , z spatial coordinates.

The simplicity of this code not only makes the execution fast, but also opens up the possibility of easily replacing the shown numpy numerical vector algebra package with computational graph based auto-differentiating backends that have risen to great popularity in recent years due to their machine learning applications.

```

1 import numpy as np
2
3 def beamParams2LRcentralMoments(params1, params2, beam1Pos, beam2Pos):
4
5     I = np.zeros(len(beam1Pos))
6     I1x = np.zeros(len(beam1Pos))
7     I2x = np.zeros(len(beam1Pos))
8     I1y = np.zeros(len(beam1Pos))
9     I2y = np.zeros(len(beam1Pos))
10    I1z = np.zeros(len(beam1Pos))
11    I2z = np.zeros(len(beam1Pos))
12
13    for p1 in params1:
14        for p2 in params2:
15
16            V1, mu1x, mu1y, mu1z, s1x, s1y, s1z = p1
17            V2, mu2x, mu2y, mu2z, s2x, s2y, s2z = p2
18
19            mu1xs = mu1x + beam1Pos[:,0]
20            mu2xs = mu2x + beam2Pos[:,0]
21            mu1ys = mu1y + beam1Pos[:,1]
22            mu2ys = mu2y + beam2Pos[:,1]
23            mu1zs = mu1z + beam1Pos[:,2]
24            mu2zs = mu2z + beam2Pos[:,2]
25
26            A1 = V1 / ((2*np.pi)**1.5 * s1x * s1y * s1z)
27            A2 = V2 / ((2*np.pi)**1.5 * s2x * s2y * s2z)
28
29            sx = (1/(s1x**2 + s2x**2))**0.5
30            sy = (1/(s1y**2 + s2y**2))**0.5
31            sz = 0.5*(s1z**2 + s2z**2)**0.5
32
33            tempX1 = s1x**2*mu2xs + s2x**2*mu1xs
34            tempX2 = s1x**2+s2x**2
35            muX = tempX1/tempX2
36            tempY1 = s1y**2*mu2ys + s2y**2*mu1ys
37            tempY2 = s1y**2+s2y**2
38            muY = tempY1/tempY2
39            tempZ1 = (s1z**2 + s2z**2)/(4*s1z**2*s2z**2)
40            tempZ2 = ((s2z**2 - s1z**2)/(s1z**2 + s2z**2))*(s1z**2*mu2zs - s2z**2*mu1zs) + s1z**2*mu2zs + s2z**2*mu1zs
41            muZ = tempZ1 * tempZ2
42
43            cx = - ( s1x**2*mu2xs**2 + s2x**2*mu1xs**2 - tempX1**2/tempX2) / (2*s1x**2*s2x**2)
44            cy = - ( s1y**2*mu2ys**2 + s2y**2*mu1ys**2 - tempY1**2/tempY2) / (2*s1y**2*s2y**2)
45            cz = tempZ1 * tempZ2**2
46            cz -= (s1z**2*mu2zs**2 + s2z**2*mu1zs**2 - (s1z**2*mu2zs - s2z**2*mu1zs)**2/(s1z**2 + s2z**2))
47            cz /= (2*s1z**2*s2z**2)
48            A = A1*A2
49            A *= np.sqrt( (s1z*s2z)**2 / (s1z**2 + s2z**2) )*(2*np.pi)**0.5
50            A *= np.exp( cx + cy + cz )
51
52            V = A*((2*np.pi)**1.5 * sx * sy * sz)
53
54            I1x += V * (muX)
55            I2x += V * (sx**2 + muX**2)
56            I1y += V * (muY)
57            I2y += V * (sy**2 + muY**2)
58            I1z += V * (muZ)
59            I2z += V * (sz**2 + muZ**2)
60            I += V
61
62            I1x /= I
63            I2x /= I
64            I1y /= I
65            I2y /= I
66            I1z /= I
67            I2z /= I
68
69            w2X = I2x-I1x**2
70            w2Y = I2y-I1y**2
71            w2Z = I2z-I1z**2
72            return I, I1x, I1y, I1z, w2X, w2Y, w2Z

```

Listing 2: A vectorized implementation of the simplified luminous region parameter computation.

Appendix D

Derivation of the emittance scan based linearity modification factor

Let $SBIL_0$ and $SBIL_m$ be the true and measured value of the single bunch instantaneous luminosity, the second as measured by a detector with a non-zero linearity α according to

$$SBIL_m = SBIL_0(1 - \alpha \cdot SBIL_0). \quad (D.1)$$

Traditionally, in physics fills SBIL is in the range of 4-10 Hz/ μ b depending on the Era, and α is a few times $\pm 0.1\%$ /(Hz/ μ b) up to about $\pm 1\%$ /(Hz/ μ b).

Since the emittance scans are performed using normalized rates, it is more convenient to use this quantity in our computations. The normalized rate measured by the luminometer is

$$r_m = SBIL_m \cdot \sigma_{vis}^0 / (J_1 J_2), \quad (D.2)$$

where J_i indicate the bunch proton counts and σ_{vis}^0 is the unbiased σ_{vis} measured in vdM conditions. A true rate can also be associated to $SBIL_0$ analogously

$$r_0 = SBIL_0 \cdot \sigma_{vis}^0 / (J_1 J_2). \quad (D.3)$$

With this, we can write

$$r_m = r_0(1 - \beta r_0), \quad (D.4)$$

where

$$\beta = \frac{\alpha J_1 J_2}{\sigma_{vis}}. \quad (D.5)$$

Suppose a factorizable Gaussian bunch overlap shape. The true rate a detector with

perfect linearity ($\alpha = 0$) would measure is:

$$r_0(x, y) = R_0 \exp\left(-\frac{x^2}{2\Sigma_x^2} - \frac{y^2}{2\Sigma_y^2}\right). \quad (\text{D.6})$$

The associated calibration constant according to the vdM method is

$$\sigma_{\text{vis}}^0 = \int r_0 \, dx \, dy = 2\pi R_0 \Sigma_x \Sigma_y. \quad (\text{D.7})$$

After applying the non-linearity, we can compute the biased σ_{vis} value using a two-dimensional integral:

$$\begin{aligned} \sigma_{\text{vis}}^\alpha &= \int r_m \, dx \, dy = \int r_0(1 - \beta r_0) \, dx \, dy \\ &= \int r_0 \, dx \, dy - \beta \int r_0^2 \, dx \, dy = \sigma_{\text{vis}}^0 - \beta \int r_0^2 \, dx \, dy. \end{aligned} \quad (\text{D.8})$$

Now,

$$r_0^2(x, y) = R_0^2 \exp\left(-2\frac{x^2}{2\Sigma_x^2} - 2\frac{y^2}{2\Sigma_y^2}\right) = R_0^2 \exp\left(-\frac{x^2}{2(\Sigma_x/\sqrt{2})^2} - \frac{y^2}{2(\Sigma_y/\sqrt{2})^2}\right), \quad (\text{D.9})$$

thus

$$\int r_0^2 \, dx \, dy = 2\pi R_0^2 (\Sigma_x/\sqrt{2})(\Sigma_y/\sqrt{2}) = 2\pi R_0^2 \Sigma_x \Sigma_y / 2 = R_0 \sigma_{\text{vis}}^0 / 2, \quad (\text{D.10})$$

which leads to

$$\sigma_{\text{vis}}^\alpha = \sigma_{\text{vis}}^0 (1 - \beta R_0 / 2). \quad (\text{D.11})$$

Let us now write back β from Eq. D.5:

$$\sigma_{\text{vis}}^\alpha = \sigma_{\text{vis}}^0 \left(1 - \frac{\alpha J_1 J_2}{\sigma_{\text{vis}}} R_0 / 2\right) = \sigma_{\text{vis}}^0 \left(1 - \frac{\alpha R_0 J_1 J_2}{2 \sigma_{\text{vis}}}\right). \quad (\text{D.12})$$

Notice that

$$\frac{R_0 J_1 J_2}{\sigma_{\text{vis}}} = \text{SBIL}_{\text{peak}}, \quad (\text{D.13})$$

which gives

$$\boxed{\sigma_{\text{vis}}^\alpha = \sigma_{\text{vis}}^0 \left(1 - \frac{\alpha}{2} \cdot \text{SBIL}_{\text{peak}}\right)}. \quad (\text{D.14})$$

So, based on an analytical calculation, if we are able to extract σ_{vis} using a two-dimensional integral, the slope of the measured σ_{vis} as a function of peak SBIL only amounts to half of the true non-linearity. The exact coefficient is of course also subject to variation based on the validity of the assumption in Eq. D.5. Based on this, the modification factor m is defined for any quantity q_0 for which the effect of non-linearity can be approximated as

$$q_\alpha = q_0 (1 - m\alpha \cdot \text{SBIL}_{\text{peak}}). \quad (\text{D.15})$$

In case a two-dimensional integral is not an option, the measured calibration constant $\sigma_{\text{vis}}^m = 2\pi R_m \Sigma_x^m \Sigma_y^m$ is also impacted by the non-linearity induced non-factorization. Examine the one-dimensional integrals measured in this case (using the vdM method):

$$I_x^m = \int r_m(x, 0) dx \quad (\text{D.16})$$

$$= \int r_0(x, 0)(1 - \beta r_0(x, 0)) dx \quad (\text{D.17})$$

$$= \int r_0(x, 0) dx - \beta \int r_0^2(x, 0) dx. \quad (\text{D.18})$$

Introducing

$$I_x^0 = \int dx r_0(x, 0), \quad (\text{D.19})$$

similarly to Eq. D.9, we get

$$\int r_0^2(x, 0) dx = \sqrt{2\pi} R_m \Sigma_x^m / \sqrt{2} = I_x^0 / \sqrt{2}, \quad (\text{D.20})$$

and, therefore,

$$\sigma_{\text{vis}}^m = 2\pi R_m \Sigma_x^m \Sigma_y^m = \frac{I_x^m I_y^m}{R_m} = \frac{I_x^0 I_y^0 (1 - \beta R_0 / \sqrt{2})^2}{R_0 (1 - \beta R_0)}, \quad (\text{D.21})$$

and finally

$$\sigma_{\text{vis}}^m = \sigma_{\text{vis}}^0 \frac{(1 - \beta R_0 / \sqrt{2})^2}{1 - \beta R_0}. \quad (\text{D.22})$$

Performing an approximation to a first order in $\beta R_0 = \alpha \cdot \text{SBIL}_{\text{peak}} \ll 1$, we get

$$\begin{aligned} \frac{(1 - \beta R_0 / \sqrt{2})^2}{1 - \beta R_0} - 1 &= \frac{(1 - \beta R_0 / \sqrt{2})^2 - (1 - \beta R_0)}{1 - \beta R_0} \\ &= \frac{-\sqrt{2}\beta R_0 + \beta^2 R_0^2 / 2 + \beta R_0}{1 - \beta R_0} \\ &\approx (1 - \sqrt{2})\beta R_0 \\ &\approx -0.414\beta R_0. \end{aligned} \quad (\text{D.23})$$

With this, and rounding the modification factor to two decimal places, a practical formula taking the non-linearity-induced non-factorisation into account is obtained:

$$\boxed{\sigma_{\text{vis}}^m \approx \sigma_{\text{vis}}^0 (1 - 0.41\alpha \cdot \text{SBIL}_{\text{peak}})}. \quad (\text{D.24})$$

This result is conditioned to the true bunch overlap shape being a factorizable single Gaussian. The effect of this assumption was studied by computing the factor numerically for radially parametrized super- and q-Gaussian true bunch overlap shapes. The calculation probes three α values and uses a maximum head-on SBIL of 8 Hz/ μb . The results are shown in Figure D.1. Both Eq. D.14 and D.24 are verified in the numerical

calculation (see $p = q = 1$). The results obtained using a two-dimensional integral do not depend on the value of α , while the results using one-dimensional fits do due to approximations made in Eq. D.23 losing their validity. The divergence, however, is small with respect to the overall effect using these parameters. With the p and q values obtained in emittance scan fits known, it is possible to estimate an uncertainty on this factor based on these curves.

All results in this section are also conditioned on perfect one- and two-dimensional fits as well. The accuracy of modelling beam overlap shapes distorted by non-linearity is a possible direction for future research.

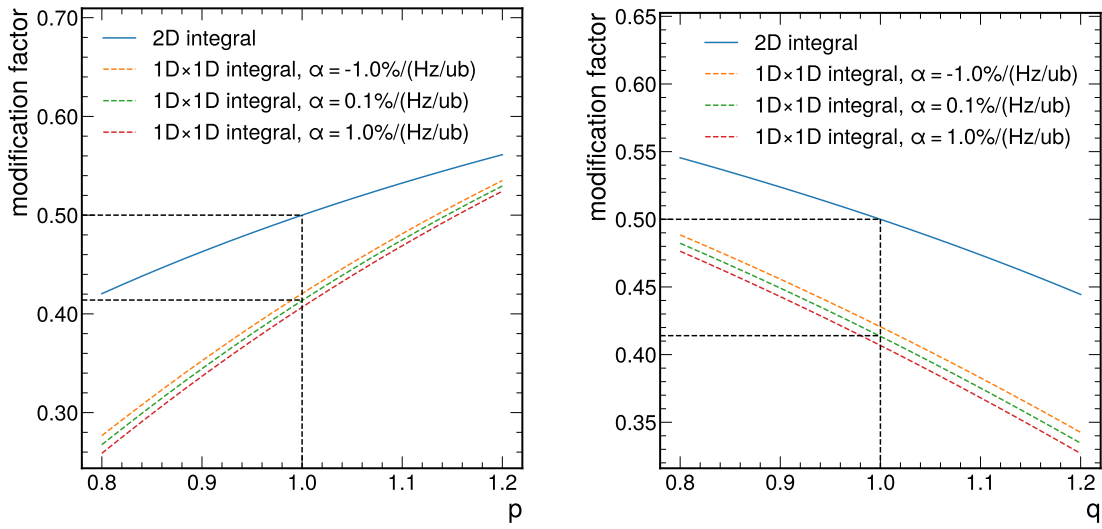


Figure D.1: The modification factor of the linearity derived in the emittance scan analysis for a super-Gaussian (left) and a q -Gaussian (right) true beam overlap shape. The 2D integral method automatically takes the non-linearity-induced non-factorisation effect into account, while the $1D \times 1D$ method, which follows the standard vdM strategy, does not, hence, results in a different factor. The vertical dashed line at $p = q = 1$ indicates the analytically explored case in this section, and the horizontal lines indicate the obtained analytical solutions, respectively. To demonstrate the effect of the approximations, the $1D \times 1D$ method is plotted using three different α values. A good agreement between the analytic and the numerical solution is seen. The simulation was done using a maximal head-on SBIL of $8 \text{ Hz}/\mu\text{b}$.

D.1 Effect on the beam overlap width

In this section only the standard vdM case is considered, i.e., the workflow using two one-dimensional fits. Since using the current approximations

$$2\pi R_m \Sigma_x^m \Sigma_y^m = \sigma_{\text{vis}}^0 (1 - (\sqrt{2} - 1)\alpha \cdot \text{SBIL}_{\text{peak}}), \quad (\text{D.25})$$

$$2\pi R_m \Sigma_x^m \Sigma_y^m = 2\pi R_0 \Sigma_x^0 \Sigma_y^0 (1 - (\sqrt{2} - 1)\alpha \cdot \text{SBIL}_{\text{peak}}), \quad (\text{D.26})$$

and, since $R_m = R_0(1 - \alpha \cdot \text{SBIL}_{\text{peak}})$, we get

$$R_0(1 - \alpha \cdot \text{SBIL}_{\text{peak}}) \Sigma_x^m \Sigma_y^m = R_0 \Sigma_x^0 \Sigma_y^0 (1 - (\sqrt{2} - 1)\alpha \cdot \text{SBIL}_{\text{peak}}), \quad (\text{D.27})$$

$$\Sigma_x^m \Sigma_y^m \approx \Sigma_x^0 \Sigma_y^0 (1 - (\sqrt{2} - 1)\alpha \cdot \text{SBIL}_{\text{peak}})(1 + \alpha \cdot \text{SBIL}_{\text{peak}}), \quad (\text{D.28})$$

$$\Sigma_x^m \Sigma_y^m \approx \Sigma_x^0 \Sigma_y^0 (1 - (\sqrt{2} - 2)\alpha \cdot \text{SBIL}_{\text{peak}}), \quad (\text{D.29})$$

$$\Sigma_x^m \approx \Sigma_x^0 \left(1 - \frac{\sqrt{2} - 2}{2} \alpha \cdot \text{SBIL}_{\text{peak}}\right), \quad (\text{D.30})$$

where $\alpha \cdot \text{SBIL} \ll 1$ was extensively used to omit non-first order terms and the symmetry in the x and y widths was exploited. This relationship is useful when the uncorrected beam overlap width is compared between multiple detectors with varying non-linearities.

Rounding the modification factor to two decimal places, a practical formula is

$$\boxed{\Sigma_x^m \approx \Sigma_x^0 (1 + 0.29\alpha \cdot \text{SBIL}_{\text{peak}})}. \quad (\text{D.31})$$

The validity of this formula is verified numerically for super and q-Gaussian true beam overlap shapes. The calculation probes three α values and uses a maximum head-on SBIL of $8 / (\text{Hz}/\mu\text{b})$. The results are shown in Figure D.2. A reasonable agreement is seen between the $\alpha = 0.1\% / (\text{Hz}/\mu\text{b})$ line and the analytically derived factor at $p = q = 1$ but for larger values of non-linearity the approximations in equations D.28-D.30 lose validity and the results diverge from Eq. D.31.




Figure D.2: The modification factor for the beam overlap shape to be used for non-linear detectors for a super-Gaussian (left) and a q -Gaussian (right) true beam overlap shape. The vertical dashed line at $p = q = 1$ indicates the analytically explored case in this section and the horizontal line indicates the obtained analytical solution. To demonstrate the effect of the approximations made, the $1D \times 1D$ method is plotted using three different α values. A good agreement between the analytic and the numerical solution at the smallest value of α . The simulation was done using a maximal head-on SBIL of $8 \text{ Hz}/\mu\text{b}$.

Bibliography

- [1] Michael E. Peskin. *An Introduction To Quantum Field Theory*. CRC Press, 2018. ISBN: 9780429972102. DOI: 10.1201/9780429503559.
- [2] Sheldon L. Glashow. "Partial-symmetries of weak interactions". In: *Nuclear Physics* 22.4 (1961), pp. 579–588. ISSN: 0029-5582. DOI: 10.1016/0029-5582(61)90469-2.
- [3] Steven Weinberg. "A Model of Leptons". In: *Physical Review Letters* 19.21 (1967), pp. 1264–1266. ISSN: 0031-9007. DOI: 10.1103/physrevlett.19.1264.
- [4] Roman Jackiw. "Physical Formulations: Elementary Particle Theory . Relativistic Groups and Analyticity. Proceedings of the eighth Nobel Symposium, Aspenäs-gården, Lerum, Sweden, May 1968. Nils Svartholm, Ed. Interscience (Wiley), New York, and Almquist and Wiksell, Stockholm, 1969." In: *Science* 168.3936 (1970), pp. 1196–1197. ISSN: 1095-9203. DOI: 10.1126/science.168.3936.1196.b.
- [5] G Zweig. *An SU₃ model for strong interaction symmetry and its breaking*. en. 2002. DOI: 10.17181/CERN-TH-412.
- [6] M. Breidenbach et al. "Observed Behavior of Highly Inelastic Electron-Proton Scattering". In: *Physical Review Letters* 23.16 (1969), pp. 935–939. ISSN: 0031-9007. DOI: 10.1103/physrevlett.23.935.
- [7] G.'t Hooft. "Renormalizable Lagrangians for massive Yang-Mills fields". In: *Nuclear Physics B* 35.1 (1971), pp. 167–188. ISSN: 0550-3213. DOI: 10.1016/0550-3213(71)90139-8.
- [8] F.J. Hasert et al. "Search for Elastic ν_e Electron Scattering". In: *Physics Letters* B46 (1973), pp. 121–124. DOI: 10.1016/0370-2693(73)90494-2.
- [9] F.J. Hasert et al. "Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment". In: *Physics Letters* B46 (1973), pp. 138–140. DOI: 10.1016/0370-2693(73)90499-1.
- [10] G. Amison et al. "Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV". In: *Physics Letters* B122 (1983), pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.

- [11] M. Banner et al. "Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider". In: *Physics Letters B*122 (1983), pp. 476–485. DOI: 10.1016/0370-2693(83)91605-2.
- [12] G. Amison et al. "Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/c² at the CERN SPS Collider". In: *Physics Letters B*126 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.
- [13] P. Bagnaia et al. "Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN anti-p p Collider". In: *Physics Letters B*129 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.
- [14] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Physical Review Letters* 13.9 (1964), pp. 321–323. ISSN: 0031-9007. DOI: 10.1103/physrevlett.13.321.
- [15] Peter W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Physical Review Letters* 13.16 (1964), pp. 508–509. ISSN: 0031-9007. DOI: 10.1103/physrevlett.13.508.
- [16] The CMS Collaboration. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Phys. Lett. B* 716 (2012) 30 716.1 (31, 2012), pp. 30–61. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].
- [17] The ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys.Lett. B*716 (2012) 1-29 716.1 (31, 2012), pp. 1–29. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [18] Vernon Barger, Danny Marfatia, and Kerry Lewis Whisnant, eds. *The physics of neutrinos*. Princeton: Princeton University Press, 2012. 224 pp. ISBN: 1400845599.
- [19] U. Sarkar. *Particle and astroparticle physics*. Series in high energy physics, cosmology, and gravitation. New York: Taylor & Francis, 2008. 524 pp. ISBN: 9781584889328.
- [20] Thomas Mannel. "Theory and Phenomenology of CP Violation". In: *Nuclear Physics B - Proceedings Supplements* 167 (2007), pp. 170–174. ISSN: 0920-5632. DOI: 10.1016/j.nuclphysbps.2006.12.083.
- [21] X. Fan et al. "Measurement of the Electron Magnetic Moment". In: *Physical Review Letters* 130.7 (2023), p. 071801. ISSN: 1079-7114. DOI: 10.1103/physrevlett.130.071801.
- [22] P A Zyla et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8 (2020). ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104.

- [23] Sz. Borsanyi et al. "Leading hadronic contribution to the muon magnetic moment from lattice QCD". In: *Nature* 593.7857 (2021), pp. 51–55. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03418-1.
- [24] Gerard 't Hooft et al. *Recent Developments in Gauge Theories*. Springer US, 1980. ISBN: 9781468475715. DOI: 10.1007/978-1-4684-7571-5.
- [25] A. Zee. *Fearful Symmetry: The Search for Beauty in Modern Physics*. Princeton University Press, 2007. ISBN: 9781400874507. DOI: 10.1515/9781400874507.
- [26] Stephen P. Martin. "A Supersymmetry Primer". In: (16, 1997), pp. 1–98. ISSN: 1793-1339. DOI: 10.1142/9789812839657_0001. arXiv: hep-ph/9709356 [hep-ph].
- [27] Horace W. Babcock. "The rotation of the Andromeda Nebula". In: *Lick Observatory Bulletins* 19 (1939), pp. 41–51. ISSN: 0075-9317. DOI: 10.5479/ads/bib/1939licob.19.41b.
- [28] Katherine Garrett and Gintaras Duda. "Dark Matter: A Primer". In: *Adv.Astron.2011 968283,2011* 2011 (2010), pp. 1–22. ISSN: 1687-7977. DOI: 10.1155/2011/968283. arXiv: 1006.2483 [hep-ph].
- [29] S. Perlmutter et al. "Measurements of Omega and Lambda from 42 High-Redshift Supernovae". In: *Astrophys.J.*517:565-586,1999 517.2 (8, 1998), pp. 565–586. ISSN: 1538-4357. DOI: 10.1086/307221. arXiv: astro-ph/9812133 [astro-ph].
- [30] E. Wigner. "On the Consequences of the Symmetry of the Nuclear Hamiltonian on the Spectroscopy of Nuclei". In: *Physical Review* 51.2 (1937), pp. 106–119. ISSN: 0031-899X. DOI: 10.1103/physrev.51.106.
- [31] Hironari Miyazawa. "Spinor Currents and Symmetries of Baryons and Mesons". In: *Physical Review* 170.5 (1968), pp. 1586–1590. ISSN: 0031-899X. DOI: 10.1103/physrev.170.1586.
- [32] Sidney Coleman and Jeffrey Mandula. "All Possible Symmetries of the SMatrix". In: *Physical Review* 159.5 (1967), pp. 1251–1256. ISSN: 0031-899X. DOI: 10.1103/physrev.159.1251.
- [33] Sergio Ferrara, ed. *Supersymmetry and supergravity '82. Proceedings of the Trieste Sept. 1982 school*. Singapore: World Scientific, 1983. 334 pp. ISBN: 9971950685.
- [34] Yu. A. Gol'fand and E. P. Likhtman. "Extension of the algebra of Poincare group generators and violation of P-invariance". In: *Supergravities in Diverse Dimensions*. World Scientific Publishing Company, 1989, pp. 20–23. ISBN: 9789814542340. DOI: 10.1142/9789814542340_0001.

- [35] D. V. Volkov and V. P. Akulov. “Possible universal neutrino interaction”. In: *Supersymmetry and Quantum Field Theory*. Springer Berlin Heidelberg, pp. 383–385. ISBN: 9783540646235. DOI: 10.1007/bfb0105270.
- [36] P. Fayet. “Spontaneously broken supersymmetric theories of weak, electromagnetic and strong interactions”. In: *Physics Letters B* 69.4 (1977), pp. 489–494. ISSN: 0370-2693. DOI: 10.1016/0370-2693(77)90852-8.
- [37] Luc Pape and Daniel Treille. “Supersymmetry facing experiment: much ado (already) about nothing (yet)”. In: *Reports on Progress in Physics* 69.11 (2006), R01. ISSN: 1361-6633. DOI: 10.1088/0034-4885/69/11/r01.
- [38] G. F. Giudice and R. Rattazzi. “Theories with Gauge-Mediated Supersymmetry Breaking”. In: *Phys.Rept.* 322 (1999) 419-499 322.6 (12, 1998), pp. 419–499. ISSN: 0370-1573. DOI: 10.1016/S0370-1573(99)00042-3. arXiv: hep-ph/9801271 [hep-ph].
- [39] Patrick Meade, Nathan Seiberg, and David Shih. “General Gauge Mediation”. In: *Prog.Theor.Phys.Suppl.* 177:143-158,2009 177 (2008), pp. 143–158. ISSN: 0375-9687. DOI: 10.1143/ptps.177.143. arXiv: 0801.3278 [hep-ph].
- [40] I. -A. Melzer-Pellmann and P. Pralavorio. “Lessons for SUSY from the LHC after the first run”. In: *EPJC* 74 (2014) 2801 74.5 (28, 2014). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-014-2801-y. arXiv: 1404.7191 [hep-ex].
- [41] M. Miura. “Search for Proton Decay via $p \rightarrow e^+\pi^0$ and $p \rightarrow \mu^+\pi^0$ in 0.31 megaton-years exposure of the Super-Kamiokande Water Cherenkov Detector”. In: *Phys. Rev. D* 95, 012004 (2017) 95.1 (12, 2016), p. 012004. ISSN: 2470-0029. DOI: 10.1103/physrevd.95.012004. arXiv: 1610.03597 [hep-ex].
- [42] D. J. Kapner et al. “Tests of the Gravitational Inverse-Square Law below the Dark-Energy Length Scale”. In: *Phys.Rev.Lett.* 98:021101,2007 98.2 (14, 2006), p. 021101. ISSN: 1079-7114. DOI: 10.1103/physrevlett.98.021101. arXiv: hep-ph/0611184 [hep-ph].
- [43] *Extra dimensions, gravitons, and tiny black holes*. URL: <https://home.cern/about/physics/extra-dimensions-gravitons-and-tiny-black-holes> (visited on 06/22/2024).
- [44] Anthony Zee. *Quantum field theory in a nutshell*. Second edition. Princeton: Princeton University Press, 2010. 1576 pp. ISBN: 9781400835324.
- [45] V.A. Rubakov and M.E. Shaposhnikov. “Do we live inside a domain wall?” In: *Physics Letters B* 125.2–3 (1983), pp. 136–138. ISSN: 0370-2693. DOI: 10.1016/0370-2693(83)91253-4.

- [46] Krzysztof A. Meissner and Hermann Nicolai. “Conformal Symmetry and the Standard Model”. In: *Phys.Lett.B648:312-317,2007* 648.4 (15, 2006), pp. 312–317. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2007.03.023. arXiv: hep-th/0612165 [hep-th].
- [47] Erick J. Weinberg. “Radiative Corrections as the Origin of Spontaneous Symmetry Breaking”. In: (2005). DOI: 10.48550/ARXIV.HEP-TH/0507214. arXiv: hep-th/0507214 [hep-th].
- [48] Christopher T. Hill and Elizabeth H. Simmons. “Strong Dynamics and Electroweak Symmetry Breaking”. In: *Phys.Rept.381:235-402,2003; Erratum-ibid.390:553-554,2004* 381.4–6 (7, 2002), pp. 235–402. ISSN: 0370-1573. DOI: 10.1016/s0370-1573(03)00140-6. arXiv: hep-ph/0203079 [hep-ph].
- [49] M. Milgrom. “A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis”. In: *The Astrophysical Journal* 270 (1983), p. 365. ISSN: 1538-4357. DOI: 10.1086/161130.
- [50] Indranil Banik and Hongsheng Zhao. “From galactic bars to the Hubble tension: weighing up the astrophysical evidence for Milgromian gravity”. In: *Symmetry*, 14, 1331 (2022) 14.7 (13, 2021), p. 1331. ISSN: 2073-8994. DOI: 10.3390/sym14071331. arXiv: 2110.06936 [astro-ph.CO].
- [51] LIGO Scientific Collaboration et al. “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *ApJL*, 848:L12, 2017 848.2 (16, 2017), p. L12. ISSN: 2041-8213. DOI: 10.3847/2041-8213/aa91c9. arXiv: 1710.05833 [astro-ph.HE].
- [52] Tobias Mistele et al. “Indefinitely Flat Circular Velocities and the Baryonic Tully-Fisher Relation from Weak Lensing”. In: *ApJL* 969 L3 (2024) 969.1 (14, 2024), p. L3. ISSN: 2041-8213. DOI: 10.3847/2041-8213/ad54b0. arXiv: 2406.09685 [astro-ph.GA].
- [53] Indranil Banik et al. “Strong constraints on the gravitational law from Gaia DR3 wide binaries”. In: *Monthly Notices of the Royal Astronomical Society* 527.3 (2023), pp. 4573–4615. ISSN: 1365-2966. DOI: 10.1093/mnras/stad3393.
- [54] Peter R. Hobson. “The Large Hadron Collider: A Marvel of Technology, edited by Lyndon Evans: Scope: edited book. Level: general reader and non-specialists”. In: *Contemporary Physics* 52.2 (2011), pp. 172–173. ISSN: 1366-5812. DOI: 10.1080/00107514.2010.540674.
- [55] L. Evans. “The Large Hadron Collider, a personal recollection”. In: *Progress of Theoretical and Experimental Physics* 2014.2 (2014), 2A102–. ISSN: 2050-3911. DOI: 10.1093/ptep/ptt117.

- [56] Tai Sakuma and Thomas McCauley. “Detector and Event Visualization with SketchUp at the CMS Experiment”. In: *J. Phys. Conf. Ser.* 513 (2014). Ed. by D. L. Groep and D. Bonacorsi, p. 022032. DOI: 10.1088/1742-6596/513/2/022032. arXiv: 1311.4942 [physics.ins-det].
- [57] The CMS Collaboration. “Particle-flow reconstruction and global event description with the CMS detector. Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003. arXiv: 1706.04965. URL: <https://cds.cern.ch/record/2270046>.
- [58] The CMS Collaboration. “Uniformity and stability of the CMS RPC detector at the LHC”. In: *Journal of Instrumentation* 8.03 (2013), P03017–P03017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/8/03/p03017.
- [59] The CMS Collaboration. “Commissioning of the CMS Magnet”. In: *IEEE Transactions on Applied Superconductivity* 17.2 (2007), pp. 1185–1190. ISSN: 1051-8223. DOI: 10.1109/tasc.2007.897754.
- [60] The CMS Collaboration. *The CMS tracker system project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. ISBN: 9789290831242. URL: <https://cds.cern.ch/record/368412>.
- [61] W. Adam et al. “The CMS Phase-1 pixel detector upgrade”. In: *Journal of Instrumentation* 16.02 (2021), P02027–P02027. ISSN: 1748-0221. DOI: 10.1088/1748-0221/16/02/p02027.
- [62] The CMS Collaboration. *The CMS electromagnetic calorimeter project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. ISBN: 92908311221. URL: <https://cds.cern.ch/record/349375>.
- [63] The CMS Collaboration. “Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV”. In: *JINST* 8 (2013), P09009. DOI: 10.1088/1748-0221/8/09/P09009. arXiv: 1306.2016. URL: <https://cds.cern.ch/record/1554142>.
- [64] The CMS Collaboration. *The CMS hadron calorimeter project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. ISBN: 9789290831099. URL: <https://cds.cern.ch/record/357153>.
- [65] The CMS Collaboration. *The CMS muon project: Technical Design Report*. Technical design report. CMS. Geneva: CERN, 1997. ISBN: 9789290831211. URL: <https://cds.cern.ch/record/343814>.

- [66] CMS Collaboration. *CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems*. Technical design report. CMS. ISBN: 9290831102. URL: <https://cds.cern.ch/record/706847>.
- [67] A Tapper and Darin Acosta. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Tech. rep. CERN-LHCC-2013-011, CMS-TDR-12. 2013. URL: <https://cds.cern.ch/record/1556311>.
- [68] The CMS Collaboration. *Commissioning of the Particle-Flow reconstruction in Minimum-Bias and Jet Events from pp Collisions at 7 TeV*. Tech. rep. CMS-PAS-PFT-10-002. Geneva: CERN, 2010. URL: <https://cds.cern.ch/record/1279341>.
- [69] The CMS Collaboration. “The CMS Particle Flow Algorithm”. In: *EPJ Web Conf.* 191 (2018), p. 02016. DOI: 10.1051/epjconf/201819102016. URL: <https://cds.cern.ch/record/2678077>.
- [70] The CMS Collaboration. *Performance of reconstruction and identification of tau leptons in their decays to hadrons and tau neutrino in LHC Run-2*. Tech. rep. CMS-PAS-TAU-16-002. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2196972>.
- [71] The CMS Collaboration. *In preparation: Precision luminosity measurement at CMS in proton-proton collisions at 13 TeV in 2017 and 2018*. Tech. rep. CMS-LUM-20-001. Geneva: CERN, 2024.
- [72] Peter Skands, Stefano Carrazza, and Juan Rojo. “Tuning PYTHIA 8.1: the Monash 2013 tune”. In: *Eur. Phys. J. C* 74 (2014), p. 3024. DOI: 10.1140/epjc/s10052-014-3024-y. arXiv: 1404.5630 [hep-ph].
- [73] M. Cepeda et al. “Report from Working Group 2: Higgs Physics at the HL-LHC and HE-LHC”. In: *CERN Yellow Rep. Monogr.* 7 (2019). Ed. by Andrea Dainese et al., pp. 221–584. DOI: 10.23731/CYRM-2019-007.221. arXiv: 1902.00134 [hep-ph].
- [74] The CMS Collaboration. *The Phase-2 Upgrade of the CMS Beam Radiation Instrumentation and Luminosity Detectors*. Tech. rep. CERN-LHCC-2021-008, CMS-TDR-023. Geneva: CERN, 2021. URL: <https://cds.cern.ch/record/2759074>.
- [75] Gabriella Pásztor. *Exploring the secrets of elementary particles at the CERN LEP and LHC colliders*. MTA Doctoral Dissertation. Budapest: MTA, 2023.
- [76] Albert M Sirunyan et al. “Measurements of the W boson rapidity, helicity, double-differential cross sections, and charge asymmetry in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 102 (2020), p. 092012. DOI: 10.1103/PhysRevD.102.092012. arXiv: 2008.04174 [hep-ex].

- [77] Armen Tumasyan et al. “Measurement of the mass dependence of the transverse momentum of lepton pairs in Drell-Yan production in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 83.7 (2023), p. 628. DOI: 10.1140/epjc/s10052-023-11631-7. arXiv: 2205.04897 [hep-ex].
- [78] Albert M Sirunyan et al. “Combined measurements of Higgs boson couplings in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 79.5 (2019), p. 421. DOI: 10.1140/epjc/s10052-019-6909-y. arXiv: 1809.10733 [hep-ex].
- [79] Armen Tumasyan et al. “Measurement of differential $t\bar{t}$ production cross sections in the full kinematic range using lepton+jets events from proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 104.9 (2021), p. 092013. DOI: 10.1103/PhysRevD.104.092013. arXiv: 2108.02803 [hep-ex].
- [80] The CMS Collaboration. “Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS”. In: *Eur. Phys. J. C* 81.9 (2021), p. 800. DOI: 10.1140/epjc/s10052-021-09538-2. arXiv: 2104.01927. URL: <https://cds.cern.ch/record/2759951>.
- [81] The CMS Collaboration. *CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. CMS-PAS-LUM-17-004. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2621960>.
- [82] The CMS Collaboration. *CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. CMS-PAS-LUM-18-002. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2676164>.
- [83] CMS BRIL Collaboration. “The Pixel Luminosity Telescope: A detector for luminosity measurement at CMS using silicon pixel sensors”. In: *Eur. Phys. J. C* 83 (2023) 673 83.7 (17, 2022). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-11713-6. arXiv: 2206.08870 [physics.ins-det].
- [84] Fejes Milan. “Kísérleti adatminőség meghatározása az LHC CMS kísérletében gépi tanulási módszerrel”. Eötvös Loránd University, Faculty of Sciences, 2020.
- [85] Fejes Milan. “Kísérleti adatminőség meghatározása az LHC CMS kísérletében gépi tanulási módszerrel”. Eötvös Loránd University, Faculty of Informatics, 2022.
- [86] J. MacQueen. “Some methods for classification and analysis of multivariate observations”. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*. 1967, p. 281.
- [87] S van der Meer. *Calibration of the effective beam height in the ISR*. Tech. rep. CERN-ISR-PO-68-31, ISR-PO-68-31. Geneva: CERN, 1968. URL: <https://cds.cern.ch/record/296752>.

- [88] C Barschel et al. *Results of the LHC DCCT Calibration Studies*. CERN-ATS-Note-2012-026 PERF. CERN, 2012. URL: <https://cds.cern.ch/record/1425904>.
- [89] David Belohrad et al. “Upgrade of the LHC Bunch by Bunch Intensity Measurement Acquisition System”. In: *Proceedings, 5th International Beam Instrumentation Conference (IBIC 2016), Barcelona, Spain, September 11–15, 2016*. 2017, MOPG39. DOI: 10.18429/JACoW-IBIC2016-MOPG39.
- [90] Michal Krupa and Marek Gasior. “The Wall Current Transformer - a New Sensor for Precise Bunch-by-Bunch Intensity Measurements in the LHC”. In: *Proceedings, 5th International Beam Instrumentation Conference (IBIC 2016), Barcelona, Spain, September 11–15, 2016*. 2017, WEAL02. DOI: 10.18429/JACoW-IBIC2016-WEAL02.
- [91] A. Jeff et al. “Longitudinal density monitor for the LHC”. In: *Phys. Rev. ST Accel. Beams* 15 (2012), p. 032803. DOI: 10.1103/PhysRevSTAB.15.032803.
- [92] “Measurement of the offline integrated luminosity for the CMS proton-proton collision dataset recorded in 2023”. In: (2024). URL: <https://cds.cern.ch/record/2904808>.
- [93] M Gasior et al. “First operational experience with the LHC Diode ORbit and OScillation (DOROS) system”. In: *Proceedings of the 5th International Beam Instrumentation Conference (IBIC 2016)*. 2017, 43 (MOPG07). DOI: 10.18429/JACoW-IBIC2016-MOPG07.
- [94] M Gasior, J Olexa, and RJ Steinhagen. “BPM electronics based on compensated diode detectors—results from development systems”. In: *Proceedings of the 2012 Beam Instrumentation Workshop (BIW’12)*. 2012, 44 (MOPG010). URL: <https://cds.cern.ch/record/1476070>.
- [95] W Kozanecki, T Pieloni, and J Wenninger. *Observation of beam-beam deflections with LHC orbit data*. CERN Accelerator Note CERN-ACC-NOTE-2013-0006. CERN, 2013. URL: <https://cds.cern.ch/record/1581723>.
- [96] A. Babaev et al. “Impact of beam–beam effects on absolute luminosity calibrations at the CERN Large Hadron Collider”. In: *The European Physical Journal C* 84.1 (2024). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-12192-5.
- [97] The CMS Collaboration. *CMS Luminosity Based on Pixel Cluster Counting - Summer 2013 Update*. Tech. rep. CMS-PAS-LUM-13-001. Geneva: CERN, 2013. URL: <https://cds.cern.ch/record/1598864>.
- [98] The CMS Collaboration. *CMS luminosity measurement for the 2015 data-taking period*. Tech. rep. CMS-PAS-LUM-15-001. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2138682>.

- [99] The CMS Collaboration. *CMS Luminosity Measurements for the 2016 Data Taking Period*. Tech. rep. CMS-PAS-LUM-17-001. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2257069>.
- [100] Markus Klute, Catherine Medlock, and Jakob Salfeld-Nebgen. “Beam Imaging and Luminosity Calibration”. In: *Journal of Instrumentation* 12.03 (11, 2016), P03018–P03018. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/03/p03018. arXiv: 1603.03566 [hep-ex].
- [101] Joscha Knolle. “Measuring luminosity and the $t\bar{t}Z$ production cross section with the CMS experiment”. Hamburg U., 2020. URL: <https://cds.cern.ch/record/2744549>.
- [102] The CMS Collaboration. *Offline luminosity measurement for the 2022 pp collisions at 13.6 TeV data set at CMS*. CMS Physics Analysis Summary CMS-PAS-LUM-22-001. 2024. URL: <https://cds.cern.ch/record/2890833>.
- [103] *CMS luminosity measurement using nucleus-nucleus collisions at $\sqrt{s_{NN}} = 5.02$ TeV in 2018*. Tech. rep. CMS-PAS-LUM-18-001. Geneva: CERN, 2022. URL: <https://cds.cern.ch/record/2809613>.
- [104] “Nonfactorization in Van der Meer scans in Run 2”. In: (2019). URL: <https://cds.cern.ch/record/2681801>.
- [105] Andrew Kuhl. “Search for supersymmetry using a photon, b-jets, and E_T^{miss} final state with the ATLAS detector in proton-proton collisions at 8 TeV center-of-mass energy”. PhD thesis. UC Santa Cruz, 2015.
- [106] Mateusz Dyndal. *Non factorisation in ATLAS and ALICE vdM scans: experience in Run 2, impact on luminosity calibration*. Tech. rep. ATL-DAPR-PROC-2019-002. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2701733>.
- [107] Nóra Rab. “Beyond the standard model at the LHC”. MA thesis. Eötvös Loránd University, Faculty of Sciences, 2024.
- [108] G Papotti et al. “Longitudinal Beam measurements at the LHC: The LHC Beam Quality Monitor”. In: CERN-ATS-2011-220 (2011). URL: <https://cds.cern.ch/record/1399087>.
- [109] José C. Pinheiro and Douglas M. Bates. “Unconstrained parametrizations for variance-covariance matrices”. In: *Statistics and Computing* 6.3 (1996), pp. 289–296. ISSN: 1573-1375. DOI: 10.1007/BF00140873. URL: <https://doi.org/10.1007/BF00140873>.

- [110] F James and M Roos. “MINUIT: a system for function minimization and analysis of the parameter errors and corrections”. In: *Comput. Phys. Commun.* 10.6 (1975), pp. 343–67. DOI: 10.1016/0010-4655(75)90039-9. URL: <https://cds.cern.ch/record/310399>.
- [111] CMS Collaboration. *Stairway to discovery: a report on the CMS programme of cross section measurements from millibarns to femtobarns*. 2024. arXiv: 2405.18661 [hep-ex]. URL: <https://arxiv.org/abs/2405.18661>.
- [112] The CMS Collaboration. *CMS luminosity measurement for nucleus-nucleus collisions at $\sqrt{s_{NN}} = 5.02$ TeV in Run 2*. CMS Physics Analysis Summary CMS-PAS-LUM-20-002. 2024. URL: <https://cds.cern.ch/record/2905093>.
- [113] The CMS Collaboration. *CMS luminosity measurement using 2016 proton-nucleus collisions at nucleon-nucleon center-of-mass energy of 8.16 TeV*. Tech. rep. CMS-PAS-LUM-17-002. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2628652>.
- [114] ATLAS Collaboration. “Search for photonic signatures of gauge-mediated supersymmetry in 8 TeV pp collisions with the ATLAS detector”. In: *Phys. Rev. D* 92, 072001 (2015) (2015). DOI: 10.1103/PhysRevD.92.072001. arXiv: 1507.05493 [hep-ex].
- [115] ATLAS Collaboration. “Search for photonic signatures of gauge-mediated supersymmetry in 13 TeV pp collisions with the ATLAS detector”. In: *Phys. Rev. D* 97, 092006 (2018) 97.9 (9, 2018), p. 092006. ISSN: 2470-0029. DOI: 10.1103/physrevd.97.092006. arXiv: 1802.03158 [hep-ex].
- [116] The CMS Collaboration. “Search for new physics in multijet events with at least one photon and large missing transverse momentum in proton-proton collisions at 13 TeV”. In: *JHEP* 2310 (2023), p. 046. DOI: 10.1007/JHEP10(2023)046. arXiv: 2307.16216. URL: <https://cds.cern.ch/record/2866291>.
- [117] Andy Buckley et al. “General-purpose event generators for LHC physics”. In: *Phys. Rept.* 504 (2011), pp. 145–233. DOI: 10.1016/j.physrep.2011.03.005. arXiv: 1101.2599 [hep-ph].
- [118] Simone Alioli et al. *Monte Carlo event generators for high energy particle physics event simulation*. MCnet-19-02. 2019. arXiv: 1902.01674 [hep-ph].
- [119] J. M. Campbell et al. “Event Generators for High-Energy Physics Experiments”. In: *Snowmass 2021*. CP3-22-12, DESY-22-042, FERMILAB-PUB-22-116-SCD-T, IPPP/21/51, JLAB-PHY-22-3576, KA-TP-04-2022, LA-UR-22-22126, LU-TP-22-12, MCNET-22-04, OUTF-22-03P, P3H-22-024, PITT-PACC 2207, UCI-TR-2022-02. 2022. arXiv: 2203.11110 [hep-ph].

- [120] S. Agostinelli et al. “GEANT4 — a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), p. 250. DOI: 10.1016/S0168-9002(03)01368-8.
- [121] LHC SUSY Cross Section Working Group. *SUSY cross sections*. 2024. URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections> (visited on 08/28/2024).
- [122] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- k_t jet clustering algorithm”. In: *JHEP 0804:063,2008* 2008.04 (8, 2008), pp. 063–063. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].
- [123] Emil Bols et al. “Jet Flavour Classification Using DeepJet”. In: *JINST 15 (2020) P12012* 15.12 (24, 2020), P12012–P12012. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/12/p12012. arXiv: 2008.10519 [hep-ex].
- [124] Huilin Qu and Loukas Gouskos. “ParticleNet: Jet Tagging via Particle Clouds”. In: *Phys. Rev. D 101, 056019 (2020)* 101.5 (22, 2019), p. 056019. ISSN: 2470-0029. DOI: 10.1103/physrevd.101.056019. arXiv: 1902.08570 [hep-ph].
- [125] A L Read. “Modified frequentist analysis of search results (the CL_s method)”. In: CERN-OPEN-2000-205 (2000). DOI: 10.5170/CERN-2000-005.81. URL: <https://cds.cern.ch/record/451614>.
- [126] The CMS Collaboration. *The CMS statistical analysis and combination tool: Combine*. Tech. rep. CMS-CAT-23-001, CERN-EP-2024-078, CMS-CAT-23-001-003. Geneva: CERN, 2024. DOI: 10.48550/arXiv.2404.06614. arXiv: 2404.06614. URL: <https://cds.cern.ch/record/2895097>.
- [127] Peter Major. “Precision luminosity measurement of proton-proton collisions at the CMS experiment in Run 2”. In: *Proceedings of European Physical Society Conference on High Energy Physics — PoS(EPS-HEP2019)*. Vol. 364. 2020, p. 198. DOI: 10.22323/1.364.0198.
- [128] Peter Major. “Studies of the factorisation of proton densities in van der Meer scans and its impact on precision luminosity measurements for CMS”. In: *Proceedings of European Physical Society Conference on High Energy Physics — PoS(EPS-HEP2019)*. Vol. 364. 2020, p. 199. DOI: 10.22323/1.364.0199.
- [129] A. Fehérkuti, P. Major, G. Pásztor. “XY factorization bias in luminosity measurements”. In: *Submitted to Int. J. of Modern Phys. A*. CMS-CR-24-093 (2024).
- [130] The CMS Collaboration. *Measurement of the offline integrated luminosity for the CMS proton-proton collision dataset recorded in 2023*. Tech. rep. CERN, 2024. URL: <https://cds.cern.ch/record/2904808>.

- [131] Aram Hayrapetyan et al. “Measurement of the inclusive cross sections for W and Z boson production in proton-proton collisions at $\sqrt{s} = 5.02$ and 13 TeV”. In: CMS-SMP-20-004, CERN-EP-2024-134 (2024). DOI: 10.48550/arXiv.2408.03744. arXiv: 2408.03744 [hep-ex].
- [132] Aram Hayrapetyan et al. “Stairway to discovery: a report on the CMS programme of cross section measurements from millibarns to femtobarns”. In: CMS-SMP-23-004, CERN-EP-2024-117 (2024). DOI: 10.48550/arXiv.2405.18661. arXiv: 2405.18661 [hep-ex].
- [133] Peter Major. “Search for supersymmetry in final states with photons and missing transverse momentum in pp collisions at 13 TeV using the CMS detector”. In: *Proceedings of XXVII International Workshop on Deep-Inelastic Scattering and Related Subjects — PoS(DIS2019)*. Vol. 352. 2019, p. 098. DOI: 10.22323/1.352.0098.
- [134] Debarghya Ghoshdastidar, Ambedkar Dukkipati, and Shalabh Bhatnagar. “Newton based Stochastic Optimization using q-Gaussian Smoothed Functional Algorithms”. In: *Automatica* 50.10 (10, 2013), pp. 2606–2614. DOI: 10.1016/j.automatica.2014.08.021. arXiv: 1311.2296 [math.OA].