Sapienza
Università di Roma

Indirect on-shell constraint on the Higgs boson width via signal–background interference in the diphoton channel at CMS

Scuola di Dottorato in Scienze Astronomiche, Chimiche, Fisiche e Matematiche 'Vito Volterra'

Dottorato di Ricerca in Fisica (XXXVIII cycle)

Candidate

Ruben Gargiulo
ID number 1824800

Thesis Advisors

Dr. E. Di Marco, Dr. F. Errico

Indirect on-shell constraint on the Higgs boson width via signal–background interference in the diphoton channel at CMS
Ph.D. thesis. Sapienza – University of Rome

This thesis has been typeset by LaTeX and the Sapthesis class.

Author's email: ruben.gargiulo@uniroma1.it

Abstract

The Standard Model Higgs boson with a mass of 125 GeV is predicted to have a decay width of 4.1 MeV. Direct measurements of the Higgs width using on-shell Higgs boson production are limited by the experimental resolution, which is about 1 GeV in the diphoton and four-lepton final states. This thesis presents a constraint on the Higgs boson width obtained from the diphoton invariant mass distribution in on-shell Higgs decays, using the interference between the amplitudes of the gluon-gluon fusion to Higgs resonance, with subsequent diphoton decay, and the continuum QCD gluon-gluon to diphoton process. The study was performed using proton-proton collision data at a center-of-mass energy of 13 TeV, collected by the CMS experiment during LHC Run 2, corresponding to an integrated luminosity of 137 inverse femtobarns. The observed (expected) upper limit on the Higgs boson width is 92 MeV (138 MeV) at the 95% confidence level.

Contents

Introduction

The most important goal of experimental particle physics is to uncover how matter and the fundamental forces behave at the smallest and most energetic scales of nature. Within this field, the Standard Model (SM) of particle physics provides a theoretical description of all known elementary particles and of three of the four fundamental interactions: electromagnetic, weak, and strong. Over the past several decades, the SM has undergone extensive and precise experimental validation. Despite its success, however, both experimental evidence and theoretical considerations point to the fact that the SM is an incomplete picture of the physical world. It fails, for instance, to account for the observed asymmetry between matter and antimatter, the existence and nature of dark matter, the origin of neutrino masses and fermion generations, and the emergence of its several free parameters.

One of the pillars of modern particle physics is the Large Hadron Collider (LHC) at CERN (European Organization for Nuclear Research), which is currently the most powerful and complex particle accelerator ever built. Among the LHC's four main experiments, the Compact Muon Solenoid (CMS) detector serves as a general-purpose apparatus designed to reconstruct and study the particles emerging from high-energy proton–proton collisions. The CMS physics program focuses on two main goals: the precision study of the Standard Model, particularly the Higgs sector, and the exploration of possible phenomena that could reveal physics beyond the SM. The Higgs boson, being linked to several of the SM's fundamental parameters, represents an especially sensitive probe for such new effects.

Since the discovery of the Higgs boson in 2012, the CMS collaborations have performed increasingly precise measurements of its properties to assess their consistency with the SM predictions. Because of its numerous couplings to Standard Model particles, the Higgs boson provides an ideal testing ground for possible deviations that might indicate new physics.

Over the first two data-taking periods of the LHC, between 2011 and 2018, several of the Higgs boson's properties have been accurately determined using CMS data. In particular, its mass has been precisely measured in the two original discovery channels, $H \to ZZ^*$ and $H \to \gamma\gamma$. Nonetheless, one of the remaining open experimental challenges is the determination of the Higgs boson's total decay width, which is directly related to its lifetime. The Standard Model predicts this quantity to be approximately 4.1 MeV; any measurable deviation from this value could hint at physics beyond the Standard Model, either through altered Higgs couplings to known particles or via interactions with undiscovered ones.

The present thesis focuses on the measurement of the Higgs boson's total decay width in the $H \to \gamma\gamma$ decay channel. The analysis exploits a novel technique based on quantum interferometry, an approach first discussed in theoretical works published around 2012. Although the diphoton final state provides a clean and well-reconstructed signature, the measurement remains challenging because the intrinsic Higgs width is roughly three orders of magnitude smaller than the detector's mass resolution, which is about 1.5 GeV. The method relies on the quantum interference between the resonant Higgs production and the non-resonant diphoton continuum background. When two processes share the same initial and final states, interference occurs and induces small distortions in the diphoton invariant-mass spectrum. The magnitude of these distortions carries direct information about the total width of the Higgs boson.

The analysis presented here is based on the complete Run 2 dataset collected by the CMS detector between 2016 and 2018, corresponding to an integrated luminosity of about 137 fb^{-1} at a centre-of-mass energy of 13 TeV. The following sections outline the overall structure of this thesis and summarise my personal contributions to the research, as a member of the CMS collaboration.

Thesis structure

The thesis is divided into two parts. Part I describes the broader scientific and experimental context in which this research was conducted and introduces the experimental tools and methodologies used in the analysis. Part II focuses on my original work carried out for the measurement of the Higgs boson decay width in the diphoton channel. The datasets used for this analysis were collected during Run 2 (2016–2018), and most experimental inputs - such as efficiencies and corrections - were already established prior to my contribution.

- **Chapter 1** introduces the Higgs boson within the Standard Model and provides an overview of the current experimental status of Higgs physics;
- **Chapter 2** discusses theoretical aspects related to the measurement of Higgs couplings and decay width, summarizing the main available techniques along with their advantages and limitations;
- **Chapter 3** presents the LHC accelerator complex and the CMS detector, with particular emphasis on the electromagnetic calorimeter;
- **Chapter 4** describes the reconstruction, identification, and energy calibration of photons and diphotons in CMS, which are crucial steps for the $H \to \gamma\gamma$ analysis;
- **Chapter 5** outlines the analysis strategy developed to measure the Higgs boson width in the diphoton decay channel;
- **Chapter 6** details the fitting procedures and modelling techniques applied to extract the Higgs boson decay width;
- **Chapter 7** presents the final results and discusses the constraints obtained on the Higgs boson total width.

Within the CMS collaboration, this analysis was conducted by me and another analyzer, each focusing on different approaches using a shared event categorization. The work described in this thesis covers the most sensitive of the two methods, for which I was the only one involved, except for the part related to event categorization. A preliminary version of this analysis has been published as a CMS Physics Analysis Summary [1] and presented at two international conferences, as of November 2025. A journal paper summarizing the technique and the results is currently in preparation and will represent the first published measurement using this interferometry-based technique by either the ATLAS or CMS collaborations.

Part I

Higgs boson physics and CMS experiment

Chapter 1

Standard Model and Higgs boson

The Standard Model (SM) of particle physics provides the foundational theoretical framework describing all known elementary particles and their interactions. This chapter offers an overview of the Standard Model, briefly summarizes the current state of Higgs boson measurements, and sets the stage for the analysis presented in this thesis. Section 1.1 outlines the SM's particle content and interactions, including the mechanism responsible for mass generation. Section 1.2 provides an overview of Higgs boson phenomenology and the latest experimental results related to its properties.

1.1 Standard Model framework

The Standard Model of particle physics is the theoretical framework that describes three of the four known fundamental forces in nature—electromagnetic, weak, and strong interactions—excluding gravity. It also provides a classification scheme for all known elementary particles. Developed over the second half of the 20th century through the contributions of scientists worldwide, the Standard Model reached its current form in the mid-1970s, following experimental validation of the existence of quarks. Since then, further discoveries—such as the top quark in 1995, the tau neutrino in 2000, and the Higgs boson in 2012—have provided strong confirmation of the model's predictive power. Additionally, it has accurately predicted properties of weak neutral currents as well as the masses and interactions of the W and Z bosons.

As a paradigm of quantum field theory (QFT), it showcases phenomena such as spontaneous symmetry breaking, gauge anomalies, and non-perturbative effects. While it serves as a robust foundation, it also provides a launching point for more speculative theories that seek to address its limitations - such as models including supersymmetry, extra dimensions, and new particles that could explain dark matter and neutrino mass.

Despite its success and internal theoretical consistency, the Standard Model is not a complete theory of fundamental interactions. It fails to explain several key phenomena, including the observed matter–antimatter asymmetry, and it does not incorporate gravity as described by general relativity. One of the approaches used to search for SM extensions is the investigation of Higgs boson properties, looking for possible deviation from theory predictions.

1.1.1 Symmetry groups and particle content

The Standard Model is a relativistic Lorentz-invariant quantum field theory based on gauge invariance of the Lagrangian density under specific Lie groups. In the framework of SM, each physical particle is described as an excitation of its quantum field and possesses specific of quantum numbers in certain representations of the symmetry groups, that dictate the relations and the dynamical properties of the particles.

The SM symmetry groups are: $SU(3)_C$ for quantum chromodynamics (QCD), and $SU(2)_I \times U(1)_{Y_W}$ for the electroweak interaction, which unifies electromagnetism and the weak force. Here, the indices represent the charge conserved by each group: C represents the colour charge associated with SU(3), I denotes the weak isospin, and Y_W denotes the weak hypercharge, as it will be detailed in the next paragraphs.

Fundamental particles are classified as fermions, which have semi-integer spin and are the components of matter (electrons, quarks that constitute nucleons, and neutrinos), and bosons, which have integer spin and mediate the interactions between fermions. In the SM, matter is only consisting of fermions, because the only stable boson in the physical realization of the SM is the photon, which is massless.

The interactions between fermions are mediated by bosons, in the sense that the Lagrangian contains terms coupling fermions to bosons and no fermion-fermion coupling is present, so fermions are only coupled indirectly (see next paragraphs), with different coupling constants for the three groups.

Fermions are divided in two types: quark and leptons. Quarks are in the triplet representation of $SU(3)_C$, i.e. they possess a definite colour charge (being autostates of colour) with three values, while leptons are in the singlet representation of $SU(3)_C$, i.e. they do not possess any colour charge. Fermions are in the doublet representation of $SU(2)_I$, i.e. they are autostates of the weak isospin T_3, and consequently both quarks and leptons are divided in other two categories, i.e. up/down-like quarks and neutrinos/charged leptons, depending on their weak isospin value. Finally, for each fermion with certain quantum numbers, three version of it exist, called families. Lepton families are electron, muon and tau, with the corresponding neutrinos, while quarks are divided as up/down, charm/strange, top/bottom. Quark types are also called flavours. Note that the fermion quantum numbers for the two electroweak groups depend on the chirality state (left or right) of the particle, as detailed in Table 1.1, therefore symmetries are automatically not conserved if fermions are massive, because massive particles do not have definite chirality states. On the other hand, the electric charge, defined as $Q = Y_W + T_3$, is the same for the two chiral states.

Type	**Left-chiral fermions**				**Right-chiral fermions**			
		Q	T_3	Y_W		Q	T_3	Y_W
Leptons	ν_e, ν_μ, ν_τ	0	$+\frac{1}{2}$	$-\frac{1}{2}$	ν_R	0	0	0
	e^-, μ^-, τ^-	-1	$-\frac{1}{2}$	$-\frac{1}{2}$	e^-_R, μ^-_R, τ^-_R	-1	0	-1
Quarks	u, c, t	$+\frac{2}{3}$	$+\frac{1}{2}$	$+\frac{1}{6}$	u_R, c_R, t_R	$+\frac{2}{3}$	0	$+\frac{2}{3}$
	d, s, b	$-\frac{1}{3}$	$-\frac{1}{2}$	$+\frac{1}{6}$	d_R, s_R, b_R	$-\frac{1}{3}$	0	$-\frac{1}{3}$

Table 1.1: Quantum numbers of left- and right-chiral fermions under electroweak groups.

There are two configuration of the SM, before and after spontaneous electroweak symmetry breaking (SSB). Before the SSB, all particles are massless, except for the Higgs fields (with imaginary masses), all three gauge symmetries (C, W, Y_W) are valid, and electroweak quantum numbers are conserved.

The vector bosons in the SM before SSB are all massless, and are divided into:

- gluons: the carriers of $SU(3)_C$ (in its octet representation), invariant under electroweak groups and not participating in the symmetry breaking

- $W^{1,2,3}$: the carriers of $SU(2)_W$ in their triplet representation (representing the generators of the group), invariant under color and weak hypercharge. They are not autostates of weak isospin and thus do not have definite electric charge
- B^0: the neutral carrier of $U(1)_{Y_W}$ in its adjoint representation, and invariant under other groups

Other than fermions and vector bosons (force carriers), the SM before SSB contains the Higgs fields, a scalar weak isospin doublet (H^*, H), with $T^3 = \pm 1/2$ and $Y_W = +1/2$. Due to a complex field potential the masses of the two components are imaginary, and the corresponding particles are tachyonic, which physically means that the potential is unstable when the field is null.

It is important to note that before SSB the symmetries are allowed to be conserved because all fields are massless. Already with non-null fermion masses, the chiral states with different quantum numbers would mix, and therefore weak isospin and hypercharge would not be conserved.

After SSB, which is made possible via the Higgs mechanism (see next paragraphs), the fermions become massive, and the boson structure is rearranged, thus generating:

- $W^{\pm} = 1/\sqrt{2}(W^1 \mp iW^2)$: a weak isospin doublet with $T_3 = \pm 1$, degenerate in mass with $M_W \sim 81$ GeV, and spin 1
- $Z^0 = \cos\theta_W W^3 - \sin\theta_W B$: a neutral massive vector boson with $M_Z \sim M_W / \cos\theta_W$, invariant under weak isospin
- $\gamma = \sin\theta_W W^3 + \cos\theta_W B$: the photon, with zero mass and invariant under weak isospin

Here θ_W is the weak mixing angle, defined as the smallest angle of a right-angled triangle with the catets proportional to the $SU(2)_W$ and $U(1)_{Y_W}$ coupling constants g and g' (Weinberg triangle), as sketched in Figure 1.1. Both weak isospin and weak hypercharge are not conserved anymore, while the electric charge is still a valid symmetry and is associated to the photon, which takes the role of a Goldstone boson and remains massless.

The weak isospin and hypercharge pattern after the SSB is shown in Figure 1.2.

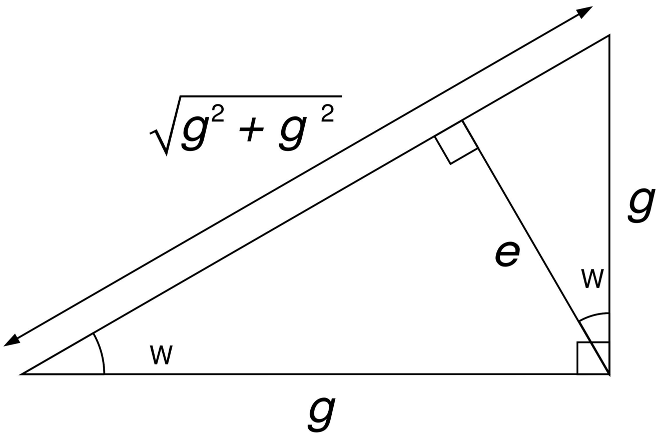

Figure 1.1: Sketch of the Weinberg triangle.

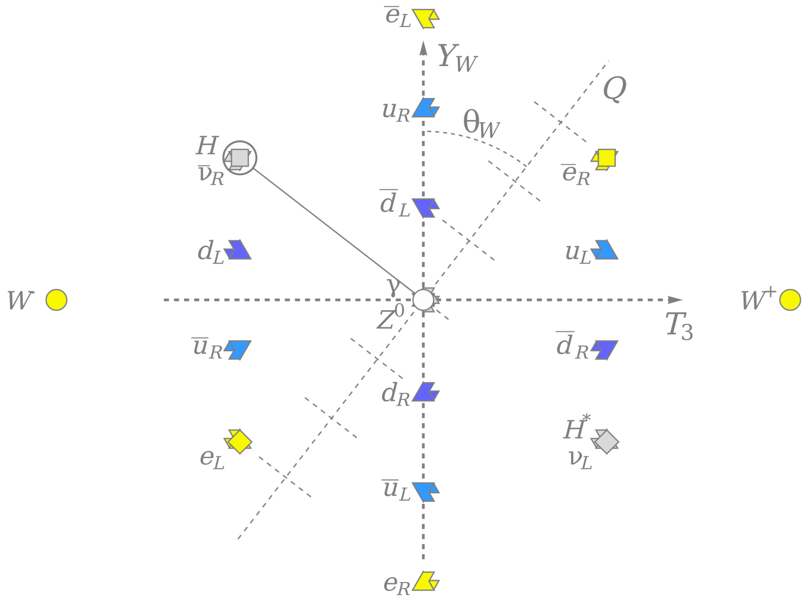

Figure 1.2: Pattern of weak isospin, weak hypercharge and electic charge for the SM particles after spontaneous symmetry breaking.

The Higgs field after SSB reduces to a single component and becomes massive at $m_h = 125$ GeV, as shown in Figure 1.3 together with a sketch of the changes before and after SSB.

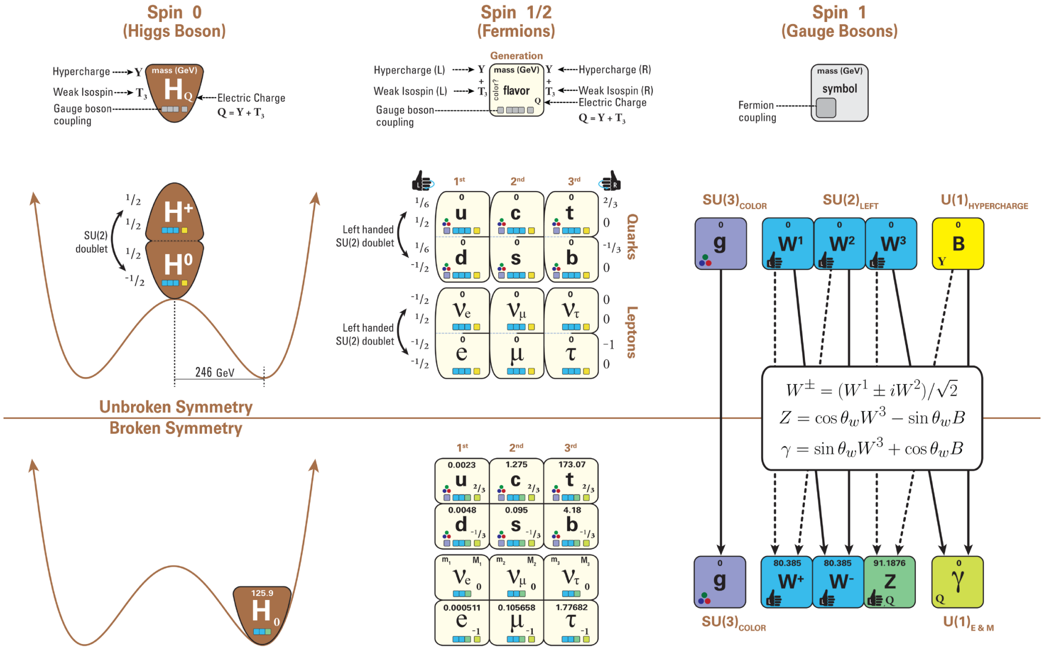

Figure 1.3: Sketch of the SM particle content before and after spontaneous symmetry breaking.

Quantum electrodynamics

Quantum electrodynamics (QED) was the first quantum field theory to be developed and it is described by the simplest gauge group, $\mathrm{U}(1)_Q$. The $U(1)_Q$ is the residual conserved symmetry group of the original electroweak $SU(2)_W \times U(1)_{Y_W}$ after symmetry breaking, reflecting the fact that in the physical realization of the SM only the electric and color charges are conserved. QED describes the interactions of charged particles, which arise by requiring local gauge invariance under transformations of the $\mathrm{U}(1)_Q$ group. To illustrate how this works, the Lagrangian describing a freely propagating electron (or any fermion in general) can be written as:

$$\mathcal{L}_e = i\bar{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\bar{\psi}\psi, \tag{1.1}$$

where m is the mass of the electron, ψ represents the electron field, γ^μ are the gamma matrices, and 'freely' indicates the potential term is zero. Local gauge invariance for QED means that the equations of motion do not change under the following transformation:

$$\psi \to \psi' = e^{iq\theta(x)}\psi \tag{1.2}$$

where $\theta(x)$ is a spacetime dependent phase and q is the coupling strength. To make the Lagrangian invariant under such transformations, additional terms must be introduced. More specifically, the derivative must transform as a covariant derivative, that is $\partial_\mu \to D_\mu = \partial_\mu - iqA_\mu$, introducing a vector boson field (A_μ), which represents a massless spin-1 boson such as the photon, which transforms as:

$$A_\mu \to A_\mu + \frac{1}{q}\partial_\mu\theta(x) \tag{1.3}$$

for any function $\theta(x)$. The result is that in this theory not only there is a vector field, but it also couples to the fermion with strength q. Physically, this can be interpreted as the electric charge of the fermion ψ.

Substituting the above equations (1.2) and (1.3) into equation (1.1) the interaction Lagrangian for QED is obtained:

$$\mathcal{L}_{\mathrm{QED}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \bar{\psi}(i\not{D} - m)\psi \tag{1.4}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the field strength tensor for electrodynamics, and the covariant derivative is defined as $\not{D} = \gamma^\mu(\partial_\mu - iqA_\mu)$ in units where $\hbar = c = 1$. It is important to note that there are no mass terms of the form $\frac{1}{2}m^2 A_\mu A^\mu$ present for the gauge field because such a term would break local gauge invariance, thus the force mediator of QED is required to be massless. Maxwell's equations for electromagnetic waves are recovered from (1.4) in the absence of any sources, then $\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$. The QED Lagrangian is not part of the original SM; it emerges following the spontaneous symmetry breaking of the $\mathrm{SU}(2)_L\times\ \mathrm{U}(1)_{Y_W}$ group.

An example of a QED interaction is given by the Feynman diagram in Figure 1.4.

Quantum chromodynamics

Quantum chromodynamics (QCD) describes the interactions of coloured particles, where the prefix chromo derives from the Ancient Greek *khroma* mean-

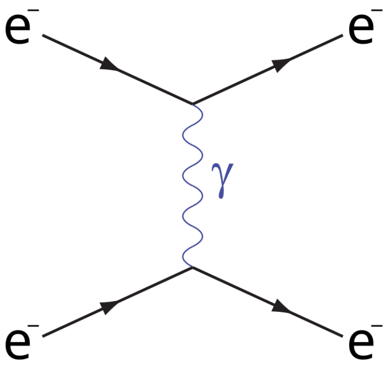

Figure 1.4: A Feynman diagram depicting an example QED process; electron scattering via the exchange of a photon, known as Møller scattering.

ing colour. Generally speaking, QCD describes the interactions of quarks and gluons, which are the only particles in the SM that have colour charge.

QCD is based on the non-Abelian $SU(3)$ gauge symmetry. Quarks carry one of the three colour charges labelled red (R), blue (B) or green (G), and likewise antiquarks carry an anticolour charge, antired ($\bar{R}$), antiblue ($\bar{B}$) and antigreen ($\bar{G}$). Any observed bound state is required to be colourless, which is obtained by combining colour-anticolour components such that the overall net charge is zero. The strong force is mediated by eight massless force carriers known as gluons, which along with quarks, are the only particles in the SM that carry colour charge. The QCD Lagrangian is defined by:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a + i\bar{\psi}\gamma^\mu(\partial_\mu - ig_s G^a_\mu T^a)\psi, \tag{1.5}$$

where ψ is the quark fermion field, the gluon field is defined in terms of the field strength tensor $G^a_{\mu\nu}$, g_s is the coupling constant for $SU(3)_C$ defined by $g_s = \sqrt{4\pi\alpha_s}$, and T^a are the set of matrices representing the group generators. The index a refers to the gluon colour states and is summed from 1 to 8.

The force carriers of QCD themselves carry a colour charge; a consequence of this feature is the ability to self-interact via both cubic and quartic couplings. Such self-interactions amongst the gluons result in a dynamically changing coupling constant for the strong force, which depends on the energy scale. As can be seen in equation (1.6), the value of the strong coupling constant varies as a function of the energy scale of any interaction, parameterised by μ^2. This is known as the running of the coupling:

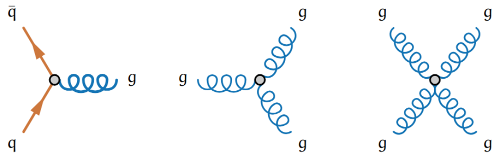

Figure 1.5: The three elementary QCD Feynman vertices, valid in the perturbative expansion.

$$\alpha_s(\mu^2) = \frac{12\pi}{(33 - 2n_f)\ln(\mu^2/\Lambda^2_{\mathrm{QCD}})} \tag{1.6}$$

where n_f is the number of quark flavours contributing at a given energy scale, and $\Lambda_{\mathrm{QCD}} \sim 200$ MeV quantifies the energy scale where the perturbative QCD coupling diverges.

The perturbative expansion is valid only in the energy regime where $\mu^2 > \Lambda_{QCD}$. The three elementary QCD Feynman vertices, which are valid in the perturbative expansion, include the quark-gluon coupling, and cubic and quartic couplings between gluons, and are shown in Figure 1.5.

Two unique and fundamental properties of QCD are uncovered when considering either the very high energy, or the very low energy extremes. The general trend of equation 1.6 as μ^2 increases is that α_s decreases. In the low energy case, when μ^2 is very small, α_s becomes exceedingly large. This behaviour is known as confinement, and is responsible for bounding quarks together inside hadrons. Hadrons, which refers to a colourless composite particle made of quarks, when consisting of two or three quarks, are known as mesons or baryons, respectively. As of yet, no free quarks or free gluons have been found to exist. As a consequence of confinement, any high-energy process which produces a free quark in the final state will radiate gluons and subsequently hadronise, producing a collimated shower of energetic hadrons known as a jet. The proton is one example of a hadron, more specifically a baryon, as it is comprised of two up-type quarks and one down-type quark, held together by gluons. Hadrons in general are comprised of two main components: the valence quarks and the so-called sea of virtual quarks and gluons. Valence quarks are defined by the quark content of the hadron, and carry all relevant quantum numbers: in the proton for example the three valence quarks are thus uud.

In the extreme case where $\mu^2 \to \infty$, $\alpha_s \to 0$, which is known as asymp-

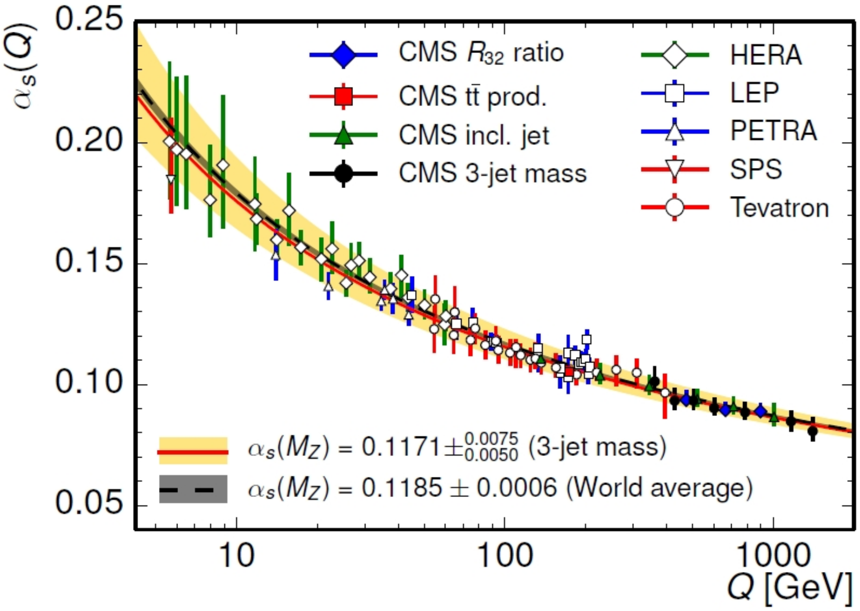

Figure 1.6: Running of the strong coupling constant α_s with respect to scale [2].

totic freedom. This property physically manifests itself when considering ultra high-energies, for example in high energy hadron collisions, where the partons are essentially free and interact independently from the other partons in the nucleons. The dependance of α_s on the energy scale μ^2 has been experimentally verified to agree with QCD predictions, as shown in Figure 1.6. These peculiar features of QCD are of particular importance when conducting physics analyses with the LHC, given the colliding beams are made of protons.

Weak interactions

The weak interaction is one of the four fundamental forces in nature, responsible for processes like beta decay and neutrino interactions. It acts on all leptons and quarks and is essential in nuclear physics and particle astrophysics. What distinguishes it from the electromagnetic and strong forces is its ability to change the flavor of fundamental particles and its violation of certain symmetries, such as parity (P) and charge conjugation (C). It is also extremely short-ranged due to the large masses of its mediating bosons (W and Z).

Historically, the first successful description of the weak force came from Enrico Fermi in 1933, who formulated a theory to explain beta decay. He proposed a four-fermion pointlike interaction involving the neutron, proton, electron, and neutrino. This theory assumed that the weak force acted instantaneously at a point, similar to how electromagnetic interactions were once viewed. The effective interaction Fermi wrote down had the structure of a current-current interaction, where a "hadronic current" (involving nucleons) interacts with a "leptonic current" (involving electrons and neutrinos), both coupled through a universal coupling constant, G_F, now known as the Fermi constant. The theory took the form of a contact term in the Lagrangian:

$$\mathcal{L}_F = -\frac{G_F}{\sqrt{2}} \left(\bar{p}\gamma^\mu n\right) \left(\bar{e}\gamma_\mu \nu_e\right)$$

This approach was very successful at low energies, but it lacked internal consistency at higher energies due to its non-renormalizable nature. It also could not explain the structure of weak interactions beyond beta decay, nor could it accommodate observed parity violation.

The breakthrough came in the 1960s and 1970s, with the development of the electroweak theory by Glashow, Salam, and Weinberg. They proposed a unified framework that merged the weak and electromagnetic forces into a single gauge theory, based on the symmetry group $SU(2)_L \times U(1)_Y$. In this theory, weak interactions are mediated by the exchange of massive $W^\pm$ and Z^0 bosons. The weak force, therefore, is no longer described by a contact interaction but instead by the exchange of real or virtual gauge bosons, as shown in Figure 1.7.

Fermi's four-fermion interaction reappears in this framework as a low-energy approximation to the full theory. At energies much lower than the mass of the W boson, the propagator of a virtual W can be expanded as:

$$\frac{-ig^{\mu\nu}}{q^2 - M_W^2} \approx \frac{ig^{\mu\nu}}{M_W^2} \quad \text{for } q^2 \ll M_W^2$$

This approximation effectively "shrinks" the interaction to a point, reproducing Fermi's result with the identification:

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8M_W^2} \tag{1.7}$$

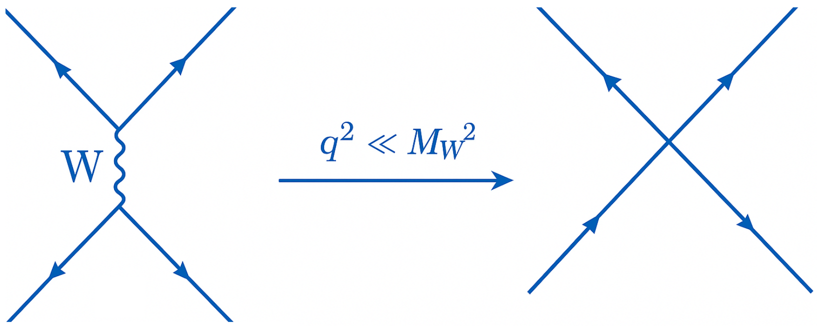

Figure 1.7: Feynman diagrams for charged-current interaction in the perturbative SM and in the Fermi theory.

where g is the SU(2) gauge coupling constant, and M_W is the mass of the W boson. This relation connects the phenomenological Fermi constant to the underlying parameters of the electroweak theory.

One of the key concepts in the electroweak theory is the idea of weak currents. There are two types of them: *charged currents*, which involve the exchange of a $W^{\pm}$ boson, and *neutral currents*, mediated by the Z^0 boson (not included in Fermi theory). Charged current interactions involve a change in electric charge and are responsible for processes like $\mu^- \to e^- \bar{\nu}_e \nu_\mu$ or $u \to de^+ \nu_e$, so they are able to change particle flavour. These interactions couple only to left-handed fermions and are maximally parity violating. Neutral currents, on the other hand, do not change the charge of the particles involved not the flavour, as in $\nu_\mu + e^- \to \nu_\mu + e^-$ for instance, and were first observed experimentally in 1973.

The charged current takes the form:

$$J^{\mu}_{\text{charged}} = \bar{u}_f \gamma^{\mu} (1 - \gamma^5) u_i$$

This structure reflects the vector-minus-axial-vector (V–A) nature of weak interactions, emphasizing that only left-handed fermions (and right-handed antifermions) participate to it. This chiral asymmetry is built into the gauge theory itself through the SU(2) group structure.

This arrangement ensures that only left-handed particles feel the SU(2) weak force, as they are assigned to doublets under the $\mathrm{SU}(2)_L$ gauge group, while right-handed particles are singlets and thus do not couple to the $\mathrm{SU}(2)_L$ gauge fields. The gauge principle then dictates the form of the Lagrangian through minimal coupling: the kinetic term for fermions is modified by replacing the ordinary derivative with a covariant derivative that includes the gauge fields.

For an SU$(2)_L$ × U$(1)_Y$ gauge theory, the covariant derivative acting on a left-handed lepton doublet $L = \begin{pmatrix} \nu_e \\ e \end{pmatrix}_L$ takes the form:

$$D_\mu = \partial_\mu - ig\frac{\tau^a}{2}W_\mu^a - ig'\frac{Y}{2}B_\mu,$$

where τ^a are the Pauli matrices (generators of SU(2)), W_μ^a are the SU(2) gauge fields, B_μ is the U$(1)_Y$ gauge field, and Y is the weak hypercharge of the field. The SU(2) coupling constant is g, and g' is the U(1) coupling, as shown in Figure 1.1. The full kinetic term in the Lagrangian for the left-handed doublet is then:

$$\mathcal{L}_{\text{kin}} = \bar{L}i\gamma^\mu D_\mu L.$$

Expanding this explicitly shows how the gauge fields couple to the fermions, and the SU(2) structure ensures that weak isospin interactions are mediated by the triplet W_μ^a. The charged current interaction arises specifically from the off-diagonal components W^1 and W^2, while the neutral current comes from W^3 and its mixing with the hypercharge field B_μ.

The SU(2) symmetry constrains the form of all interaction terms involving the gauge fields and fermions: gauge invariance requires that these terms are constructed using the covariant derivative and the field strength tensors of the gauge fields. The Lagrangian includes the pure gauge field terms:

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

where $W_{\mu\nu}^a$ and $B_{\mu\nu}$ are the field strength tensors for SU(2) and U(1), respectively. Together with the fermion kinetic terms and their gauge interactions, these define the gauge-invariant electroweak Lagrangian before spontaneous symmetry breaking.

Thus, the SU$(2)_L$ symmetry not only explains why the weak interactions couple only to left-handed fermions, but also dictates the full structure of the weak interaction Lagrangian. The theory's predictive power, including the precise form of charged and neutral current interactions, arises directly from these symmetry principles.

1.1.2 Higgs mechanism

In order to understand the Higgs mechanism in a simplified framework, the case of SSB in an abelian U(1) gauge theory is discussed, to demonstrate how the mass of the corresponding gauge boson (an analogue of the photon) comes about [3]. The abelian example will then be generalized to the electroweak Standard Model [4].

Spontaneous symmetry breaking in a U(1) theory

The U(1) gauge invariant kinetic term of the photon is given by

$$\mathcal{L}_{\text{kin}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{1.8}$$

where

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \tag{1.9}$$

That is, $\mathcal{L}_{\text{kin}}$ is invariant under the transformation: $A_\mu(x) \to A_\mu(x) - \partial_\mu \eta(x)$ for any η and x.

If a mass term for the photon is added to the Lagrangian,

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{1}{2}m^2 A_\mu A^\mu, \tag{1.10}$$

this violates the local gauge symmetry.

In order to avoid introducing manifestly symmetry-violating to the Lagrangian, the model is extended by introducing a complex scalar field with charge $(-e)$ that couples both to itself and to the photon:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + (D_\mu\phi)^\dagger(D^\mu\phi) - V(\phi), \tag{1.11}$$

where $D_\mu = \partial_\mu - ieA_\mu$ and $V(\phi) = -\mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$. In this case the Lagrangian is invariant under the gauge transformations:

$$A_\mu(x) \to A_\mu(x) - \partial_\mu\eta(x), \tag{1.12}$$

$$\phi(x) \to e^{ie\eta(x)}\phi(x). \tag{1.13}$$

If $\mu^2 < 0$, the state of minimum energy will be the one with $\phi = 0$, as shown in Figure 1.8, and the potential will preserve the symmetries of the Lagrangian, therefore the theory would be equal to scalar QED, with a massless photon and a charged scalar field ϕ with mass μ.

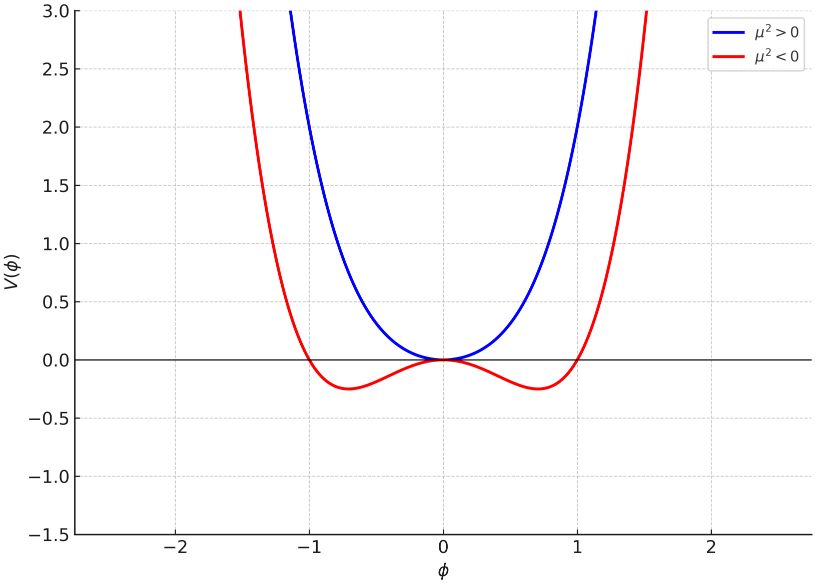

Figure 1.8: Higgs potential in one dimension, with the two cases depending on the sign of μ^2.

However, if $\mu^2 > 0$, the minimum at $\phi = 0$ is unstable and therefore the vacuum state will move to the global minimum, and the field ϕ will acquire a vacuum expectation value (VEV):

$$\langle\phi\rangle = \sqrt{\mu^2/2\lambda} \equiv \frac{v}{\sqrt{2}}, \tag{1.14}$$

breaking spontaneously the global U(1) gauge symmetry, as a consequence.

It is convenient to parameterize ϕ as

$$\phi = \frac{v+h}{\sqrt{2}} e^{i\chi/v}, \tag{1.15}$$

where h and χ, which are referred to as the Higgs boson and the Goldstone boson, respectively, are real scalar fields which have no VEVs. Substituting Equation 1.15 back into the Lagrangian, this expression is found:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - evA_\mu\partial^\mu\chi + \frac{e^2v^2}{2}A_\mu A^\mu + \frac{1}{2}(\partial_\mu h\partial^\mu h - 2\mu^2h^2)+ \tag{1.16}$$

$$\frac{1}{2}\partial_\mu\chi\partial^\mu\chi + (\text{h}, \chi\,\text{interactions}) \tag{1.17}$$

This Lagrangian now describes a theory with a photon of mass $m_A = ev$, a Higgs boson h with $m_h = \sqrt{2\mu} = \sqrt{2\lambda}v$, and a massless Goldstone χ. The χ–A_μ mixing can be removed by making the following gauge transformation:

$$A_\mu \rightarrow A'_\mu = A_\mu - \frac{1}{ev}\partial_\mu \chi. \tag{1.18}$$

The gauge choice with the transformation above is called the unitary gauge. The Goldstone χ will then completely disappear from the theory and its degree of freedom will be used to give mass to the photon.

It is instructive to count the degrees of freedom (DOF) before and after SSB has occurred, to verify the matching and understand their evolution. Originally there were a massless photon (2 DOF) and a complex scalar field (2 DOF), for a total number of 4 DOF. Indeed, massless vector bosons (like the photon in the SM) possess 2 polarization DOF, both transverse, while massive vector bosons possess all three DOF. After the SSB, there is one massive photon (3 DOF) and a real scalar field h (1 DOF), again for a total of 4 DOF.

Spontaneous symmetry breaking in the SM

In the SM, the Higgs field Φ is a complex scalar field transforming as a doublet under $SU(2)_L$ with hypercharge $Y = 1$:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix},$$

where ϕ^+ and ϕ^0 are complex scalar components.

The covariant derivative acting on the Higgs doublet is defined as

$$D_\mu = \partial_\mu - ig\frac{\sigma^a}{2}W^a_\mu - ig'\frac{Y}{2}B_\mu,$$

where σ^a are the Pauli matrices, generators of $SU(2)_L$, and $Y = 1$ is the hypercharge of Φ.

Explicitly,

$$D_\mu\Phi = \partial_\mu\Phi - ig\frac{\sigma^a}{2}W^a_\mu\Phi - ig'\frac{1}{2}B_\mu\Phi.$$

The Higgs sector Lagrangian is

$$\mathcal{L}_{\text{Higgs}} = (D_\mu\Phi)^\dagger(D^\mu\Phi) - V(\Phi),$$

where the Higgs potential is

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

with $\lambda > 0$ and $\mu^2 < 0$.

The negative μ^2 triggers spontaneous symmetry breaking by making the symmetric vacuum unstable.

The configuration of the stable vacuum is found by minimizing the potential:

$$\frac{\partial V}{\partial \Phi} = 0.$$

Since V depends only on $\Phi^\dagger \Phi$, the minimum satisfies

$$\Phi^\dagger \Phi = -\frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2},$$

where $v = \sqrt{-\frac{\mu^2}{\lambda}}$.

By a gauge transformation, the VEV can be chosen in the neutral component, spontaneously breaking the gauge symmetry:

$$\boxed{\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}.}$$

After SSB the Higgs field can be written as a fluctuation around the VEV:

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix},$$

where $h(x)$ is the physical Higgs boson field.

The gauge boson mass terms are found by looking at the covariant derivatives after SSB when $h(x)$ is 0:

$$D_\mu \langle \Phi \rangle = \partial_\mu \langle \Phi \rangle - ig \frac{\sigma^a}{2} W_\mu^a \langle \Phi \rangle - ig' \frac{1}{2} B_\mu \langle \Phi \rangle.$$

Since $\langle \Phi \rangle$ is constant, $\partial_\mu \langle \Phi \rangle = 0$, therefore:

$$\frac{\sigma^a}{2} W_\mu^a \langle \Phi \rangle = \frac{1}{2} \left(\sigma^1 W_\mu^1 + \sigma^2 W_\mu^2 + \sigma^3 W_\mu^3 \right) \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}.$$

Calculating explicitly using the Pauli matrices (SU(2) generators):

$$D_\mu\langle\Phi\rangle = -i\frac{v}{2\sqrt{2}}\begin{pmatrix} g(W_\mu^1 - iW_\mu^2) \\ -gW_\mu^3 + g'B_\mu \end{pmatrix}.$$

The gauge boson mass term is:

$$\mathcal{L}_{\text{mass}} = (D_\mu\langle\Phi\rangle)^\dagger(D^\mu\langle\Phi\rangle) = \frac{v^2}{8}\left[g^2(W_\mu^1W^{\mu 1} + W_\mu^2W^{\mu 2}) + (-gW_\mu^3 + g'B_\mu)^2\right].$$

Defining the charged W bosons as:

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

the charged boson mass term is

$$\mathcal{L}_{W^\pm\ \text{mass}} = \frac{g^2v^2}{4}W_\mu^- W^{\mu +}.$$

and the charged W boson mass is:

$$\boxed{m_W = \frac{gv}{2}.} \tag{1.19}$$

The neutral boson mass terms come from

$$\frac{v^2}{8}(-gW_\mu^3 + g'B_\mu)^2 = \frac{v^2}{8}\begin{pmatrix} W_\mu^3 & B_\mu \end{pmatrix}\begin{pmatrix} g^2 & -gg' \\ -gg' & g'^2 \end{pmatrix}\begin{pmatrix} W^{\mu 3} \\ B^\mu \end{pmatrix}.$$

This symmetric matrix can be diagonalized by an orthogonal transformation, defining the weak mixing angle θ_W:

$$\tan\theta_W = \frac{g'}{g}.$$

Defining therefore the physical fields as linear combinations:

$$\begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & -\sin\theta_W \\ \sin\theta_W & \cos\theta_W \end{pmatrix}\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix},$$

the mass matrix becomes diagonal with eigenvalues

$$m_Z^2 = \frac{v^2}{4}(g^2 + g'^2), \quad m_A^2 = 0.$$

Therefore the photon, with zero mass, and coupling to charges $e = g \sin\theta_W$, is recovered and the Z boson is generated, with a mass given by:

$$\boxed{m_Z = \frac{v}{2}\sqrt{g^2 + g'^2} = \frac{m_W}{\cos\theta_W},}$$

The relation between g, m_W, and G_F is given in Equation 1.7. By using the experimentally measured value of G_F from muon decay [5], and determining θ_W from the ratio of charged- to neutral-current interactions, the weak coupling g can be calculated via

$$e = g\ \sin\theta_W.$$

This, in turn, allows one to predict m_W and, consequently, m_Z.

On the other hand, to generate fermion masses, the Standard Model introduces Yukawa couplings between fermions and the Higgs field.

For a generic fermion f, the left-handed doublet L_f and right-handed singlet R_f are written as:

$$L_f = \begin{pmatrix} \nu_f \\ f \end{pmatrix}_L, \quad R_f = f_R.$$

The Yukawa interaction term is

$$\mathcal{L}_{\text{Yukawa}} = -y_f \bar{L}_f \Phi R_f + \text{h.c.},$$

where y_f is the Yukawa coupling constant for fermion f.

After spontaneous symmetry breaking, the Higgs field Φ is replaced by its VEV:

$$\Phi \to \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v \end{pmatrix}.$$

This gives a mass term for the fermion:

$$\mathcal{L}_{\text{mass}} = -\frac{y_f v}{\sqrt{2}} \bar{f}_L f_R + \text{h.c.} = -m_f \bar{f} f,$$

where

$$\boxed{m_f = \frac{y_f v}{\sqrt{2}}.} \tag{1.20}$$

The fermion mass is therefore proportional to the Higgs VEV and its Yukawa coupling.

Understanding the Higgs mechanism requires careful accounting of the physical degrees of freedom present in the theory before and after spontaneous symmetry breaking. This counting not only confirms the internal consistency of the mechanism but also provides insight into how the gauge bosons acquire mass through the absorption of Goldstone bosons degrees of freedom. Prior to spontaneous symmetry breaking, the Higgs field Φ is an $SU(2)_L$ complex doublet with hypercharge $Y = 1$. Being a complex doublet means that Φ has two complex components, which correspond to four real scalar degrees of freedom in total, written as:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix},$$

where $\phi_1, \phi_2, \phi_3, \phi_4$ are real scalar fields, and each contributes one real degree of freedom. Thus,

$$\text{Number of scalar degrees of freedom} = 4.$$

In addition to the scalar fields, the gauge sector of the electroweak theory consists of four gauge bosons:

- Three gauge bosons W_μ^a associated with the $SU(2)_L$ symmetry ($a = 1, 2, 3$),
- One gauge boson B_μ associated with the $U(1)_Y$ symmetry.

Before symmetry breaking, all these gauge bosons are massless. A massless vector boson in four-dimensional spacetime possesses only two physical polarization states, corresponding to two transverse degrees of freedom. Therefore, the total degrees of freedom in the gauge boson sector before symmetry breaking are:

$$3 \times 2 = 6 \quad \text{for } W_\mu^a, \quad \text{and} \quad 2 \quad \text{for } B_\mu,$$

which sum to

$$6 + 2 = 8.$$

Hence, before the Higgs field acquires a vacuum expectation value, the combined total of physical degrees of freedom in the scalar and gauge sectors is

$$4 + 8 = 12.$$

When the Higgs field acquires a vacuum expectation value (VEV), the electroweak symmetry $SU(2)_L \times U(1)_Y$ is spontaneously broken down to the electromagnetic $U(1)_{\text{EM}}$. This symmetry breaking has profound implications for the physical content of the theory.

The crucial insight of the Higgs mechanism is that three of the four original scalar degrees of freedom in the Higgs doublet do not appear as physical scalar particles after symmetry breaking. Instead, these three degrees of freedom become the longitudinal polarization modes of the gauge bosons that acquire mass, specifically the charged $W^{\pm}$ bosons and the neutral Z boson.

The massive vector bosons have three polarization states each: two transverse and one longitudinal. The absorption of the scalar degrees of freedom as longitudinal components allows the gauge bosons to become massive without explicitly breaking gauge invariance.

Below the explicitly count of the degrees of freedom:

- The two $W^{\pm}$ and Z bosons carry three polarization states, so together: $W^+/W^-/Z : 3 \times 3 = 9\,\text{degrees of freedom}$
- The photon A_μ is massless with only two transverse degrees of freedom: $\gamma :\ 2\,\text{degrees of freedom}$
- the Higgs boson h, which corresponds to fluctuations around the vacuum expectation value, has one degree of freedom (being a real scalar)

Therefore, after symmetry breaking, the total number of physical degrees of freedom in the scalar and gauge sectors is:

$$9(W^+/W^-/Z) + 2(\gamma) + 1(h) = 12$$

which matches the total number of degrees of freedom before the SSB.

1.2 Higgs boson production modes and decay channels

The discovery in 2012 of a scalar resonance with a mass near 125 GeV by the ATLAS [6] and CMS [7] collaborations marked a milestone in the understanding of high-energy particle physics. Using the full dataset from LHC Run 1, consisting in about 30 fb^{-1}, collected at center-of-mass energies of 7 and 8 TeV between 2010 and 2012, the experimental evidence established the first consistent picture of electroweak symmetry breaking within the framework of the Standard Model (SM).

The Run 2 of the LHC, conducted from 2015 to 2018 at a center-of-mass energy of 13 TeV, provided a significantly larger dataset of about 138 fb^{-1} and allowed for more precise measurements. These data confirmed the compatibility of the newly discovered boson with the SM Higgs boson, disfavoring many alternative scenarios [8]. As of 2025, the Run 3 is underway and expected to continue through 2026, aiming to deliver an integrated luminosity of approximately 350 fb−1 per experiment [4].

Across all observed production mechanisms and decay modes, the measured cross sections and differential distributions are in agreement - within current experimental and theoretical uncertainties - with SM predictions [8, 9].

In high-resolution final states such as four-leptons and diphotons, the Higgs boson mass has been measured with sub-per-mille precision, yielding a combined value of 125.11 ± 0.11 GeV for ATLAS and 125.04 ± 0.12 GeV for CMS in the four-lepton channel [10, 11, 12]. More recently, there has been evidences of Higgs boson decay to dimuons [13, 14]. However, direct evidence for couplings to first-generation fermions remains beyond the current experimental reach: this reflects the inherent challenges posed by these interactions, which are predicted to be extremely small in the SM. For instance, searches of $H \to e^+e^-$ by CMS yielded limits at 95% corresponding to about 60000 times the SM prediction [15].

Despite the growing data statistics, the nature of the Higgs boson is still not fully probed. A more complete Higgs sector, possibly comprising additional scalar states, is still possible within current experimental constraints. An example is the the Two-Higgs-Doublet Model (2HDM), which extends the Standard Model with an additional Higgs doublet, leading to a rich phenomenology including extra neutral and charged Higgs bosons and modified couplings. Searches of 2HDM were performed by CMS, as shown in Figure 1.9, with no deviations with respect to SM predictions.

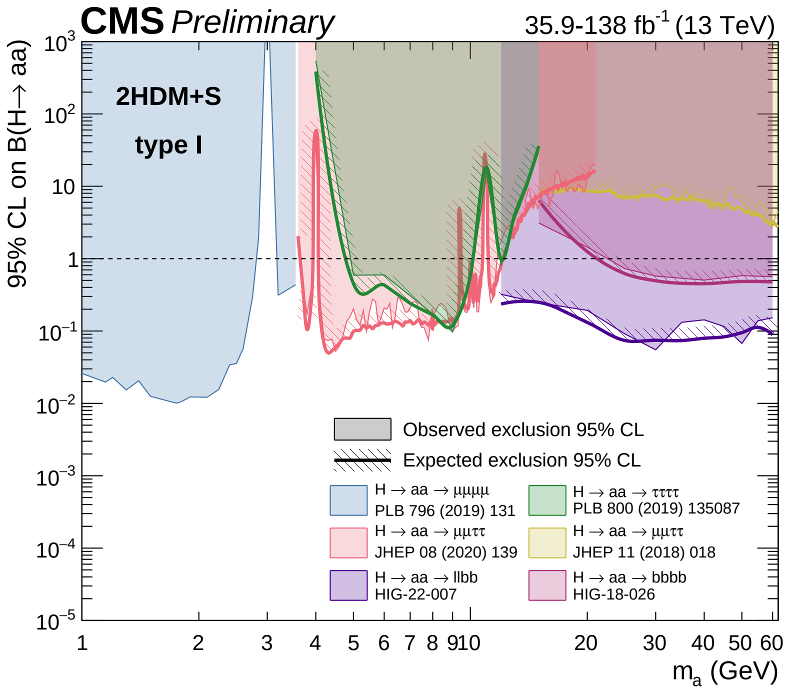

Figure 1.9: 95% CL on $B(h \to aa)$ in the 2HDM+S type-1 scenario for exotic h decay searches performed with data collected at 13 TeV center-of-mass energy. References are summarized in the plot legend.

Moreover, it has not yet been determined whether the Higgs is a fundamental scalar or a composite object, unlike any previously observed particle. These open questions continue to drive a vibrant research program in both theoretical and experimental high-energy physics [16].

The existence of the Higgs boson is essential for the internal consistency of the SM. In its absence, perturbative unitarity would be violated at high energies, particularly in longitudinal vector boson scattering processes. Furthermore, the theory would suffer from problematic divergences in radiative corrections to electroweak observables. The Higgs mechanism resolves these issues by ensuring a renormalizable and unitary framework. Consequently, the SM can, in principle, be extrapolated to energies beyond the electroweak scale without the need of additional degrees of freedom. Nevertheless, the Higgs sector introduces a conceptual tension known as the hierarchy problem. As a fundamental scalar, the Higgs mass is quadratically sensitive to ultraviolet (UV) physics. In the absence of a protective symmetry, quantum corrections drive its mass toward the highest energy scale in any theory that UV-completes the SM, implying a need for unnatural fine-tuning to stabilize the electroweak scale. This theoretical issue has motivated many scenarios that extend the SM near the TeV scale[17].

In the thirteen years since its discovery, the Higgs boson has emerged as a central topic in particle physics and a particular window into physics beyond the Standard Model. Its properties continue to be tested with increasing precision, and the study of Higgs boson phenomenology remains at the frontier of particle physics.

1.2.1 Higgs boson production modes

At hadron colliders such as the LHC, several mechanisms contribute to the production of Higgs bosons. The dominant production channels in the Standard Model (SM) are gluon fusion (ggH), vector boson fusion (VBF), associated production with electroweak gauge bosons (denoted as VH, where $V = W, Z$), and associated production with top quarks ($t\bar{t}H$). Representative Feynman diagrams for these processes are shown in Figure 1.10.

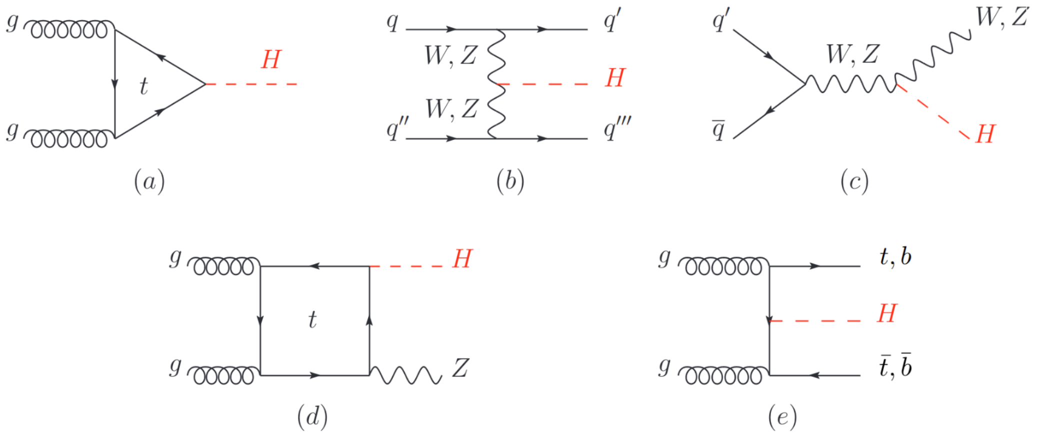

(a) (b) (c)

(d) (e)

Figure 1.10: Feynman diagrams for the main Higgs boson production mechanisms at LHC: (a) gluon fusion, (b) vector boson fusion, (c) associated production with W/Z bosons, (d) associated production with top quark pairs.

Theoretical predictions for the production cross sections in each of these channels have reached a high level of precision, incorporating up to next-to-leading-order (NLO) and next-to-next-to-leading-order (NNLO) QCD and electroweak corrections, reaching N3LO for inclusive gluon-gluon fusion [18]. The energy dependence of the production cross sections in proton-proton (pp) collisions, along with theoretical uncertainty bands, is illustrated in Figure 1.11 [4, 19] for $m_H = 125$ GeV.

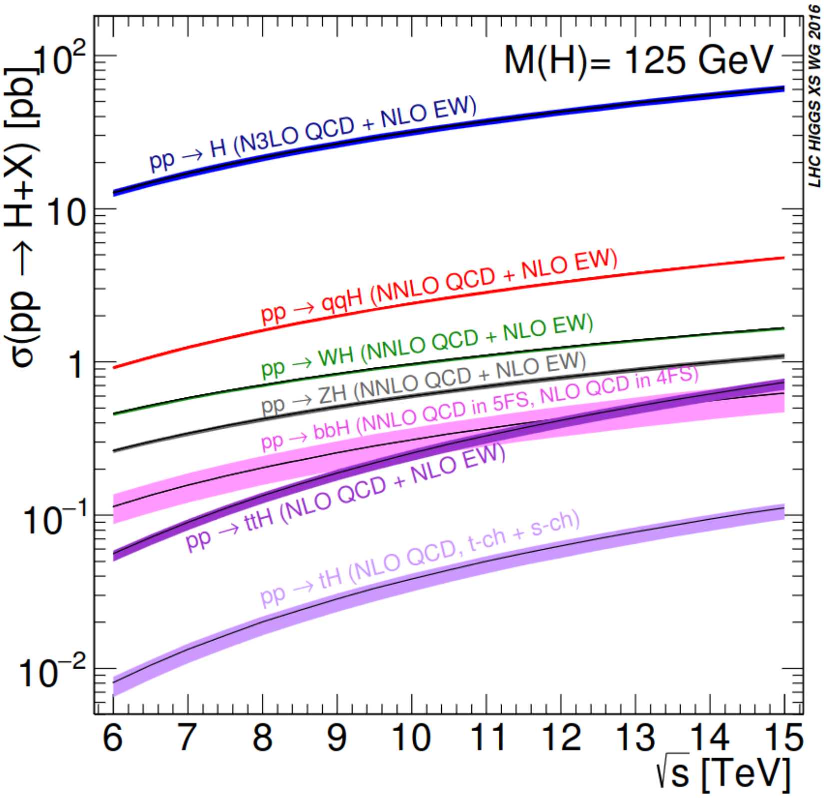

Figure 1.11: SM Higgs boson production cross sections as a function of the center of mass energy for pp collisions, for $m_H = 125$ GeV [4, 19].

Gluon-gluon fusion (ggH)

The gluon-gluon fusion process, $gg \rightarrow H + X$, shown in Figure 1.10(a), is the dominant Higgs production mechanism at the LHC due to the high gluon-gluon effective luminosity in the proton parton distribution functions (PDF) at the large LHC center-of-mass energy. Although the Higgs bosons does not couple directly to gluons at tree level in the SM, because gluons are massless, this process proceeds via a top-quark loop, which mediates an effective ggH vertex. The large top-quark mass enhances the loop-induced coupling, making ggH the leading production mode over a wide energy range.

Vector boson fusion (VBF)

Vector boson fusion is the second most significant production mode at the LHC. The process $qq \to qqH$, shown in Figure 1.10(b), occurs via t- or u-channel exchange of a virtual W or Z boson, with the Higgs radiated from the exchanged gauge boson. The final-state quarks typically emerge as two high-p_T jets in opposite hemispheres of the detector, resulting in a large rapidity gap between them, as shown in Figure 1.12, and suppressed hadronic activity in the central region.

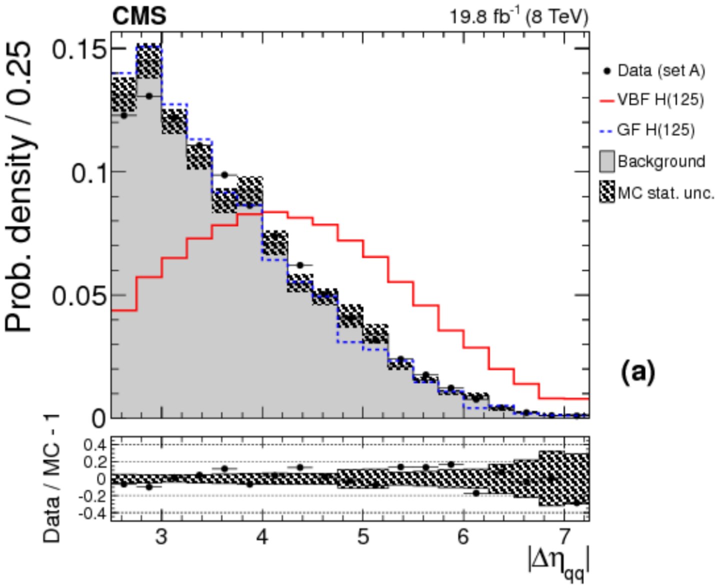

Figure 1.12: Distributions of rapidity gap between jets in events with two high-energy jets in the bb channel, for different production modes [20].

This characteristic topology arises from the color-singlet nature of the weak bosons fusion, which reduces QCD radiation and can be exploited to discriminate VBF events from backgrounds. When proper event selection cuts are applied, the VBF process stands out as a clean environment allowing Higgs boson precision coupling measurements, particularly in channels with otherwise low S/B ratio or no way to trigger signal events, for example $H \to$ invisible searches [21].

Associated production with weak gauge bosons (VH)
The associated production of a Higgs boson with a W or Z boson, $q\bar{q}' \to WH$ or $q\bar{q} \to ZH$, shown in Figure 1.10(c-d), also contributes significantly to the Higgs production rate. While subdominant at the LHC, the presence of a vector boson in the final state, especially when decaying leptonically, provides an experimental signature that allows sufficient background suppression, particularly in channels with otherwise large S/B ratio, as for measurements targeting the $H \to bb$ decay [22].

Associated production with top quarks pairs ($t\bar{t}H$)
The production of a Higgs boson in association with top quarks is of particular interest because it allows for direct measurement of the top-Higgs Yukawa coupling. In the $pp \to t\bar{t}H$ process, shown in Figure 1.10(e), the Higgs boson is radiated off a top-quark line, and the complex final state can be exploited to suppress the background and observe the signal despite its small cross section, and to directly measure the top-Higgs boson coupling, including its CP properties.

1.2.2 Higgs boson decay channels

A comprehensive understanding and interpretation of the experimental results requires the precise calculation of all relevant Higgs boson branching ratios. The Higgs boson mass, of approximately 125 GeV, enables the investigation of its couplings to a wide range of SM particles, as shown in Figure 1.13. The dominant decay channels are $H \to b\bar{b}$ and $H \to WW^*$, followed by $H \to gg$, $H \to \tau^+\tau^-$, $H \to c\bar{c}$, and $H \to ZZ^*$. Decays into $\gamma\gamma$, γZ, and $\mu^+\mu^-$ occur with significantly smaller branching fractions. Notably, the loop-induced decays into gluons, diphotons, and $Z\gamma$ provide indirect sensitivity to the Higgs boson couplings to WW, ZZ, and top quarks, allowing access to combinations of these couplings that are otherwise difficult to probe directly.

The sensitivity of a given decay channel depends on several factors: the Higgs production cross section, the decay branching fraction, the mass resolution achievable through reconstruction, the event selection efficiency, and the background contamination in the final state.

At the LHC, five decay channels are particularly important for a Higgs boson with $m_H \approx 125$ GeV. The $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ channels allow for precise measurements of all final-state particles, resulting in excellent reconstructed mass resolutions of approximately 1–2%. The $H \to$

$W^{\pm}W^{\mp}(*) \rightarrow \ell^{+}\nu\ell^{-}\bar{\nu}$ channel benefits from a relatively large branching fraction and clean final state, but the presence of neutrinos escaping detection degrades the mass resolution, as the W mass must be inferred indirectly from variables such as the transverse mass, i.e. the invariant mass computed only in the transverse plane to the beams direction. The $H \rightarrow b\bar{b}$ and $H \rightarrow \tau^{+}\tau^{-}$ channels face large backgrounds, but also feature large signal yields, and have intermediate mass resolutions of roughly 10% and 15%, respectively.

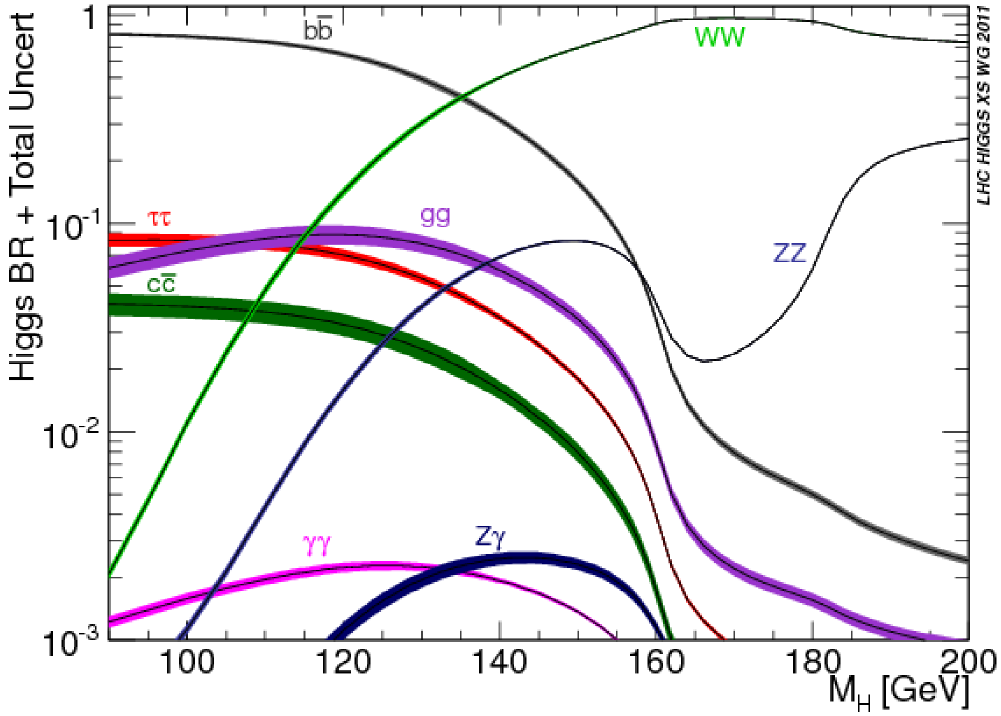

Figure 1.13: Branching ratios of the Higgs boson as a functino of its mass [23].

$\mathbf{H} \rightarrow \gamma\gamma$ decay

In the diphoton channel, it is possible to measure a narrow peak over a smoothly falling background in the invariant mass distribution of two high energy photons, as shown in Figure 1.14.

The background is significant and originates from prompt $pp \rightarrow \gamma\gamma$ processes, which constitute irreducible backgrounds, as well as from $\gamma +$jet and

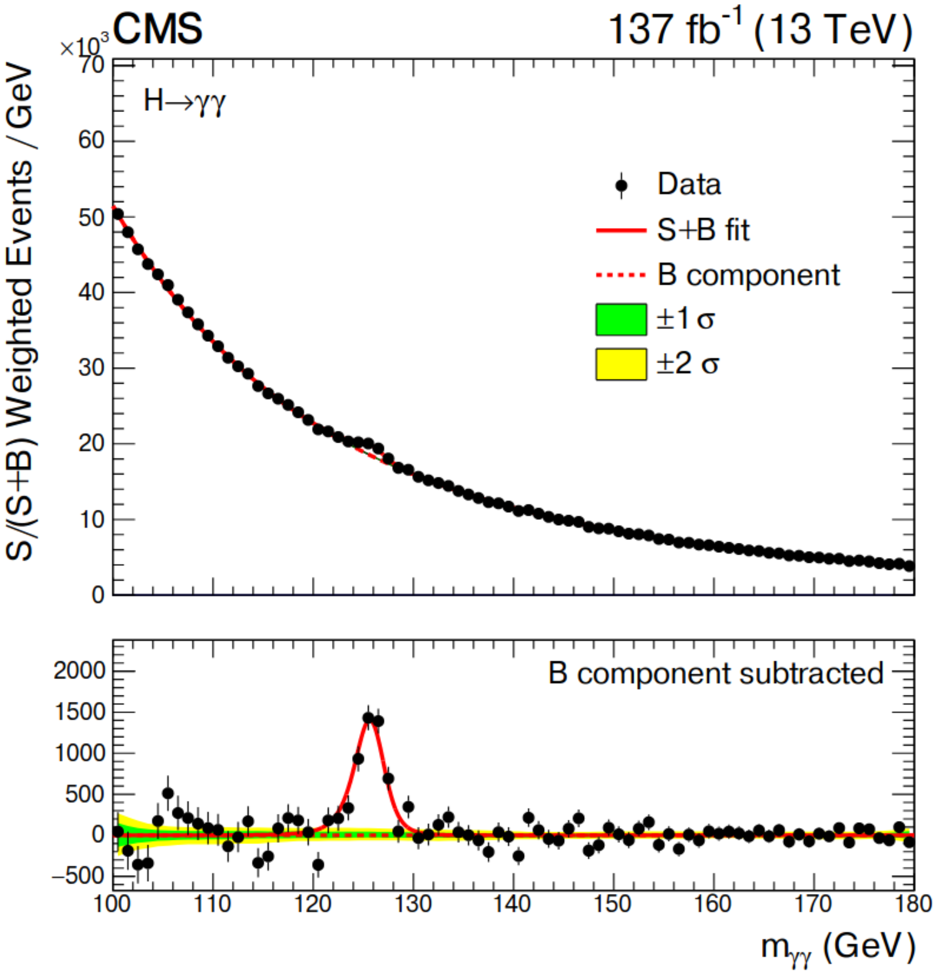

Figure 1.14: Best fit signal-plus-background model, overlaid on the S/(S+B)-weighted distribution of the data points (black), for the sum of all the categories. S and B represent the fitted number of Higgs boson candidates and background events in the mass peak region [24].

dijet processes where one or two jet typically fragments into a high-energy π^0, misidentified as a photon, representing reducible backgrounds.

Both ATLAS and CMS have performed precise calibrations of the photon energy response, using control samples such as $Z \to e^+e^-$ and $Z \to \mu^+\mu^-\gamma$ [25, 26] as detailed in Section 4.2.1. This calibration is employed to correctly simulate the signal invariant mass line-shapes. Parametric signal models are then fitted to these line-shapes to define the functional form of the signal. The background shapes are modeled by fitting the data $m_{\gamma\gamma}$ distributions in each category, using simple analytical functions. Categories, with different S/B ratio, are fitted simultaneously to extract the signal yield or other properties for a given Higgs boson mass hypothesis, as shown in Figure 1.14. This channel is the one employed in this thesis, and will therefore discussed in detail in the following chapters.

$\mathbf{H \to ZZ^* \to 4\ell}$ decay

The so-called golden channel involves a search for a narrow mass peak over a small continuous background, dominated by non-resonant ZZ^* production from $q\bar{q}$ annihilation and gluon-gluon fusion processes, as shown in Figure 1.15. The irreducible background shape and normalization are taken from simulation, unlike the diphoton case. This is possible because reducible backgrounds, including $Z + b\bar{b}$, $t\bar{t}$, and $Z +$ jets events, are suppressed by requirements on lepton isolation and vertex impact parameters, and their yields can be estimated from control samples in data.

To discriminate the Higgs boson signal from the dominant non-resonant ZZ^* background, both ATLAS [27] and CMS [28] employ multivariate techniques and matrix element likelihood approaches, constructing kinematic discriminants based on the ratio of leading-order matrix elements $\left|M^2_{\text{sig}}/M^2_{\text{bkg}}\right|$ [28]. These discriminants are calculated for each four-lepton event under the signal hypotheses (e.g., $gg \to H \to 4\ell$, $q\bar{q} \to VH(\to 4\ell)$, $q\bar{q} \to qqH(\to 4\ell)$) and the main background hypothesis $q\bar{q} \to ZZ \to 4\ell$.

Since the four-lepton mass resolutions and reducible background levels differ among the sub-channels 4μ, $4e$, and $2e2\mu$, these are analyzed separately and statistically combined to produce the final results, as shown in Figure 1.15.

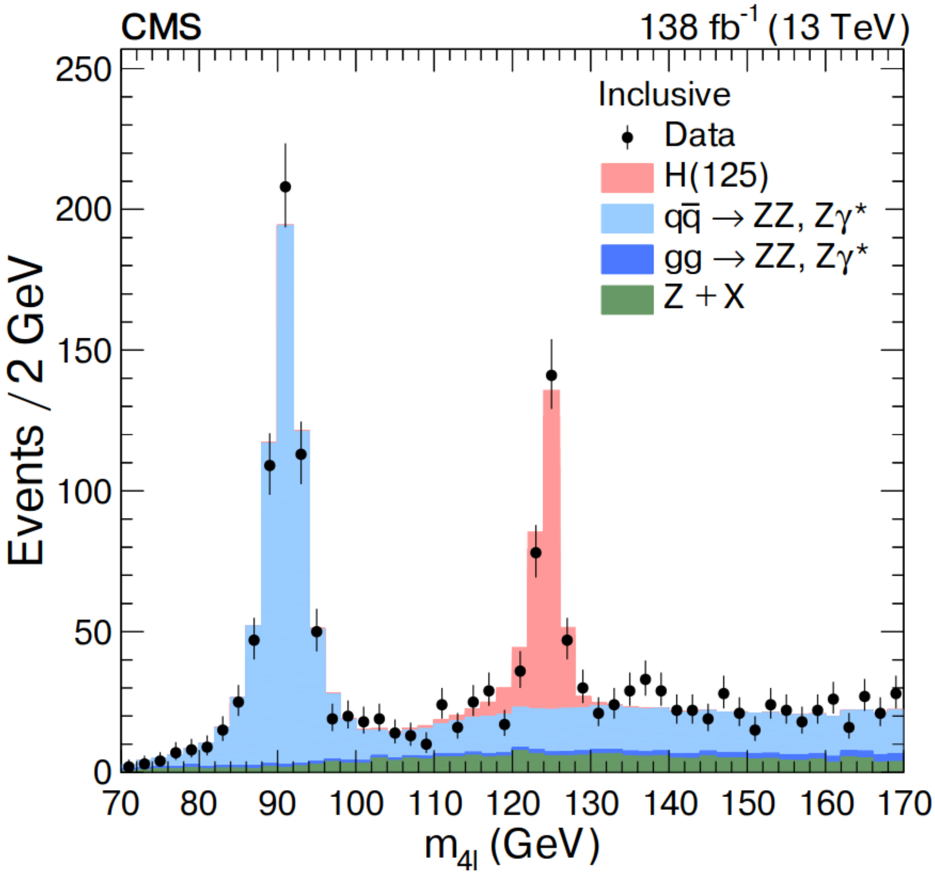

Figure 1.15: Observed (points) and predicted (stacked histograms) m_{4l} distributions in the inclusive final states. The predictions for the Higgs boson signal and the three main backgrounds are given by the different colors [28].

$\mathbf{H \to W^+W^- \to \ell^+\nu\ell^-\bar{\nu}}$ decay

The $H \to W^+W^- \to \ell^+\nu\ell^-\bar{\nu}$ channel is experimentally challenging due to the presence of neutrinos in the final state. Searches focus on events with two oppositely charged leptons (electrons or muons) accompanied by missing transverse energy and/or jets [29].

Events are tipically categorized by lepton flavor combination (e^+e^-, $\mu^+\mu^-$, and $e^\pm\mu^\mp$) and jet multiplicity ($N_{\text{jet}} = 0, 1, \geq 2$).

There are many different backgrounds, depending on the category. For opposite-flavor leptons with no jets, the dominant background is non-resonant WW production, while same-flavor leptons have large Drell-Yan contamination. On the other hand, $t\bar{t}$, tW, and $W+\text{jets}$ (where jets can be misidentified as leptons) contribute across all categories. Smaller contributions arise from non-resonant WZ, ZZ, and $W\gamma^*$ processes. Selecting events with large missing transverse energy (E_T^{miss}) can help reducing Drell-Yan and multijet backgrounds, as well as vetoing events with dilepton invariant mass near the Z-boson mass (in same-flavor categories) and events with identified b-jets, (to reduce $t\bar{t}$). Requiring very tight isolation for the leptons, on the other hand, reduce the $W+\text{jets}$ background from jets mis-identified as leptons. Due to the scalar nature of the Higgs boson and the $V-A$ structure of the W-boson decay, the two charged leptons tend to be emitted with small angular separation. Thus, the dilepton invariant mass $m_{\ell\ell}$ and the azimuthal angle difference $\Delta\phi_{\ell\ell}$ are used to discriminate signal from background. The transverse mass, defined as

$$m_T = \sqrt{2p_T^{\ell\ell}E_T^{\text{miss}}(1-\cos\Delta\phi_{E_T^{\text{miss}},p_T^{\ell\ell}})},$$

where p_T^{ll} is the transverse momentum of the dilepton system, works as an good proxy to suppress backgrounds, but has relatively poor mass resolution, as shown in Figure 1.16. In this channel background rates are estimated from data-driven control samples with floating normalization parameters.

$\mathbf{H} \to \tau^+\tau^-$ decay

In the search for $H \to \tau^+\tau^-$ decays, tau leptons decaying into electrons, muons , and hadrons are considered [30]. The reconstruction of the $\tau^+\tau^-$ invariant mass ($m_{\tau\tau}$) relies on a kinematic fit that incorporates both the visible decay products and the missing transverse energy in the event.

In contrast to the WW channel, the $\tau\tau$ channel allows for better mass reconstruction techniques despite the fact that the final states have the same number of neutrinos, because the $\tau\tau$ decay kinematics can be more tightly constrained. In particular, due to the relatively low mass of the τ lepton ($m_\tau = 1.777$ GeV), the visible decay products tend to be nearly collinear with the neutrinos when the τ leptons are highly boosted, a common occurrence in Higgs decays, enabling the use of the so-called collinear approximation, or more advanced likelihood-based techniques such as the Missing Mass Calculator (MMC) [31] and SVfit [32, 33], which scan over the unknown neutrino momenta to provide an estimator of the invariant mass of the $\tau\tau$ system. These methods significantly improve the mass resolution compared to the transverse mass used in the WW channel, but it is nonetheless limited to approximately 10–15%, necessitating a search for a broad excess over background in the $m_{\tau\tau}$ distribution, as shown in Figure 1.17.

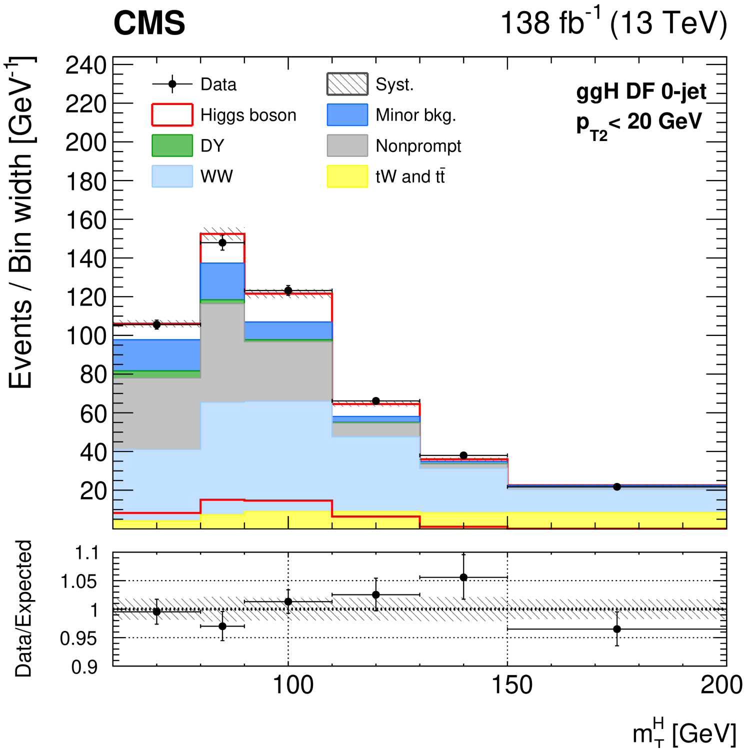

Figure 1.16: Observed distributions of the Higgs boson candidated transverse mass in an example category with 0 reconstructed jets of the CMS Run2 measurement in the $H \to WW$ channel. [29].

The dominant backgrounds in this channel include Drell–Yan processes ($Z \to \tau^+\tau^-$, $Z \to e^+e^-$), W+jets, top-quark pair production ($t\bar{t}$), and multijet events. Events are tipically categorized based on the number and kinematic properties of additional jets, the reconstructed Higgs transverse momentum, and the angular separation ($\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$) between the tau candidates. Categories enriched in vector boson fusion (VBF) events, characterized by the presence of two energetic forward jets with large pseudorapidity separation, typically offer the best signal-to-background ratios and sensitivity.

Both ATLAS and CMS observed Higgs boson decays in ditaus with more than 5 σ [30], and performed precision cross-sections measurements, as well as and CP-odd couplings searches in this channel [34] [35].

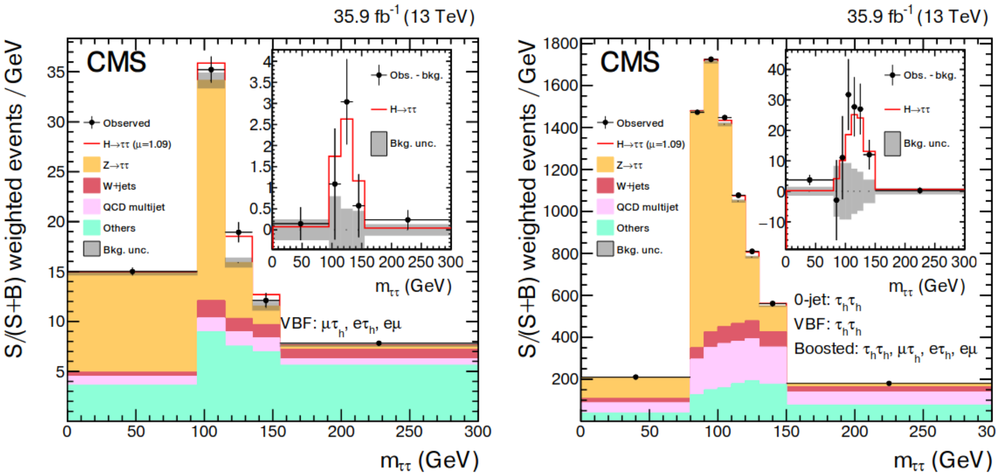

Figure 1.17: Combined observed and predicted $m_{\tau\tau}$ distributions for the $H \to \tau\tau$ channel. The left pane includes the VBF category of the muon-hadronic, electron-hadronic and muon-electron channels, and the right pane includes all other channels [30].

$\mathbf{H \to b\bar{b}}$ decay

The $b\bar{b}$ decay channel suffers from huge QCD backgrounds, therefore the search for $H \to b\bar{b}$ decays is most sensitive in production modes that include an associated vector boson (VH), which facilitates triggering and suppresses QCD backgrounds. W bosons are reconstructed via their leptonic decays ($W \to \ell\nu$, with $\ell = e, \mu, \tau$), and Z bosons via $Z \to e^+e^-$, $\mu^+\mu^-$, or $\nu\bar{\nu}$. The Higgs candidate is reconstructed from two b-tagged jets, and the dijet invariant mass ($m_{b\bar{b}}$) provides the primary discriminating observable. either used as a fit variable, or as an input to a multi-variate discriminant. Backgrounds include V+jets with both light and heavy flavors, $t\bar{t}$, dibosons, and multijet events.

At the LHC, ATLAS and CMS employ tipically a boosted VH topology to suppress background, focusing on events where the vector boson and Higgs candidate have large transverse momentum and are back-to-back in azimuth.

Events are divided into $p_T(V)$ bins, and MVA classifiers are trained on various kinematic and b-tagging features. With Run2 data, this decay channel has been observed with more than 5 σ [22] [36], as shown in Figure 1.18.

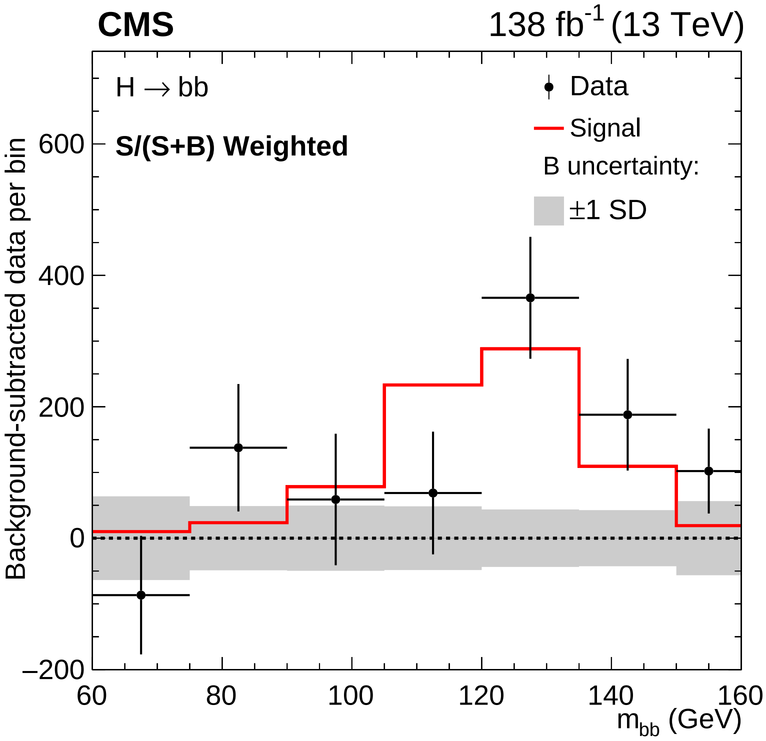

Figure 1.18: The background-subtracted diparticle invariant mass distribution targeting the study of the decay channel $H \to bb$ [22].

Chapter 2

Higgs boson coupling and width

2.1 Coupling measurements

With increasing dataset sizes, measurements in the most sensitive channels are now performed also in exclusive categories defined by distinct production and decay characteristics, that allow extracting the Higgs boson couplings to various particles. Indeed, differential cross-section measurements are performed in combinations of production processes and decay channels, as beyond SM effects are expected to be more sizeable in particular kinematical regions.

Most experimental categories are dominated by signal events from a single Higgs decay mode but generally include a mixture of production processes. For example, a typical VBF selection requires the presence of two energetic jets (with transverse momentum $p_T \geq 30\,\text{GeV}$) featuring a large dijet invariant mass ($m_{jj} \geq 400\,\text{GeV}$) and a substantial rapidity gap ($\Delta\eta_{jj} \geq 3.5$). Although such a category is enriched in VBF-produced Higgs bosons, significant contamination from ggH and VH processes remains. Consequently, the measured signal rate within the VBF category does not correspond directly to the pure VBF production cross section, since the ggH and VH contributions cannot be fully disentangled. Detailed simulations are employed to estimate the relative contributions of the various Higgs production modes within each event category, which are disentangled by means of simultaneous fits.

For each decay mode, multiple exclusive sub-channels are defined to target different production mechanisms and kinematic regimes. These sub-channels are then statistically combined to maximize sensitivity to overall couplings to all SM particles. A natural extension of this strategy is to perform a global combination across all decay channels, enhancing sensitivity to production and decay modes, as well as allowing precise measurements of the coupling strengths.

At a hadron collider like the LHC, it is not possible to measure the total Higgs production cross section independently of its decay branching ratios. Consequently, the total Higgs width cannot be directly determined from cross-section measurements only. However, a combined analysis of many exclusive categories—each with different sensitivities to production and decay processes—allows for detailed measurements of the Higgs boson's production rates, decay rates, and couplings, albeit under certain assumptions.

A key assumption in any combination is that all channels refer to the same Higgs boson mass hypothesis. This is ensured within individual experiments, and was enforced across experiments during Run 1 by first performing a combined ATLAS+CMS mass measurement, yielding 125.09 ± 0.24 GeV [37], before combining signal strength (measured cross-section divided by its SM value) results.

The statistical combination of all Higgs boson channels, is performed via a simultaneous fit of signal and background models to data across hundreds of exclusive categories. In the ATLAS+CMS Run1 combination [38], around 600 categories were used, spanning different center-of-mass energies, decay modes, control regions, and phase space regions. Combination of channels have been performed also for Run 2 by the single experiments [8] [39].

The core of the combination is a formula which describes the expected number of signal events n_s^c in category c as:

$$n_s^c = \left(\sum_{i,f} \mu_i \, \sigma_i^{\text{SM}} \cdot A_{if}^c \cdot \epsilon_{if}^c \cdot \mu_f \, \text{BR}_f^{\text{SM}} \right) \cdot \mathcal{L}^c, \tag{2.1}$$

where:

- i and f index the production ($i \in \{\text{ggH}, \text{VBF}, ZH, WH, t\bar{t}H, tH, \dots\}$) and decay modes ($f \in \{\gamma\gamma, WW, ZZ, b\bar{b}, \tau\tau, \mu\mu, Z\gamma\}$),
- σ_i^{SM} and BR_f^{SM} are the SM predictions for the production cross section and decay branching ratio,
- A_{if}^c and ϵ_{if}^c are the acceptance and efficiency in category c for a given (i, f) process,

- $\mathcal{L}^c$ is the integrated luminosity in that category,
- μ_i and μ_f are signal strength modifiers for production and decay, respectively [4].

The so-called κ framework, described in detail in [40, 19], facilitates the extraction of Higgs coupling properties in terms of a series of Higgs coupling strength modifier parameters κ_i, defined as the ratios of the couplings of the Higgs boson to each particle i relative to their corresponding Standard Model (SM) values. The κ framework assumes a single narrow resonance such that the zero-width approximation can be used to decompose the cross section as a product of two factors characterising the production and decay of the Higgs boson. The κ parameters are introduced by expressing each of these factors as their SM expectation multiplied by the square of a coupling strength modifier for the corresponding process at leading order:

$$(\sigma \cdot \mathrm{BR})(i \to H \to f) = \sigma_i^{\mathrm{SM}} \kappa_i^2 \cdot \frac{\Gamma_f^{\mathrm{SM}} \kappa_f^2}{\Gamma_H^{\mathrm{SM}} \kappa_H^2} \quad \to \quad \mu_{i,f} \equiv \frac{\sigma \cdot \mathrm{BR}}{\sigma^{\mathrm{SM}} \cdot \mathrm{BR}^{\mathrm{SM}}} = \frac{\kappa_i^2 \kappa_f^2}{\kappa_H^2}, \tag{2.2}$$

where μ_i^f is the rate relative to the SM expectation and κ_H^2 adjusts the SM Higgs width to account for modifications induced by the altered Higgs boson couplings, i.e.

$$\kappa_H^2 = \sum_f \mathrm{BR}_f^{\mathrm{SM}} \, \kappa_f^2.$$

When all $\kappa_i = 1$, the SM is exactly reproduced [4]. The above formula also shows the degeneracy between couplings and total decay width, if both are constrained only from cross-section measurements. For loop-induced processes, e.g. $H \to \gamma\gamma$, one may either resolve the coupling strength modifier in terms of its SM constituents, i.e., $\kappa_\gamma(\kappa_t, \kappa_W)$, or treat κ_γ as an effective coupling strength parameter.

In this formalism, several assumptions are explicitly made:

1. The observed signals in different search channels originate from a single narrow resonance with defined mass
2. The narrow width approximation is assumed to allow factorisation of signal yields into production and decay components.
3. The tensor structure of the couplings is assumed identical to that of the SM Higgs boson, implying the observed state is a CP-even scalar as in the SM.

Global fits are performed expressing the μ_i and μ_f parameters in terms of a limited set of κ_j parameters or their ratios, under various assumptions.

Moreover, measuring the Higgs boson couplings to Standard Model (SM) fields requires making certain assumptions about either the total Higgs boson width (forcing it to be equal to the SM one), the Higgs coupling to vector bosons, or applying a direct experimental constraint on the Higgs boson width. More specifically, a degeneracy arises because at the LHC the measured observables are only production cross sections multiplied by branching fractions, as shown in Equation 2.1, resulting in an intrinsic degeneracy between all coupling modifiers κ_i and the total width modifier κ_H.

The Higgs boson branching fractions are generally divided into visible, invisible, and undetected final states. The 'visible' fraction corresponds to decays into SM particles which can be detected experimentally, while the 'invisible' fraction (which is non-zero even in the SM because of the $H \to ZZ \to 4\nu$ decay with 0.12% BR) is constrained by direct invisible decay searches (using particular topology and missing energy tagging) [21], and the 'undetected' fraction, represents exotic decay modes escaping current detection [4].

As it can be noted, in order to maintain agreement with the SM values for all couplings, any deviation in the total Higgs width κ_H must be exactly compensated by a corresponding scaling of all κ_i. This is known as a "flat direction," where all couplings scale by a factor α and the total width scales by α^2, because the signal in any channel is proportional to $\kappa_i \kappa_f / \Gamma_H$. A concrete model realizing such a flat direction is presented in [41], which introduces a scalar field that decays into light hadrons, rendering its decays 'undetected' by current experiments. This scalar cannot be 'invisible,' as that scenario is constrained by direct searches for invisible Higgs decays using VBF or $t\bar{t}H$ topologies and missing energy. Moreover, this model requires Higgs couplings to SM fields to be larger than in the SM and the Higgs coupling to the new scalar to maintain the observed rates constant, a situation considered exotic and finely tuned.

As already said, this flat direction can be broken in two main ways. First, by imposing $\kappa_V \leq 1$, which naturally constrains κ_H through processes like VBF production with decays into two vector bosons, whose rates scale as κ_V^4/κ_H^2. Second, the Higgs width constraint can be experimentally derived from off-shell Higgs production measurements, or, **as it will be explained in this thesis, by on-shell production bounds with indirect techniques**. BSM scenarios with $\kappa_V > 1$ are considered exotic, because the

Higgs field would no longer cancel the divergences in the vector boson scattering amplitudes correctly. This typically would lead to a breakdown of perturbative unitarity at relatively low energies, unless some new physics (such as additional Higgs bosons or other particles) enters to restore unitarity. The coupling measurements with three possible assumptions, i.e. SM width (BSM branching ratio equal to 0), $k_V < 1$ and experimentally constrained width (through off-shell measurements), are shown in Figure 2.2 for the case of CMS Run 2 combination of all differential cross-sections in all the measured channels [8]. For all the couplings, is it possible to show their strength as a function of the mass of the SM particle the Higgs boson couples, assuming the expected relation and showing graphically the striking agreement between data and SM in this sector (see Figure 2.1).

2.2 Width measurements

Approaches to determine the Higgs boson width Γ_H, equal to 4.1 MeV in the SM [19], are commonly classified as either direct or indirect, according to whether the method measures the lineshape of the resonance or instead relies on alternative observables. Indirect methods, in the Higgs boson physics at LHC, are further categorized depending on the phase-space where the measurement is performed. If the kinematic region is the one around the Higgs boson resonant production, it is on-shell, while it is called off-shell otherwise.

2.2.1 Direct methods

Neglecting detector resolution effects, the invariant mass distribution of an unstable particle follows a relativistic Breit-Wigner distribution, defined as:

$$f(m_f,\, m,\, \Gamma) \propto \frac{1}{(m_f^2 - m^2)^2 + m^2\Gamma^2} \tag{2.3}$$

where m_f is the invariant mass of the final state, m is the pole mass of the particle and Γ is its decay width, related to the lifetime τ by $\Gamma = \hbar/\tau$. The relativistic Breit-Wigner distribution derives from the squared module of the kinematic part of the propagator for unstable particles in the momentum space:

$$\Delta_F(p) = \frac{i}{q^2 - m^2 + im\Gamma} \tag{2.4}$$

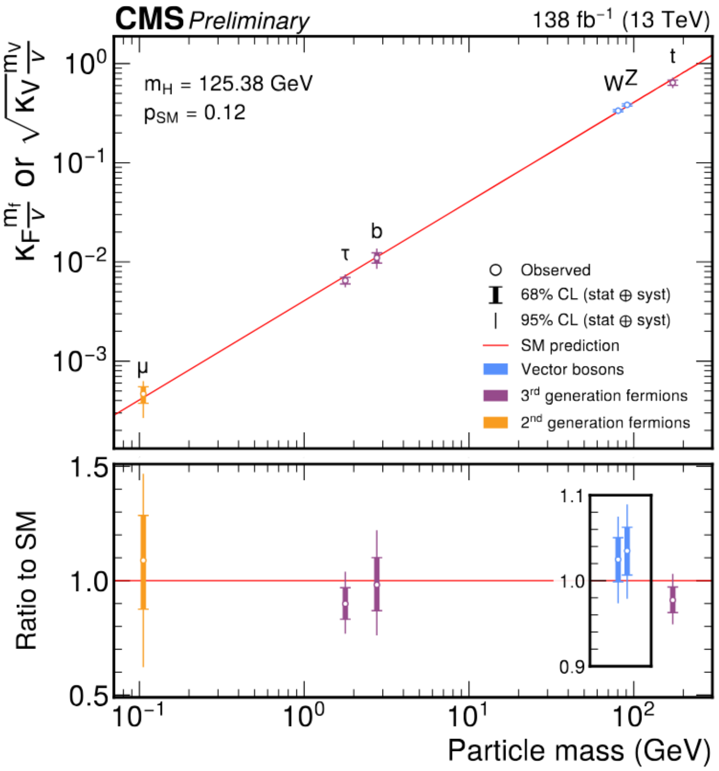

Figure 2.1: The measured coupling modifiers of the Higgs boson to fermions and gauge bosons, multiplied by fermion masses, of, for vector bosons, for the square root of their masses, as a function of the masses themselves [42]. This linear proportionality is predicted by the SM, as shown in Equation 1.20 and 1.19.

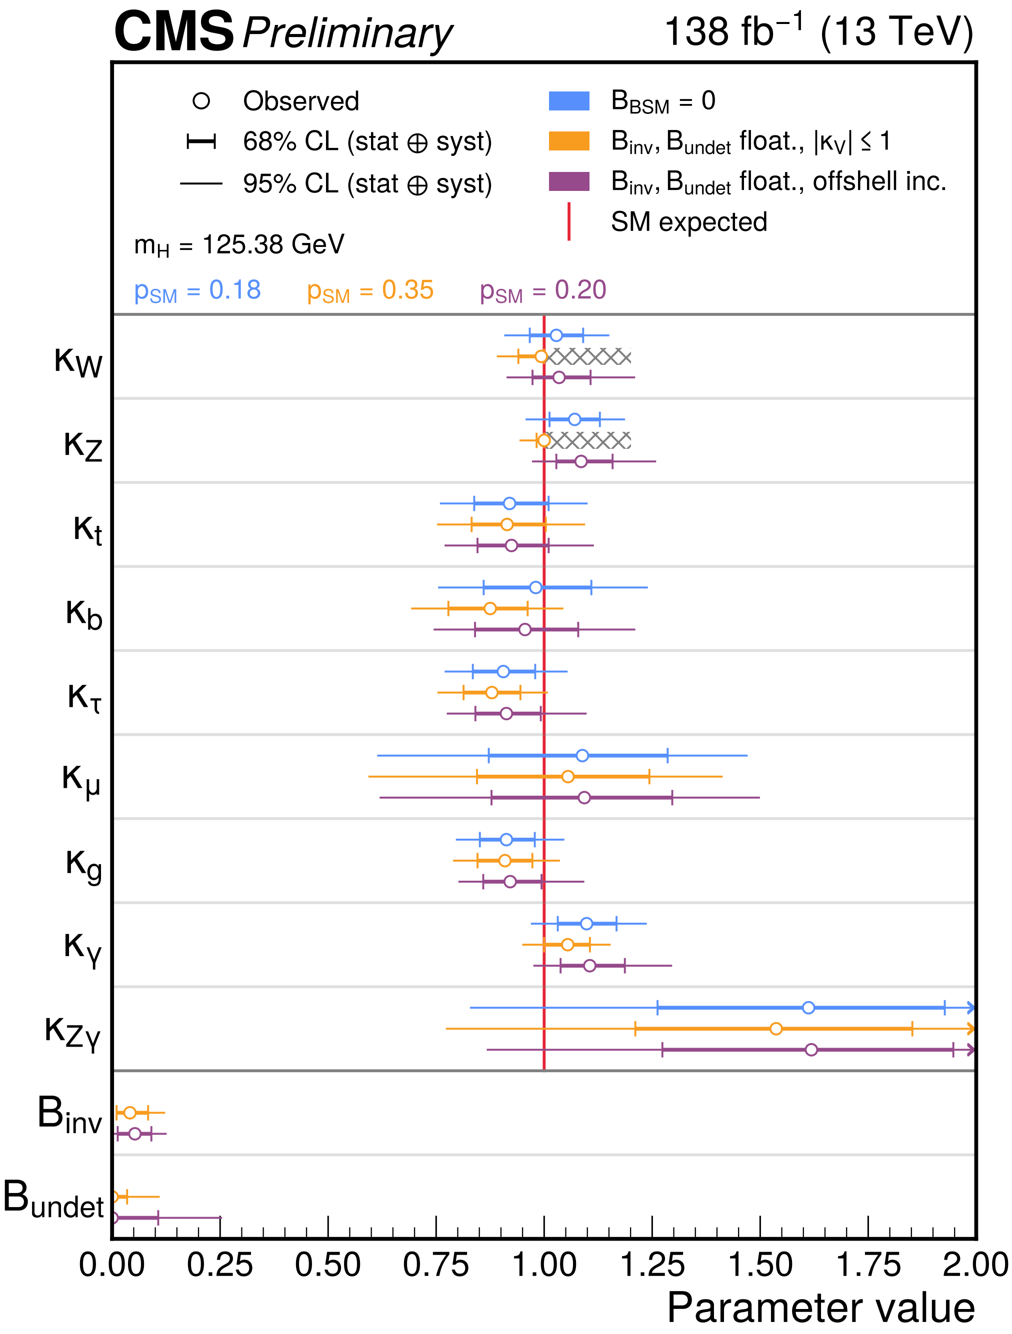

Figure 2.2: The measured Higgs boson coupling modifiers in the κ framework. The best-fit values, 68%, and 95% CL intervals are shown for each of the models considered. The results where no additional BSM contributions to the Higgs boson total decay width are shown in blue. The results for the model which introduces BSM contributions but places an external constraint (hatched boxes) on vector bosons couplings, namely $k_V \leq 1$, are shown in orange. The results for the fit in which the width measurements from off-shell production are included are shown in purple [42].

where q is the four momentum of the unstable particle. If the distribution of m_f is experimentally observed, the width of a resonant peak will be given by the total decay width of the particle, convoluted with the detector resolution of the final state particles.

The most intuitive method to measure the decay width of the Higgs boson requires therefore fitting the invariant mass distribution of the Higgs boson decay products, correcting for the experimental resolution, and extracting the Breit-Wigner width.

The measurement following this approach has been performed by CMS in the $H \to ZZ^* \to 4l$ channel with the full Run 2 data set [28].

The employed signal models are a convolution of double Crystall Ball (DCB) functions to model the experimental resolution [43], i.e. distributions with a Gaussian core and power low tails, with a relativistic Breit-Wigner with a width Γ_H, as shown in Figure 2.3.

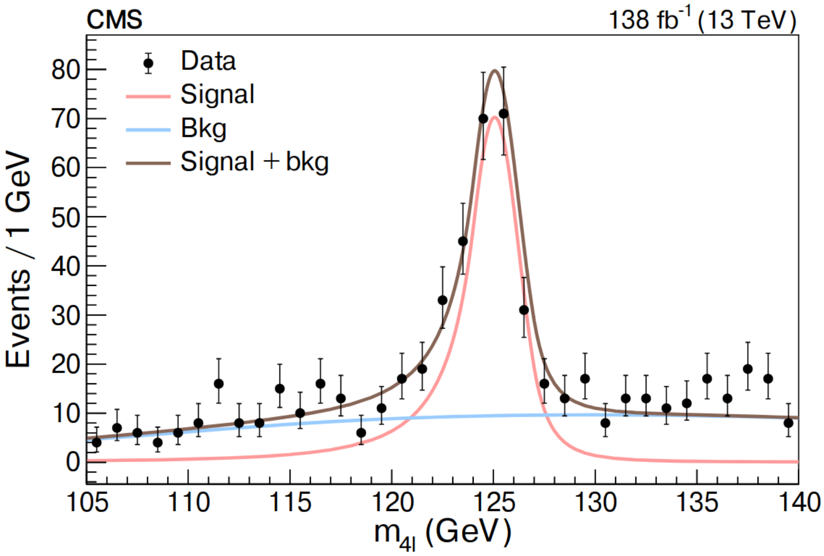

Figure 2.3: Illustration of how the on-shell analytical model is constructed, combining the m_{4l} distributions from all data-taking years and all final states for the $H \to ZZ^* \to 4l$ channel. The red, blue, and brown lines show the results of the fit to the signal, background, and their sum, respectively. The solid black points with vertical bars show the data and the associated statistical uncertainties [28]

Because the predicted Higgs boson width is extremely small — effectively setting a strict lower bound — the Feldman–Cousins procedure [44] is employed to determine confidence intervals and limits. Confidence levels (CLs) are calculated for a range of width hypotheses using distributions derived from simulated pseudo-experiments. The resulting observed (expected) upper limits on Γ_H are 50 (320) MeV at 68% CL and 330 (640) MeV at 95% CL. Figure 2.4 presents the corresponding $1 - \mathrm{CL}$ distribution as a function of Γ_H. The measurement precision is primarily limited by statistical uncertainty ($\sigma_{stat} \propto \sigma_{mass}/\sqrt{N_{signal}}$), with a smaller contribution from systematic effects (with the current data statistics), the largest of which arises from the lepton momentum resolution.

An alternative approach to measuring the Higgs boson width involves determining its decay lifetime. The average lifetime of Higgs boson candidates, denoted $\langle \Delta t \rangle$, can be expressed as

$$\langle \Delta t \rangle = \tau_H \equiv \frac{1}{\Gamma_H}, \tag{2.5}$$

where Δt represents the lifetime of an individual Higgs boson candidate. Because the lifetime and total decay width are inversely related, measuring the lifetime directly provides information about the width.

The Higgs boson lifetime can be inferred by reconstructing the decay particles' flight distance within the detector. The distribution of measured lifetimes depends on the invariant mass of the decay products and the transverse momentum of the reconstructed Higgs boson. The Standard Model predicts an expected lifetime of approximately $\tau_H \approx 48\,\mathrm{fm}/c$.

CMS performed an analysis of this type using Run 1 data in the $H \to ZZ^* \to 4l$ channel. The distribution of the reconstructed flight distances, $c\Delta t$, is displayed in Figure 2.5, where the expected Standard Model signal is indicated in red. The resulting upper limit on the Higgs boson lifetime was set at

$$\tau_H < 1.9 \times 10^{-13}\,\mathrm{s} = 190\,\mathrm{fs}$$

at the 95% confidence level, which is roughly nine orders of magnitude larger than the Standard Model prediction $\tau_H^{\mathrm{SM}} = 1.6 \times 10^{-22}\,\mathrm{s}$ [45]. This upper limit on the lifetime corresponds to a lower bound on the total decay width of

$$\Gamma_H > 3.5 \times 10^{-9}\,\mathrm{MeV}.$$

A fundamental limitation of this technique arises from the fact that the predicted flight distance of the Higgs boson is far too small to be resolved

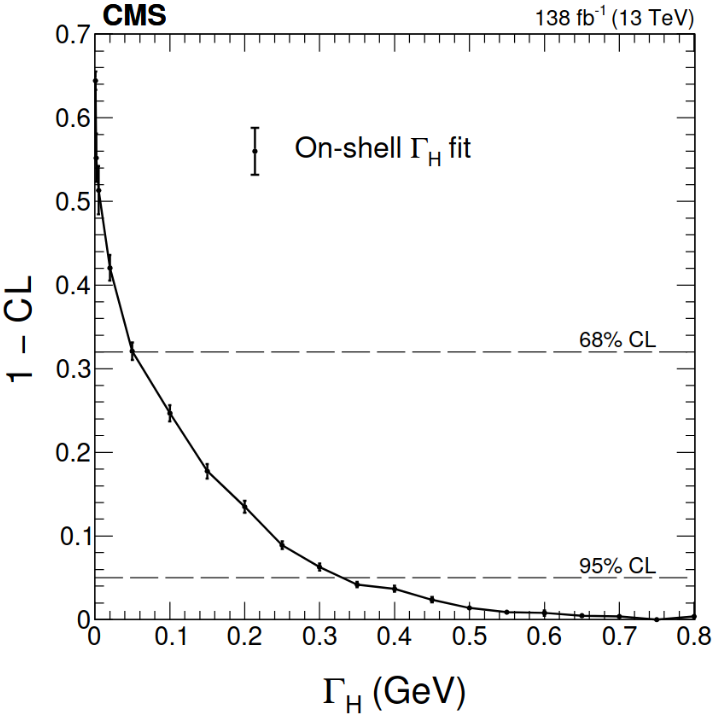

Figure 2.4: Distribution of (1 - CL) vs. Γ_H from the fit in the measurement of the Higgs boson width using on-shell production in the $H \to ZZ^* \to 4l$ channel [28]. The CL values shown by the points are extracted using the Feldman-Cousins approach. The vertical bars on the points represent the spread of the simulated pseudo-experiment values. The 68 and 95% CL values are shown by the dashed horizontal lines.

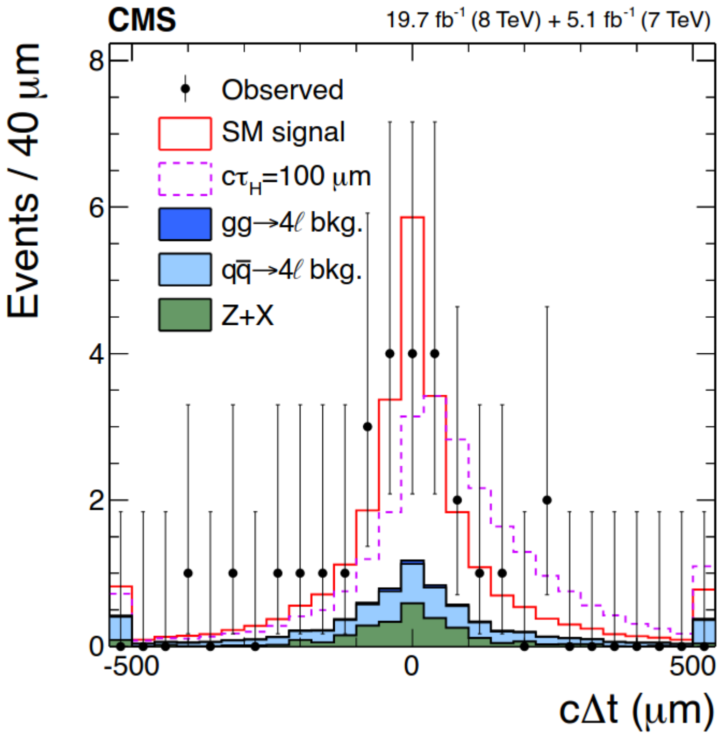

Figure 2.5: The $c\delta t$ distribution used to place an upper limit of $\tau_H < 1.910^{-13}$ s on the lifetime of the Higgs boson at the 95% confidence level [45]. The predicted SM signal is indicated by the red line. A possible alternate non-SM signal with $c\tau_H = 100$ μm is shown by the dashed purple line.

with the current vertex displacement resolution of the ATLAS and CMS detectors, which is on the order of $10\,\mu$m. This lack of sensitivity has prevented subsequent lifetime-based measurements of the Higgs boson.

Due to these challenges, namely relatively large experimental resolution in the invariant mass and flight distance, direct methods have historically been unable to measure the Higgs width with precision close to the Standard Model expectation.

In order to overcome the difficulties, in 2013, two innovative indirect approaches were proposed [46, 47]:

- The *off-shell* method, which compares the ratio of Higgs boson production cross-sections in the off-shell and on-shell invariant mass regions, typically using $H^* \to ZZ$ or $H^* \to WW$ decays.

- The *interferometry* technique, which analyzes distortions in the on-shell invariant mass distribution of the Higgs boson caused by interference effects with non-resonant background processes, notably in the $H \to \gamma\gamma$ channel. While interference effects also occur in the $H \to ZZ^*$ channel, their impact on the measured Higgs mass is minor (less than 10 MeV) provided the total width is close to the Standard Model value. By contrast, interference in the diphoton channel is expected to produce more pronounced effects because of the different background yield. The measurement of this interference is the subject of this thesis and will be discussed in detail later.

2.2.2 Indirect off-shell methods

The off-shell method provides a way to extract the Higgs boson width by studying the decay channels $H^* \to ZZ$ and $H^* \to WW$. Here the focus will be on the $H^* \to ZZ$ channel, as it offers the greatest sensitivity because of the clean four-lepton final state. This final state is experimentally preferrable compared to the $H^* \to WW$ channel, which involves neutrino, but conceptually the off-shell technique applies equally to both channels and, indeed, the measurement performed by CMS in this channel yielded a value of $\Gamma_H = 3.9^{+2.7}_{-2.2}$ MeV. Although not discussed further, the measurement was performed also in the $H^* \to ZZ \to \ell\ell\nu\nu$ decay channel [28].

Figure 2.6(left) illustrates the leading-order Feynman diagram for the gluon fusion production of the Higgs boson followed by its decay into two

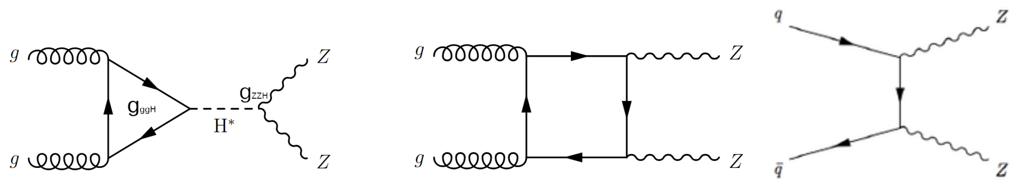

Figure 2.6: Feynman diagrams for the (left) $gg \to H^* \to ZZ$ non-resonant signal process, (center) the $gg \to ZZ$, and (right) $q\bar{q} \to ZZ$ background processes. The couplings g_{ggH} and g_{HZZ} correspond, respectively, to the effective gluon-Higgs vertex via quark loop and to the Z-Higgs coupling.

Z bosons ($gg \to H^* \to ZZ$). In this process, the gluon fusion production proceeds through loop diagrams involving mainly top quarks. Because the Higgs boson coupling scales with mass, the dominant fermion loops typically involve top (or bottom) quarks. The quark loop can be approximated by an effective vertex characterized by a coupling strength g_{ggH}. The coupling governing the decay of the Higgs boson to two Z bosons is denoted g_{HZZ}.

One of the backgrounds to this signal arises from the continuum process $gg \to ZZ$, depicted in 2.6(right). This background has a production cross-section approximately six times larger than the $H^* \to ZZ$ signal process, and interferes with the signal itself, making its precise understanding crucial for the off-shell width measurement, while the main background in terms of yield is actually $q\bar{q} \to ZZ$.

The differential cross-section in the ZZ (or ZZ^* if the Higgs boson is on-shell) invariant mass for the full $gg \to H \to ZZ$ process can be expressed as:

$$\frac{d\sigma(gg \to H \to ZZ)}{dm_{ZZ}} = \frac{g^2_{ggH} g^2_{HZZ}}{(m^2_{ZZ} - m^2_H)^2 + m^2_H \Gamma^2_H} f(m_{ZZ}, m_H) \tag{2.6}$$

where $f(m_{ZZ}, m_H)$ contains the residual dependence on m_{ZZ} of the product of the Higgs boson production cross-section and of the branching ratio $BR(H \to ZZ)$. In the on-shell case, thanks to the narrow width approximation, $f(m_{ZZ}, m_H)$ is constant and the differential cross-section is a pure relativistic Breit-Wigner. In the off-shell case not only the production cross-section evolves with m_{ZZ}, but also the branching ratio increases largely, above twice the mass of the Z boson, because the available phase-space becomes larger with the opening of the decay channel to two on-shell Z bosons. A similar effect happens at the production side with the $t\bar{t}$ threshold, because of the gluon-gluon fusion loop, which gets enhanced if $m_{ZZ} > 2m_t$, since the top quarks can be produced on-shell, removing the suppression factor associated with the off-shell production in the $m_H < 2m_t \sim 350$ GeV regime.

If the Higgs boson signal differential cross-section is integrated in the two regions, i.e. on-shell and off-shell, the numerical terms $f(m_{ZZ}, m_H)$ can be re-absorbed, making clearer the dependence of the rates on the couplings and the width:

$$\sigma_{on-shell} \propto \frac{g_{ggH}^2 g_{HZZ}^2}{\Gamma_H},\ \sigma_{off-shell} \propto g_{ggH}^2 g_{HZZ}^2 \tag{2.7}$$

where the presence of Γ_H at the $\sigma_{on-shell}$ denominator is due to the integral over the Breit-Wigner function.

If one considers the ratio of these two cross-sections, the coupling factors ideally cancels out, enabling a direct determination of the width. This cancellation holds true only if the couplings g_{ggH} and g_{HZZ} remain the same across the two energy regimes under consideration, but, on the contrary, the strong and electroweak couplings vary with the energy scale at which they are evaluated, i.e. $\sim$ 125 GeV on-shell, and, in the case of the ZZ measurement, 180-500 GeV. To address this, coupling terms are expressed via coupling modifiers κ, by normalizing the coupling g_X to the Standard Model expectation g_X^{SM}, such that $\kappa_X = g_X/g_X^{SM}$ for any given Higgs coupling X.

By normalizing each cross-section to its Standard Model predictions, the couplings g_X cancel out, leaving only the coupling modifiers and additional terms. This leads to the definition of the signal strength parameter μ, which for the on-shell region can be written as:

$$\mu_{\text{on-shell}} = \left.\frac{\sigma_{pp\to H\to ZZ}}{\sigma^{SM}_{pp\to H\to ZZ}}\right|_{\text{on-shell}} = \frac{\kappa_g^{2,\,\text{on-shell}} \cdot \kappa_Z^{2,\,\text{on-shell}}}{\Gamma_H/\Gamma_H^{SM}}, \tag{2.8}$$

while in the off-shell region it becomes:

$$\mu_{\text{off-shell}} = \left.\frac{\sigma_{pp\to H\to ZZ}}{\sigma^{SM}_{pp\to H\to ZZ}}\right|_{\text{off-shell}} = \kappa_g^{2,\,\text{off-shell}} \cdot \kappa_Z^{2,\,\text{off-shell}}. \tag{2.9}$$

This results in the relation:

$$\frac{\mu_{\text{off-shell}}}{\mu_{\text{on-shell}}} = \frac{\Gamma_H}{\Gamma_H^{SM}}. \tag{2.10}$$

It is important to note that such cancellation of coupling modifiers only occurs if $\kappa_{g,\text{off-shell}} = \kappa_{g,\text{on-shell}}$ and $\kappa_{Z,\text{off-shell}} = \kappa_{Z,\text{on-shell}}$. These conditions imply that the ratios of on-shell and off-shell couplings match the Standard Model predictions exactly and any deviation is to attributed to deviations in the width. Consequently, this approach restricts the analysis to SM-like couplings and reduces sensitivity to potential Beyond Standard Model (BSM) physics effects. Given that current LHC measurements are consistent with SM expectations for cross-sections, this assumption is reasonable; however, it is important to remember that this underlying assumption is intrinsic in this off-shell technique. Despite the theoretical importance of this assumption, analyses have been performed while allowing also some kind of off-shell BSM contributions, and the resulting bounds are indeed relaxed with respect to the nominal ones, but the sensitivity is only impacted by around a factor 2 [28].

There are three sources of events in the off-shell region: one is off-shell production of the Higgs boson (signal), the second are the background $gg, q\bar{q} \rightarrow ZZ$, and the third is the interference between signal and $gg \rightarrow ZZ$ background. The presence of such interference results in a deficit of continuum ZZ background events, which overall increases the sensitivity of the analysis. The differential cross-section for the combined signal (S), background (B) and interference (I) contributions is given by:

$$\frac{d\sigma_{incl.}}{dm_{ZZ}} = \mu_{off-shell}\frac{d\sigma_{sig}}{dm_{ZZ}} + \sqrt{\mu_{off-shell}}\frac{d\sigma_{int}}{dm_{ZZ}} + \frac{d\sigma_{bkg}}{dm_{ZZ}} \tag{2.11}$$

where it can be noted that the interference term scales as a function of $\sqrt{\mu_{off-shell}}$. Note also that in the SM scenario ($\mu_{off-shell} = 1$), the contribution to the cross-section due to the interference is larger compared to the signal, in the off-shell region, because it is multiplied by the larger background cross-section and it not suppressed by the narrow Higgs boson width, as it would be in the on-shell case.

Figure 2.7 shows the differential cross-section in a wide range including both on-shell and off-shell regions.

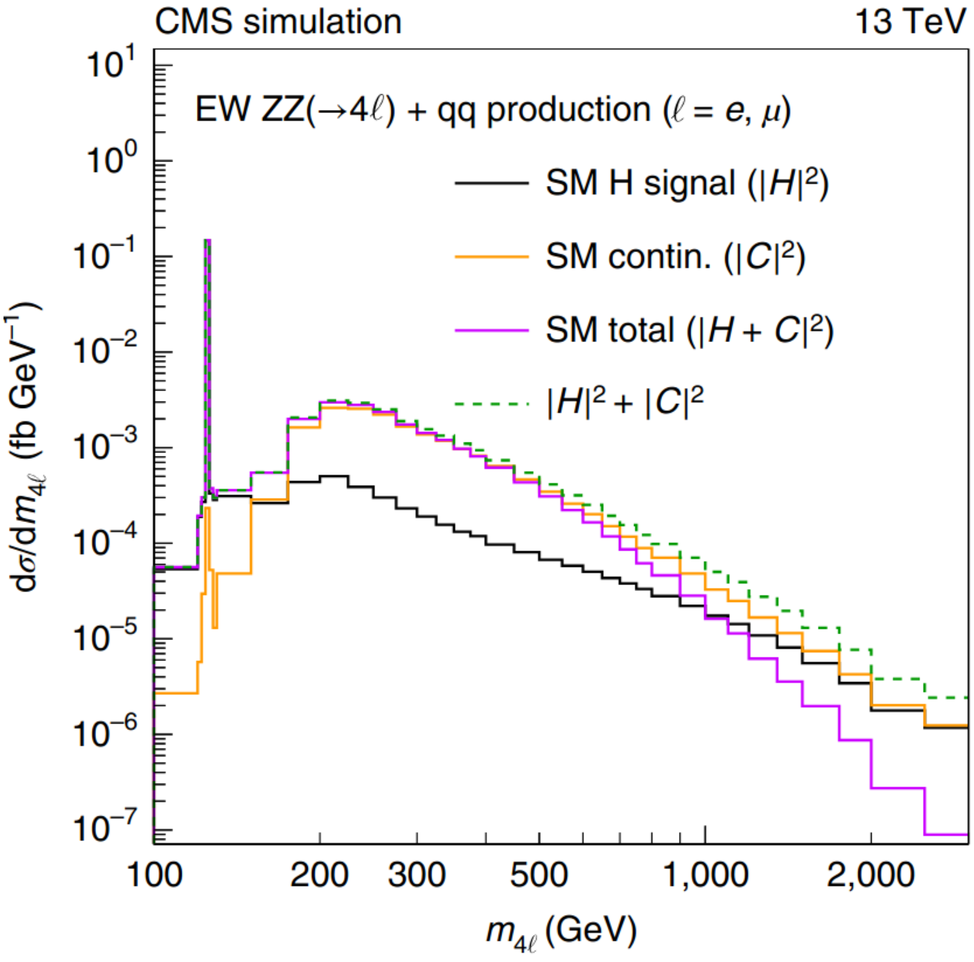

Figure 2.7: Distributions for the $4l$ invariant mass. The dashed green curve represents the direct sum without interference of Higgs boson production and continuum amplitudes, and the solid magenta curve represents the sum with interference included. Note that the interference is destructive, and its magnitude grows as the mass increases [28].

Both ATLAS [48] and CMS [28] performed such measurements of Γ_H. In the CMS case, focusing on the $4l$ case, after pre-selection of events, two kinematical discriminant able to distinguish between signal and background and to tag interference is evaluated event-by-event, and the 3D distribution of the four-lepton mass and this two discriminants is fitted to extract Γ_H. In Figure 2.8 the $4l$ mass distribution for a large region in the discriminants plane enhancing the Higgs boson signal, together with the likelihood scan for Γ_H are shown.

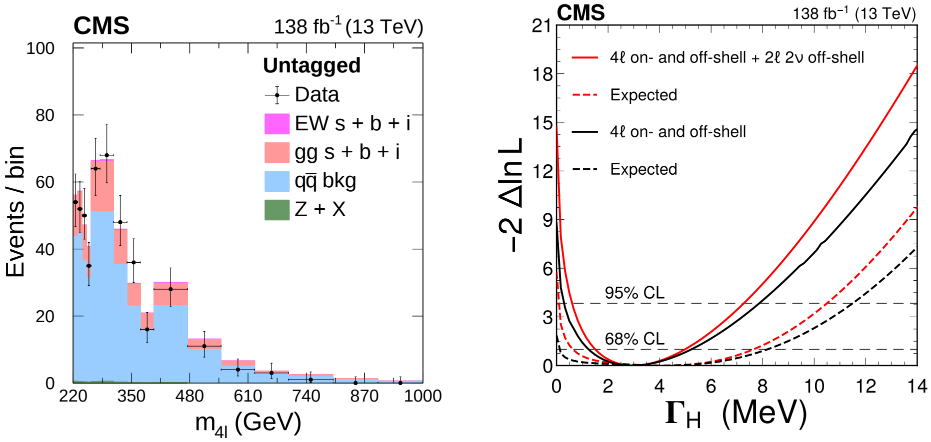

Figure 2.8: Distributions of $4l$ invariant mass, after applying a cut on the kinematic background discriminants, to enhance the Higgs boson signal contribution (left), and (right) observed (solid) and expected (dashed) one-parameter likelihood scans over Γ_H [28].

Notably, a 3.6σ significance for the off-shell Higgs boson production is obtained. Including also events from the $2l2\nu$ channel, a value of $\Gamma_H = 3.0^{+2.0}_{-1.5}$ MeV is obtained, with 95% CL limits between [0.6, 7.3] MeV.

2.2.3 Indirect on-shell strategy in $H \to \gamma\gamma$ via interferometry

The constraint given in the off-shell regime is very strong, but requires to assume the equality of on-shell and off-shell couplings. This motivates indirect measurements using only the on-shell phase-space, but not limited by the experimental resolution as in direct lineshape contraints. A process with the required characteristics, i.e. a cross-section scaling differently from $g_i^2 g_f^2/\Gamma_H$, is represented by any type of interference between the resonant production and continuum diagrams without Higgs bosons. Indeed, quantum interference arises whenever two or more diagrams share the same initial and final states; such interference can produce observable effects. In the case of gluon-fusion Higgs production with decay into two photons, $gg \to H \to \gamma\gamma$, the interference with the direct continuum diphoton background, $gg \to \gamma\gamma$, is particularly relevant because of the large cross-section and relatively clean final state. Exploiting these interference patterns to probe the Higgs boson decay width is the center of this thesis. This section outlines the theoretical framework and possible analysis strategy motivated by the theory.

Signal-background interference in gluon-gluon fusion

At the LHC, the SM Higgs boson production mode proceeds mainly through gluon-gluon fusion. The resonant Higgs-mediated process, $gg \to H \to \gamma\gamma$, and the continuum background process, $gg \to \gamma\gamma$, share the same initial (two gluons) and final (two photons) states, leading naturally to interference.

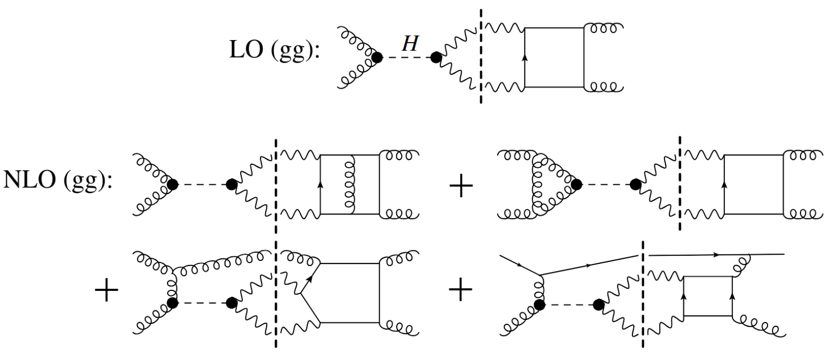

Figure 2.9: Representative Feynman diagrams for gg-initiated interference at leading and next-to-leading order in QCD. The dashed vertical lines represent the separation between signal (left) and background (amplitude), highlighting the interfering final states.

The inclusive amplitude for the $gg \to \gamma\gamma/gg \to H \to \gamma\gamma$ process can be expressed as

$$\mathcal{M}_{incl} = \frac{\mathcal{M}_{\text{sig}}}{m_{\gamma\gamma}^2 - m_H^2 + i\Gamma_H m_H} + \mathcal{M}_{\text{bkg}}, \tag{2.12}$$

where $m_{\gamma\gamma}$ is the invariant mass of the diphoton system, m_H is the Higgs boson mass, Γ_H its total width, and $\mathcal{M}_{\text{bkg}}$ is the amplitude of the $gg \to \gamma\gamma$ process [49]. Note that in the $\gamma\gamma$ final state, the prompt diphoton production $pp \to \gamma\gamma$, of which $g \to \gamma\gamma$ is a small contribution, is not the only experimental background, because of contaminations from γ+jets and multi-jet processes where one or two jets are misidentified as photons. Squaring the amplitude yields

$$\begin{aligned} |\mathcal{M}_{incl}|^2 = &\underbrace{\frac{|\mathcal{M}_{\text{sig}}|^2}{(m_{\gamma\gamma}^2 - m_H^2)^2 + \Gamma_H^2 m_H^2}}_{:=|\mathcal{M}_{\text{S}}|^2} + |\mathcal{M}_{\text{bkg}}|^2 \\ &+ 2\,\text{Re}\left[\frac{\mathcal{M}_{\text{sig}}}{m_{\gamma\gamma}^2 - m_H^2 + i\Gamma_H m_H}\mathcal{M}_{\text{bkg}}^\dagger\right]. \end{aligned} \tag{2.13}$$

Separating the signal and background amplitudes into real and imaginary parts,

$$\mathcal{M}_{\text{sig,bkg}} = \text{Re}\,\mathcal{M}_{\text{sig,bkg}} + i\,\text{Im}\,\mathcal{M}_{\text{sig,bkg}}, \tag{2.14}$$

one can rewrite the squared amplitude as [50]:

$$\begin{aligned} &|\mathcal{M}_{gg\to\gamma\gamma}|^2 = |\mathcal{M}_{\text{S}}|^2 + |\mathcal{M}_{\text{bkg}}|^2 \\ &+ \frac{2}{(m_{\gamma\gamma}^2 - m_H^2)^2 + \Gamma_H^2 m_H^2}\left[(m_{\gamma\gamma}^2 - m_H^2)\cdot(\text{Re}\,\mathcal{M}_{\text{bkg}}\,\text{Re}\,\mathcal{M}_{\text{sig}} + \text{Im}\,\mathcal{M}_{\text{bkg}}\,\text{Im}\,\mathcal{M}_{\text{sig}})\right] \\ &+ \frac{2}{(m_{\gamma\gamma}^2 - m_H^2)^2 + \Gamma_H^2 m_H^2}\left[\Gamma_H m_H\cdot(\text{Re}\,\mathcal{M}_{\text{bkg}}\,\text{Im}\,\mathcal{M}_{\text{sig}} - \text{Im}\,\mathcal{M}_{\text{bkg}}\,\text{Re}\,\mathcal{M}_{\text{sig}})\right]. \end{aligned} \tag{3.14}$$

The signal and interference cross-sections for gluon-gluon fusion (ggH) can then be written schematically in the narrow-width approximation as [46]:

$$d\sigma_s^{ggH}/dm_{\gamma\gamma} = \frac{c_{g\gamma}^2 S}{(m_{\gamma\gamma}^2 - m_H^2)^2 + m_H^2\Gamma_H^2} \tag{2.15}$$

$$d\sigma_i^{ggH}/dm_{\gamma\gamma} = c_{g\gamma}\frac{(m_{\gamma\gamma}^2 - m_H^2)R + m_H\Gamma_H I}{(m_{\gamma\gamma}^2 - m_H^2)^2 + m_H^2\Gamma_H^2} \tag{2.16}$$

where S, R, I absorb, respectively, the numerical coefficients for signal, real, and imaginary part of the interference, and $c_{g\gamma} = \kappa_{gH}\kappa_{\gamma H}$ is the product of gluon-Higgs and photon-Higgs effective couplings modifiers, i.e. the ratios between the physical coupling and the coupling in the SM.

The total signal cross-section, after integrating the Breit-Wigner (BW) function, is proportional to $\mu_{ggH} = c_{g\gamma}^2/\Gamma_H$, similarly to the $H \to ZZ \to 4l$ case, while the interference cross-section is proportional to $\sqrt{\mu_{ggH}\Gamma_H}$. The different dependence on Γ_H of signal and interference is the key handle to constrain Γ_H itself. The real part of the interference term changes sign at $m_{\gamma\gamma} = m_H$, and, when summed with the signal component, shifts the peak of the resonance and modify the shape of the differential cross-section as a function of $m_{\gamma\gamma}$. On the other hand, the imaginary part is symmetric and alters the total cross section by around 2%, without shifting the peak, as shown in Figure 2.10. In the SM, the imaginary component is suppressed at leading order and becomes more relevant only at higher orders in QCD.

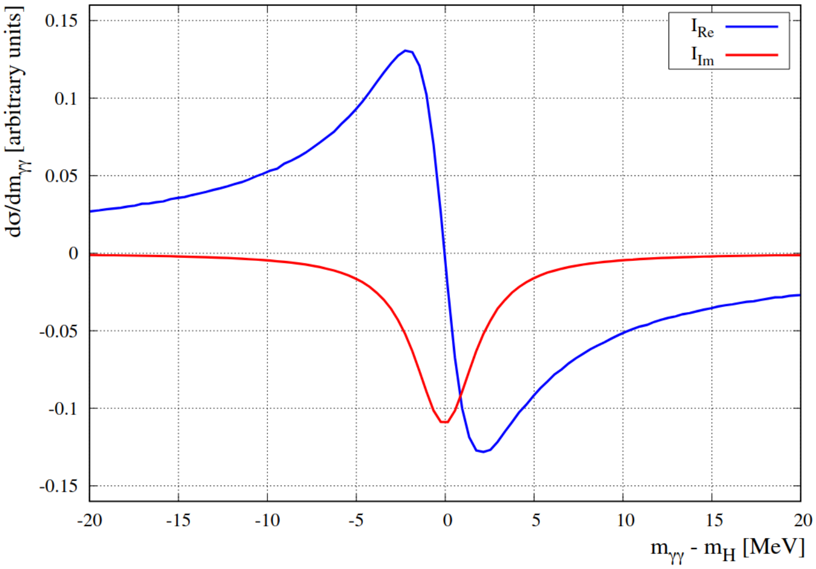

Figure 2.10: Predicted diphoton invariant mass shape for real and imaginary interference components at $\sqrt{s} = 13.6$ TeV [51].

Detector Effects

In a real detector, finite photon energy resolution smears the invariant mass distribution, modifying also the interference lineshape. In order to model the experimental smearing of the invariant mass reconstruction, a Gaussian convolution with a standard deviation matching the expected resolution is used in theory predictions, as a first approximation [46]. The smearing does not cover the peak-dip features with very long tails, as shown in Figure 2.11. On the contrary, this effect, when summing the interference with the signal component, generates an apparent mass shift, as shown in Figure 2.12. A more realistic treatment, with respect to the first attemps in the theoretical predictions published in literature [46, 50, 51], requires parametric or full detector simulation, as it will be detailed in Part 2 of this thesis.

The distortion in the Higgs boson mass spectrum due to the interference can thus be exploited by measuring precisely the shape of the distribution including both ggH signal and interference. The size of these distortions is directly related to the real part of the interference amplitude. As a result, the Higgs boson width Γ_H can be extracted by measuring the amount of deviation from the pure BW lineshape induced by the interference with the $gg \to \gamma\gamma$ continuum, because its differential cross section is proportional to $c_{g\gamma} = \sqrt{\mu_{\gamma\gamma}\Gamma_H}$.

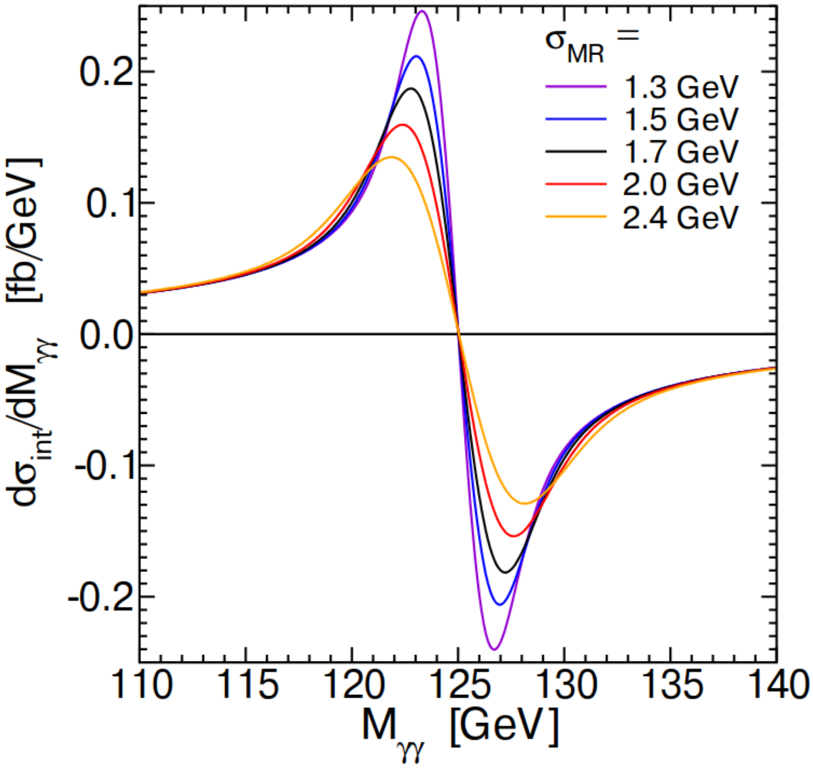

Figure 2.11: Illustration of detector resolution effects on the interference lineshape, with a parametric Gaussian convolution. The cross-section is evaluated at $\sqrt{s} = 8$ TeV [50].

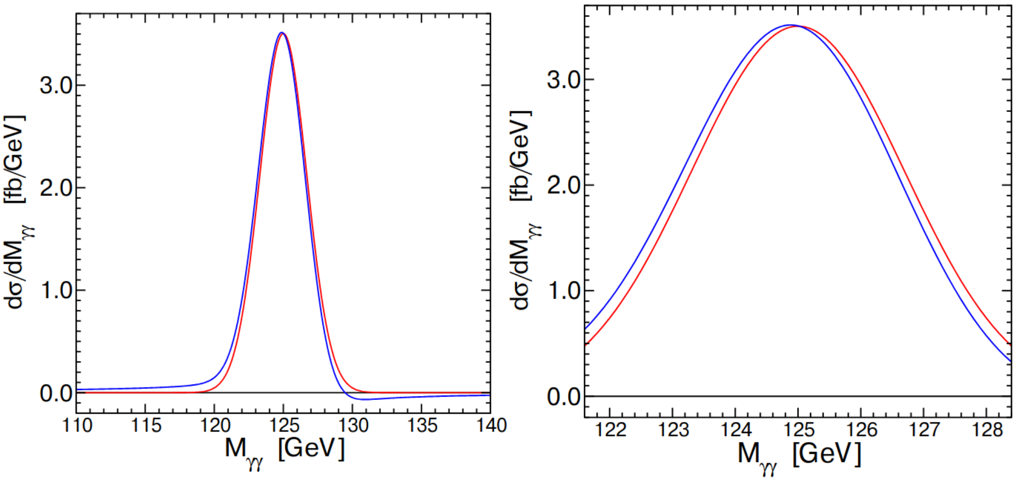

Figure 2.12: Diphoton invariant mass distributions with a Gaussian mass resolution at $\sigma = 1.7$ GeV. In each panel, the right-shifted (red) curve includes only the Higgs contribution without interference, and the left-shifted (blue) curve also includes the interference contribution. The cross-section is evaluated at $\sqrt{s} = 8$ TeV [50].

Sub-dominant interference contributions

Higgs bosons can also be produced in processes initiated by quark-gluon scattering, such as $qg \to q\gamma\gamma$ at tree level, which acts as a background to $qg \to Hq \to \gamma\gamma q$. These channels have smaller rates than the gluon-fusion process due to the reduced parton luminosity.

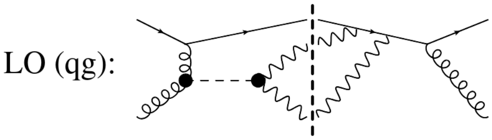

Figure 2.13: Representative diagrams for the sub-dominant $qg \to (H) \to \gamma\gamma$ interference contributions at the lowest order in QCD.

The qg interference has a similar lineshape to the gg interference but with an opposite phase, shifting the resonance peak to slightly higher $m_{\gamma\gamma}$ values [46]. When detector resolution is included, both gg and qg interference contributions can be compared; the gg term remains dominant, but the qg term partially counteracts its mass-shift effect.

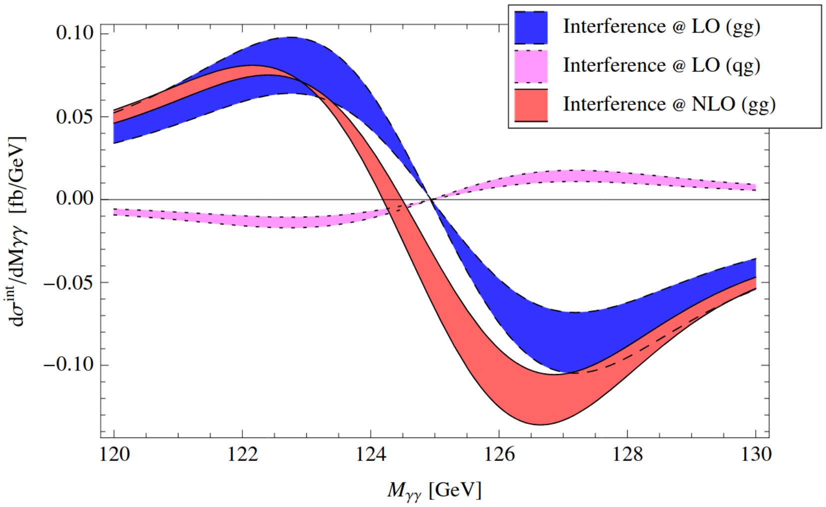

Figure 2.14: Comparison of dominant and sub-dominant interference effects, including the experimental resolution. The differential cross-section for gg production, computed at NLO in QCD, is also shown. Cross-sections evaluated at $\sqrt{s} = 7$ TeV [46]. The shift in the y-axis from the LO to NLO in the gg production is due to the imaginary component of the amplitude, which is suppressed at LO.

Mass shift and interference-sensitive observables

In the original proposals of width measurements using interference in $H \to \gamma\gamma$, focusing on the theoretical aspects [46, 50, 51], the main focus was on the apparent mass shift as the handle to constrain Γ_H. Nevertheless, the mass shift is only a simple way to quantify the distortions due to interference in different segments of the kinematical phase spaces, Higgs boson transverse momentum (p_T) and number of jets in the final state. A possible implementation of this method is with a simple integrated estimator, i.e. by defining an invariant mass window $m_H - \epsilon < m_{\gamma\gamma} < m_H + \epsilon$ and evaluating the integral mean $\langle m \rangle_\epsilon$, by integrating $m_{\gamma\gamma}[d\sigma_s^{ggH}/dm_{\gamma\gamma} + d\sigma_i^{ggH}/dm_{\gamma\gamma}]$ over the mass window:

$$\langle m \rangle_\epsilon = \int_{m_H-\epsilon}^{m_H+\epsilon} m_{\gamma\gamma}[d\sigma_s^{ggH}/dm_{\gamma\gamma} + d\sigma_i^{ggH}/dm_{\gamma\gamma}]dm_{\gamma\gamma} \tag{2.17}$$

The mass shift is then defined as $\Delta M_\epsilon = \langle m \rangle_\epsilon - m_H$, which effectively is equal to the integral mean of $d\sigma_i^{ggH}/dm_{\gamma\gamma}$. If the detector resolution can be approximated as a Gaussian distribution, then $d\sigma_s^{ggH}/dm_{\gamma\gamma}$ is symmetric with respect to its peak position, and ΔM_ϵ is proportional to the interference cross-section, because of integral linearity, in particular to the coefficient R and to the coupling product $c_g\gamma$, therefore $\Delta M_\epsilon \propto \sqrt{\mu\Gamma_H/\Gamma_H^{SM}}$.

As shown in Figure 2.15 for small ϵ ($< 1GeV$), the mass shift corresponds only to the shift in the maximum point of the distribution, which experimentally would correspond to the peak position of the function used to fit the background-subtracted data. If, instead, a wider window is included, the shift is not only larger, but increases approximately linearly with ϵ for all $\epsilon > 2\sigma$, due to the very long positive (negative) tail at lower (higher) $m_{\gamma\gamma}$ [50].

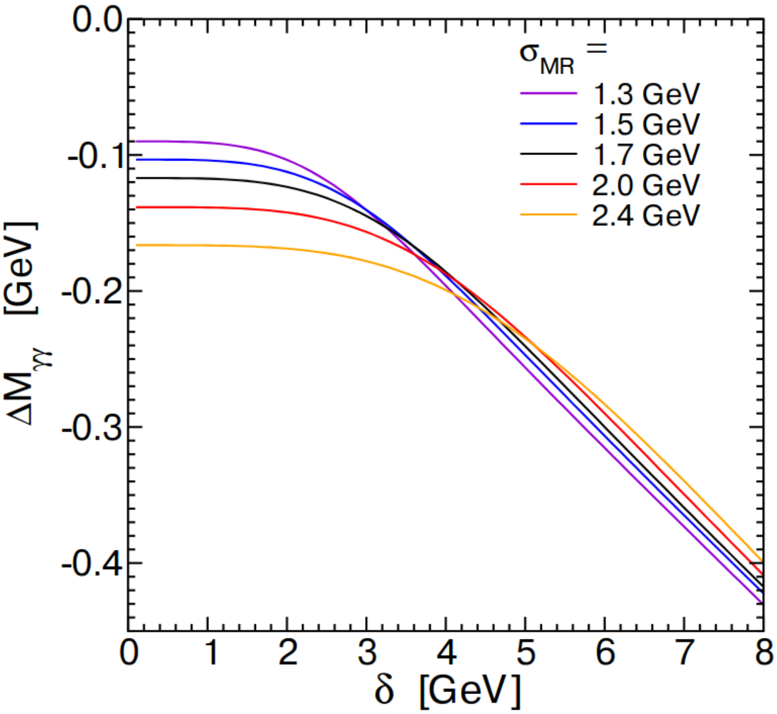

Figure 2.15: Dependence of the mass shift on the window ϵ [50].

For this reason, to have better experimental sensitivity using the mass shift, one must extract it using some kind of mean over the signal distribution, which is not the way experimental analyses usually proceeds, that is fitting the signal distribution with a dedicated function. Moreover, this tail-dependent mass-shift is difficult to define in the presence of data-driven background, as is the case for $H \to \gamma\gamma$. For these reasons, in this thesis, the mass shift is not employed as an observable, but this tail-dependent effect is partially recovered by fitting the full signal shape, in particular the ”shoulder” regions, that have far better signal-to-background ratio than the tails, using templates able to quantify the interference cross-section. Nonetheless, the mass-shift is used to understand the dependence of interference distortions on key observables and finally to choose the features used for categorization. The most sensitive variable is diphoton transverse momentum, as shown in Figure 2.16, where at large $p_{T,\gamma\gamma}$ the mass-shift vanishes. It is also important to note that, for $\gamma\gamma + 2$ jets events, the mass-shift is very small or even going in the opposite direction [52].

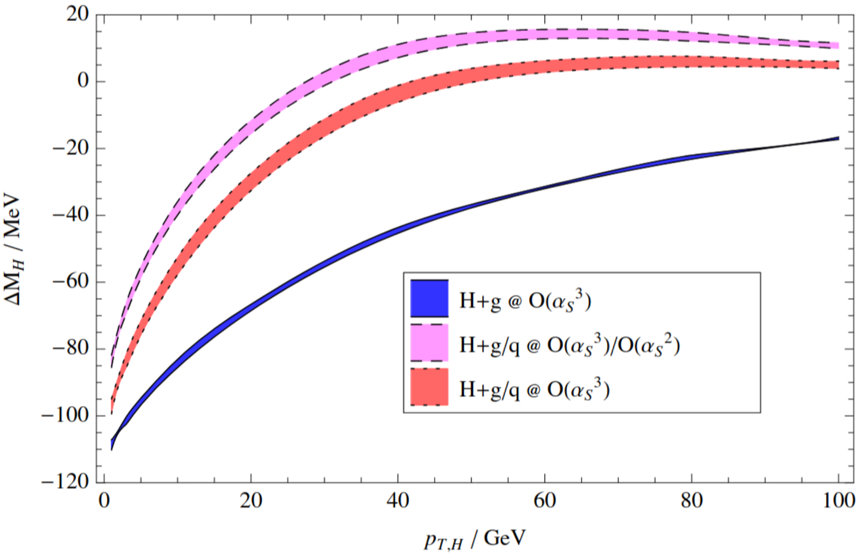

Figure 2.16: Dependence of mass-shift from a lower cut on $p_{T,\gamma\gamma}$ [46].

Chapter 3

CMS experiment at the LHC

3.1 Hadronic collisions

Hadronic collisions are central to the study of high-energy particle physics and nuclear dynamics. At high energies, hadrons are not viewed as elementary particles but as bound states of quarks and gluons. Not all the nucleons or partons are involved in the hard scattering, but usually only one parton per hadron contributes. The partons that do not directly participate in the collision are referred to as spectators, and usually have a very small scattering angle and therefore are not detected. Initial-state partons interact by QCD, QED of electroweak force, while partons in the final state result in the production of jets, which are collections of baryons, resulting from the hadronization of final state partons, when multiple interactions create particles in a shower-like fashion.

The hadronization starts with a phase called parton shower, where, after the hard process, the partons created by the hard collision emit gluon radiation (governed by perturbative QCD), and continue doing so until a low virtuality scale $O(1\text{ GeV})$ is reached. To describe this perturbative evolution phase, quantum coherence effects are incorporated, which connect the probabilities of radiation off different partons in the event, to ensure color neutrality of the final physical particles. Once the low virtuality scale is reached, the memory of the hard-process phase has been lost due to the very different time scales, and the final phase of hadronization starts. At this point, non-perturbative QCD takes over and nearby partons merge into colour-singlet clusters, which create physical hadrons. Because of the de-coupling from the hard-process phase, the hadronization is assumed to be independent of the initial hard process, and its parametrization, tuned to the observables of some reference process, can then be used in other hard interactions. Note that the evolution is not entirely independent of what happens in the hard event, because after the collision, at least colour quantum numbers must be exchanged between final state partons to guarantee the overall color neutrality and conservation of baryon number [53].

In generic hadronic collisions, the incoming hadrons can be thought of as a collection of partons with specific momentum distributions, called parton distribution functions (PDF). The interaction between these partons can be described using a cross-section, calculated as an integral over the parton distributions of the incoming hadrons. The cross-section for a given hadronic process is given by:

$$\sigma_{AB} = \int dx_1\, dx_2\, f_A(x_1) f_B(x_2) \hat{\sigma}_{ab}(x_1, x_2, \hat{s}),$$

where $f_A(x_1)$ and $f_B(x_2)$ are the parton distribution functions (PDFs) for the hadrons A and B, respectively, and $\hat{\sigma}_{ab}(x_1, x_2, \hat{s})$ is the parton-level cross-section, which depends on the momentum fractions x_1 and x_2 of the partons and the center-of-mass energy squared $\hat{s}$. The PDFs are not fixed and evolve with Q^2, the momentum transfer in the interaction. This evolution is described by the DGLAP equations, which are a set of partial differential equations governing the scale dependence of the PDFs:

$$\frac{d}{d \ln Q^2} f_a(x, Q^2) = \sum_b \int_x^1 \frac{dz}{z} P_{ab}(z) f_b\left(\frac{x}{z}, Q^2\right),$$

where $P_{ab}(z)$ are the so-called splitting functions, that describe the probability of a parton of type b radiating a parton of type a with momentum fraction z. The DGLAP equations describe the evolution of PDFs as the energy scale Q^2 changes, and they are crucial for accurately modeling parton distributions in high-energy collisions.

Simulating hadronic collisions involves complex dynamics that cannot be solved analytically. Monte Carlo (MC) methods are therefore employed to generate events and study the various processes involved in hadronic collisions. MC event generators simulate the hard scattering, parton showering, hadronization, and final-state particle production. Some well-known MC generators include:

- **PYTHIA** [54]: A versatile MC event generator for simulating high-energy collisions, including hadronization, parton showers, and multi-parton interactions.
- **HERWIG** [55]: Another MC generator that focuses on parton showers and hadronization, with a different approach compared to PYTHIA.
- **SHERPA** [56]: A general-purpose MC generator capable of simulating hadronic collisions with detailed models for parton showers, hadronization, and multi-parton interactions.

These generators allow for the simulation of event kinematics, particle spectra, and cross-sections, providing a statistical description of the expected outcomes of hadronic collisions.

The process of generating events in high-energy collisions involves several key steps:

1. **Hard Scatterings:** Partons are selected from the incoming hadrons according to their PDFs, depending on the energy scale with the DGLAP equations. The hard scattering process is modeled by calculating the parton-level cross-section for the interaction.
2. **Parton Showers:** The outgoing partons from the hard scattering interact by emitting additional partons (gluon radiation), forming a parton shower.
3. **Hadronization:** After the partons have undergone sufficient evolution, they undergo hadronization, where quarks and gluons combine to form physical hadrons.

3.2 Large Hadron Collider

The Large Hadron Collider (LHC) [57] is the world's largest and most powerful particle accelerator, located at CERN on the Franco-Swiss border near Geneva. Built inside a circular underground tunnel of 27 km circumference, the LHC provides collisions of protons and heavy ions at unprecedented energies, enabling studies of the fundamental interactions of nature. The LHC is the latest stage of a long series of accelerators at CERN, which together form a complex system for producing, accelerating, and injecting particle beams into the collider ring. Since its first collisions in 2009, the LHC has delivered a wealth of data that has been essential for discoveries such as the Higgs boson and continues to push the frontiers of high-energy physics.

3.2.1 Accelerator complex

The LHC relies on a sophisticated accelerator chain to bring protons to the required energy before they are injected into the main collider ring. The process begins with a hydrogen source, from which protons are extracted and accelerated in the linear accelerator (LINAC). The most recent upgrade, LINAC4, replaces the older LINAC2 and injects protons into the subsequent stages at an energy of 160 MeV. From there, the beam enters the Proton Synchrotron Booster (PSB), which increases the energy to 2 GeV.

The next stage is the Proton Synchrotron (PS), a circular machine with a circumference of about 628 m. The PS has been in operation since the 1950s and it increases the proton energy to 26 GeV before sending them to the Super Proton Synchrotron (SPS). The SPS, with a circumference of 7 km, further accelerates the beam to 450 GeV, the injection energy of the LHC. Both the PS and the SPS, when not filling LHC, extract the beams for various purposes, including physics experiments (NA62, NA64, AMBER, SHINE) and test beams for detectors R&D.

Once injected into the LHC, the beams circulate in opposite directions in two separate vacuum pipes. Superconducting dipole magnets, cooled with liquid helium to about 1.9 K, bend the beams along the circular trajectory, while quadrupole magnets focus them to maintain beam stability. The LHC accelerates the beams up to a design energy of 7 TeV per proton, corresponding to a center-of-mass energy of 14 TeV in proton–proton collisions. Collisions occur at four interaction points, each hosting a major experiment: ATLAS, CMS, ALICE, and LHCb.

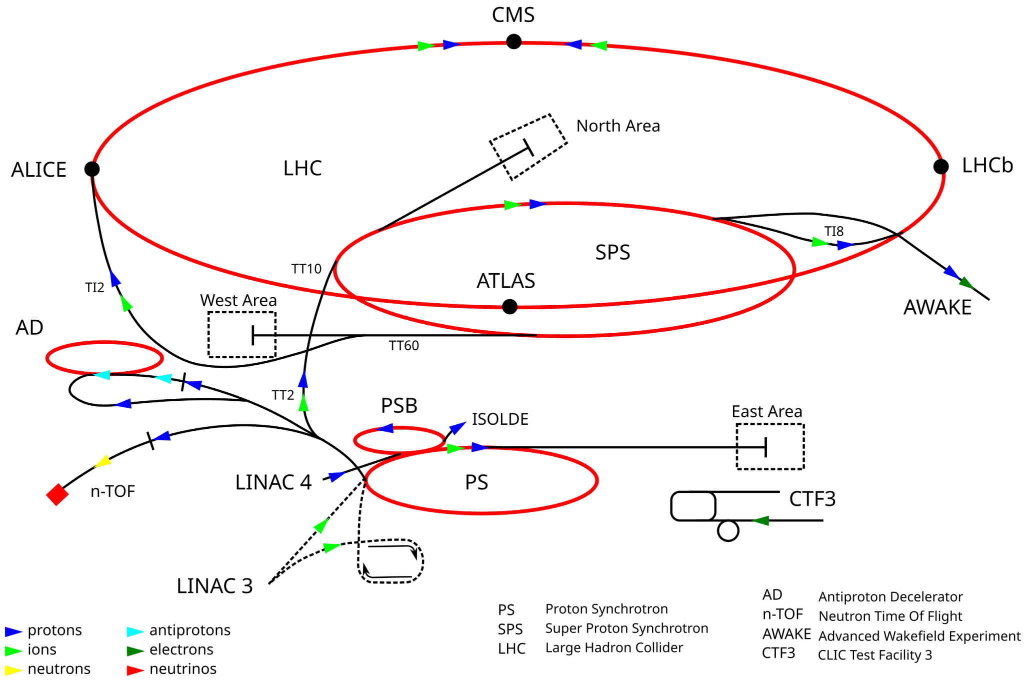

Figure 3.1: CERN accelerator complex scheme.

3.2.2 Luminosity and pile-up

A key performance parameter of the collider is luminosity, defined as a measure of the number of collisions produced per unit time and per unit cross-sectional area. The instantaneous luminosity L can be expressed as

$$L = \frac{N_b^2 n_b f_{\text{rev}} \gamma_r}{4\pi \epsilon_n \beta^*} F, \tag{3.1}$$

where N_b is the number of protons per bunch, n_b is the number of bunches, f_{rev} is the revolution frequency, γ_r is the relativistic Lorentz factor, ϵ_n is the normalized transverse beam emittance, β^* is the beta function at the interaction point, and F accounts for the reduction in overlap due to the crossing angle of the beams. The integrated luminosity L_{int}, obtained by integrating the instantaneous luminosity over time, determines the statistical power of physics analyses. For any process with a cross-section σ, the number of events is the product of the σ and L_{int}. Higher luminosity operation increases the number of recorded collisions but also leads to multiple softer interactions per bunch crossing, known as "pile-up." The average number of such interactions, μ, is proportional to the instantaneous luminosity and the total proton–proton cross section. While pile-up increases detector occupancy and complicates event reconstruction, modern analysis techniques mitigate these effects, for instance, through pile-up subtraction in jets and isolation variables, or by exploiting high-granularity detector information.

During Run 2 (2015–2018), the LHC routinely operated at an instantaneous luminosity above 1.5×10^{34} cm^{-2}s^{-1}, with an average pile-up of 30–50 interactions per bunch crossing, as shown in Figure 3.2. The High-Luminosity LHC (HL-LHC) upgrade, planned to start in 2031, aims to increase the peak luminosity by about a factor of five, leading to an average pile-up of up to 200. This will present significant challenges for detectors and reconstruction algorithms but will also greatly enhance the dataset available for precision studies and searches for new physics.

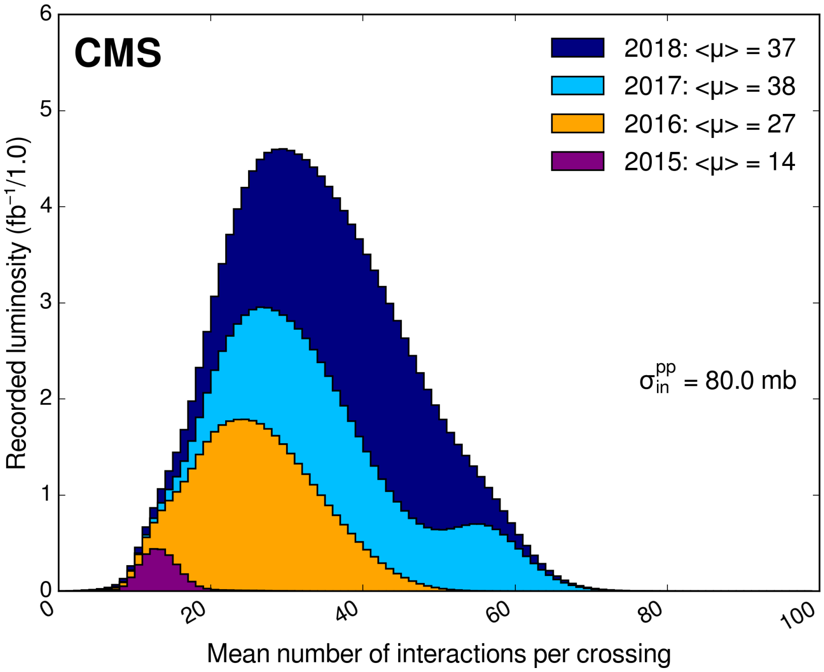

Figure 3.2: Pileup distribution in Run 2 [58].

3.2.3 Data-taking periods

The LHC operation is organized into runs of various years, separated by long shutdowns (LS) for upgrades and maintenance. Each run corresponds to a distinct data-taking period with specific beam conditions, detector configurations, and delivered luminosity.

Run 1 (2010–2012) marked the first operational phase of the LHC. Proton–proton collisions were delivered at center-of-mass energies of 7 TeV in 2010 and 2011, and 8 TeV in 2012. The integrated luminosity delivered to the ATLAS and CMS experiments was about 25 fb^{-1}, and the highlight of Run 1 was the discovery of the Higgs boson in 2012.

Run 2 (2015–2018) followed the first long shutdown (LS1), during which the accelerator and detectors underwent significant upgrades. Proton–proton collisions were delivered at 13 TeV, nearly doubling the energy of Run 1. The LHC reached new records in luminosity, delivering over 100 fb^{-1} to the general-purpose experiments. Run 2 data provided the basis for precision measurements of the Higgs boson properties and extensive searches for physics beyond the Standard Model.

Run 3 started in 2022 after the completion of the second long shutdown (LS2). It features proton–proton collisions at 13.6 TeV and improved detector performance.

Looking ahead, the **HL-LHC** project, planned to begin operation after a long shutdown for the upgrade of the machine and of the detectors, starting in 2026 and ending in 2031, will extend the physics reach of the collider significantly. By accumulating up to 3000 fb^{-1} of data, the HL-LHC will enable precision measurements of Higgs boson couplings at the percent level and expand the sensitivity to new physics scenarios well beyond the current limits.

3.3 CMS experiment

The Compact Muon Solenoid (CMS) experiment is a general-purpose detector. It is formed by different detector layers, concentric and roughly cylindrical, in a typical ”onion” shape structure. Different layers are specialized in different measurements or sensitive to a specific category of particles. Most detectors are naturally split into a central, barrel section, which has a tubular geometry around the beam line, and two forward regions, the endcaps, that increase the detector acceptance to the regions in the immediate sur-

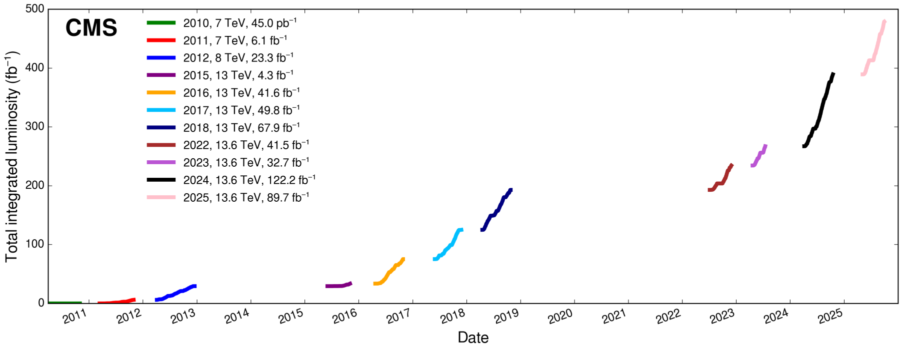

Figure 3.3: Integrated luminosity in CMS during Run1-3 [59].

roundings of the beam line, far from the interaction region. A schematic view of the CMS detector and a section with particles passing through it are in Figure 3.4.

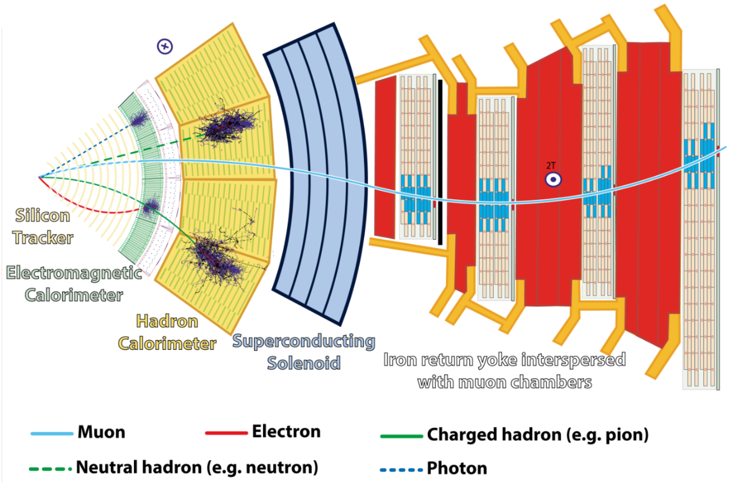

Figure 3.4: Slice showing CMS sub-detectors and how particles interact with them [60].

Solenoid Magnet

The CMS central magnet is a superconducting solenoid designed to generate a strong and uniform magnetic field. It measures 12.5 m in length and

6 m in diameter, and is surrounded by a massive iron return yoke that both guides the magnetic flux and provides bending power in the outer detector regions. The solenoid encloses the tracking system and calorimeters and produces a 3.8 T field aligned with the beam axis. This field is essential for curving the trajectories of charged particles, enabling the determination of their charge sign and the reconstruction of their transverse momentum via the Lorentz force.

Tracker

The innermost detector surrounds the beam pipe with a cylindrical structure of radius up to about 130 cm. Operating in the highest particle-density environment, it is optimized to measure charged-particle trajectories with excellent spatial resolution and minimal material budget, i.e. amount of material normalized by its radiation length, to minimize scattering. The fine segmentation of the silicon sensors provides hit positions to be measured with a precision of the order of 10 μm, allowing charged-particle tracks to be reconstructed across multiple layers with high accuracy. Thanks to the magnetic field, tracks are fitted to evaluate their curvature in the transverse plane, in order to measure the transverse momentum p_T.

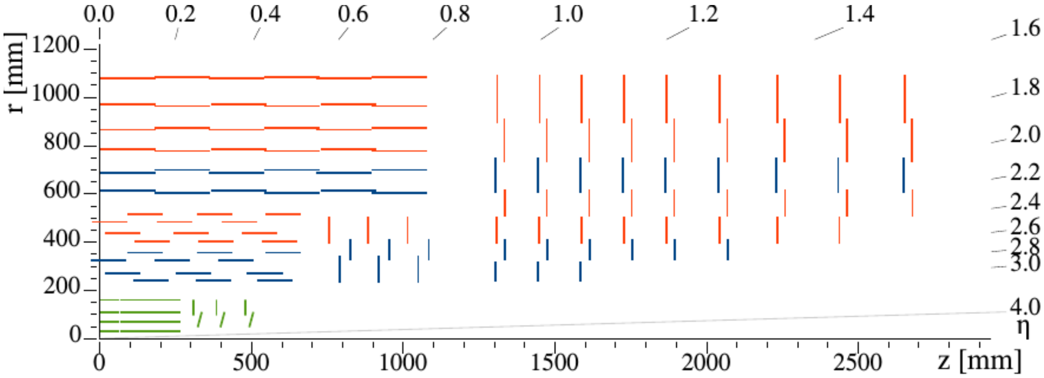

Figure 3.5: Schematic drawing of the CMS tracker [61].

The tracking system is entirely based on silicon sensors and consists of two main technologies: hybrid pixel detectors and silicon strip detectors. The pixel system comprises four concentric barrel layers together with three disks on each endcap. Each pixel cell has an area of $100 \times 150\,\mu\text{m}^2$, with the outermost barrel layer reaching a radius of 16 cm and the endcap disks positioned within 50 cm from the interaction point. The strip system extends tracking coverage up to 1.2 m in radius and $|\eta| < 2.5$. Strip thicknesses range from 320 μm to 500 μm with pitches between 122 μm and 205 μm. Despite being optimized to minimize material budget, it reaches 2 radiation lengths for some polar angle values, as shown in Figure 3.7 therefore photons and electrons, except at $\eta \sim 0$, often have hard interactions, i.e. conversions or bremmstrhalung, upstream of the electromagnetic calorimeter.

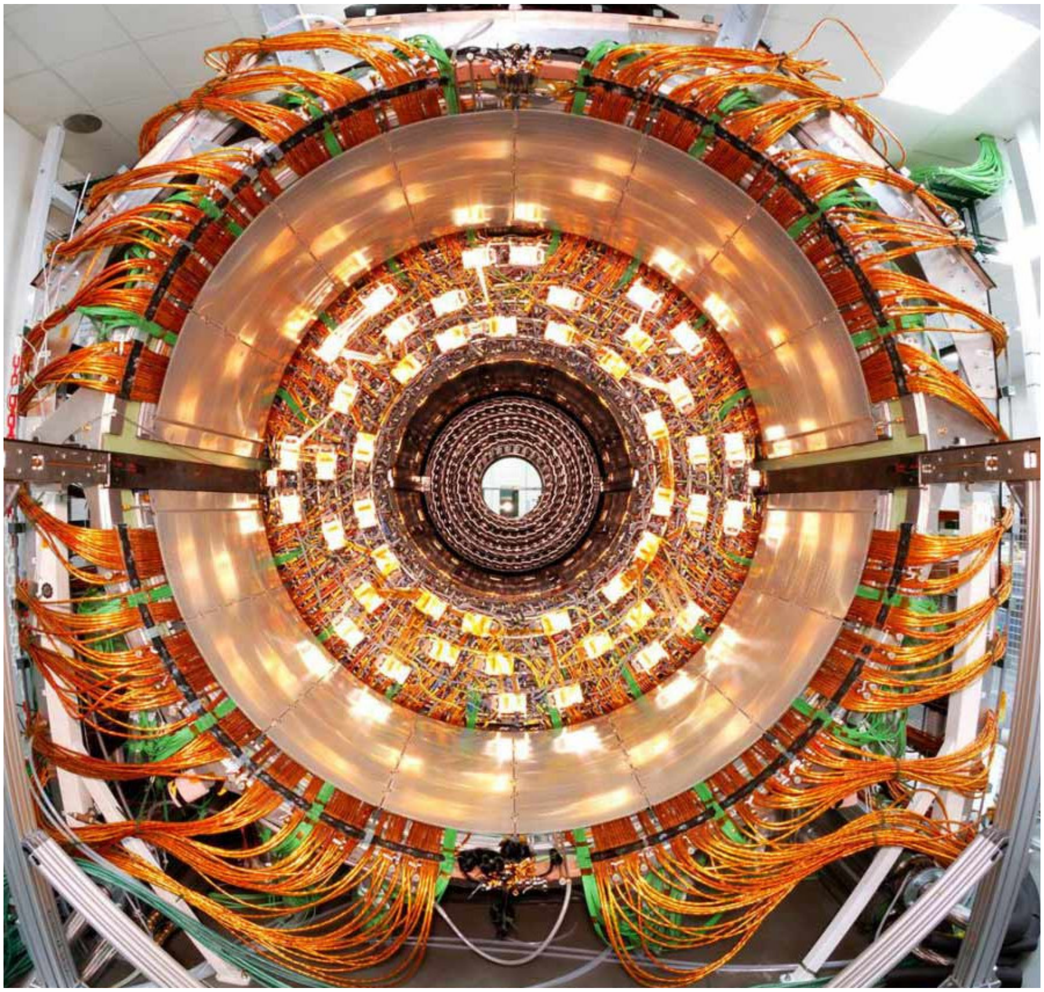

Figure 3.6: Picture of the CMS Tracker outer barrel.

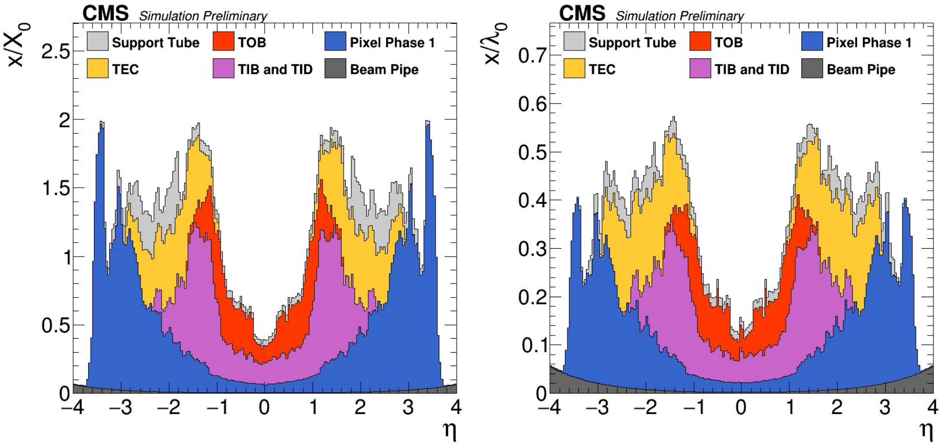

Figure 3.7: Material budget simulation of the tracker. Shown is the breakdown of the total material budget into sub-detectors, i.e. tracker endcap (TEC), tracker outer barrel (TOB), tracker inner barrel (TIB), tracker inner detector (TID) and pixels. The beam pipe and support tube are also included in the material budget. 100,000 single neutrino gun events are used to compile the plots. The passage of the neutrino through the material of the simulated CMS detector is tracked in Geant4, imaging the material distribution of the detector by plotting, for each step of the neutrino's spatial trajectory, the fraction of the total radiation or hadronic interaction length of its trajectory, traversed in that step [62, 63].

Electromagnetic Calorimeter (ECAL)

Surrounding the tracking detectors, the electromagnetic calorimeter (ECAL) measures the energies and positions of electrons and photons with high precision. It is a homogeneous and hermetic detector composed of about 75,000 lead tungstate ($PbWO_4$) crystals, arranged in a barrel section ($|\eta| < 1.479$) and two endcaps ($1.479 < |\eta| < 3.0$). Crystals sizes are about 2.2×2.2 cm in the barrel and 2.9×2.9 cm or 3×3 cm in the endcaps, with lengths of about 23 cm. The shape of the crystals depend on η, such that their axis always point towards the interaction point, with a tilt of about 1° to avoid losing efficiency for particles traversing the space between crystals or modules. Lead tungstate is well suited for compact calorimetry due to its short radiation length ($X_0 \approx 0.89\,\text{cm}$), small Molière radius ($R_M = 1.96\,\text{cm}$), and fast scintillation response ($\sim 15\,\text{ns}$). Although the light yield is relatively low (around 100 photons/MeV), the high density and fast response compensate for this, enabling excellent energy resolution in a compact volume.

The scintillation light is detected through avalanche photodiodes (APDs) in the barrel and vacuum phototriodes (VPTs) in the endcaps, chosen for their radiation tolerance and performance in high magnetic fields.

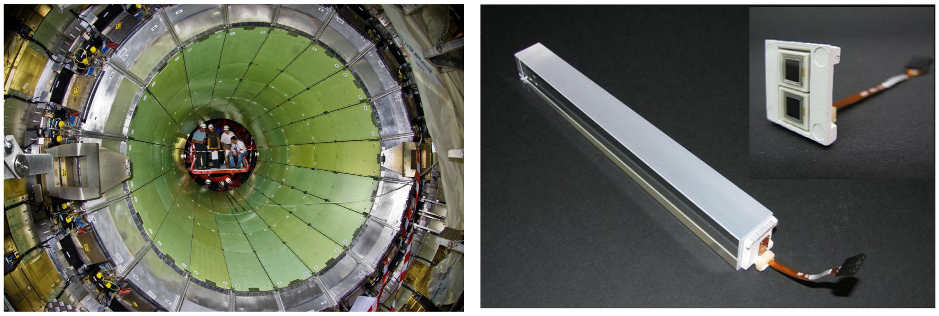

Figure 3.8: Pictures of ECAL (left) mounted inside CMS, and of one crystal in the barrel with its APD capsule [64].

Temperature and high-voltage stability are critical for maintaining calibration, contributing less than 0.2% to the constant term of the resolution. Uniformity in light collection along the crystal length, which could otherwise degrade performance, is improved by controlled depolishing of one lateral face. Transparency variations due to radiation-induced color centers are continuously monitored and corrected using a laser light injection system, ensuring stability of the response during data taking, as detailed in Section 3.4.

Before installation, ECAL modules were characterized in electron beam tests, confirming the expected performance of energy resolution of the ECAL, parameterized as:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \tag{3.2}$$

where $S = 2.8\%$ represents the stochastic term, $N = 12\%$ corresponds to electronic noise, and $C = 0.3\%$ denotes the constant term. In real operation, radiation damage and ageing can slightly worsen the resolution, depending on the detector region and on data-taking eras, as shown in Figure 3.9.

During LHC operation, in situ calibrations are refined using well-understood physics processes such as $\pi_0 \rightarrow \gamma\gamma$, $Z \rightarrow ee$, ensuring uniform and stable response over time.

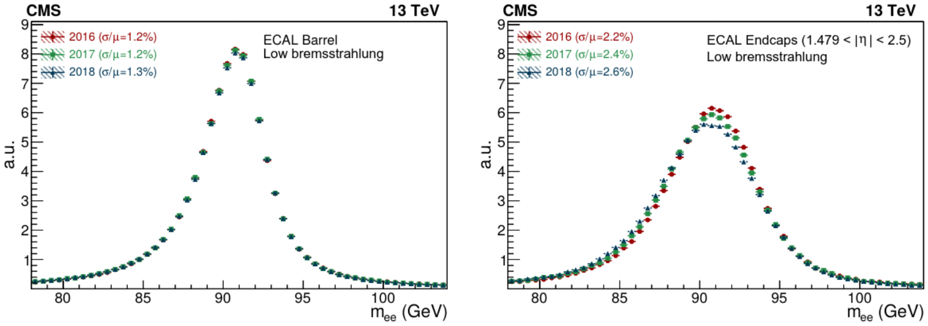

Figure 3.9: Invariant mass spectrum of $Z \to ee$ candidates in data reconstructed using ECAL, in different data-taking years [25].

In addition to energy measurement, the ECAL also provides accurate position information for photons by reconstructing the energy-weighted centroid of the crystal deposits, which is essential for photon identification and π^0 reconstruction. A detailed discussion of ECAL clustering and reconstruction is presented in Section 3.4.

In order to reject high-energy neutral pions producing diphotons that shower in the same ECAL crystal, it was planned to have high-granularity silicon-tungsten pre-shower detectors in both barrel and endcap, featuring 2 X_0 thickness and the capability to distinguish, with geometrical information, between a double photon cluster from a π^0 and a single cluster from a prompt photon. Only the endcap preshower was finally mounted in CMS, with somewhat limited capability of π^0 veto due to ageing and irradiation, and therefore its shower shape information is not fully employed for photon identification algorithms.

Hadronic Calorimeter (HCAL)

The hadronic calorimeter (HCAL) measures the energy of hadrons using a sampling design, with alternating layers of brass absorber and plastic scintillator. The detector is about 1 meter thick, and is built to be hermetic, ensuring that the full energy of jets is accounted for and that missing transverse momentum can be reliably reconstructed. The resolution at beam tests for pions, after corrections using the energy deposited in ECAL is:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + C^2, \tag{3.3}$$

with $S = 85\%$ (stochastic term) and $C = 7\%$ (constant term) [65].

Muon System

The muon spectrometer surrounds the calorimeters and is integrated with the steel return yoke of the magnet. In the barrel region, muons are measured with drift tubes, while in the endcaps Cathode Strip Chambers (CSCs) and Resistive Plate Chambers (RPCs) are used. The return yoke provides a field of about 1.5 T, bending muon tracks sufficiently to allow momentum determination from the curvature of their trajectories. The combination of multiple technologies ensures both precise spatial measurements and fast triggering, particularly in the forward regions.

Trigger System

The LHC provides collisions at a frequency of 40 MHz, corresponding to about 40 million events per second. Given that a typical event occupies about 1 MB, recording every collision is unfeasible. Moreover, most events arise from low-energy interactions that are not of primary interest. To reduce the data rate to a manageable level of roughly 2 kHz, CMS employs a two-stage trigger system. The first stage is the Level-1 (L1) trigger, implemented in custom electronics, while the second stage is the High-Level Trigger (HLT), a software-based system that runs advanced reconstruction and selection algorithms.

3.4 CMS ECAL object reconstruction

3.4.1 ECAL clustering

Accurate localization and measurement of energy deposits in the ECAL is a non-trivial task. The electromagnetic shower initiated by an electron often produces radiated photons, which in turn generate several secondary clusters that must be correctly associated with the primary electron-induced cluster. The electron bremmstrhalung also required dedicated tracking algorithm like Gaussian Sum filters (GSF) [66], able to fit the kinks in the tracks due to photon emission, instead of the standard Kalman filter [67]. The electron bremmstrhalung, and also photon conversions, require a generalization of the standard clustering of calorimetric hits and the definition of the so-called super-cluster, containing a collection of "base" clusters created with geometric procedures, as shown in Figure 3.10. For precise energy measurements of electron and photon, it is needed to take into account, respectively, photon clusters from electron bremmstrhalung, and electrons and positrons clusters from photon conversions.

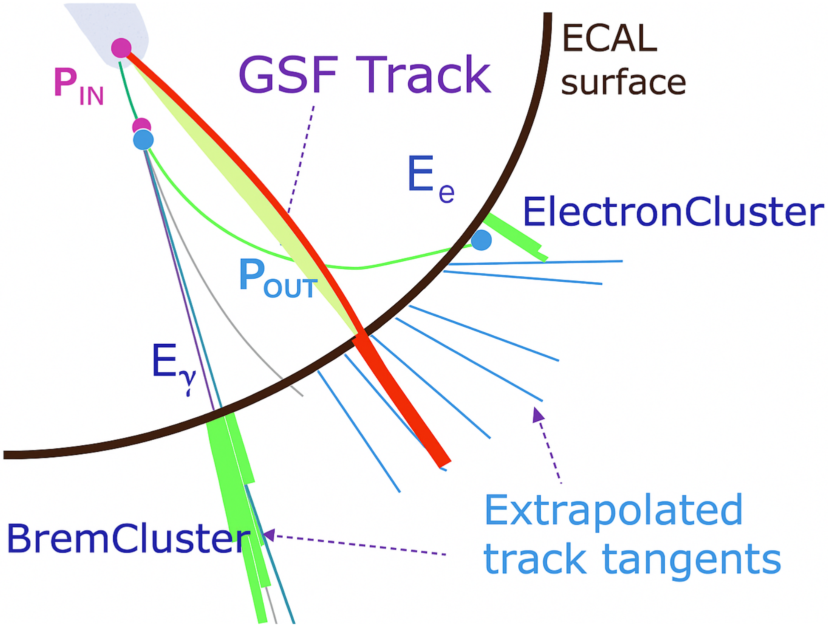

Figure 3.10: Scheme of a super-cluster in CMS ECAL.

Two algorithms are primarily employed for ECAL super-clustering. The first is the *mustache algorithm*, which is optimized for low-energy deposits [25]. Starting from a seed cluster, it associates additional clusters lying in the (η, ϕ) plane around the seed. This region is elongated along ϕ due to the CMS solenoidal magnetic field, which spreads bremsstrahlung photons preferentially in that direction rather than in η, and the width along η depends slightly on the ϕ coordinate. The characteristic "mustache" shape originates from the interplay between magnetic bending and the cylindrical geometry of the ECAL. As the electron traverses the tracker, it curves in the (ϕ) direction due to the solenoidal magnetic field, while bremsstrahlung photons emitted along its trajectory continue straight toward the calorimeter. Each photon is radiated from a different point along the curved path, and since the ECAL surface is cylindrical rather than planar, photons emitted at different radii intersect the ECAL at slightly different pseudorapidities (η). When the ECAL energy deposits are projected into the η–ϕ plane, they therefore appear spread out in ϕ (because of magnetic bending), while the spread in η varies with ϕ, because of the projection geometry. The resulting envelope is not rectangular but curved, giving rise to the distinctive "mustache" pattern used to associate bremsstrahlung clusters with the parent electron, as shown in Figure 3.11. The size of the mustache window depends on the particle's transverse energy E_T: higher-energy electrons and photons are less deflected by the magnetic field, resulting in narrower spreads.

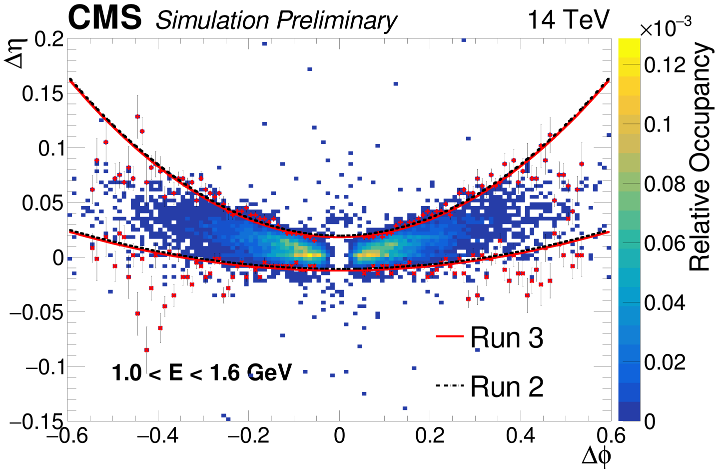

Figure 3.11: Distribution of $\Delta\eta = \eta_{\text{seed-cluster}} - \eta_{\text{cluster}}$ versus $\Delta\phi = \phi_{\text{seed-cluster}} - \phi_{\text{cluster}}$ for simulated electrons. The red curve indicates approximately the set of clusters selected by the mustache algorithm [25].

The second algorithm, referred to as the *refined clustering algorithm*, improves upon the mustache reconstruction by incorporating additional information from tracking and photon-conversion identification. It combines the mustache output with track-to-cluster matching, GSF tracks, and dedicated algorithms that reconstruct bremsstrahlung photon clusters tangent to the electron trajectory. A boosted decision tree (BDT) is also employed to identify both single- and double-leg photon conversions. Based on this combined information, the refine algorithm adjusts the supercluster composition by adding or removing satellite clusters.

3.4.2 Energy reconstruction

The energy of an electromagnetic supercluster in the ECAL is reconstructed according to:

$$E_{e,\gamma} = \sum_i \left[A_i(t) \cdot L_i(t) \cdot C_i(t)\right] \cdot G(\eta) \cdot F_{e,\gamma} + E_{\text{preshower}}, \tag{3.4}$$

where the sum runs over all crystals i included in a supercluster (SC). Here $A_i(t)$ denotes the pulse amplitude measured by the photosensors, while $L_i(t)$ and $C_i(t)$ are time-dependent corrections accounting respectively for radiation-induced transparency loss, evaluated using the laser-based monitoring system, as shown in Figure 3.12 and for inter-crystal response equalization. The factor $G(\eta)$ represents the absolute calibration converting ADC counts into GeV units and is derived from calibrations using Z decays, while $F_{e,\gamma}$ represents particle-dependent corrections compensating for clustering inefficiencies, geometry, and upstream material. Finally, the energy deposited in the two preshower planes is added in the endcap region.

Signal amplitude reconstruction

Single-crystal energy measurements rely on the extraction of the APD pulse amplitude. Amplitudes are extracted from waveforms, which are sampled at 40 MHz, i.e. one sample / bunch-crossing (25 ns), using the Multifit algorithm. Multifit treats overlapping signals in the same crystal by fitting the waveform with one in-time pulse and up to nine out-of-time pulses, whose amplitudes and are free parameters, while the pulse shape itself remains fixed, as shown in Figure 3.13 [68].

Both the signal pulse shape and pedestal offset are essential for this procedure and must be monitored regularly. Radiation damage alters crystal transparency and therefore changes the shape of collected signals. For this reason, during Run 2, pulse templates were recomputed every 3–4 fb^{-1} to track these changes. The electronic noise, resulting in pedestal fluctuations, on the other hand, depend mostly on the photosensors. A gradual increase of about 40 MeV per year in the pedestal fluctuation was observed throughout LHC operation. Pedestal runs are taken approximately every 40 minutes to monitor this behavior, using the same laser system employed for transparency studies.

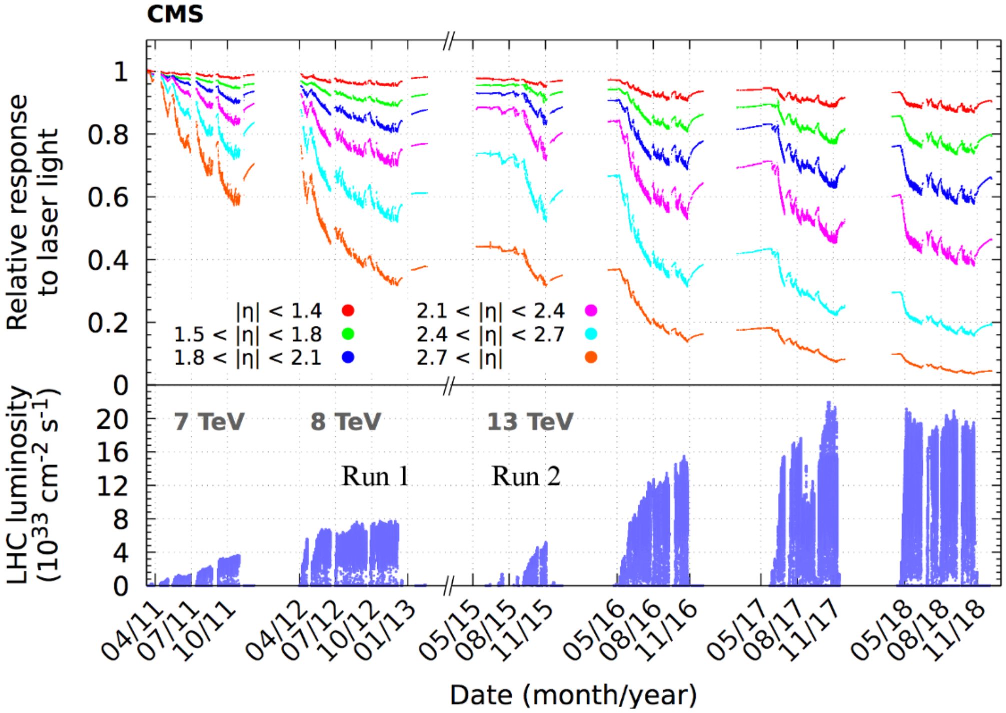

Figure 3.12: Laser transparency corrections throughout Run 2, divided in $|\eta|$ regions.

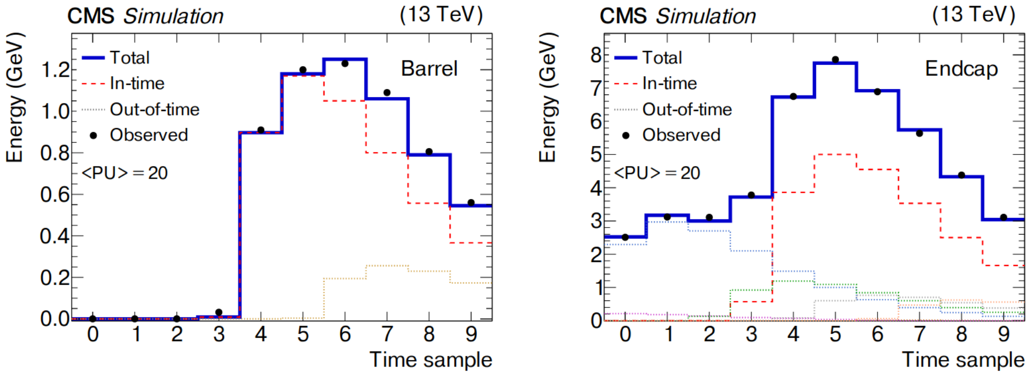

Figure 3.13: Two examples of fitted pulses for simulated events with 20 average pileup interactions and 25 ns bunch spacing. Signals from individual crystals are shown. They arise from a pT = 10 GeV photon shower in the barrel (left) and in an endcap (right). In the left panel, one OOT pulse, in addition to the IT pulse, is fitted. In the right panel, six OOT pulses, in addition to the IT pulse, are fitted [68].

Crystal Transparency Corrections

High radiation fluxes in LHC collisions cause significant degradation of crystal transparency, particularly at large $|\eta|$ [69]. This reduces the light output and distorts pulse shapes on short time scales. Although partial recovery occurs during no-beam periods, continuous corrections are required. CMS employs a laser+LED monitoring system [69] that probes each crystal every 40 minutes during data-taking. The relation between the crystal response to electromagnetic showers (S) and to laser light (L) can be parameterized as:

$$\frac{S}{S_0} = \left(\frac{L}{L_0}\right)^{\alpha}, \tag{3.5}$$

where S_0 and L_0 are reference responses at the beginning of data-taking, and the exponent α is measured to be ~ 1.5 in the barrel and 0.6–1.1 in the endcaps. Figure 3.14 shows the normalized π^0 mass reconstructed with (green) and without (red) laser corrections, demonstrating their crucial impact on stability. Residual long-term response losses are observed due to radiation damage in laser reference diodes and fibers. These effects are corrected using the ratio E/p between ECAL energy and track momentum for electrons from W and Z decays.

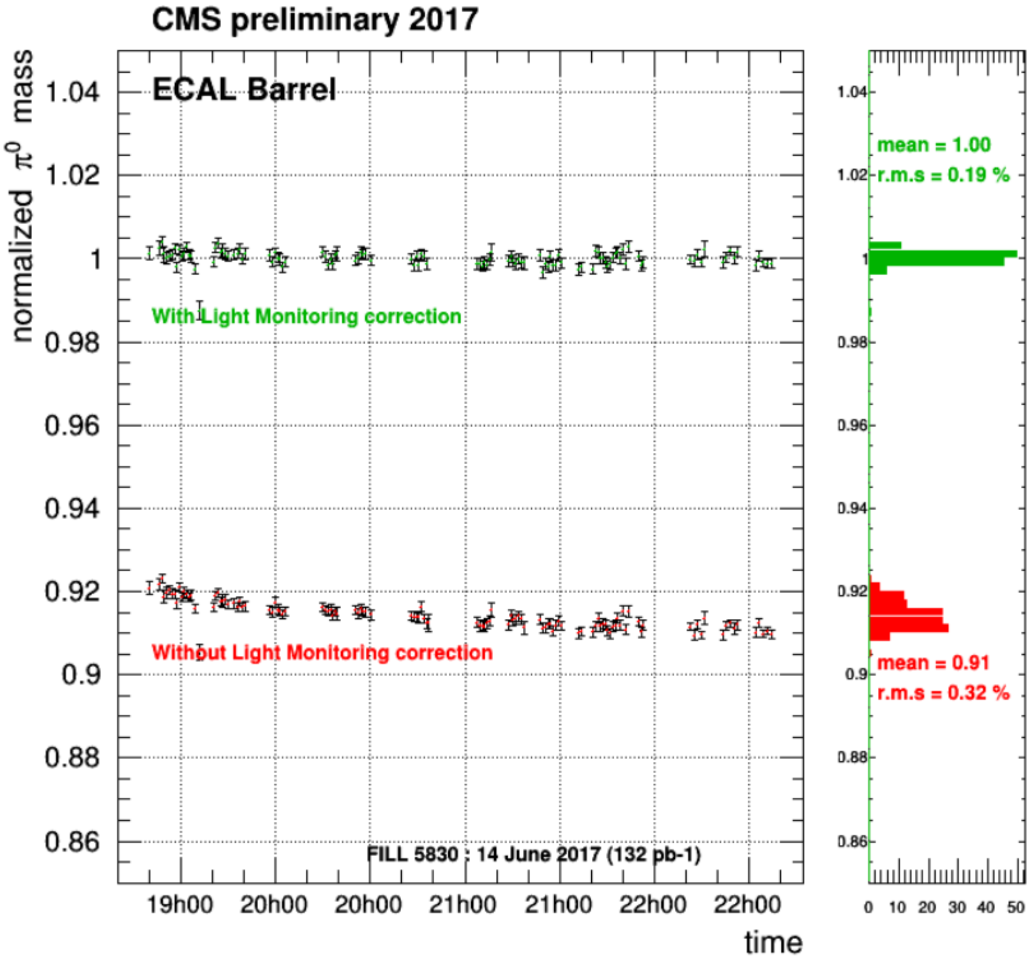

Figure 3.14: Normalized $\pi^0 \to \gamma\gamma$ invariant mass over a few hours data-taking period with (green) and without (red) laser corrections applied.

Inter-Crystal Calibrations

Inter-calibration ensures uniform detector response by equalizing the output of individual crystals. CMS employs three complementary techniques:

- $\pi^0 \to \gamma\gamma$: the invariant mass is reconstructed for photon pairs, as shown in Figure 4.8 and deviations from the nominal π^0 mass are used to iteratively correct responses in fixed ϕ rings;
- E/p: the energy-to-momentum ratio for electrons from W and Z decays is constrained to unity, providing crystal-level corrections;
- $Z \to e^+e^-$: available from Run 2 thanks to large datasets, this method calibrates crystal responses by maximizing a likelihood, comparing reconstructed and simulated Z mass distributions.

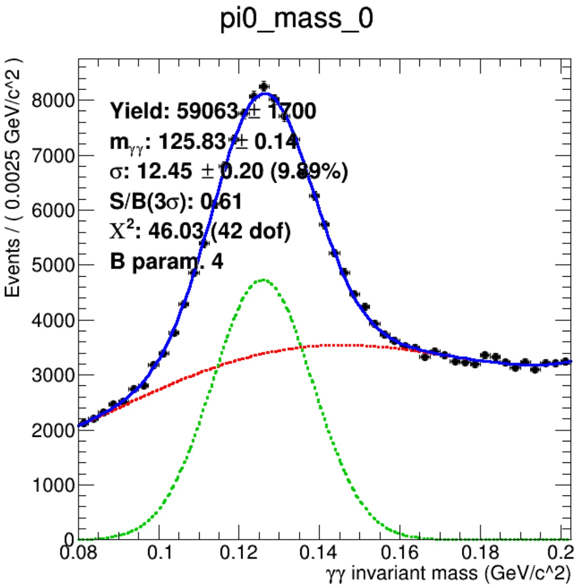

Figure 3.15: Invariant mass of photon pairs in the $\pi^0 \to \gamma\gamma$ intercalibration sample, selecting photons in the barrel.

The three methods are combined to obtain a final per-crystal calibration, exploiting their different and systematic uncertainty as a function of η, as shown in Figure 3.16.

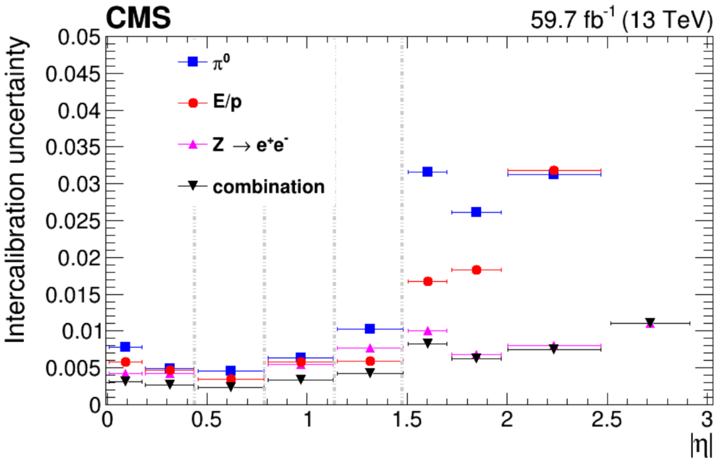

Figure 3.16: The statistical uncertainties in the different intercalibration methods for data collected in 2018, and their combination. The vertical dotted lines mark the boundary between the ECAL modules in the EB and the EB/EE transition.

3.4.3 Energy Regression

A semi-parametric Boosted Decision Tree (BDT) regression technique is employed to mitigate potential performance degradation arising from electromagnetic shower leakage, inactive detector regions, as well as material and geometrical effects. The regression makes use of several variables describing the shower shape, and, in the case of electrons, the bremsstrahlung energy loss, together with the angular coordinates of the track associated with the ECAL supercluster. The isolation is measured by summing the energy of other particles candidates in a cone centered on the supercluster.

The method employs a series of boosted decision tree (BDT) regressions trained on simulated samples, where the true energy of each electromagnetic object is known. The overall goal is to correct the energy using the most probable value of the ratio between the reconstructed and the true energy, which is learned using a semi-parametric approach by means of a probability density of the energy ratio $y = E_{\mathrm{reco}}/E_{\mathrm{true}}$. Each regression step is implemented as a gradient-boosted decision tree optimized through a log-likelihood loss function. For a given set of electromagnetic (e/γ) objects in Monte Carlo simulation, the likelihood function is expressed as:

$$\mathcal{L} = - \sum_{\text{MC } e/\gamma \text{ objects}} \ln p(y|\vec{x}), \tag{3.6}$$

where $p(y|\vec{x})$ represents the probability density for an object with input variables $\vec{x}$ to have the observed energy ratio $y = E_{\text{reco}}/E_{\text{true}}$. The regression learns this probability density by fitting to simulated samples in which the true and reconstructed energies are known.

The probability density function adopted for this regression is a double-sided Crystal Ball (DCB) function, which has a Gaussian core with power-law tails on both sides, defined as:

$$\mathrm{DCB}(y) = \begin{cases} N\, e^{-\xi^2/2}, & \text{if } -\alpha_L \le \xi(y) \le \alpha_R, \\ N\, e^{-\frac{1}{2}\alpha_L^2} \left(\frac{n_L}{\alpha_L} \left(\frac{n_L}{\alpha_L} - \alpha_L - \xi(y) \right) \right)^{-n_L}, & \text{if } \xi(y) < -\alpha_L, \\ N\, e^{-\frac{1}{2}\alpha_R^2} \left(\frac{n_R}{\alpha_R} \left(\frac{n_R}{\alpha_R} - \alpha_R + \xi(y) \right) \right)^{-n_R}, & \text{if } \xi(y) > \alpha_R, \end{cases} \tag{3.7}$$

where $\xi(y) = (y - \mu)/\sigma$, μ and σ denote the mean and width of the Gaussian core, and $\alpha_{L,R}$ and $n_{L,R}$ are the shape parameters governing the left and right tails. The normalization constant N ensures that the total probability integrates to unity. During the training, the regression predicts the DCB parameters $\mu(\vec{x})$, $\sigma(\vec{x})$, $\alpha_{L,R}(\vec{x})$, and $n_{L,R}(\vec{x})$ as functions of the input features $\vec{x}$. The most probable value μ represents the energy correction factor for a given e/γ candidate, while σ corresponds to the estimated per-object energy resolution.

The full electron and photon energy corrections are derived from three successive regressions. The first regression (step 1) applies a correction to the supercluster (SC) energy to account for energy lost in inactive material or module gaps of the ECAL. The regression inputs include the SC energy, its position, and shower-shape observables such as the lateral spread in η and ϕ, the number of saturated crystals, and the ratio R_9, defined as the sum of the energy of the 3×3 crystal matrix centered on the most energetic crystal divided by the total SC energy. An estimate of the transverse energy density in the calorimeter, due to all clusters and tracks in the event, is also used to mitigate effects from pileup.

The second regression (step 2) estimates the energy resolution of the SC. It uses the same set of input variables as step 1, but the regression target is now the ratio between the true and reconstructed energy. The output is an estimate of the per-object energy uncertainty $\sigma_E(\vec{x})$. This correction improves the modeling of the ECAL resolution, especially in the presence of intercalibration variations among crystals. The training employs two simulated samples: one assuming ideal intercalibration constants ("ideal IC") and one incorporating expected random variations ("real IC"), allowing the regression to learn how to mitigate calibration-related fluctuations.

A third regression step (step 3) is applied only to electrons and combines the ECAL-based energy measurement with the momentum estimate obtained from the tracker. Because the ECAL and tracker provide independent information on the electron momentum, a weighted combination of the two measurements is constructed as:

$$E^{\mathrm{reco}}_{\mathrm{combined}} = \frac{E_{\mathrm{ECAL}}/\sigma_E^2 + p_{\mathrm{tracker}}/\sigma_p^2}{1/\sigma_E^2 + 1/\sigma_p^2}, \tag{3.8}$$

where E_{ECAL} and σ_E are the corrected ECAL energy and its estimated resolution, and p_{tracker} and σ_p are the momentum and its resolution measured in the tracker. This combination improves the overall energy estimate, particularly at low transverse energies ($E_T > 15$ GeV) where the tracker momentum resolution exceeds that of the ECAL, as shown in Figure 3.17.

The three regressions are therefore applied sequentially:

1. Step 1 corrects the raw SC energy using the ideal IC simulation.
2. Step 2 estimates the energy resolution based on the real IC simulation.
3. Step 3 combines the ECAL and tracker estimates for electrons only, producing the final corrected momentum and its uncertainty.

Photon energies are corrected using only the first two steps, since there is no corresponding tracker momentum measurement. The two regressions employed for electrons and photons are trained indipendently, in order to take into account their differences in the shower shape variables distributions. The regressions significantly enhance both the energy scale and resolution of electrons and photons, as shown in Figure 3.17. The improvements are most pronounced in regions of the detector affected by material transitions or intercalibration fluctuations, and lead to a better uniformity of the energy response across η and p_T.

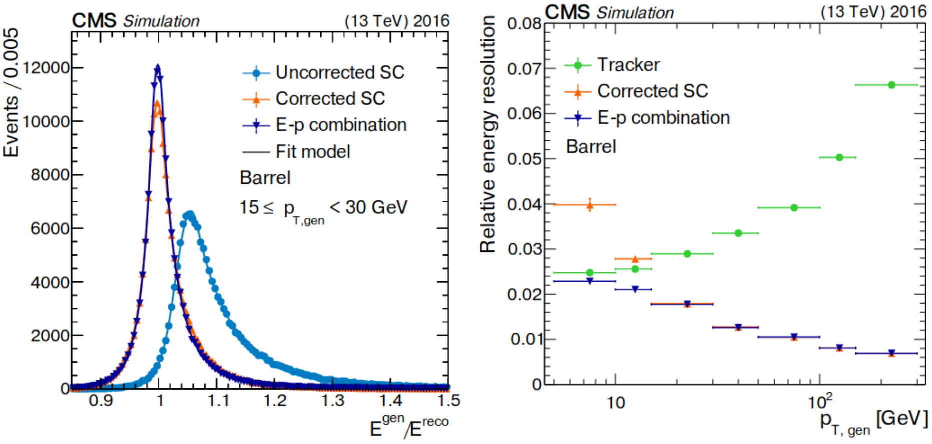

Figure 3.17: Ratio of the true to the reconstructed electron energy (left), in the p_T range 15–30 GeV with and without regression corrections, with a DCB function fit overlaid, and relative electron resolution versus electron p_T (right), as measured by the ECAL (“corrected SC”), by the tracker, and in the combination after the step 3 regression, as found in 2016 MC samples for barrel electrons.

Chapter 4

Diphoton reconstruction for $H \to \gamma\gamma$

4.1 Photon selection

Photons selected for all CMS $H \to \gamma\gamma$ analyses have to satisfy a plethora of requirements, such as trigger, preselection, electron-veto, and multi-variate photon identification. These selection are outlined in the following sections [70, 71].

4.1.1 Trigger

L1 trigger

The diphoton High-Level Trigger (HLT) path employed in this analysis is initiated by at least one electromagnetic candidate identified at the Level-1 (L1) trigger. In the ECAL Barrel (EB), the geometry is organized such that five strips of five crystals each (along the azimuthal direction ϕ) form a *trigger tower* (TT), corresponding to a 5×5 array of crystals in (η, ϕ). In the ECAL Endcap (EE), a similar structure is implemented, though the geometry is more complex due to the X–Y arrangement of crystals. The transverse energy (E_T) from all crystals within a TT is summed by the front-end electronics into a *trigger primitive* (TP), which is transmitted via optical fibres off-detector [72]. The off-detector electronics combines pairs of TPs into L1 trigger candidates over regions of size 4×4 TTs. These are forwarded to the *Global Calorimeter Trigger* (GCT), which selects the four most energetic candidates and transmits them to the *Global Trigger*

(GT). The GT produces the final L1 decision by applying transverse energy thresholds. The electron/photon (e/γ) algorithm operates on a 3×3 TT sliding window:

- The candidate E_T is defined as the sum of the central TP and the largest energy deposit among its four adjacent TTs.
- Electromagnetic (EM) showers exhibit a compact lateral spread. Candidates must satisfy the so-called fine grain requirements, i.e. a central tower with two adjacent strips that together carry at least $\sim 90\%$ of the tower E_T, as shown in Figure 4.1.

In addition, the hadronic calorimeter (HCAL) energy contribution must satisfy:

$$\frac{H}{E} < 5\% \quad \text{for the central tower and candidates with} E_T > 2\text{GeV}.$$

For isolated candidates, an additional isolation condition is imposed: at least five of the eight nearest neighbouring TTs must each have

$$E_T < 3.5\text{GeV}.$$

Once an L1 seed is present, the HLT reconstructs ECAL clusters using the readout units within a rectangular window centered on the L1 candidate. The thresholds and corresponding seed names for the lowest unprescaled single and double L1 seeds are summarized in Table 4.1.

High-level trigger

The HLT selection strategy remains consistent across the three data-taking years, with the exception of the transverse energy requirement on the unseeded leg, set to 18 GeV in 2016 and 22 GeV in 2017 and 2018. A set of general requirements is applied to all photon candidates, while each candidate must additionally satisfy either isolation-based or a high-$R9$ selection [70]. The variables used for selection are H/E, $\sigma_{i\eta i\eta}$, $\text{Iso}_{ph}^{\text{HLT}}$, $\text{Iso}_{track}^{\text{HLT}}$, which represent, respectively, the energy reconstructed in the HCAL TT above the ECAL seed cluster, divided by the photon energy, the width of the photon cluster (measured as number of crystals in the η direction), isolation with respect to other photons or tracks. The isolation of a photon γ with respect to ohher particles is generally defined as the sum of other particles p_T in a cone centered on the photon γ. The summary of HLT requirements is shown in Table 4.2.

	E_T threshold [GeV]
2016	
Single Electron	40
Single Electron (isolated)	30
Double Electron	23, 10
2017	
Single Electron	40
Single Electron (isolated)	32
Single Electron (isolated, $\|\eta\| < 2.1$)	30
Double Electron	25, 14
2018	
Single Electron	40
Single Electron (isolated)	32
Single Electron (isolated, $\|\eta\| < 2.1$)	30
Double Electron	25, 14

Table 4.1: Lowest unprescaled E_T thresholds for L1 trigger seeds for electromagnetic candidates and corresponding seed names for single and double L1 electron candidates in 2016–2018.

	H/E	$\sigma_{i\eta i\eta}$	R_9	Iso_{ph}^{HLT} [GeV]	Iso_{track}^{HLT} [GeV] (unseeded)
EB; $R_9 > 0.85$	< 0.12	–	–	–	–
EB; $R_9 \leq 0.85$	< 0.12	< 0.015	> 0.5	$< (6.0 + 0.012\,E_T)$	$< (6.0 + 0.002\,E_T)$
EE; $R_9 > 0.90$	< 0.1	–	> 0.8	–	–
EE; $R_9 \leq 0.90$	< 0.1	< 0.035	> 0.8	$< (6.0 + 0.012\,E_T)$	$< (6.0 + 0.002\,E_T)$
Other trigger requirements					
HLT s. $E_T > 30$ GeV		HLT us. $E_T > 18(22)$ GeV in 2016(17-18)			
$m_{\gamma\gamma} > 90$ GeV					

Table 4.2: Summary of HLT photon selection requirements, including calorimeter identification, isolation, and other trigger conditions. Photon legs where L1 fired or not are outlined as seeded (s.) or unseeded (us.)

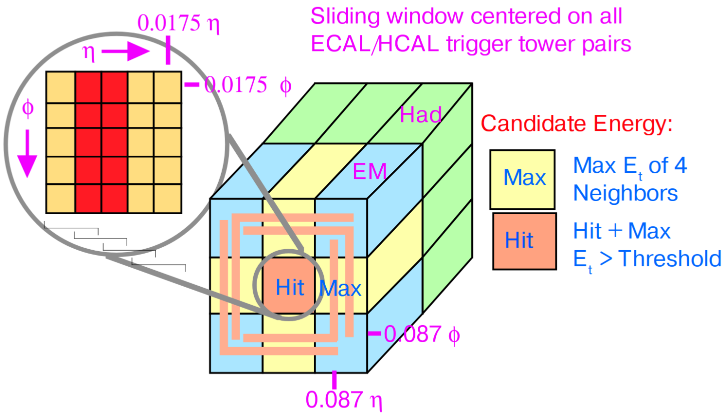

Figure 4.1: Scheme of the L1 electron/photon trigger algorithm. The candidate ET is the sum of the central TP (orange) and highest ET TT from the 4 broadside neighbours (yellow). The Fine Grain profile (left box) and the ratio with HCAL TT energy H/E (green) are used as vetoes while the quiet corners (orange L-shapes) are used to separate isolated from non-isolated candidates. The HCAL TT are aligned with the ECAL TT in pseudorapidity [72].

Trigger efficiency measurements in data

Trigger efficiency of the diphoton HLT path was measured on data using the tag-and-probe technique with $Z \to ee$ events, that work as a control sample in $H \to \gamma\gamma$ analyses because of the high statistic and good understanding of the physics process, together with the high purity of the sample. This is due to the similarity between electron and photon responses in ECAL, and the close energy scales between $Z \to ee$ and $H \to \gamma\gamma$, but comes with non-trivial subtleties, for instance the different shower shape variables distributions between electrons and photon and the different $p - T$ spectra [70]. The idea of tag-and-probe is to use a clean sample of resonant decays where one decay product (the "tag") is required to pass very tight identification and isolation criteria and to fire a single-electron trigger, ensuring that the event is recorded. The other decay product (the "probe") is then unbiased with respect to the trigger under study and can be used to test its efficiency. In this case, the tag is one electron from the Z decay, while the probe is the other electron reconstructed as a photon candidate, mimicking the response of a real photon in the electromagnetic calorimeter. The trigger efficiency is computed as the fraction of probe candidates passing the diphoton trigger, measured separately in bins defined by the probe shower-shape variable R_9, its pseudorapidity $|\eta|$ and the offline transverse energy, as shown in Figure 4.2. Moreover, the offline and HLT electron candidate are matched geometrically.

Since electrons from $Z \to ee$ and photons from $H \to \gamma\gamma$ do not have identical R_9 and $|\eta|$ distributions, the $Z \to ee$ sample is reweighted in these variables so that the probe kinematics and shower properties match those of simulated Higgs boson photons with $m_H = 125$ GeV. This procedure makes the extracted efficiency directly applicable to $H \to \gamma\gamma$ analyses.

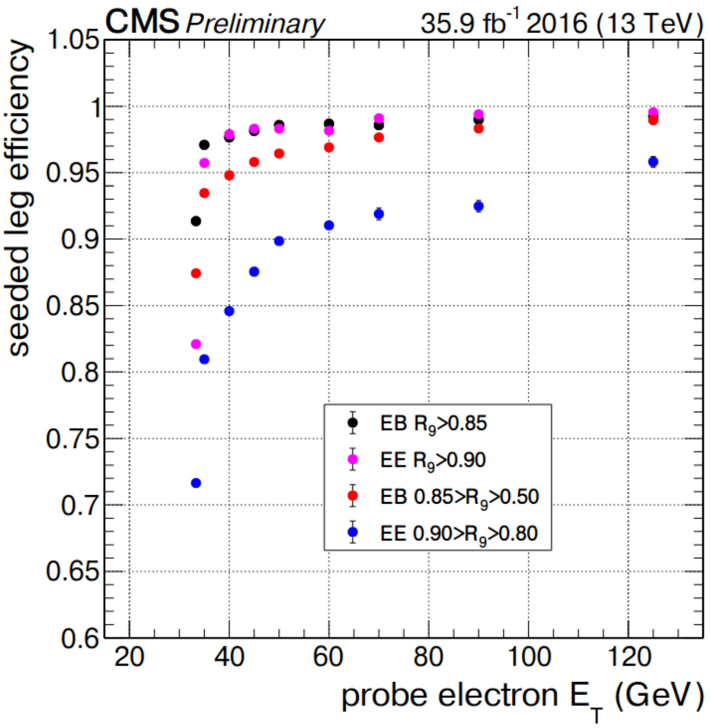

Figure 4.2: Diphoton trigger efficiency measured on 2016 data for $Z \to ee$ events using tag and probe method. Efficiency with respect to offline probe ET, shown in 4 analysis categories defined according to probe R9 and $|\eta|$. The efficiency of the seeded HLT leg is represented, which requires the L1 has been fired.

4.1.2 Photon preselection

The HLT online selection is required both in data and in simulation. However, detectors conditions evolve during time, while the MC samples are generated only once per data-taking era. For this reason the trigger turn-on, as a function of (E_T, R_9, η), is not identical in data and MC, and an additional offline preselection is needed. This offline selection is designed to be slightly more stringent than the HLT requirements, so that the analysis is performed on the plateau region of the trigger efficiency as measured in data. Moreover, the offline selection requires photon candidates within the geometrical acceptance ($|\eta| < 2.5$), excluding the barrel–endcap transition region $1.4442 < |\eta| < 1.566$).

The photon with the largest transverse momentum (leading photon) is required to have $p_T > 35$ GeV, while the subleading photon must have $p_T > 25$ GeV. Both photons are required to pass a "conversion-safe" electron veto, employing the pixel detector to veto electrons without removing photons undergoing conversions. The preselection criteria are summarized in Table 4.3. The selected diphoton pair is finally required to satisfy $m_{\gamma\gamma} > 100$ GeV.

Region	H/E	$\sigma_{i\eta i\eta}$	R_9	Iso_γ [GeV]	Iso_{track} [GeV]
EB, $R_9 > 0.85$	< 0.08	–	> 0.5	–	–
EB, $R_9 \leq 0.85$	< 0.08	< 0.015	> 0.5	< 4.0	< 6.0
EE, $R_9 > 0.90$	< 0.08	–	> 0.8	–	–
EE, $R_9 \leq 0.90$	< 0.08	< 0.035	> 0.8	< 4.0	< 6.0

Table 4.3: Offline photon preselection requirements for barrel (EB) and endcap (EE) regions. The conditions are split according to the R_9 category.

As shown in Figure 4.2, the trigger efficiencies are very steep around 35 GeV, which, together with the preselection, has the effect of decreasing the combined trigger-preselection efficiency of $gg \to H \to \gamma\gamma$ diphotons to about 84%, 57% 77%, when photons are, respectively, both in the barrel, one in barrel and one in endcaps, both in endcaps, with a global inclusive 74% efficiency.

4.1.3 Multivariate photon identification

The goal of the final photon identification (photonID) is to distinguish *prompt photons* from *non-prompt photons*. Non-prompt photons mainly originate from high-momentum neutral mesons, inside jets, decaying into two photons, where both photons are reconstructed as a single photon candidate because they share the same supercluster. There are also contributions from neutral hadrons or charged hadrons (with failed track association) that deposit most of their energy in ECAL. On the other hand, photons produced directly from hard interactions are labeled as prompt photons. To separate these two classes of photons, a BDT is applied after the preselection. The BDT is trained on $H \to \gamma\gamma$ and photon+jets simulated datasets at $\sqrt{s} = 13$ TeV, taking as signal the $H \to \gamma\gamma$ photons, and as background the non-prompt ones in the γ+jets sample which arise from jets, and most importantly, meson decays. In order to make the training independent of the photon kinematic variables, a two-dimensional reweighting is performed as a function of the supercluster pseudorapidity η_{SC} and the photon candidate transverse momentum p_T. This procedure matches the kinematic distributions of signal prompt photons to those of background non-prompt photons. Separate trainings are performed for the 2016, 2017, and 2018 ultra-legacy datasets and for the barrel (EB) and endcap (EE) regions. The input variables include kinematic, shower-shape, and isolation observables, as well as a few additional ones:

- **Kinematic variables:** η_{SC};
- **Shower-shape variables:** $\sigma_{i\eta i\eta}$, $\mathrm{COV}_{i\eta i\phi}$, $E_{2\times2}/E_{5\times5}$, R_9, σ_η, and σ_ϕ;
- **Isolation variables:** photon isolation and track isolation;
- **Additional variables:** ρ_{PU}, the supercluster raw energy E_{RAW} (used in both EB and EE);
- **Endcap-only variables:** preshower spread σ_{RR}, and $E_{\mathrm{ES}}/E_{\mathrm{RAW}}$ (the energy deposited in the preshower divided by the supercluster raw energy),

where $\mathrm{COV}_{i\eta i\phi}$ is the covariance of the single crystal η and ϕ in terms of crystal cells within the 5x5 crystal matrix centred on the supercluster seed, $\sigma_\eta(\sigma_\phi)$ is the logarithmic energy weighted standard deviation of single crystal

η within the supercluster, and ρ is the average energy density in ECAL due to pileup interactions, which is useful to make the pdotonID efficiency independent on pileup.

The distribution of two example variables, i.e. R_9, track isolation and $\sigma_{i\eta i\eta}$, is shown in Figure 4.4. Track isolation and $\sigma_{i\eta i\eta}$, in particular, have different distributions for prompt photons and non-prompt photons, due to, respectively, association with hadron tracks, and non-zero distance in η between the two prompt photons coming from neutral mesons.

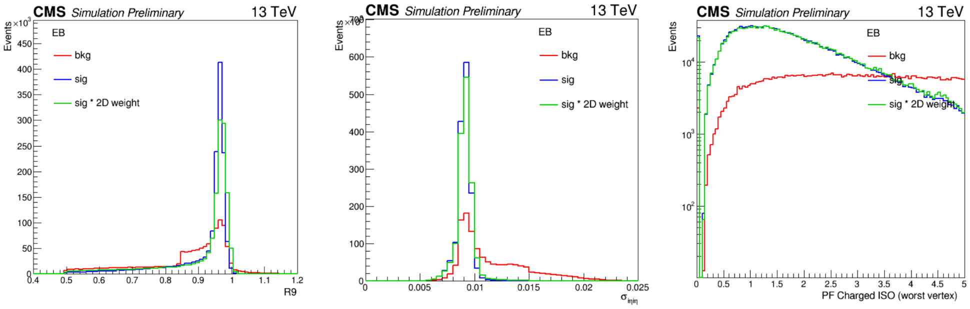

Figure 4.3: Distribution of R_9, track isolation and $\sigma_{i\eta i\eta}$ for prompt and non-prompt photons.

The BDT score distribution for signal ($H \to \gamma\gamma$) and unmatched reconstructed photons from the $\gamma+$ jets MC sample is shown in Figure 4.4.

The data/MC agreement of the photon ID BDT score was validated using both $Z \to ee$, as shown in Figure 4.5 and $Z \to \mu\mu\gamma$ events.

The selection on the single photon employed in $H \to \gamma\gamma$ Run 2 analyses is $\text{BDT}_\gamma > -0.9$, i.e. very loose, because, after the photon selection, another BDT, labelled "diphoton BDT", distinguish between signal and background diphotons, leveraging the better discrimination provided by the composite diphoton object, and employing as an input the photon ID score. After the photonID loose selection, the background consists mainly in γ+jets events, multijet events, both with unmatched reconstructed photons, and $pp \to \gamma\gamma$ events (prompt diphotons), where, before the diphoton BDT application , the γ+jets events contribution is much larger than the prompt diphoton contribution.

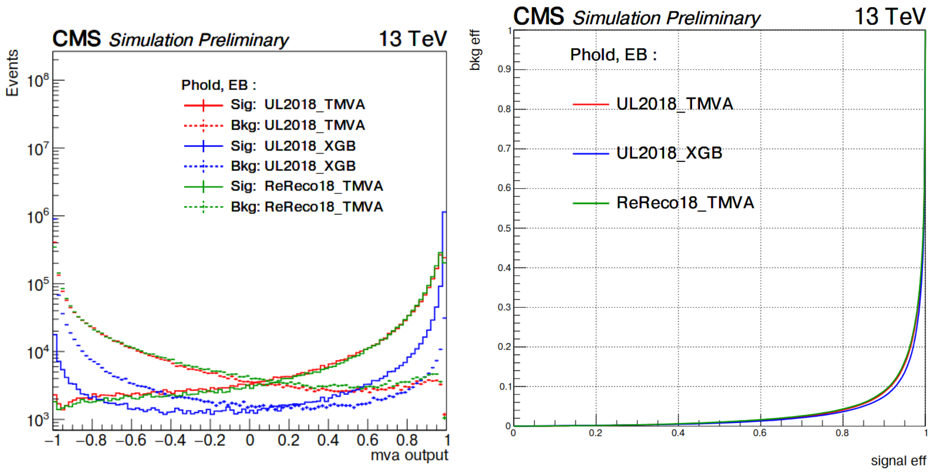

Figure 4.4: BDT score distributions for signal and background photons in the barrel. The chosen algorithm is Xgboost (XGB) trained on 2018 best calibration ”UltraLegacy” (UL) conditions MC samples (red).

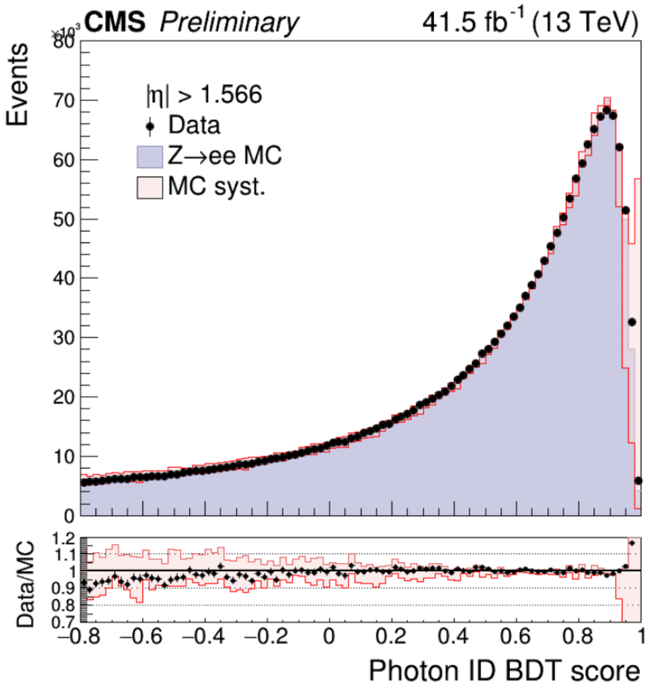

Figure 4.5: Comparison between photon ID BDT output scores distribution between data and MC for the 2017 $Z \rightarrow ee$ sample

Note that the simulation of the unmatched reconstructed photons from multijet and γ+ jet processes is not completely reliable, because their cross-sections are very large, but the jet into photon mis-identification probability is very small $O(10^{-4})$, therefore both the statistical and systematic component of the MC uncertainties are too large for a reliable estimation of this background contribution, which is estimated from data.

Non-prompt photons sources in γ+jets

It is interesting to understand why it is difficult to separate prompt from non-prompt photons with only calorimetric information, by using a small sample of simulated γ + jets events. Reconstructed photons are selected based on photonID, and diphotons are selected requiring $|\eta| < 2.5$, $p_T >$ 35 GeV for the leading and $>$ 25 GeV for the sub-leading photon. In a sample with about 100×10^3 γ+jets generated events, a set of 37.8×10^3 events contain a diphoton. Within this set, only about 7000 have a generated particle matched to both reconstructed photons. Due to the nature of the sample, these diphotons always contain a non-prompt photon, and, given their matching to a generated particle, it is possible to investigate their source by looking at their mother in the Monte Carlo generated particle list of the event. It is found that, out of these 7000 events, 93% contain a non-prompt photon whose mother is a meson, as shown in Figure 4.6. Out of these mesons, the most frequent, as could be expected, are π^0, η and others with known radiative decay pathways.

It is also possible to investigate the source of the non-prompt photons for the 38×10^3 events without generator matching, by looking at unstable generated particles in a cone $DR = 0.1$ centered on the non-prompt photon position. Out of 37.8×10^3 events, A fraction of 76% of the candidates contains at least one unstable generated particle in the cone, with a particle type distribution similar to the previous case, as shown in Figure 4.7.

In particular, 19.2×10^3 events have a π^0 in the cone, and for those events the distribution of π^0 p_T versus minimum opening angle is plotted, as shown in Figure 4.8. The typical opening values are around ΔR=0.01-0.02, while a single cell of ECAL, except in the most inner regions of EE, has a granularity $\delta\eta \times \delta\phi = 0.0175 \times 0.0175$, which explains how difficult is to reject these type of background. This also explains the importance of employing isolation variables for photonID.

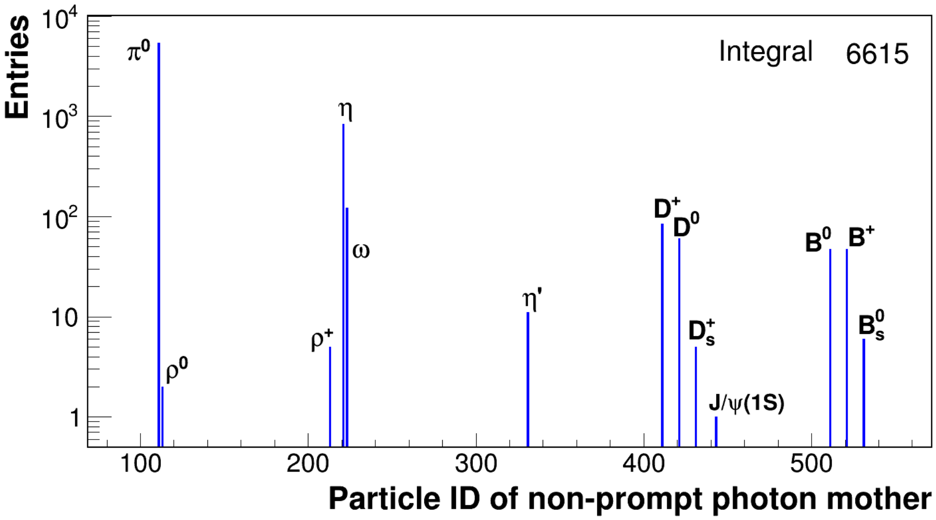

Figure 4.6: Particle ID distribution of the mother of non-prompt photon with generator matching.

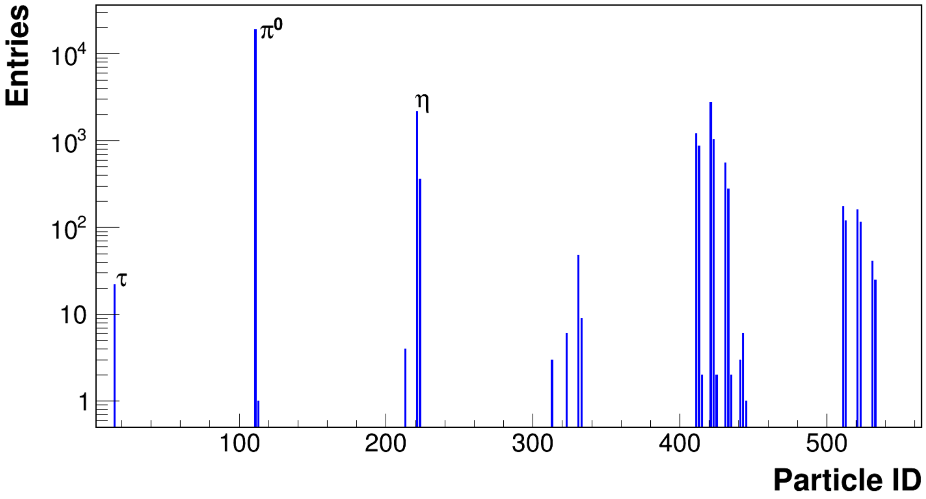

Figure 4.7: Particle ID distribution of unstable generated particles in a cone $DR = 0.1$ centered on the non-prompt photon position.

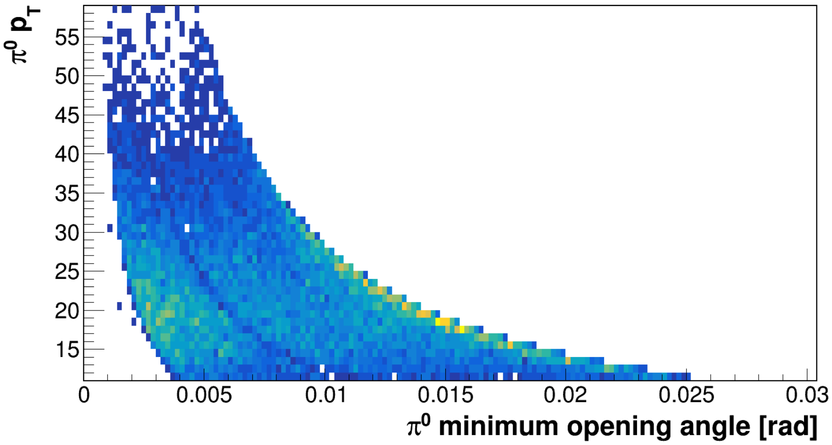

Figure 4.8: Distribution of p_T versus minimum opening angle for neutral pions in a cone $\Delta R = 0.1$ around non-prompt photons with failed generator matching.

4.2 Photon corrections

4.2.1 Photon energy scale and resolution

After applying the standard ECAL calibration procedure, residual differences remain between the energy scale and resolution of electrons and photons in data and simulation [71]. The differences in scale are corrected in the data, and the differences in resolution are corrected in MC by smearing the energy values in the simulation. The differences between photon and electron energy scales are addressed in further steps with dedicated systematics.

Two complementary techniques are employed to determine the energy scale and smearing corrections in the 2016, 2017, and 2018 datasets: the *median method* and the *global smearing method*.

Median method

The first approach relies on the median value of the reconstructed Z boson invariant mass, obtained from dielectron events where both electrons are reconstructed using supercluster energies corrected by the energy regression. The median of the invariant mass distribution is evaluated in data, and the corresponding scale correction is defined such that the median value is aligned with the nominal Z boson mass reported by the PDG. The scale corrections are derived as a function of the run number and the supercluster pseudorapidity $|\eta|$, since residual differences between data and simulation are both time-dependent and $|\eta|$-dependent due to radiation damage in the ECAL crystals.

The ECAL barrel and endcap are each divided into two pseudorapidity regions: $|\eta| < 1$ and $1 \leq |\eta| \leq 1.4442$ for the barrel, and $1.566 \leq |\eta| < 2$ and $2 \leq |\eta| \leq 2.5$ for the endcap. Run bins are defined to contain at least 10,000 dielectron events passing the selection criteria. For each run–$|\eta|$ bin, the scale correction ΔP is defined as:

$$\Delta P = \frac{m_Z}{\mathrm{median}(m_{ee}^{\mathrm{data}})}. \tag{4.1}$$

Typical correction magnitudes range from 1% to 2%, reaching up to 4% in the endcap for 2016 and 2017 data, and up to 5% in 2018.

Global smearing method

Once the median energy scale corrections have been applied, residual differences in scale and resolution are extracted simultaneously using the *global smearing method*. This approach targets both the remaining scale offsets and the ECAL energy resolution. The method relies on a maximum-likelihood fit between data and simulation. By using the simulated Z line shape as a reference, all detector effects, such as reconstruction efficiency and kinematic acceptance, are properly taken into account.

Residual discrepancies between data and simulation are parameterized by a small additional energy scale shift and by a Gaussian smearing term that accounts for differences in resolution. One advantage of the global smearing approach is the possibility of performing a simultaneous fit across several electron categories, deriving per-electron corrections, by minimizing simultaneously the corrections on single electrons in multiple categories. With N single-electron categories, $N(N + 1)/2$ dielectron invariant mass spectra are constructed in both data and simulation. A simultaneous likelihood fit is performed across all categories as a function of $2N$ parameters, $(\Delta P, \Delta\sigma)$. In the fit, electron energies in data are scaled by a factor $(1 + \Delta P)$, while MC energies are multiplied by a random factor sampled from a Gaussian distribution centered at 1 with width $\Delta\sigma$.

The scale corrections are determined in five $|\eta|$ bins and between 10 and 13 R_9 bins, depending on the $|\eta|$ region. This fine binning in R_9 ensures that the corrections are independent of the effective R_9 distributions in the analysis categories, and removes, in first approximation, the difference in scale between electrons and photons, which have different shower shape variables distributions. The smearing terms are derived in coarser bins, with four $|\eta|$ bins and two R_9 bins. In addition, an extra set of corrections for the electron energy scale is derived as a function of $|\eta|$ and E_T using the same global smearing procedure. These corrections ensure good agreement between data and simulation, as shown in Figure 4.9 for an example category. The assigned systematics on the smearing is a flat uncertainty of 0.05% on the relative mass resolution, regardless of the $R_9/|\eta|$ bin, except for 2016 and 2017, where an inflated uncertainty of 0.8% is applied for all photons with $p_T > 50$ GeV, to cover observed non-closures.

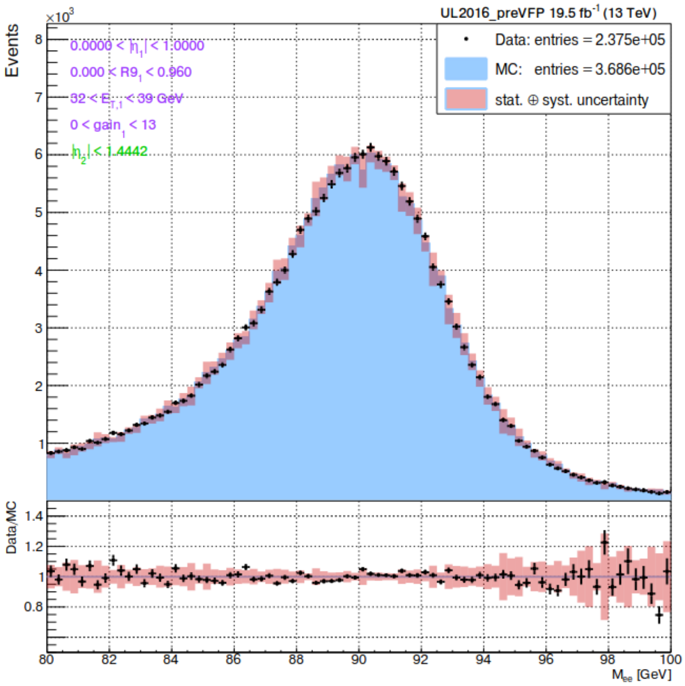

Figure 4.9: Data to MC agreement for $Z \to ee$ events after scale and smearing corrections for an example category in the 2016 dataset.

Frontal non-uniformity

An additional correction is applied to account for the impact of radiation damage on the light collection of electrons and photons in the ECAL.

When a particle interacts with a $PbWO_4$ crystal, it deposits energy along its path, producing scintillation light proportional to the local energy deposition, with an average number of optical photons per deposited energy of 100γ/MeV. This scintillation light propagates through the crystal (with around 22mm^2 face area and > 22 cm thickness), undergoing absorption, scattering, and partial reflection at the crystal surfaces. Only a small fraction of the light produced by scintillation in the crystals (100 γ/MeV) reaches the two $5 \times 5\text{mm}^2$ photodetectors, generating a measurable signal in the ECAL readout, corresponding to around 4.5 photoelectrons / MeV, because of the sensitive/crystal face area ratio of 10%, APD quantum efficiency of 60%, and 75% light collection efficiency, measured on non-irradiated crystals.

The light collection efficiency can be described also as a function of z, as LCE(z), which depends on the depth z at which the light is emitted. The overall ECAL response can therefore be expressed as:

$$S = \frac{\int E_{\text{dep}}(z)\,\text{LCE}(z)\,dz}{\int E_{\text{dep}}(z)\,dz}, \tag{4.2}$$

where LCE(z) also includes the efficiency of the photodetectors. Different types of particles (photons, electrons, hadrons, jets) with the same total energy exhibit different average longitudinal energy-deposit profiles $E_{\text{dep}}(z)$, leading to different ECAL responses.

Radiation damage degrades light collection efficiency due to (i) loss of crystal transparency, (ii) reduced APD gain in the barrel (EB), and (iii) reduced photocathode sensitivity in the endcaps (EE). Crystal transparency loss is described by an induced absorption coefficient μ_{induced}, which causes exponential attenuation of light along the crystal length L:

$$T_{\text{final}} = T_{\text{initial}} \cdot e^{-(\mu_{\text{natural}}+\mu_{\text{induced}})L}, \tag{4.3}$$

where T_{initial} and T_{final} are the initial and final light transmissions, respectively, and μ_{natural} is the absorption coefficient of the undamaged crystal.

Because of differences in shower profiles between electrons and photons, and because the light collection efficiency varies over time due to radiation damage, the ECAL response to electrons and photons differs and evolves during data taking. The ECAL energy scale calibration is derived from electrons from $Z \to e^+e^-$ decays and applied to both electron and photon objects. Hence, the response difference between electrons and photons must be corrected.

Since radiation damage is not simulated, the energy regression can only correct for intrinsic shower-shape differences, not for time-dependent light-collection effects. An additional photon-specific scale correction, known as the *Front Non-Uniformity* (FNUF) correction, is therefore applied. It is defined as:

$$\frac{1}{F} = \frac{S_\gamma}{S_e} = \frac{\dfrac{\int E^{\gamma}_{\text{dep}}(z)\,\text{LCE}(z)\,dz}{\int E^{\gamma}_{\text{dep}}(z)\,dz}}{\dfrac{\int E^{e}_{\text{dep}}(z)\,\text{LCE}(z)\,dz}{\int E^{e}_{\text{dep}}(z)\,dz}}, \tag{4.4}$$

where $E^{e}_{\text{dep}}(z)$ and $E^{\gamma}_{\text{dep}}(z)$ represent the average shower profiles of electrons and photons of the same energy. This correction compensates for the fact that photons typically deposit most of their energy in the front part of the crystal (within $\sim 13X_0$), where the radiation damage affects the transparency most severely.

Assuming a constant instantaneous luminosity of 2×10^{34} cm^{-2}s^{-1}, the corrections is evaluated using FLUKA [73] simulations to predict μ_{induced} and the corresponding transparency loss R/R_0 as a function of pseudorapidity ($0 < |\eta| < 3$ in 0.1 bins) and integrated luminosity. For each luminosity scenario and $|\eta|$ bin, the corresponding LCE(z) is then simulated using S-Litrani [74], incorporating the APD and VPT gain losses.

Although the assumed luminosity scenario may differ from that of actual data taking, the correction depends only on the measured transparency ratio R/R_0, which is derived directly from laser monitoring data rather than simulation. The average measured R/R_0 values as a function of run number are shown in Figure 4.10.

Figure 4.11 illustrates the resulting LCE(z) profiles and the computed FNUF corrections $1/F$ as a function of photon energy and R/R_0.

The FNUF correction is applied on top of all other photon energy calibrations in data, using the average R/R_0 in each $|\eta|$ bin from laser monitoring. This is the reason why the correction is lower than 1, i.e. because the application of the $Z \to ee$ scale over-corrects photons, which start their shower deeper, in regions of the crystals where the integrated dose, and the transparency loss, is less pronounced.

It is applied only to high-R_9 photons ($R_9 > 0.96$), since low-R_9 photons have shower shapes more similar to those of electrons.

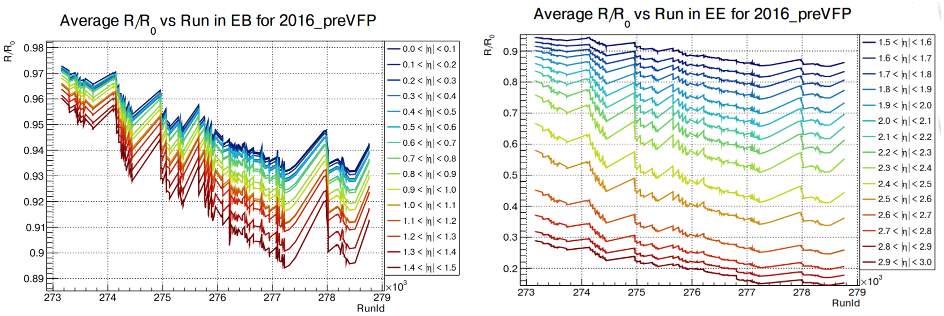

Figure 4.10: Average R/R0 from laser measurements as a function of the run, for different pseudorapidity bins in EB (left) and EE (right).

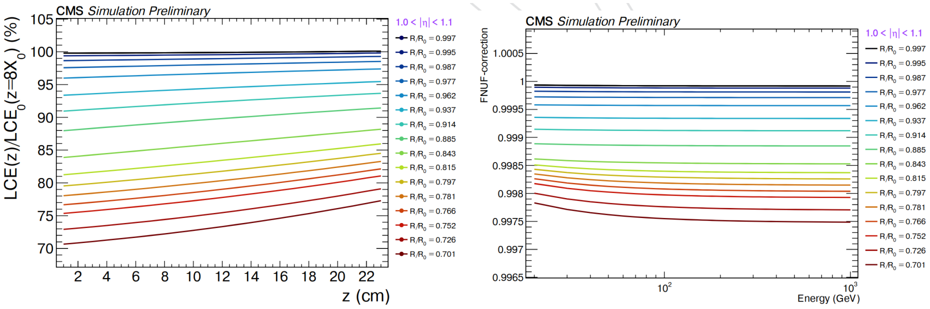

Figure 4.11: Simulated light collection efficiencies as function of crystal depth for different η regions and R/R_0 (left) and final "FNUF-correction", 1/F, to be applied to each photon as function of energy and the average R/R_0 (right), in an example $|\eta|$-bin.

Since the FNUF correction depends directly on R/R_0, any modification of R/R_0, e.g., via the corrections applied to laser amplitudes via pin-diodes (PN) monitoring the laser light output, propagates to the final energy scale. PN corrections typically increase the effective transparency loss, leading to a slightly higher FNUF correction. Their impact on the diphoton mass scale is evaluated by reapplying the FNUF correction using half of the PN corrections as input and comparing the fitted $m_{\gamma\gamma}$ peak positions.

4.2.2 Shower shape, preshower and isolation corrections

Some detector effects are not perfectly simulated and therefore require alignment with data using multivariate corrections. These effects are typically dependent on the year, but the $Z \to \mu\mu\gamma$ sample does not have enough statistics to allow it, hence it is done on the $Z \to ee$, because the discrepancies do not depend on the type of particle.

The goal of the procedure is to achieve the best possible agreement between data and simulation for the photon identification MVA output. To this end, the correction must be performed differentially in the transverse momentum p_T, pseudorapidity η, azimuthal angle ϕ, and the per-event pileup energy density ρ. In addition, correlations among the shower shape variables and among the charged isolation variables are taken into account. The distributions of the shower shape, preshower and isolation variables in simulation are then morphed so as to reproduce as closely as possible the ones observed in data.

The morphing of simulation distributions to match those observed in data is performed using the cumulative distribution functions (CDFs) of both the simulated and target (data) distributions. For each variable to be corrected, the value in simulation is replaced by the corresponding value in data that satisfies the condition that their CDFs are equal. This ensures a consistent mapping between the distributions.

This procedure is applied to all shower shape variables used as inputs to the photon identification MVA, while preserving the interdependencies among these variables and their correlations with event topology. To achieve this, a series of boosted decision trees (BDTs) is employed to model the conditional shape of the CDFs in both data and simulation.

The variables used as inputs to each BDT are selected in a chained fashion: for simulation, variables that have already been corrected are included as additional inputs when training the BDTs for subsequent variables. This method is therefore called chained quantile regression (CQR) [24]. In data, no correction is needed, so the same (uncorrected) inputs are used for each variable. This necessitates a specific correction order, which is applied consistently to both data and simulation. For each variable, 21 BDTs are trained to predict the values at specific quantile points of the CDF, i.e. 19 points from 0.05 to 0.95, spaced by 0.05, with the addition of 0.01 and 0.99. Each BDT is trained using a quantile loss function appropriate to its target quantile. To evaluate values between the trained quantile points, linear interpolation is applied. For isolation variables, the procedure must be extended to account for discontinuities in their distributions. The training of the boosted decision trees (BDTs) is performed separately for each of the three years of data-taking. The data/MC agreement before and after this correction for 2018 EB $Z \rightarrow ee$ events is shown in Figure 4.12 for two example shower-shape variables. Despite the imperfect agreement for the single variables, the agreement of the photon ID score after CQR corrections is acceptable, as shown in Figure 4.13, where the target is an impact less than 2% on the inclusive cross-section.

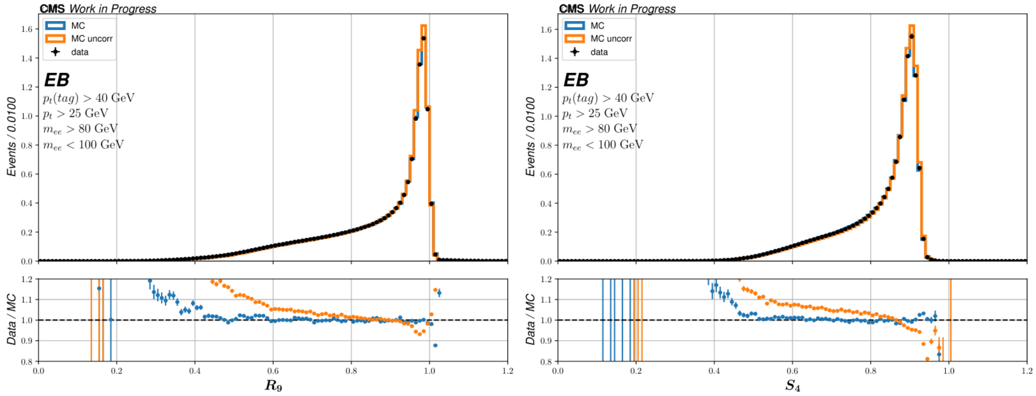

Figure 4.12: Data/simulation comparison for the shower shape variables R_9 and S_4 for 2018 EB with corrected and uncorrected distributions from simulation.

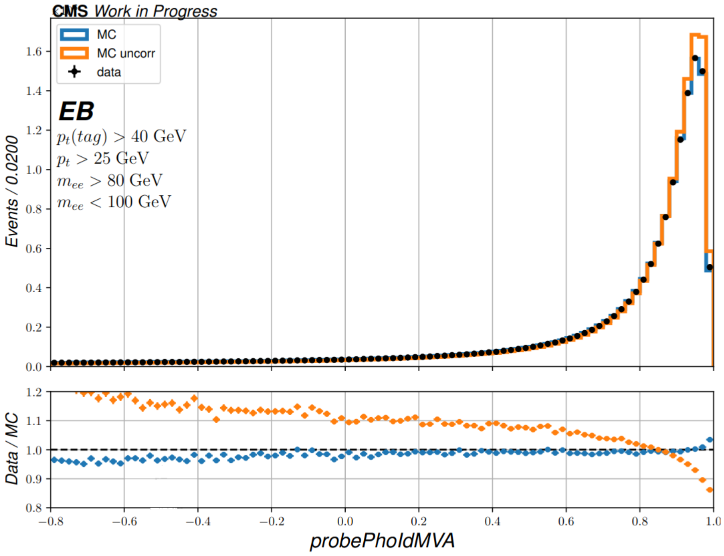

Figure 4.13: Data/simulation comparison for the photon ID score for 2018 EB with corrected and uncorrected distributions from simulation.

4.3 Primary vertex reconstruction and selection

Since events involving two photons do not necessarily contain prompt charged particles originating from the hard scattering, the precise determination of the primary interaction vertex along the beam axis is not straightforward. However, even for the dominant gluon-fusion production process, the Higgs boson is often accompanied by additional jets or, at minimum, soft hadronic activity from the underlying event. This additional activity enables the reconstruction of the vertex corresponding to the Higgs boson production in nearly all cases. Nonetheless, in the presence of a large number of pileup interactions, the identification of the correct vertex associated with the diphoton system can still be ambiguous [75]. The difficulty increases for ggH at low p_T, which is the case treated in this thesis.

The overall strategy consists of reconstructing all primary vertices in the event and subsequently selecting the one most compatible with the diphoton candidate. This selection exploits both the kinematic properties and spatial correlation of the associated hadronic activity, as well as any available information from photon conversions reconstructed in the tracking detectors.

Primary vertices are reconstructed from an inclusive set of charged-particle tracks measured in the tracking detectors. Tracks are required to satisfy quality and beamline compatibility criteria to suppress contributions from fake tracks or secondary decays. The tracks are then grouped into clusters along the z-axis of the beamline, corresponding to distinct primary vertices. The optimal division into N clusters is determined by minimizing the global χ^2 compatibility between the track z positions and uncertainties, assuming N common vertices along the beamline.

Each cluster is then fitted to a common three-dimensional vertex using an adaptive vertex fitting algorithm, which iteratively down-weights tracks that are inconsistent with the common vertex. The typical spatial resolution of this fit is of order $\mathcal{O}(50\ \mu\mathrm{m})$ in the transverse plane and $\mathcal{O}(100\ \mu\mathrm{m})$ in the longitudinal direction. Since the transverse beam width is comparable to the vertex position resolution, the known position and width of the luminous region are included as external constraints to further improve the transverse resolution. The reconstruction output consists of a list of vertices, each characterized by a fitted position, associated uncertainties, and the set of tracks assigned to it. In general, the vertex corresponding to the hard interaction exhibits larger hadronic activity than those arising from pileup. Moreover, the hadronic recoil in the transverse plane is expected to approximately balance the momentum of the diphoton system.

Use of photon conversions

When one or both photons are associated with reconstructed conversions in the tracker, their conversion vertices provide additional pointing information that can improve the determination of the production vertex [75]. The reconstructed conversion momentum vector is extrapolated from the conversion vertex back to the beamline. However, since the resolution of the conversion reconstruction depends strongly on its radial position, conversions far from the interaction region yield poorer estimates due to limited tracking measurements. In such cases, an alternative approach is used, which projects the line connecting the reconstructed conversion vertex to the corresponding ECAL supercluster position back to the beamline. The inclusion of ECAL position information significantly improves the overall pointing resolution.

An uncertainty is assigned to each conversion-based z estimate, depending on the detector region where the conversion occurred. For events with two converted photons, a weighted average of the two z estimates is computed using their uncertainties. For each reconstructed vertex, a compatibility variable is then defined:

$$\text{ConversionPull} = \frac{|z_{\text{conv}} - z_{\text{vtx}}|}{\sigma_z}, \tag{4.5}$$

where z_{vtx} is the vertex z position, z_{conv} is the z estimate from conversion pointing, and σ_z is its associated uncertainty. The uncertainty in z_{vtx} is negligible compared to that of the conversion measurement.

Multivariate selection

The final vertex selection is performed using a boosted decision tree (BDT) classifier trained to distinguish the true diphoton vertex in simulated Higgs events from pileup vertices. The input features include three so-called hadronic-recoil variables:

- **sumpt2** $= \sum p_{T,i}^2$: The sum of the squared transverse momenta of all charged tracks associated with the vertex. This variable measures the overall amount of hadronic activity,
- **ptbal** $= -\sum p_{T,i} \cdot \hat{p}_T^{\gamma\gamma}$: The sum of the transverse momenta of all charged tracks associated with the vertex, projected onto the axis of the diphoton momentum in the transverse plane. This quantifies how much the charged hadronic activity balances the transverse momentum of the diphoton system,
- **ptasym**: An asymmetry variable formed from the total transverse momentum of all charged tracks associated with the vertex and the transverse momentum of the diphoton system. It provides an additional measure of the balance between the charged hadronic activity and the diphoton system,

as well as the conversion compatibility variable described above. If no conversion is reconstructed, the corresponding input is set to a dummy value so that the classifier effectively ignores it. The vertex with the highest BDT score is then selected for use in photon kinematic reconstruction and identification.

Vertex probability estimate

Because the invariant mass resolution deteriorates significantly when the incorrect vertex is selected, due to the wrong reconstructed angles for the photons, it is important to estimate the probability that the chosen vertex is the correct one, which is an input to the diphoton BDT. A vertex is considered correctly identified if its z position lies within 1 cm of the true interaction point, a value corresponding to an error in the invariant mass tipically much smaller than the one coming from reconstructed photon energies. The per-event probability of a correct selection is estimated using an additional BDT trained to distinguish events with correctly and incorrectly selected vertices, which is an input of the diphoton BDT, that provide discrimination between signal and background.

The classifier uses the following input variables:

- Diphoton transverse momentum, $p_T^{\gamma\gamma}$: higher values correlate with increased hadronic activity and better vertex identification.
- Number of reconstructed primary vertices: events with more vertices are more prone to misassignment.
- Vertex-selection BDT scores for the three highest-ranked vertices: these contain information about the signal-likeness of each candidate.
- Distances in z between the selected vertex and the next two highest-ranked candidates.
- Number of photons associated with reconstructed conversions (0, 1, or 2): conversion information enhances vertex selection efficiency.

The BDT response is linearly transformed to yield an estimated per-event probability for correct vertex assignment, calibrated using the fraction of correctly identified vertices in simulated Higgs events. The per-event probability is shown in Figure 4.14, as a function of p_T for 2018.

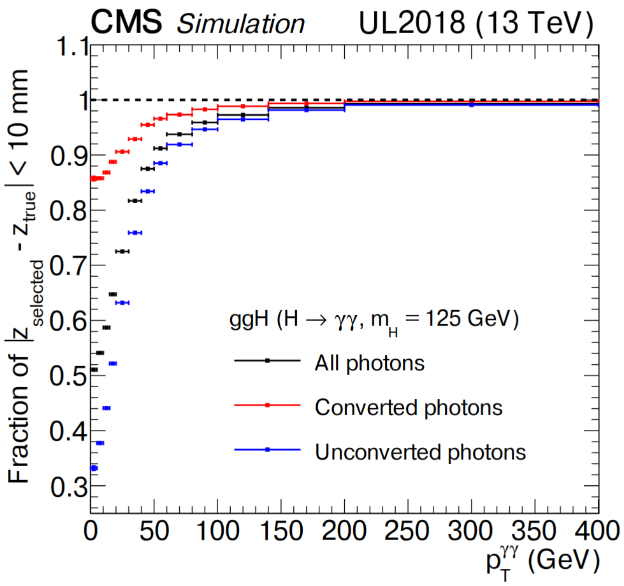

Figure 4.14: Fraction of the events with $|z_{\text{selected}} - z_{\text{true}}| < 10$ mm, i.e. the rate of the correct diphoton vertex, as a function of the diphoton pT, in 2018. Simulated events with $H \to \gamma\gamma$ are used, considering the gluon-gluon fusion production mode.

4.4 Diphoton BDT training

The main sources of background in $H \to \gamma\gamma$ analyses are QCD multijet processes (two fake photons, labelled 'fake-fake') and γ + jets events (one prompt and one fake photon, labelled 'prompt-fake'). The signal sample used for training is obtained as a cross-section-weighted mixture of the four main production mechanisms: gluon fusion (ggH), vector boson fusion (VBF), associated production with a vector boson (VH), and associated production with a top quark pair (ttH), all for a Higgs boson mass of 125 GeV [71].

Data-driven background sample

Monte Carlo simulations are known to significantly underestimate the uncertainties related to photon misidentification, particularly at low photon ID BDT scores. Therefore, the QCD and γ+jets samples exhibit large event weights and limited statistics, resulting in poor modeling of input observables, especially in the tails of distributions. To address these issues, a fully data-driven model is obtained from events that fail a minimum photon ID BDT score requirement of -0.9, defining a “low photon ID sideband” region, as shown in Figure 4.15. This data-driven sample is mixed with a MC sample of prompt diphotons ($pp \to \gamma\gamma$).

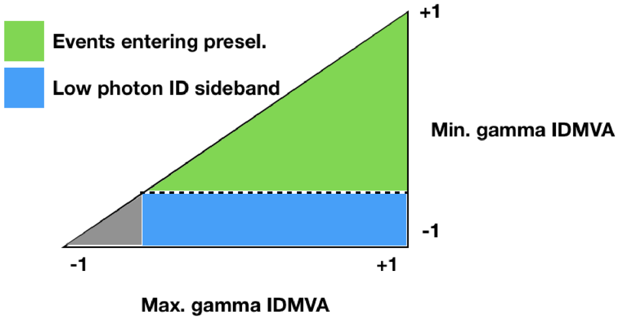

Figure 4.15: Scheme of the low photon ID sideband extraction.

A probability density function (PDF) for the minimum photon ID BDT score is derived using γ+jets Monte Carlo, where fake photons are matched to generator-level jets. This PDF is parameterized by a polynomial fit to the photon ID BDT distribution, as shown in Figure 4.16.

For each event in the sideband region, a new value of the minimum photon ID BDT score is randomly assigned within the range $[-0.9, 1.0]$ according to the fake-photon PDF, and an additional per-event weight is applied to reproduce the expected distribution of the minimum photon ID BDT score:

$$w = \frac{\int_B^A f_{\mathrm{real}}(x)dx}{\int_A^B f_{\mathrm{fake}}(x)dx}, \tag{4.6}$$

where $f_{\mathrm{fake}}(x)$ denotes the fitted PDF for the photon ID BDT score of fake photons.

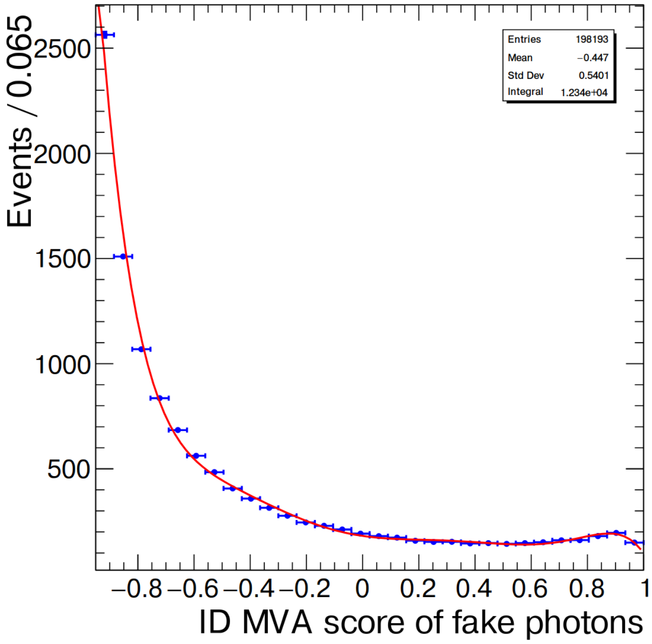

Figure 4.16: Distribution of the minimum photon ID BDT score for each diphoton, with a polynomial fit.

Since the normalization of the sideband-derived sample does not necessarily match the expected number of QCD and γ+jets events in the preselection, a simultaneous fit of the minimum and maximum photon ID BDT distributions is performed to extract normalization scale factors, such that the MC $pp \rightarrow \gamma\gamma$) sample, added to the data-driven low-ID sideband describing the QCD fake contributions, reproduces the data distribution, as shown in Figure 4.17.

Training and validation

Training events are selected using the photon preselection criteria, a loose cut on photon ID greater than -0.9, and transverse momentum thresholds corresponding to $p_T/m_{\gamma\gamma}$ values of $1/3$ and $1/4$ for the leading and subleading photons, respectively. The diphoton BDT is trained by the H$\gamma\gamma$ group using a set of ten input features. These include the kinematic properties of the diphoton system, per-event estimates of the diphoton mass resolution, and per-photon scores from the identification BDT.

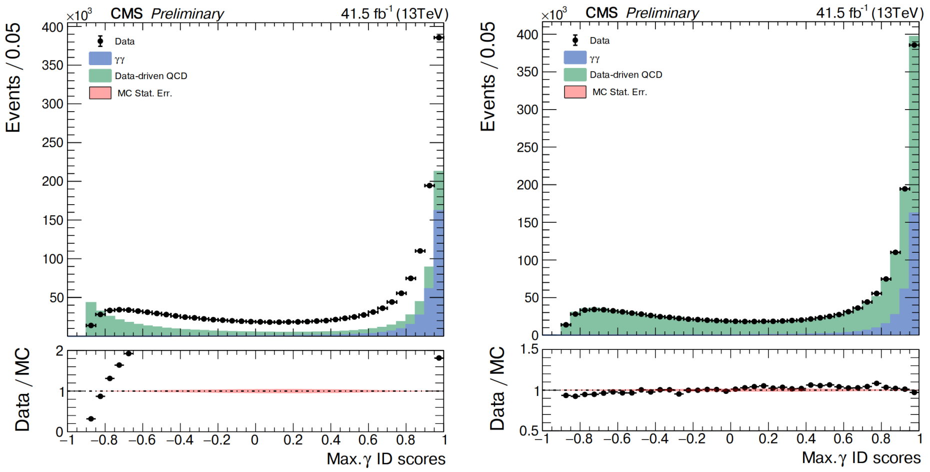

Figure 4.17: Maximum of the photon ID score of the two photons for each diphoton, before (left) and after (right) normalizing the two components.

The per-event relative mass resolution estimates are computed by propagating the photon energy resolution estimates, assuming Gaussian resolution functions. The XGBoost algorithm is employed to train the model and tune its hyperparameters. The resulting shapes of the BDT output scores on the test set are shown in Figure 4.18.

After training, diphoton data and MC events after preselection are used to validate the input variables and the output of the diphoton BDT, using events in the mass sideband regions $100 < m_{\gamma\gamma} < 115$ GeV and $135 < m_{\gamma\gamma} < 180$ GeV, as shown in Figure 4.18.

To check if the mixed sample used for the training is compatible with data, their diphoton invariant mass distributions are compared after a loose cut on the diphoton BDT score at -0.9, blinding the 115 GeV $< m_{\gamma\gamma} <$ 135 GeV region, as shown in Figure 4.19.

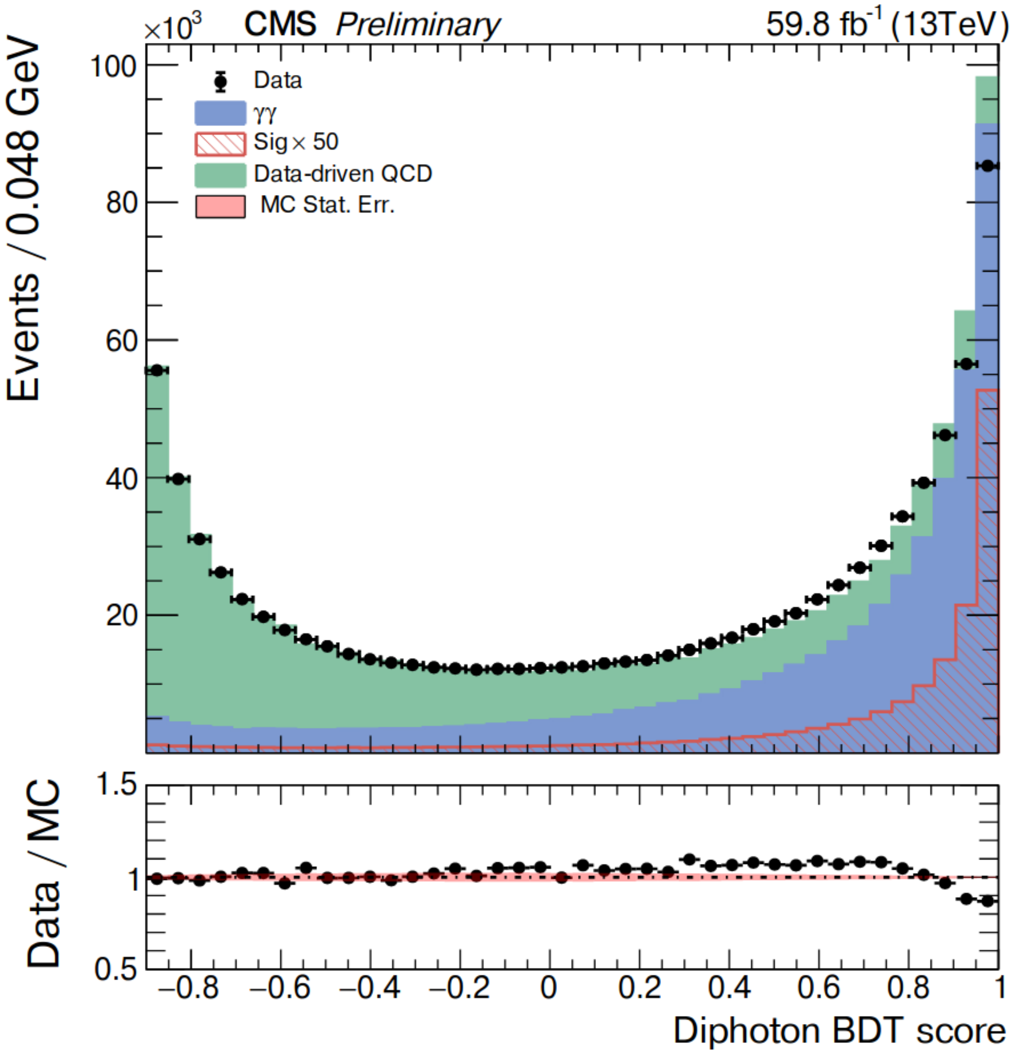

Figure 4.18: BDT score after training, applied to the data, in the invariant diphoton mass sidebands.

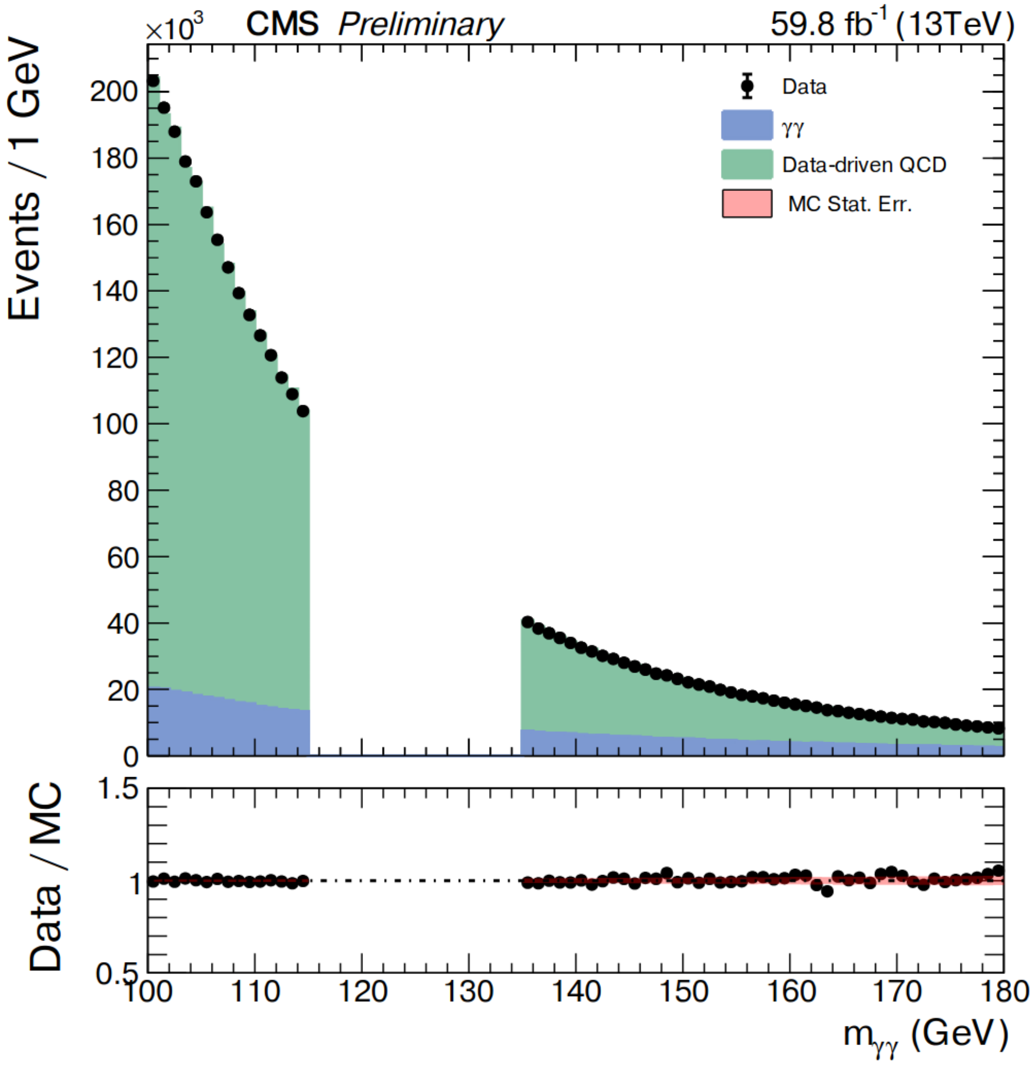

Figure 4.19: Diphoton invariant mass distribution of data and test sample after a loose cut at -0.9 on the diphoton BDT, blinding the signal region.

Part II

Contraints on Γ_H in the $\gamma\gamma$ channel at CMS

Chapter 5

Analysis strategy

5.1 Data and simulated samples

The data samples used in this analysis correspond to an integrated luminosity of 138 fb^{-1} recorded in 2016-2018 at the LHC in pp collisions at a center-of-mass energy of 13 TeV [58, 76, 77].

The signal samples for gluon–gluon fusion (ggH), vector-boson fusion (VBF), and vector-boson–associated Higgs boson production (VH) are simulated with the MadGraph5_aMC@NLO v2.4.2 matrix-element generator [78] at next-to-leading order (NLO) in QCD. For the ggH process, the simulation includes the full set of one-loop QCD diagrams describing the gluon–gluon fusion through a top-quark loop, while for the VBF and VH processes the generator includes the relevant electroweak diagrams at NLO accuracy in QCD. The generated parton-level events are interfaced with Pythia 8.2 [54] for parton showering and hadronization, using the CP5 tune [79]. The NNPDF 3.1 [80] parton distribution functions are employed consistently both in the matrix-element generation and in the shower evolution. The gluon–gluon fusion signal sample is further reweighted in the Higgs boson transverse momentum ($p_{\mathrm{T},H}$) to account for higher-order QCD corrections beyond NLO using the MiNNLOPS generator [81], which provides a consistent matching between NLO matrix elements and the parton shower. This is needed because these samples were prodduces before MiNNLOPS became available.

Samples simulating the interference between gluon-gluon fusion signal and $gg \to \gamma\gamma$ or $qg \to q\gamma\gamma$ background processes are produced separately with NLO precision using the SHERPA v.2.2.11 generator [56, 82], for Higgs boson masses of 120, 125 and 130 GeV and a total width of 4.07 MeV. Note that SHERPA is the only available generator that implements the interference process. All diagrams discussed in Reference [46] including those in Fig. 2.13 are considered for this study. The SHERPA generator also handles parton showering and hadronization. The only parton shower algorithm available is the Catani-Seymour one, which effectively corrects the generation at NLL level. A fit to the generator-level differential cross section with respect to diphoton mass, using the theoretical formula in Eq. (5.2), is shown in Figure 5.1, where the fitted parameters are the numerical coefficients R and I.

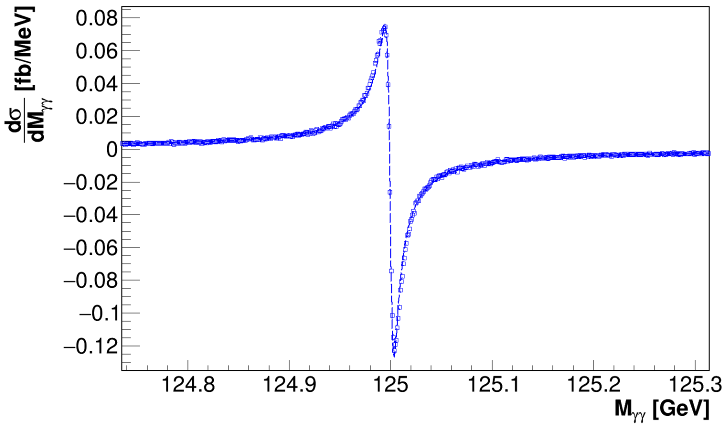

Figure 5.1: Fit to the generator-level differential cross section of the gluon-gluon interference with respect to diphoton mass, using the theoretical formula in Eq. (5.2), in a restricted phase space.

Figure 5.2 shows the comparison and compatibility betweeen the gluon-gluon interference differential cross-section at the generator level, and the one evaluated for 13.6 TeV collisions in a theoretical paper [83] in a well-defined phase-space, requiring $p_T^{\gamma} > 20$ GeV, $\sqrt{p_T^{\gamma_1} * p_T^{\gamma_2}} > 35$ GeV, $|\eta_\gamma| < 2.5$, $\Delta R(\gamma_1, \gamma_2) > 0.4$. The same comparison was not possible for 13 TeV collisions because of the lack of material in literature, and the same applied for quark-gluon interference, shown in 5.3. The inclusive cross-section value at NLO reported by SHERPA is 0.556 fb with a 4% error for the gluon-gluon interference, and 0.021 fb for the quark-gluon one. Note that these values does not have a clear physical meaning, because of the very large fraction of negative weight (50%), and the inclusion of events also at $m_{\gamma\gamma}$ very far from the Higgs boson signal region, which dominate the total cross-section integration.

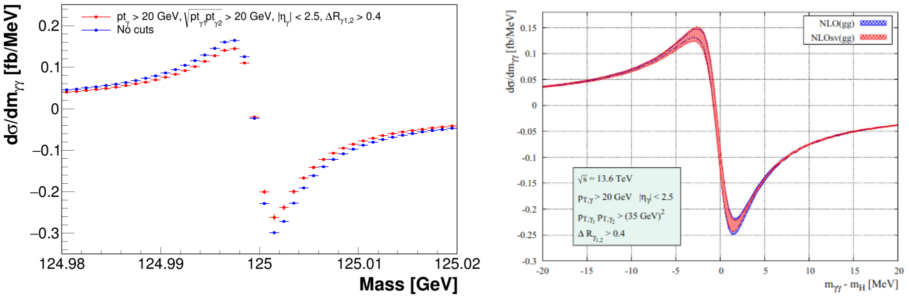

Figure 5.2: Differential cross-section of the gluon-gluon interference with respect to $m_{\gamma\gamma}$, as evaluated in the theoretical paper [83] for 13.6 TeV collisions (right), and by SHERPA NLO (with Catani-Seymour parton shower, at NLL) at 13 TeV (left). The cuts reported in the right plot have been applied to produce the left plot.

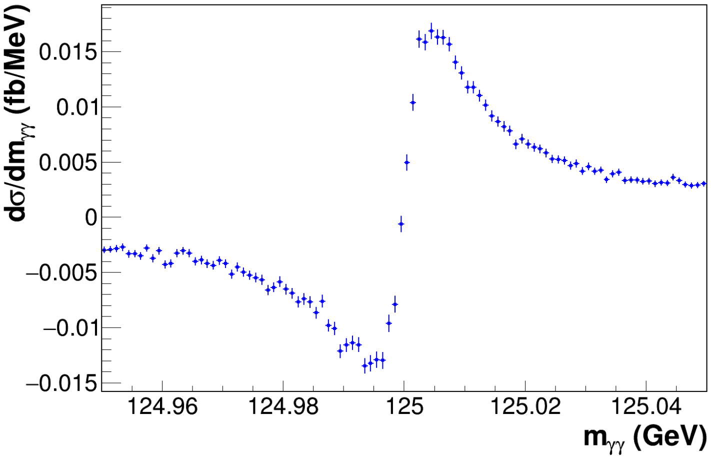

Figure 5.3: Differential cross-section of the quark-gluon interference with respect to $m_{\gamma\gamma}$, evaluated by SHERPA NLO (with Catani-Seymour parton shower, at NLL) at 13 TeV.

The interference cross-section evaluated by SHERPA in the simulated samples employed is multiplied by a global factor (K-factor), equal to $\sqrt{K_S K_B}$, where the signal K-factor was set to $K_S = 1.45$, which effectively rescales the signal prediction from SHERPA to the MiNNLOPS value used for the signal process [19], while the K-factor for the background is fixed to 1, given that there is no indication of K-factors needed.

All simulated samples are propagated through the full CMS detector simulation using the GEANT4 package [84]. Multiple nearly simultaneous proton-proton collisions happening in the same bunch crossing, termed as pileup. The mean number of pileup interactions in data for Run 2 as a whole was found to be 34. Events in the simulated samples are reweighted such that the resulting shape of the pileup distribution in simulation after this procedure matches the one in data.

5.2 Analysis strategy overview

The differential cross section as a function of the di-photon invariant mass $m_{\gamma\gamma}$, for the gluon-gluon fusion signal and interference components, can be written schematically in the NWA [46] as:

$$d\sigma_s^{ggH}/dm_{\gamma\gamma} = \frac{c_{g\gamma}^2 S}{(m_{\gamma\gamma}^2 - m_H^2)^2 + m_H^2\Gamma_H^2}, \tag{5.1}$$

$$d\sigma_i/dm_{\gamma\gamma} = c_{g\gamma}\frac{(m_{\gamma\gamma}^2 - m_H^2)R + m_H\Gamma_H I}{(m_{\gamma\gamma}^2 - m_H^2)^2 + m_H^2\Gamma_H^2}, \tag{5.2}$$

where $c_{g\gamma}$ is the product of Higgs boson couplings to gluon and photons, while S, R and I are numerical coefficients independent of the couplings or the width. The R term changes sign for $m_{\gamma\gamma}$ above m_H and introduces distortions in the $m_{\gamma\gamma}$ spectrum, while the I term affects only the normalization. The cross section of the pure ggH signal term is proportional to $c_{g\gamma}^2/\Gamma_H$. Therefore, by inverting this relation it is possible to parameterize the relationship between the couplings product and the width modifier with respect to the SM value, Γ_H/Γ_H^{SM}:

$$c_{g\gamma} = \sqrt{\mu_F\Gamma_H/\Gamma_H^{SM}} \tag{5.3}$$

where $\mu_F = \mu_{g\gamma}$ is the ratio between the measured ggH signal-only on-shell cross section and its SM value. The differential cross section for the interference $d\sigma_i/dm_{\gamma\gamma}$ has an asymmetric shape with a large negative contribution, as shown in Fig. 5.4 with and without including experimental resolution, dominated by photon energy uncertainty. When experimental resolution is taken into account, the asymmetrical shape is still visible but gets smeared along the $m_{\gamma\gamma}$ axis. The asymmetric shape is not strongly modified by the resolution effects, and then it can be experimetally exploited to constrain the total width.

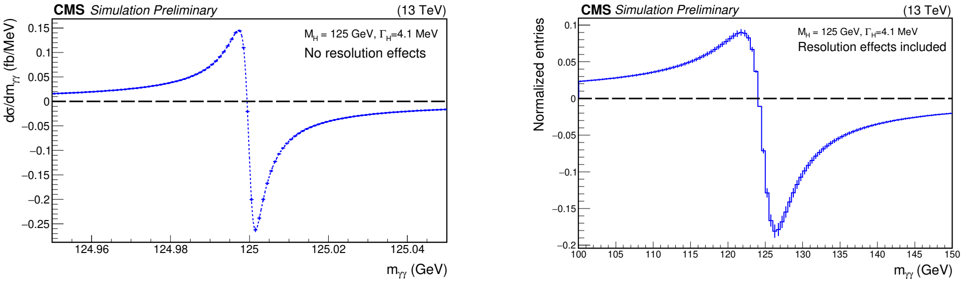

Figure 5.4: Differential cross-section as a function of $m_{\gamma\gamma}$, for the interference contribution in Eq. (5.2), without experimental resolution effects (left), and including experimental resolution (right). The experimental resolution is implemented through CMS Run 2 full simulation, for samples with $M_H = 125$ GeV and $\Gamma_H = \Gamma_H^{SM}$.

The distortion in the Higgs boson total mass spectrum due to the interference can thus be exploited by measuring precisely the shape of the $m_{\gamma\gamma}$ distribution. Indeed, as shown in Fig. 5.5, considering both $\mathcal{S}_{ggH}$ and I contributions after full detector resolution is applied, for several Γ_H values.

The distorted ggH spectrum, i.e. $d\sigma_s^{ggH}/dm_{\gamma\gamma} + d\sigma_i/dm_{\gamma\gamma}$, corresponds to a shift in the Higgs boson mass peak [46]. The size of these distortions is directly related to the R term in Eq. 5.2. As a result, the Γ_H parameter can be extracted by measuring the deviation of the shape from the one expected from $\mathcal{S}_{ggH}$-only process, induced by the interference in the Higgs boson lineshape, because its contribution is proportional to $c_{g\gamma} = \sqrt{\mu_F \Gamma_H / \Gamma_H^{SM}}$.

The ggH signal and interference differential cross sections (see Eq. (5.2)), when experimental resolution is taken into account, can be expressed in terms of the overall signal strength μ_F, i.e. the ratio of the measured and expected ggH signal yield, and explicitly on the couplings modifier, $c_{g\gamma}$, and the width,

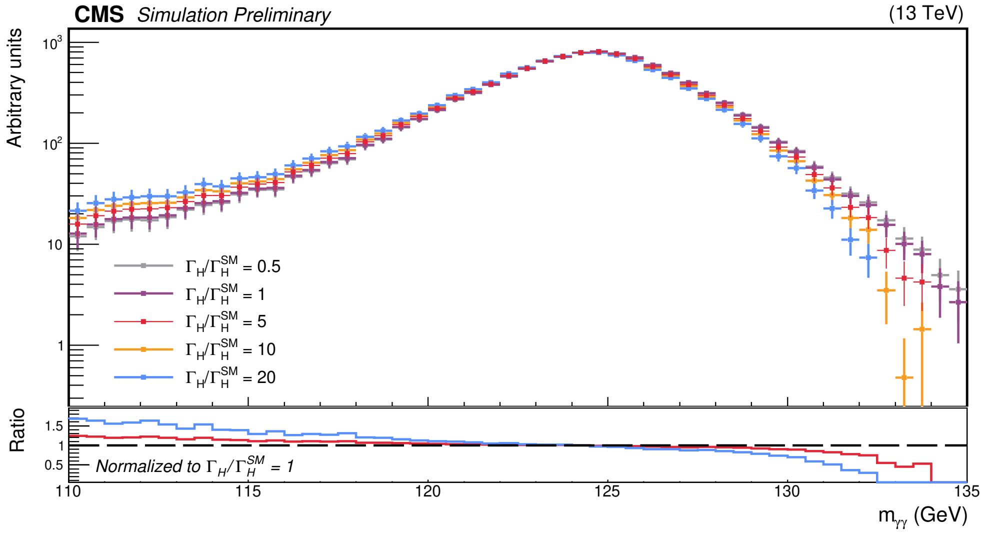

Figure 5.5: Sum of ggH signal and interference (both gg and qg) xmass spectra for different Γ_H values and M_H = 125 GeV, including resolution effects using full CMS detector simulation.

Γ_H:

$$d\sigma_s^{ggH}/dm_{\gamma\gamma} = \frac{c_{g\gamma}^2}{\Gamma_H}\Gamma_H^{SM}\mathcal{S}_{ggH}(m_{\gamma\gamma}) = \mu_F\mathcal{S}_{ggH}(m_{\gamma\gamma}), \tag{5.4}$$

$$d\sigma_i/dm_{\gamma\gamma} = c_{g\gamma}\mathcal{I}(m_{\gamma\gamma}), \tag{5.5}$$

where $\mathcal{S}_{ggH}(m_{\gamma\gamma})$ and $\mathcal{I}(m_{\gamma\gamma})$ are the probability density functions for the ggH signal and interference components, defined in such a way not to be dependent on $c_{g\gamma}$ and Γ_H. The simplification of using $\mathcal{S}_{ggH}$ and $\mathcal{I}$ templates which do not depend on Γ_H is possible because, due to experimental resolution, both $\mathcal{S}_{ggH}$ and $\mathcal{I}$ reconstructed $m_{\gamma\gamma}$ lineshapes become effectively independent of Γ_H, for Γ_H much smaller than the experimental resolution $O(1)$ GeV. From a technical point of view this simplification is very important, as it allows to produce, with the full CMS simulation, signal and interference samples at multiple Breit-Wigner Γ_H, and to only use one narrow-width ($\Gamma_H = 4$ MeV) samples, by combining linearly the signal and interference ones, as detailed below. Using Monte Carlo samples with parametric simulation, which is much less computationally intensive with respect to full

CMS detector simulation, it was verified that the impact of this assumption on the analysis is neglibible for the range of Γ_H considered in this note. These tests are detailed in Section 5.3.

The differential cross section including both ggH resonant diphoton production (signal) and interference with the continuum diphoton process can be written as:

$$d\sigma_s^{ggH}/dm_{\gamma\gamma} + d\sigma_i/dm_{\gamma\gamma} = \mu_F \mathcal{S}_{ggH} + c_{g\gamma}\mathcal{I} = (\mu_F - c_{g\gamma})\mathcal{S}_{ggH} + c_{g\gamma}(\mathcal{S}_{ggH} + \mathcal{I}) \tag{5.6}$$

This is a convenient expression for the extraction of the interference amplitude, because the $\mathcal{I}$ component is always negative above the Higgs boson mass, while $(\mathcal{S}_{ggH} + \mathcal{I})$ is positive in a large region around the Higgs boson peak, contrarily to $(\mathcal{S}_{ggH} + \mathcal{I} + \mathcal{B}_{pp\to\gamma\gamma})$, which is alway positive, where $\mathcal{B}_{pp\to\gamma\gamma}$ is the differential cross-section of the prompt diphoton background. The need for positive differential cross-sections for the signal and signal+interference processes is needed because of the statistical framework used in CMS for $H \to \gamma\gamma$ analyses, where all shapes ar positive definite.

Expressing $c_{g\gamma}$ as a function of Γ_H, the Higgs boson decay width is extracted by fitting the mass distribution in data with a linear combination of Probability Density Functions (PDF) for each category. This combination of PDFs is described by the following differential cross section expression, which includes also non-ggH Higgs boson production and background processes:

$$\frac{d\sigma_s^{ggH}}{dm_{\gamma\gamma}} + \frac{d\sigma_i}{dm_{\gamma\gamma}} + \frac{d\sigma_s^{V}}{dm_{\gamma\gamma}} + \frac{d\sigma_b}{dm_{\gamma\gamma}} =$$
$$\left(\mu_F - \sqrt{\mu_F \frac{\Gamma_H}{\Gamma_H^{SM}}}\right)\mathcal{S}_{ggH} + \sqrt{\mu_F \frac{\Gamma_H}{\Gamma_H^{SM}}}(\mathcal{S}_{ggH} + \mathcal{I}) + \mu_V \sum_{j\in\{\text{VBF, VH}\}} \mathcal{S}_{V,j} + \frac{d\sigma_b}{dm_{\gamma\gamma}}$$

where $\mu_F = \mu_{g\gamma}$ and μ_V represent, respectively, the signal strength in the ggH and electroweak bosons mediated production modes $(S_{V,j})$, i.e. VBF and VH. Templates for $\mathcal{S}_{ggH}$ and $\mathcal{S}_{ggH} + \mathcal{I}$ are constructed on the basis of ggH pure signal samples and interference samples. The background processes are mostly continuum di-photon production, γ+jet and multi-jets, with one or more jets misidentified as a photon, have no dependence on μ_F, μ_V and Γ_H. Templates for the background $(d\sigma_b/dm_{\gamma\gamma})$ are constructed by fitting the diphoton invariant mass distribution in the signal region using resonance sidebands directly from data, which include both prompt diphoton production, γ+jet, and multijet events. Note that the minor quark-gluon interference is treated together with the ggH interference and directly included in the $\mathcal{S}_{ggH} + \mathcal{I}$ templates. Note also that the ttH process is not included, because of the categorization employed in this analysis, detailed below, featuring a category targeting very loosely $\geq$ 2 jets event. Its contribution would therefore only increase by a minimal amount the event yield in this category, which is insensitive to Γ_H, and dominated by VBF and VH contributions.

The effect of the interference has a non-trivial $p_T^{\gamma\gamma}$-dependence can be quantified using the mass-shift observable, as shown in Figure 2.16. Pure NLO calculations, without NLL resummations, have unreliable interference $p_T^{\gamma\gamma}$ spectra [85]. Resummation effects are included, in SHERPA NLO generation, at NLL level, via the Catani-Seymour algorithm [56, 82], and this is shown in Figure 5.6, comparing SHERPA simulation with the theoretical predictions available in [85], which unfortunately are at 8 TeV and do not allow a one-to-one comparison. The rise above 0 in the differential cross-section as a function of $p_T^{\gamma\gamma}$, at low $p_T^{\gamma\gamma}$ values, is not completely understood, and therefore a cut is employed in the analysis, requiring $p_T^{\gamma\gamma} > 15$ GeV. Another potential mismodelling effect is due to the usage of NLO instead of NNLO, but this has an effect limited to a change of around 15% in the cross-section for the real part of the interference (see Figure 5 in [83]) and it is difficult to quantify exactly in this analysis due to the absence of dedicated generators. As it will be explained later, this effect is not a problem because the impact on the Γ_H measurement of a 7% error on the interference XS, due to QCD scale uncertainties, is below 5% and is one of the smallest systematic effects.

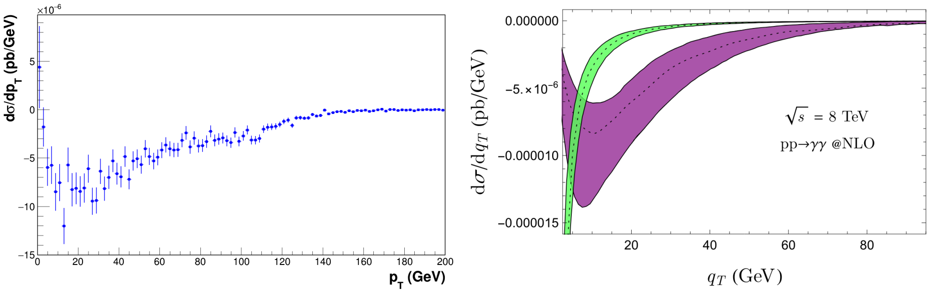

Figure 5.6: Differential cross-section of the gluon-gluon interference with respect to $p_{T,\gamma\gamma}$, evaluated by SHERPA NLO (with Catani-Seymour parton shower, at NLL) at 13 TeV (left) and in the theoretical predictions of [85], both at NLO and NLO+NLL (right).

It is not possible to verify, in a reliable way, any model-dependence possibly introduced by the Catani-Seymour parton shower model itself. Indeed, the alternative parton shower algoritm in SHERPA, so-called DIRE [86], has been deprecated and is no longer supported, because it has been shown [87] that the kinematic mappings in DIRE are not NLL safe. The distribution of $p_T^{\gamma\gamma}$ in ggH simulated signal events, after preselection cuts, is shown in Figure 5.7. On the same events, the efficiency of the $p_T^{\gamma\gamma}$ selection is 79%.

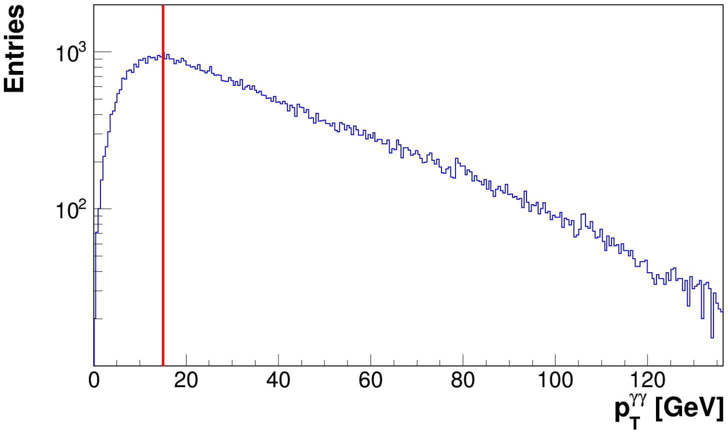

Figure 5.7: Distribution of $p_T^{\gamma\gamma}$ in ggH simulated reconstructed signal events for 2018, in arbitrary units, after preselection cuts, including generator weights. The red line shows the cut at 15 GeV.

5.3 Parametric simulations for non-SM Γ_H values

A series of preliminary studies based on generator level simulations and smearing procedures to mimic the experimental resolution were performed, in order to verify the theoretical assumptions made in the analysis strategy. In particular, the dependence of the interference spectrum on Γ_H requires confirmation.The need for generator-level studies, with a fast implementation of the detector effects, was motivated by the large computing power required to produce full simulation Monte Carlo samples at different Γ_H values. The only interference simulated samples available with full CMS detector simulation are for $\Gamma_H = \Gamma_H^{SM}$ and $M_H = 120, 125, 130$. Generator-only interference samples were generated with SHERPA (same version employed in the CMS official production), with 10^6 events for each Γ_H/Γ_H^{SM} value, ranging from 0.1 to 50. The photons from the generated samples are smeared with η dependent and R_9 dependent Gaussian resolutions. The R_9 variable is not present in generated samples, therefore only two bins in R_9 ($R_9 > 0.965$ for converted photons, and $R_9 \leq 0.965$ for unconverted photons) are simulated [88], by extracting randomly the bin using the η-dependent conversion probability for photons. The conversion probability is estimated as the η-dependent fraction of photons with $R_9 > 0.965$. The η-dependent resolution for unconverted photons is taken from $Z \to ee$ data with low R_9 (low-bremsstrahlung), while for converted photon it is taken from the inclusive $Z \to ee$ data, as shown in Figure 5.8.

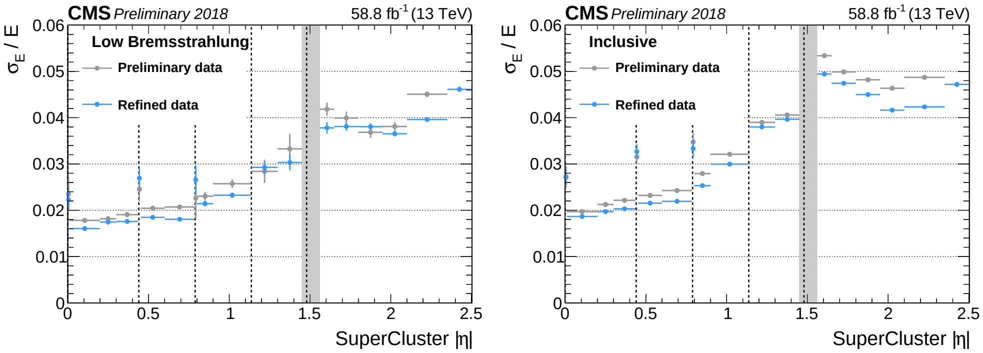

Figure 5.8: ECAL resolution from $Z \to ee$ data for the $R_9 \leq 0.956$ case, i.e. low-bremmstrhalung, (left), and the inclusive case (right).

To validate this procedure, it was applied to generator-level photons in the gluon-gluon fusion signal sample for $M_H = 125$ GeV, and the resulting diphoton mass spectrum was compared to the one obtained from full simulation, as shown in Figure 5.9. Despite the fact that the matching is not perfect, the RMS value is similar and only the left tail of the resolution function is not well described. Therefore, as a further test the smearing was applied to the privately generated interference sample for $\Gamma_H/\Gamma_H^{SM} = 1$, and the resulting diphoton mass spectrum was compared to the one obtained from full simulation, as shown in Figure 5.10.

For the interference the matching is very good, with a χ^2/NDF test statistics on the difference between the two histograms of 75.1/69, therefore the smearing procedure is assumed to be validated for the purposes of these preliminary studies.

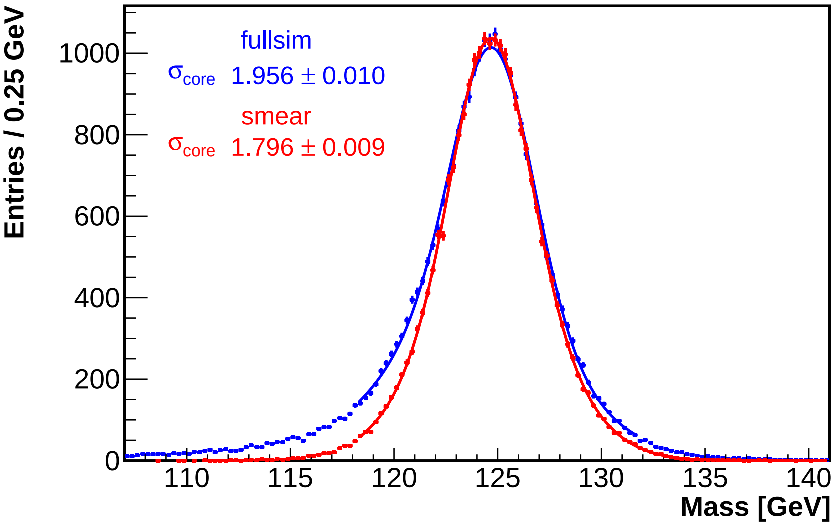

Figure 5.9: Diphoton mass spectrum obtained from generator-level quantities, convoluted with a parametric Gaussian smearing, vs. full simulation, for gluon-gluon fusion signal sample with $M_H = 125$ GeV. Preselections on the leading and subleading photons γ_1, γ_2 were applied, i.e. $|\eta| < 2.5$, $p_{T,1}/m_{\gamma\gamma} > 1/3, p_{T,2}/m_{\gamma\gamma} > 1/4$.

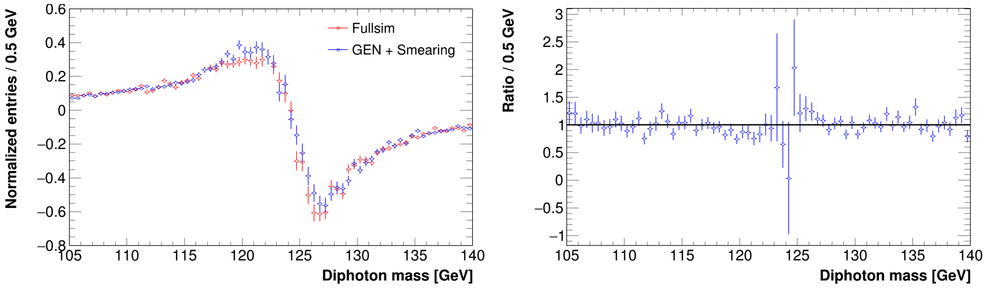

Figure 5.10: Diphoton mass spectrum obtained from generator-level quantities, convoluted with a parametric Gaussian smearing, vs. full simulation for interference sample with $M_H = 125$ GeV (left), and ratio between full-simulation and smearing-based spectra (right). Prelesections on the leading and subleading photons γ_1, γ_2 were applied, i.e. $|\eta| < 2.5$, $p_{T,1}/m_{\gamma\gamma} > 1/3, p_{T,2}/m_{\gamma\gamma} > 1/4$.

The interference distributions were first examined using samples in which only the width was varied, while the couplings remained unchanged. As shown in Figure 5.11, after the inclusion of the photon energy resolution, all the distributions with $\Gamma_H/\Gamma_H^{SM} \leq 20$ have a compatible shape. This confirms that the Γ_H dependence is solely in the interference cross-section due to the term $c_{g\gamma} = \sqrt{\mu_F \Gamma_H/\Gamma_H^{SM}}$. Indeed, it is expected that the two lobes of the interference distributions start to differ with respect to the SM Γ_H case, only when the with starts to be comparable or higher than the $m_{\gamma\gamma}$ peak resolution. This confirmation allows to employ the nominal SM Γ_H samples for both signal and interference, and to generate templates at different Γ_H by rescaling the interference cross-section as $\sqrt{\mu_F \Gamma_H/\Gamma_H^{SM}}$.

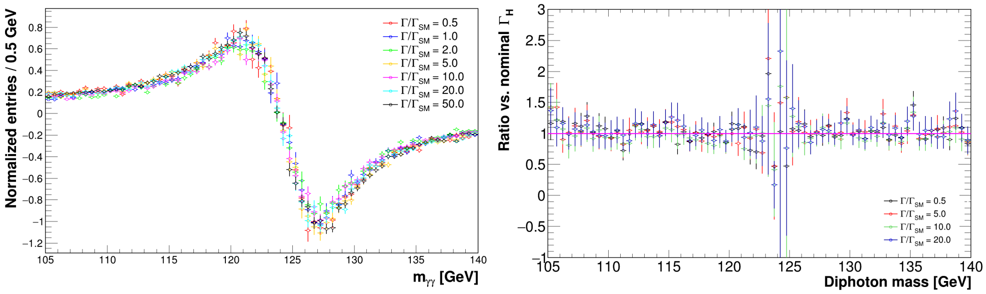

Figure 5.11: Diphoton mass spectrum obtained from generator-level quantities, convoluted with a parametric Gaussian smearing for interference samples with $M_H = 125$ GeV ad different Γ_H values (left), and ratio with respect to the nominal Γ_H case (right). Preselections on the leading and subleading photons γ_1, γ_2 were applied, i.e. $|\eta| < 2.5$, $p_{T,1}/m_{\gamma\gamma} > 1/3, p_{T,2}/m_{\gamma\gamma} > 1/4$.

It is also important to verify that the pure signal shape is independent on Γ_H, as shown in Figure 5.12, using SHERPA + smearing. For $\Gamma_H/\Gamma_H^{SM} = 25$, the corresponding oversmearing is about $\sigma^{rel}_{\Gamma_H = 100\text{MeV}}/2.35 = 42$ MeV , where 2.35 is the factor required to convert a full-width-half-maximum spread (like Γ_H) to its Gaussian equivalent. This value of oversmearing is below the error on the resolution coming from the scale and smearing procedure.

Using the same genererator-level samples employed in these tests, also the dependence of the interference and signal cross-section on Γ_H was checked, maintaining the same effective $c_{g\gamma}$ coupling, and the expected behaviour was verified, i.e. $\mathcal{I}$ independent on Γ_H, $\mathcal{S} \propto 1/\Gamma_H$. as shown in Figure 5.13.

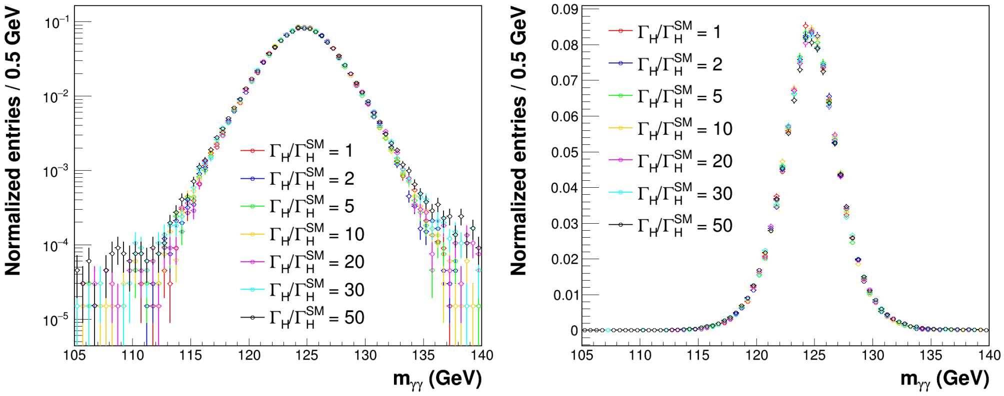

Figure 5.12: Signal mass spectrum generated with SHERPA , convoluted with parametric Gaussian smearing, for different Γ_H values.

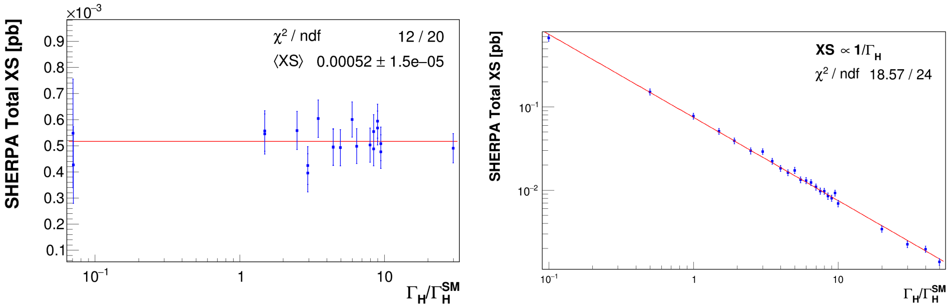

Figure 5.13: SHERPA values of cross-sections with respect to Γ_H for interference (left) and signal (right), in the case of fixed couplings. The red lines correspond to constant, and inverse proportionality fits, respectively, for the interference, and signal cross-sections. The fit quality is good in both cases.

5.4 Event categorization

Splitting the analysis into categories increase the sensitivity of $H \to \gamma\gamma$ analysis and allows to constrain Γ_H via simultaneous fits, where the likelihood is the product of single-category likelihoods, depending on the same physical parameters $(M_H, \mu_F, \mu_V, \Gamma_H)$ in all categories. In each category the background is data-driven, and this avoids complex background modelling from simulations, while the signal and interference are taken from MC and their shapes and normalizations depend on the physical parameters.

The first step of categorization uses the properties of the reconstructed diphoton and dijet systems to separate events consistent with the VBF and ggH production modes. The events passing a selection loosely targeting VBF, defined in Table 5.1, are put in a dedicated category (VBFTag), while the remaining events, predominantly ggH, are further divided in sub-categories. This is performed to mitigate as much as possible the dilution of the ggH fusion sample with events from VBF or other production modes, or events from ggH with two jets, which have smaller interference contributions [19]. The fraction of *ggH* events, after the application of the pre-selection, passing the loose VBF tag is 21%.

Loose VBF tag
$100 < m_{\gamma\gamma} < 180$ GeV
$p_T^{\gamma_1}/m_{\gamma\gamma} > 1/4$
$p_T^{\gamma_2}/m_{\gamma\gamma} > 1/5$
$\lvert\eta_\gamma\rvert < 2.5$
$p_T^{j_1} > 30$ GeV
$p_T^{j_2} > 20$ GeV
$m_{jj} > 100$ GeV

Table 5.1: Criteria for the relaxed VBF tag.

Since the sensitivity to the interference is strongly dependent on the diphoton p_T ($p_T^{\gamma\gamma}$), the events not satisfying the VBF selection are divided according to that observable. For every year, each $p_T^{\gamma\gamma}$ category c_i was divided in two sub-categories, one with high BDT score, selecting all events with BDT score (defined from -1 to 1) greater than a threshold t_i^h, and another one with medium BDT score, selecting the events with BDT score smaller than t_i^h but greater than a lower threshold t_i^m. The thresholds th and tm were subject to the optimization, as described in the following. Five $p_T^{\gamma\gamma}$ categories were used in-order to have sufficient granularity to capture the p_T^H dependence of interference effect, while maintaining sufficient statistical power and optimizing mass resolution in each category. The diphoton invariant mass distribution is constructed in each of the resulting categories, and the boundaries are then adjusted iteratively until a minimum of the following loss function is obtained. The $p_T^{\gamma\gamma}$ category boundaries are the same for all data-taking eras, while the BDT score boundaries depend on the data-taking era, since some of its input variables depend on the detector resolution, which evolved during the LHC Run 2.

$$L = \sigma_{\gamma\gamma} \times \frac{\sqrt{B}}{S} \tag{5.7}$$

where $M_{\gamma\gamma}$ and $\sigma_{\gamma\gamma}$ are the mean and the standard deviation of the diphoton invariant mass distribution respectively, calculated within a 1σ region around the peak. S and B are the signal and background yields, respectively, in the same invariant mass window. This procedure minimizes the relative diphoton mass resolution and maximizes the signal strength in each analysis category, and optimize the analysis for precision in the measurement of Higgs boson mass, which is a parameter correlated with our measurement due to the mass-shift effect. This optimization was performed using signal MCs corresponding to all four eras of Run 2 data taking, and data-driven samples for the inclusive non-resonant background, from mass sidebands. A detail of each category is shown in Table 5.2. It was verified on 2018 simulated samples, scaled at full Run 2 luminosity, that using a dedicated optimization, instead of the described procedure, targeting directly the sensitivity to Γ_H does not improve the measurement significantly over the aforementioned procedure, as shown in Section 5.4.1.

Note that the ggH and $ggH + I$ yields differ by at most 8%, depending on the category. This is coherent with the expected inclusive 2% difference in the cross-section at NLO due to the interference [46].

5.4.1 Re-optimization for Γ_H sensitivity

An alternative, ad-hoc, categorization for the shape-based method was attempted. The structure of the categories was not modified with respect to the method explained in Section 5.4, i.e., there are 5 categories in Higgs boson p_T, bound between 15 GeV and 3 TeV, each of them with 2 sub-categories based

N.	Category	$p_T^{\gamma\gamma}$ (GeV)	Expected events ggH	 ggH + I	 VBF	 VH	σ_{eff} (GeV)	S/(S+B)
0	ggH−5 − hi	≥ 77	98	95	24	23	1.2	0.28
1	ggH−5 − med		183	175	31	26	1.4	0.17
2	ggH−4 − hi	[47, 77)	87	80	9.7	4.3	1.2	0.16
3	ggH−4 − med		215	204	19	12	1.3	0.099
4	ggH−3 − hi	[31, 47)	92	88	4.4	2.0	1.0	0.13
5	ggH−3 − med		271	265	11	7.5	1.3	0.10
6	ggH−2 − hi	[21, 31)	119	117	2.7	1.4	1.1	0.10
7	ggH−2 − med		276	272	6.3	4.1	1.6	0.066
8	ggH−1 − hi	[15, 21)	96	94	1.2	0.7	1.1	0.11
9	ggH−1 − med		193	190	2.5	1.6	1.4	0.069
10	VBFTag	-	720	720	238	64	1.6	0.045

Table 5.2: Boundaries in $p_T^{\gamma\gamma}$, expected event yields for all years, mass resolution, and expected signal fraction for each category averaged over the different years. The category label (hi, med) defines the type of diphoton BDT region selected, i.e., respectively, high and medium. The mass resolution σ_{eff} is defined as the width of the region, centered on the peak, that contains 68% of the signal distribution, while the expected signal fraction $S/(S+B)$ is the ratio between signal and total events in the same region defined for the mass resolution. Note that the total events are not the sum of the events from the four processes, because ggH and (ggH+I) are combined depending on the values of μ and Γ_H. Following Eq. 5.7, in the case of $\mu = 1, \Gamma_H = 0$, only the pure ggH process would be observed, together with VBF and VH, while in the case of $\mu = 1, \Gamma_H = \Gamma_H^{SM}$, only (ggH+I), VBF and VH would be observed.

on the diphoton BDT score, bound between -1 and 1, while the boundaries have been re-optimized.

For each set of categories the figure of metric was defined, for simplicity, as the sum of the χ^2, for all categories, between the so-called Asimov histograms, at $\Gamma_H = 0$ and $\Gamma_H = 1$. Asimov histograms are first constructed by using the MC histograms in $m_{\gamma\gamma}$ (binned also in p_T and BDT score), after removal of events passing the loose VBF tag, using the formula in Equation 5.7. Then the background contribution is added to the Asimov histograms, by using an interpolation in the signal region 115-135 GeV, after fitting the data with exponential functions outside the signal region. To mimic data, the Asimov histograms are constructed with the same statistics of the data and Poissonian errors. The χ^2 sum evaluated in this way represents the square of the statistical significance of the interference process, which is a convenient figure of merit because it is linked to the analysis sensitivity. For each category the Asimov histograms were evaluated by using directly the MC datasets at $M_H = 125$ GeV for $\mathcal{S}$ ans $\mathcal{S} + \mathcal{I}$ datasets, without including the VBF and VH processes. This procedure yields slightly optimistic significance with respect of the method actually used in the analysis, because it does not allow for background shape variations, and does not include non-ggH processes.

The ten analysis categories $c_1^{\text{med-MVA}}$, $c_1^{\text{high-MVA}}$, ..., $c_5^{\text{med-MVA}}$, $c_5^{\text{high-MVA}}$ are parametrized as:

$$c_1^{\text{med-MVA}} = [(15 \text{ GeV}, pT_1), (t_1^m, t_1^h)],\ c_1^{\text{high-MVA}} = [15 \text{ GeV}, pT_1), (t_1^h, 1)] \tag{5.8}$$

$$c_2^{\text{med-MVA}} = [(pT_1, pT_2), (t_2^m, t_2^h)],\ c_2^{\text{high-MVA}} = [(pT_1, pT_2), (t_2^h, 1)] \tag{5.9}$$

$$... \tag{5.10}$$

$$c_5^{\text{med-MVA}} = [(pT_4, 3 \text{ TeV}), (t_1^m, t_5^h)],\ c_5^{\text{high-MVA}} = [(pT_4, 3 \text{ TeV}), (t_5^h, 1)] \tag{5.11}$$

By imposing that $pT_4 < 150$ GeV (to have enough statistics) and that all t^m are greater than -0.5, each set of categories is then completely determined as a function of 14 parameters $(a_1, ..., a_4), (b_1^{c_1}, b_2^{c_1}, ..., b_1^{c_5}, b_2^{c_5}))$, ranging between 0 and 1, as follows:

$$pT_0 = 15 \text{ GeV} \tag{5.12}$$

$$pT_n = pT_{n-1} + a_n(150 \text{ GeV} - pT_{n-1}) \quad \text{for } n = 1, ..., 4 \tag{5.13}$$

$$t_n^m = -0.5 + 1.5 b_1^{c_n},\ t_1^h = t_n^m + b_2^{c_n}(1 - t_n^m) \quad \text{for } n = 1, ..., 5 \tag{5.14}$$

where the $(a_1, ..., a_4)$ are used to define the pT cuts and the b ones define the cuts on the diphoton BDT score. Being the category optimization with 14 parameters a non-differentiable multi-dimensional problem, a genetic algorithm NSGA-II [89] was employed. After convergence a result better than the one obtained by default categories in Section 5.4 was found, as shown in Figure 5.14.

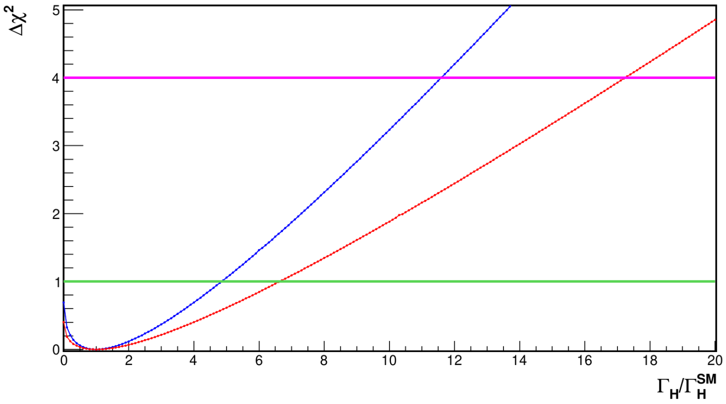

Figure 5.14: $\Delta\chi^2$ scan using the default categories (red line) or using the improved categories (blue line).

Given the relatively small improvement around 30%, the analysis was performed using the categories optimized using the metric based on mass precision.

Chapter 6

Maximum likelihood fit

6.1 Statistical modeling

As described in the previous chapter, the selected events are categorized in 5 bins of the diphoton p_T, for $p_T^{\gamma\gamma} > 15$ GeV, and in two bins of the diphoton MVA score, with an additional single category targeting VBF. To extract the Higgs boson width, a simultaneous extended maximum likelihood fit [90] to the diphoton invariant mass spectrum ($M_{\gamma\gamma}$) is performed in all analysis categories (enumerated with an index k = 0, ..., 10). Categories from 0 to 9 target ggH, while category 10 targets loosely VBF and ggH with 2 jets. The selected diphoton p_T of each pair of category decreases with the index, i.e. the first two corresponds to the highest p_T region, and so on, while categories with even index (0, 2, ...) have high diphoton BDT score, while odd index ones have medium BDT score, as shown in Table 5.2.

Systematic uncertainties are taken into account using nuisance parameters (NP) $\vec{\theta}_S, \vec{\theta}_B$, which modify the signal, $\mathcal{S}+\mathcal{I}$ and background shapes or normalizations. The likelihood model employed to fit the binned data uses parametric signal and background models. It can be written as follows:

$$\mathcal{L}_{stat}\left(\text{data} \mid \Gamma_H/\Gamma_H^{SM}, \vec{\mu}, M_H, \vec{\theta}_S, \vec{\theta}_B\right) = \tag{6.1}$$

$$= \prod_{k=1}^{n_{\text{cat}}} \prod_{i=1}^{n_{bins}} \text{Pois}\left(n_{\text{data}}^{ki} \mid N_k(i, \Gamma_H/\Gamma_H^{SM}, \vec{\mu}, M_H, \vec{\theta}_S, \vec{\theta}_B)\right) \tag{6.2}$$

where the N_k functions are the sum of signal and background models, $\vec{\theta}$ are the nuisance parameters. The mass of the Higgs boson M_H, as well as the signal strengths ($\mu_F = c^2_{g\gamma}/\Gamma_H$, μ_V) for gluon-gluon fusion and vector production modes are left floating in the fit. The signal model, i.e. the Probability Density Function for the Higgs boson resonance, is then derived for each combination of category and process from simulated samples. As explained in section 5.2, pure gluon-gluon fusion (ggH) and gluon-gluon fusion summed with the interference contribution (ggH+int) are treated as two different processes, therefore the probability density function of the signal contributions is modelled according to four production processes (ggH, ggH+int, VBF and VH). On the other hand, the background model is derived from the inclusive data, blinding the signal region in a range $115 < m_{\gamma\gamma} < 135$ GeV in each category, explained in more detail in Section 6.4. The binning employed in the analysis has a width of 0.25 GeV in $m_{\gamma\gamma}$, which is more then 5 times smaller than the mass resolution. This is motivated by the need to have sensitivity to the shape variations due to the interference, and to reduce the bias due to the background steep derivative. In the fit, for each category the expected $m_{\gamma\gamma}$ distribution is modeled as a sum of all signal processes, indexed by j, and background. Each signal process is scaled by the signal strength μ_j, multiplied by the event yield $n_{j,k}$ (see Table 5.2) and by the parametric probability density function P_{jk}, evaluated in n_{bins} bins, indexed by i, in the variable $m_{\gamma\gamma}$. Each bin i in the k-th category has the following yield:

$$N_k(i, \Gamma_H/\Gamma_H^{SM}, \vec{\mu}, M_H, \vec{\theta}) = \sum_j^{procs} \mu_j(\Gamma_H/\Gamma_H^{SM}) n_{j,k} P_{jk}^{\text{sig}}\left(x_i, M_H; \vec{\theta}\right) + P_k^{\text{bkg}}\left(x_i; \vec{\theta}\right) \tag{6.3}$$

where μ_j is equal to $\mu_{ggH} = (\mu_F - \sqrt{\mu_F \Gamma_H/\Gamma_H^{SM}})$ for the ggH process, while it is equal to $\mu_{ggH+int} = \sqrt{\mu_F \Gamma_H/\Gamma_H^{SM}}$ for the $\mathcal{S}_{ggH} + I$ process, following the formula in Equation 5.7, and is equal to μ_V (another independent parameter) for the two other production modes taken into account (VBF and VH). The estimation of the event yields and probability density functions for signal and background is described in the following sections.

6.2 Signal modeling

For each combination of category and production process, the distribution of events in $m_{\gamma\gamma}$ from the simulation is fitted using a Double Crystal Ball (DCB) function, defined in Equation 3.7. The fits are performed simultaneously using MC samples at three Higgs boson mass points (120, 125, 130 GeV), in order to have templates that can be interpolated to any M_H value. In order to do this, the DCB parameters at the different M_H values are interpolated linearly, requiring one additional parameters for each DCB parameter, while the normalization is interpolated using a parabola.

The 5 DCB parameters $(\mu, \sigma, \alpha_L, \alpha_R, n_L, n_R)$ together with the additional 5 parameters required by the linear interpolation, are fitted using the sum of the χ^2 between the DCB function and the MC histogram for all three M_H values, while the normalizations of the DCB functions are based on the integral of the histograms in the range $M_H - 15\,\text{GeV}, M_H + 10\,\text{GeV}$. This range was chosen to be inclusive on signal samples, containing a region larger than $[-10, +6]\sigma$ for a typical 1.5 GeV mass resolution, while also minimizing the overlap with the negative region due to the interference on the right part of the spectrum. The asymmetry of the range is a consequence of the asymmetric shape of $\mathcal{S}_{ggH} + I$. The fit range is between $M_H - 7\,\text{GeV}$ and $M_H + 3\,\text{GeV}$, to avoid fitting the region where the histograms for the $\mathcal{S}_{ggH} + I$ process can have negative wbins due to the interference. Before this fit procedure, a pre-fit is done only to the 125 GeV mass point, and the resulting parameters are used to initialize the final fit. Moreover, in the final fit, the value of the right-tail parameters are fixed to the pre-fit values, since they cannot be estimated reliably due to the limitation of the fit range to $M_H + 3$ GeV. Also the dependence of the signal normalization on M_H, influenced by variations in the signal cross-section σ, branching ratio B, and the product of acceptance and selection efficiency $\epsilon \times A$, as well as all other signal model parameters, are parameterized as a function of MH to allow MH variations in the final fit. Figure 6.1 illustrates example fits for two categories for the ggH+int. process, with normalization and shape parameters interpolated using second degree polynomials.

Overall, this multistep procedure defines the parametric signal model employed in the global maximum likelihood fit. All the fits performed are reported in Section A.1 A.2 in the Appendix. Figure 6.2 includes all production modes for three example categories. The complete set of plots is available in Figures A.33, A.34, A.35, A.36 in the Appendix.

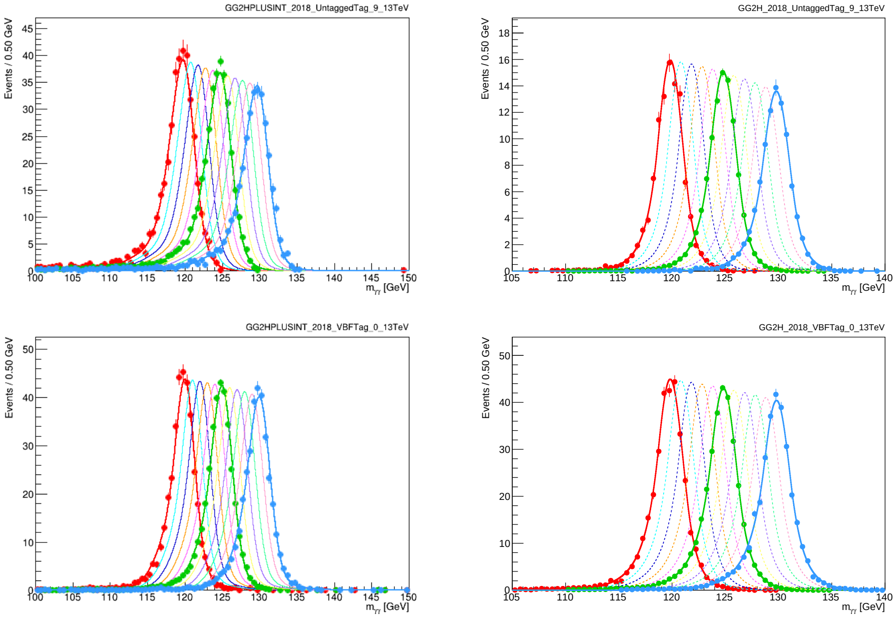

Figure 6.1: Example fits for two categories for the $\mathcal{S}_{ggH} + I$ (left) and $\mathcal{S}_{ggH}$ (right). The top plots corresponds to category 9 (lowest p_T region, high BDT score), while the bottom ones to category 10 (loose VBF tag).

Figure 6.4 presents the weighted average of all categories for the 4 different processes. The effective resolution of the signal $M_{\gamma\gamma}$ distribution, for a Higgs boson with $M_H = 125$ GeV, is also displayed. The different resolution is primarily determined by changes in the photon energy resolution between categories and processes due to their different kinematical regions.

Any mismodelling far from the peak due, for example, to the fact that the $\mathcal{S}_{ggH} + \mathcal{I}$ shape gets negative in its right tail, is covered by the very small signal/background ratio in those regions, as verified with dedicated tests, detailed in Section 6.3. In Figure 6.3 the sum of $\mathcal{S}_{ggH} + \mathcal{I}$ MC histograms for categories from 0 to 9 is shown, evaluated at the Run 2 luminosity, together with error bands calculated from the background Poissonian fluctuations. This shows that the positive and negative tails, far off the peak, which are due to the interference are in a region dominated by the background. The only missing category is the loose VBF tag one, which is richer in background because of the absence of BDT score cuts and would yield a too optimistic picture.

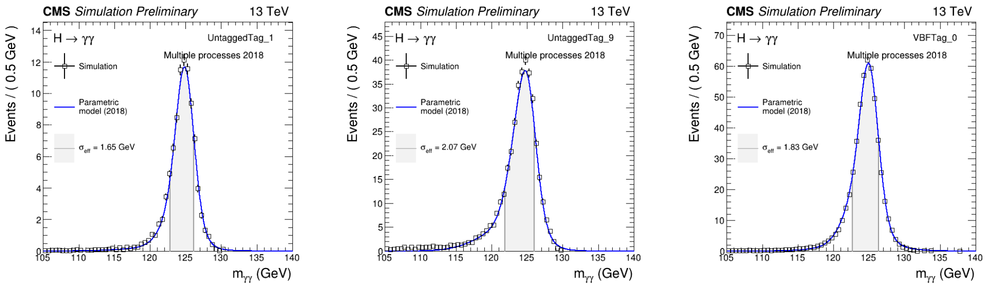

Figure 6.2: Signal model averaged over the three data-taking years including all production modes for three example categories.

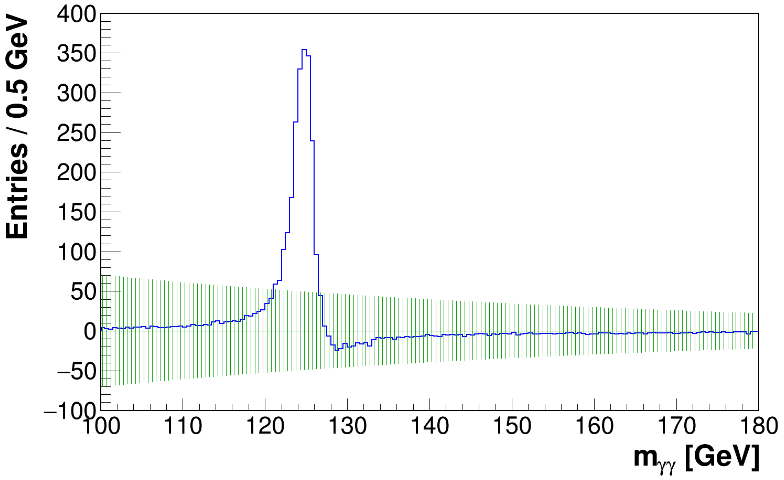

Figure 6.3: Sum of $\mathcal{S}_{ggH} + \mathcal{I}$ MC histograms for categories from 0 to 9, evaluated at the Run 2 luminosity using the samples for 2018, together with error bands calculated from the background Poissonian fluctuations. Each category is weighted with $S/(S+B)$, and the sum is normalized such that the reweighting does not affect the total event count. The only missing category is the loose VBF tag one.

In order to check for large discrepancies between MC histograms and signal models at large non-SM Γ_H, both the S+I models and the MC histograms

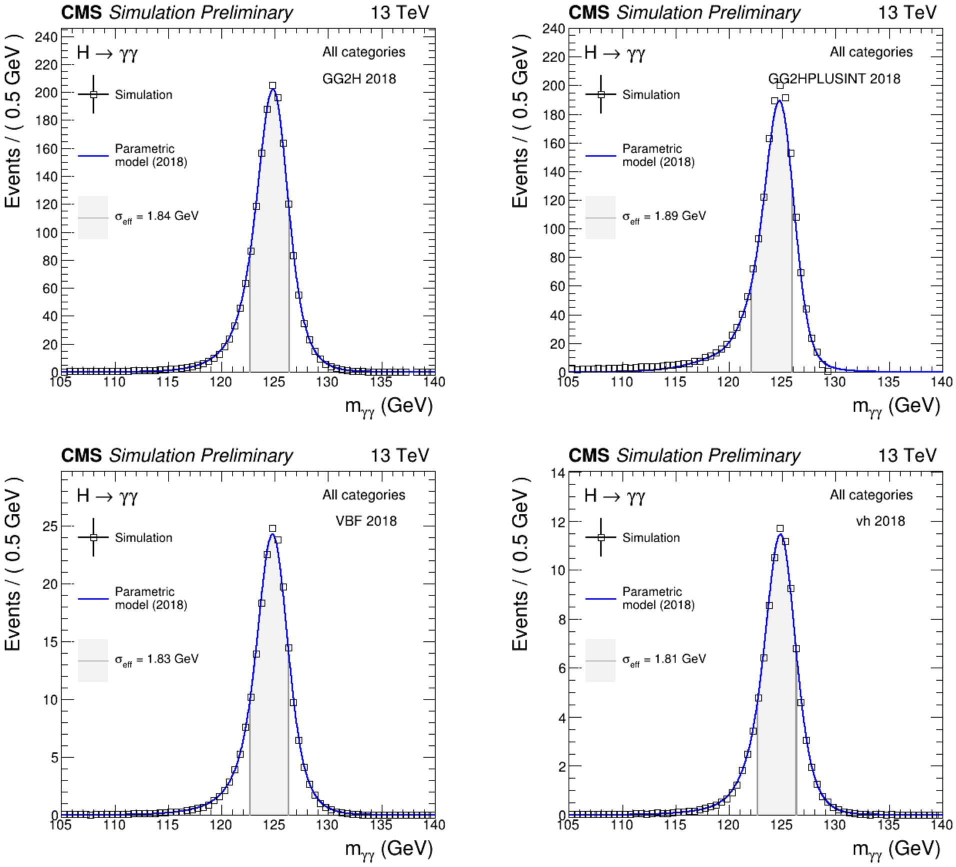

Figure 6.4: Signal models including all categories of 2018, for the four production modes. From left to right and from top to bottom: $\mathcal{S}_{ggH}$, $\mathcal{S}_{ggH} + I$, VBF, VH.

were plotted at different Γ_H values, by using the formula:

$$(\mathcal{S}_{ggH}+\mathcal{I})(m_{\gamma\gamma},\,\Gamma_H) = \left(1-\sqrt{\Gamma_H/\Gamma_H^{SM}}\right)\mathcal{S}_{ggH}(m_{\gamma\gamma},\,\Gamma_H^{SM})+\sqrt{\Gamma_H/\Gamma_H^{SM}}\,(\mathcal{S}_{ggh}+\mathcal{I})(m_{\gamma\gamma},\,\Gamma_H^{SM}), \tag{6.4}$$

including only gluon-gluon fusion signal and $\mathcal{S}_{ggh}+\mathcal{I}$ processes, for templates defined from 2018 MC samples at $\Gamma_H = \Gamma_H^{SM}$. Some examples of the best and worst cases are shown in Figure 6.5 for $M_H = 125$ GeV.

There are some discrepancies in the region of the right tail (near 128 GeV) but they occur where the signal is very low in yield and the fit is dominated by the background. They are due to the fact that on the right part of the tail there are more events in the signal than in the $\mathcal{S}_{ggh}+\mathcal{I}$ models, and therefore when the coefficient of the pure signal term is negative (when $\Gamma_H > \Gamma_H^{SM}$) a negative yield for those bins is obtained when combining the two processes. This is ultimately a consequence of the fact that the parametric signal models are by construction positive definite.

6.3 Background modeling

Parametric models to describe the background in different analysis categories are derived from data side-bands. The $M_{\gamma\gamma}$ distributions in the data side-bands, defined from 100 to 115 GeV and from 135 to 180 GeV, are fitted using various plausible background model PDFs, monotonically decreasing, including exponential functions, Bernstein polynomials, Laurent series, and power law functions and their higher order equivalents. An F-test[91] is performed with a goodness-of-fit criterion to determine a set of background PDF functions and their orders. The choice of the background PDF is considered as an additional discrete nuisance parameter to account for the uncertainty in this arbitrary function selection [92], as explained in Section 6.4.

Figure 6.6 shows the background model PDFs considered in one analysis category.

The background model in all categories are available in Figures A.37 , A.38, A.39, and A.40 in the Appendix. It is possible that a sizeable bias is introduced for large Γ^H/Γ^H_{SM} values because of the interference contributions outside the signal region (115, 135 GeV). A set of Anderson-Darling tests [93] was performed to estimate statistically the discrepancy between the best-fit PDF and the sum of the best-fit PDF and the interference at $\Gamma_H \sim 500\Gamma_H^{SM}$, for the range $M_{\gamma\gamma} > 130$ GeV. In Figure 6.7 the test score distributions for

Category VBFTag_0 (2jets events - small interference):

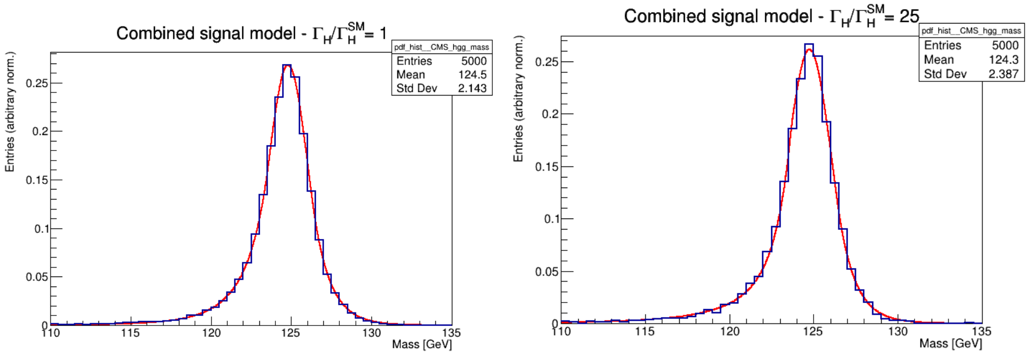

Category UntaggedTag_0 (high pT, "moderately small" interference)

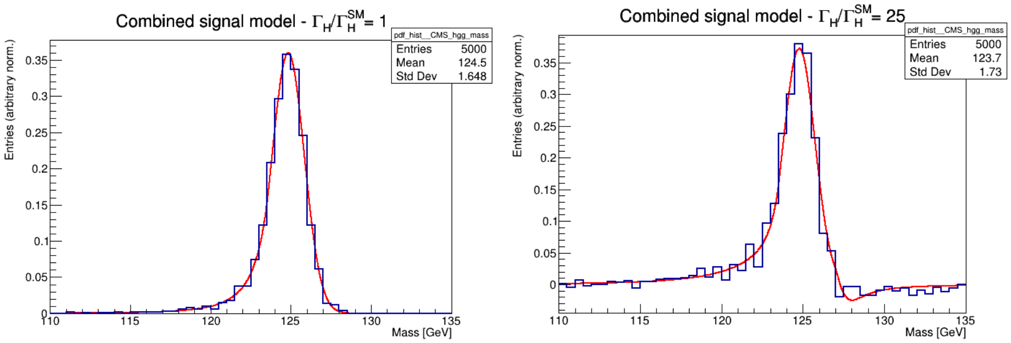

Category UntaggedTag_9 (low pt, large interference)

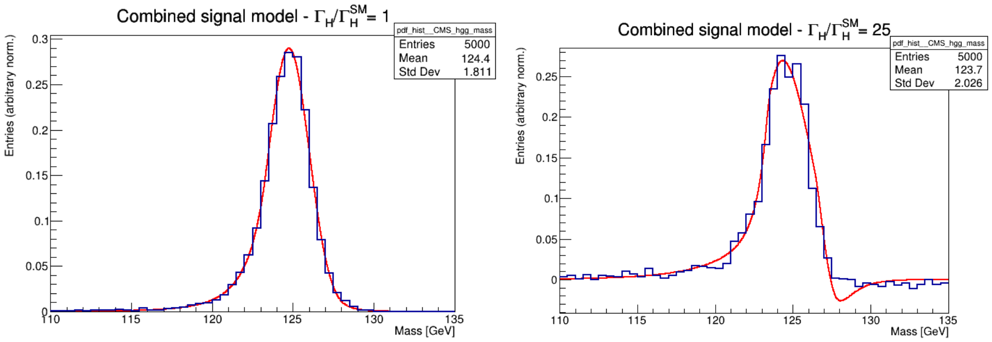

Figure 6.5: $\mathcal{S}_{ggh} + \mathcal{I}$ signal models, together with the relative histograms from MC, for $\Gamma_H/\Gamma_H^{SM} = 1, 25$. Both the signal models and the histograms have been derived from the ones at the nominal width through the proper formula used also in the physics model (6.4).

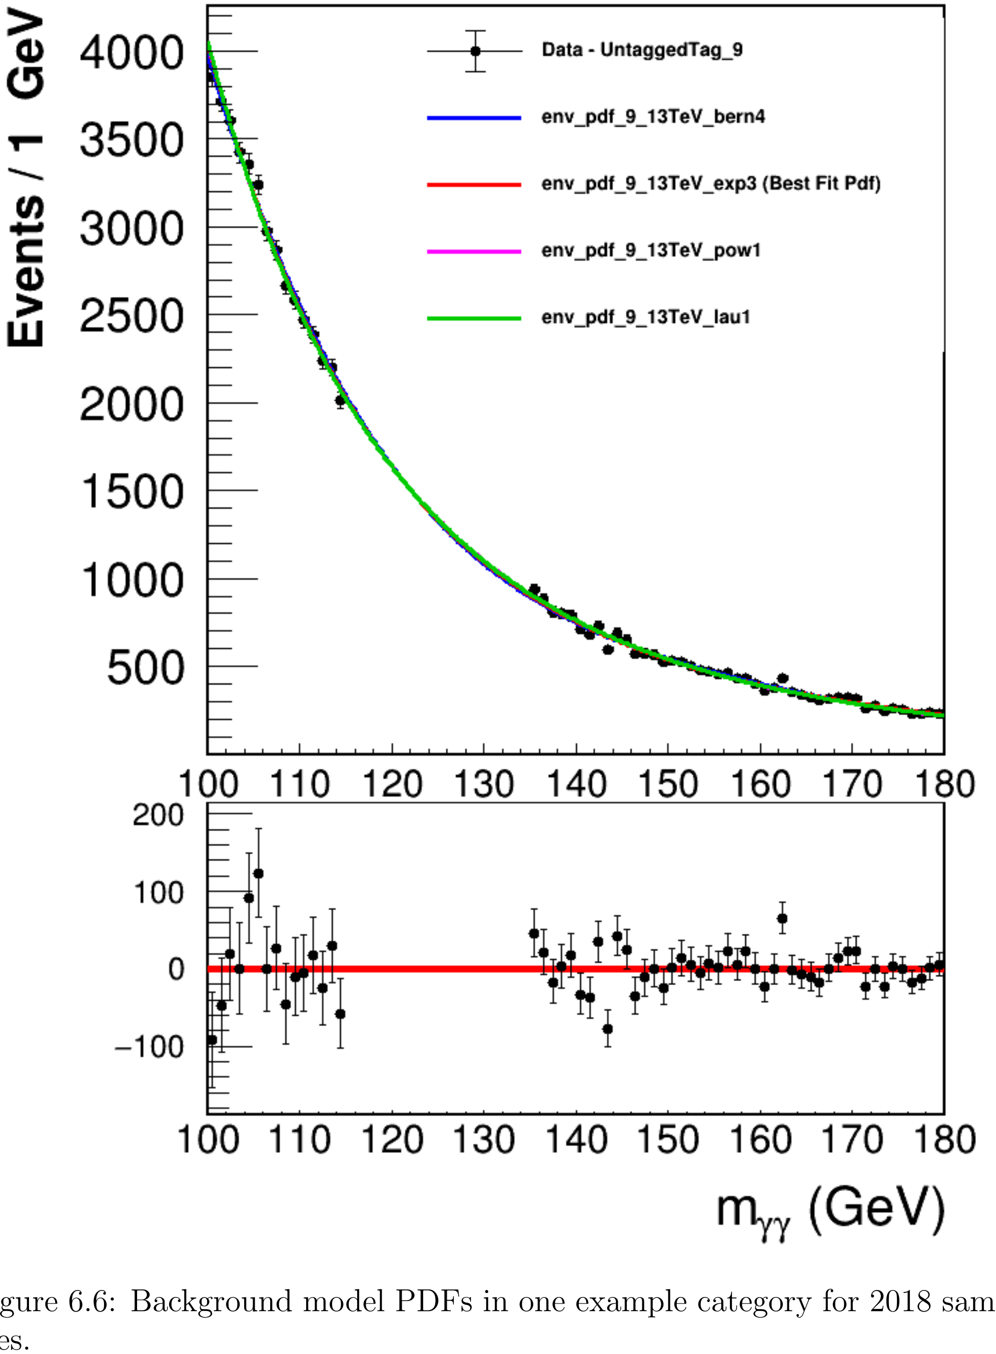

Figure 6.6: Background model PDFs in one example category for 2018 samples.

1000 toys are shown for the cases $\Gamma_H = 0$ (no interference), $\Gamma_H \sim 500\Gamma_H^{SM}$ and, as a control, $\Gamma_H \sim 5000\Gamma_H^{SM}$, showing that at $\Gamma_H \sim 500\Gamma_H^{SM}$ no significant deviations from the no-interference case are present. In particular, the similarity between the AD score distributions for the no-interference and the 500 Γ_H/Γ_H^{SM} case implies the absence of bias, unlike the case at an extremely large width value, where the difference in the distributions is evident.

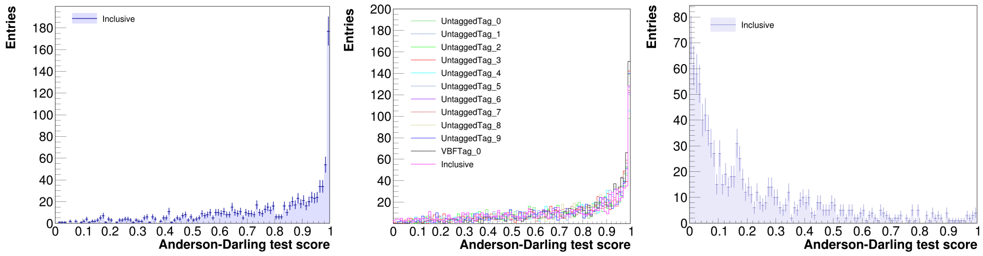

Figure 6.7: Anderson-Darling test score distributions for 1000 toys comparing the best-fit PDF and the sum of the best-fit PDF and the interference at $\Gamma_H = 0\Gamma_H^{SM}$ (left), $\Gamma_H \sim 500\Gamma_H^{SM}$ (center), $\Gamma_H \sim 5000\Gamma_H^{SM}$ (right), for the range $M_{\gamma\gamma} > 130$ GeV.

6.4 Treatment of systematic uncertainties in the fit

In the likelihood fit, systematic uncertainties on the extraction of Γ_H are implemented as nuisance parameters included in the minimization of the likelihood function. This section summarizes the treatment of nuisance parameters in likelihood-based inference, distinguishing between *discrete* and *continuous* nuisance parameters. The latter are further categorized based on their constraint modeling.

6.4.1 Discrete nuisance parameters

Let μ denote the parameter of interest and i_d a discrete nuisance parameter that can take values in a finite set $\{i_1, i_2, \ldots, i_k\}$. This represents the case of the choice of the background function, for instance. The likelihood function is denoted as $\mathcal{L}(\mu, i_d)$, and in the **discrete profiling method**, the profile likelihood is constructed by maximizing over the discrete nuisance values:

$$\mathcal{L}_{\text{discrete}}(\mu) = \max_{i=1,\dots,k} \mathcal{L}(\mu, i_i) \tag{6.5}$$

No additional constraint terms are added to the likelihood for discrete nuisance parameters. The method effectively selects, for each value of θ, the discrete configuration that yields the highest likelihood.

6.4.2 Continuous nuisance parameters

Let θ now denote a continuous nuisance parameter. In this case, the likelihood is multiplied by constraint terms to incorporate prior knowledge or auxiliary measurements. The total likelihood, if the parameter of interest is μ becomes:

$$\mathcal{L}_{\text{total}}(\mu, \theta) = \mathcal{L}_{\text{stat}}(\mu, \theta) \cdot \prod_i C_i(\theta_i) \tag{6.6}$$

where $C_i(\theta_i)$ is the constraint term for the i-th nuisance parameter. The constraints can be distinguished on the basis of the PDF used to model it.

For nuisance parameters without any prior information, no constraint term is added. These parameters are left free in the fit:

$$C(\theta) = 1 \tag{6.7}$$

Such parameters may represent unknown quantities for which no auxiliary measurement or theory uncertainty is available.

In many cases, Gaussian constraints are applied:

$$C(\theta) = \exp\left[-\frac{1}{2}\left(\frac{\theta - \theta_0}{\sigma}\right)^2\right] \tag{6.8}$$

where θ_0 is the nominal value (e.g., zero shift), and σ is the standard deviation representing the uncertainty.

For nuisance parameters affecting the normalization (e.g., luminosity or theoretical uncertainties), multiplicative uncertainties are better modeled using a log-normal distribution. The corresponding constraint takes the form:

$$C(\theta) = \frac{1}{\theta\sqrt{2\pi}\ln k} \exp\left[-\frac{1}{2}\left(\frac{\ln(\theta/\theta_0)}{\ln k}\right)^2\right] \tag{6.9}$$

where θ_0 is the nominal normalization, and k defines the relative uncertainty (e.g., $k = 1.1$ for a 10% uncertainty).

This structured approach enables the incorporation of systematic uncertainties into the inference process, ensuring robust estimation and hypothesis testing in the presence of both discrete and continuous sources of uncertainty.

6.5 Sources of systematic uncertainties

In this analysis, the systematic uncertainty related to the data-driven background estimation is addressed using the discrete profiling method, as described earlier. The uncertainties on the signal are divided in experimental and theoretical ones. They are further classified in normalization or shape uncertainties. Indeed, if, for a certain systematic uncertainty, the shape of the $m_{\gamma\gamma}$ distribution remains unaffected, the uncertainty is handled as a log-normal variation in the event yield and represents a normalization uncertainty. These include theoretical uncertainties, but also many xperimental ones, such as those influencing the BDTs used for event categorization, efficiencies, or the luminosity. On the other hand, uncertainties, both theoretical and experimental, that alter the shape of the $M_{\gamma\gamma}$ distribution are incorporated into the signal template functions as Gaussian constrained nuisance parameters, and represent shape uncertainties. In this case the DCB parameters of the models are therefore not only functions of the single parameter M_H, but depend as well on the associated NPs. For all uncertainties, their impact is evaluated individually for each analysis category. Nuisance parameters are completely correlated between different categories and processes, i.e., for each independent source of systematic uncertainty there is one nuisance parameter that is propagated in all processes and categories.

The parameters representing the peak position are modified, by adding one nuisance per systematic uncertainty, multiplied by a coefficient equal to the the 1σ variation of the mean value. The same treatment is applied to σ, a_L, and n_L, with a multiplicative instead of additive combination.

For each source of systematic uncertainty two alternative $m_{\gamma\gamma}$ distributions are produced, by varying the variable of interest (the calorimeter scale, for instance) by $\pm 1\sigma$, and applying the full analysis chain on the modified dataset. In this way changes in both selection efficiency, category migrations, and shape variations of the fit variable $m_{\gamma\gamma}$ are taken into account. The coefficients inserted in the DCB parameters are evaluated by means of additional fits to the $m_{\gamma\gamma}$ histograms from these alternative datasets, evaluating the half difference between each DCB parameter in the $+1\sigma$ and the -1σ alternative dataset fit.

As motivated earlier, the right-tail parameters are not modified because the fit range is too small to allow it, and most of the range, on the right-hand side of the peak, contains only the gaussian core of the DCB.

6.5.1 Theory systematics

Theoretical uncertainties impact both the overall cross-section prediction for a given process and the distributions of kinematic variables used in event selection and categorization. When cross-section measurements are performed, the uncertainties in the total cross-section are omitted and instead treated as uncertainties in the SM predictions. However, uncertainties related to event kinematics, which influence the efficiency and acceptance of the analysis, are still taken into account. Both the uncertainties affecting overall cross-section normalizations and those impacting event kinematics are considered when measuring signal strengths and coupling modifiers.

The sources of theoretical uncertainties considered in this analysis are:

- **QCD scale and α_S uncertainties**: The QCD scale uncertainty arises from variations in the renormalization and factorization scales used when calculating the expected SM cross-section and event kinematics, accounting for missing higher-order terms in perturbative calculations. The guidelines provided in [19] are followed. The uncertainty in the overall normalization is estimated from three sources: varying the renormalization scale by a factor of two, varying the factorization scale by a factor of two, and varying both simultaneously in the same direction. Depending on the production process, the magnitude of the normalization uncertainty ranges from approximately 0.5% for VBF production to 2% for ggH production. For the gluon-gluon fusion process, the QCD scale uncertainties are taken into account using the LHC-WG1 scheme, which includes the inclusive N3LO uncertainty [19].

  The α_S uncertainty is due to imperfect knowledge of α_S at the q^2 scale of its event, and it is evaluated for two alternative values, where the scale is fixed at m_Z: $\alpha_S^{\text{down}}(m_Z) = 0.116, \alpha_S^{\text{up}}(m_Z) = 0.120$, while the nominal value employed is $\alpha_S(m_Z) = 0.118$. The overall normalisation uncertainties are computed following the PDF4LHC prescription [19], while the uncertainties in the event kinematics are calculated from the NNPDF PDF sets, which contain sets of weights for the different choices of α_S and QCD scales.

This procedure is only applied thoroughly for the ggH, VBF, VH samples, with S component only, generated with Madgraph, while for the $\mathcal{S}_{ggH} + I$ samples (generated with SHERPA) technical problems due to the process implementation in the generator complicate the estimation of this type of errors. Indeed, there is not a reliable sets of weights in the samples generated and processed with the full CMS simulation, accounting for QCD scale and α_S variations. For this reason, 10M generator-level events were produced for each variation, in order to evaluate their differences in the differential cross section as a function of $m_{\gamma\gamma}$, around the Higgs boson mass. The ratio between the differential cross-section for each variation and the nominal sample was calculated in coarser bins and was fitted with a flat function for each variation. The maximum envelope of these ratios is determined to be around ±10%, as shown in Figure 6.8, therefore a relative Gaussian uncertainty of 7% (equivalent standard deviation of a ±10% uniform band) was associated with the ensemble of QCD scale and α_S uncertainties and was applied to the interference contribution, before summing it with the S component in the $\mathcal{S}_{ggH} + I$ PDF.

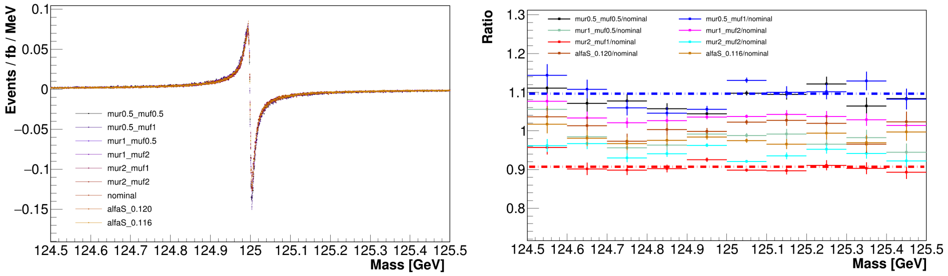

Figure 6.8: Differential cross section with respect to diphoton invariant mass around the Higgs boson mass (left) and ratio between the differential cross-section for each variation and the nominal sample in coarser bins (right).

Given the very small (2%) relative contribution of the interference in terms of cross-section [46], the uncertainty of the normalization for the ggH+int. process is taken to be the same as the ggH process.

Nonetheless, the aforementioned 7% uncertainty is taken into account by generating alternative MC distributions for ggH+int., with interference contributions scaled by 0.93 and 1.07, which are then fitted with the DCB function , as shown in Figure 6.9 for an example category. Systematic uncertainties on the relevant parameters of the $\mathcal{S}_{ggH} + I$ model, i.e. the mean, σ and left-tail DCB parameters, for each category, are evaluated by taking the differences between the fitted values and the nominal values of this parameters, as for others shape-related systematic uncertainties.

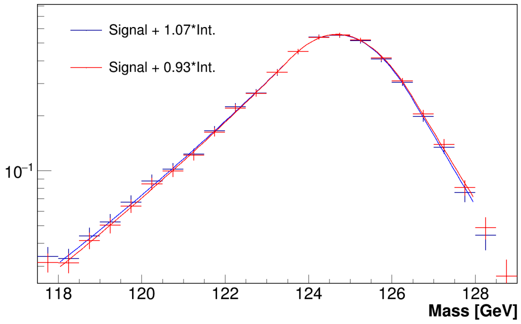

Figure 6.9: DCB fit of alternative distributions with interference contributions scaled by 0.93 and 1.07, as decribed in the text, to account for the uncertainties on the QCD modeling, muR, muF and alphaS.

- **Interference in the VBF process**: There is not much theoretical information in literature on the amount of interference for the VBF process, in particular with the categories used in this analysis, and there is no available simulation of interference in VBF. It is known that the effect should be smaller than for the gluon gluon fusion, with induced mass shifts of the order of 20 MeV, going in the opposite direction as in gluon gluon fusion [19]. For this reason a dedicated systematic uncertainty is applied, as large as the 100% of the effect (at nominal Γ_H)predicted for the gluon-gluon fusion process. This is performed by assigning systematic uncertainties on the left-tail DCB parameters for the VBF process, for each category, equal to the differences between the $\mathcal{S}_{ggh} + \mathcal{I}$ values and the signal-only values of this parameters for the gluon-gluon process. Nonetheless, the impact of this uncertainty is negligible, hence to simply the fit the relevant nuisance parameters have been pruned in the final evaluations . This has been verified by including this uncertainty on the VBF process, and estimating the expected sensitivity of the analysis.

- **PDF (parton density functions) uncertainties**: These uncertainties account for the imperfect knowledge of the momentum distributions of the parton inside protons. For signal samples, the overall normalization uncertainties are computed following the PDF4LHC prescription [19], while the uncertainties in the event kinematics are calculated from the NNPDF PDF sets. The PDF sets produce a set of 60 weights, which are combined, to simplify the fit, by using the Principal Component Analysis (PCA) algorithm, to evaluate a single PDF weight component, taking into account the main part of the variance, reducing the number of PDF events weights to 1, for all production modes [94]. This is possible because the impact of PDF uncertainties in the analysis is very limited. Principal Component Analysis (PCA) is a linear dimensionality reduction technique. It linearly transforms data onto a new coordinate system defined by principal components, which are orthonormal vectors capturing the directions of maximum variance. Formally, the principal components are a sequence of p unit vectors, where the i-th vector maximizes variance while being orthogonal to the first $i-1$ vectors. Each principal component corresponds to a direction that minimizes the average squared perpendicular distance from the data points to the line, ensuring linear uncorrelation among transformed dimensions. An important diagnostic quantity in the PCA is the *explained variance ratio*, which indicates the proportion of the total variance in the dataset that is captured by each principal component. By examining the cumulative explained variance ratio, one can determine how many components are sufficient to represent the data with minimal loss of information. For instance, if the first principal component explains most of the variance, it may be justified to retain only that component for subsequent analysis.

  In order to validate the usage of PCA, the explained variance ratio was evaluated on all samples, by performing the PCA analysis with 5 components instead of one. As shown in Figure 6.10 for the interference process, the first component contains at least 75% of the total variance, which is acceptable because the impact of PDF uncertainties is expected to be of the order of 0.1%, as in previous CMS Run 2 $H \to \gamma\gamma$

analyses [24].

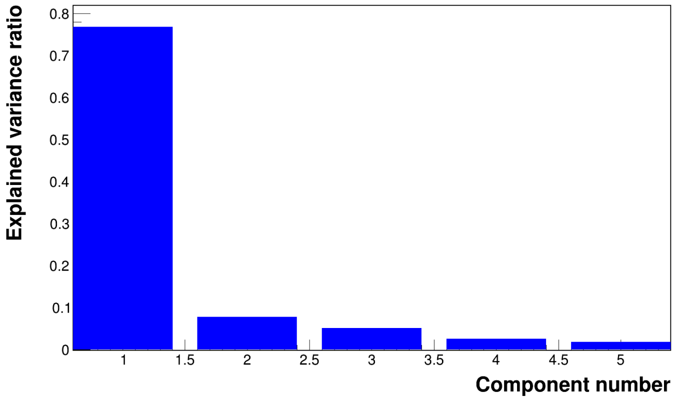

Figure 6.10: Explained variance ratio for the interference sample, after the application of the PCA algorithm with 5 components on the PDF weights.

- **Uncertainty in the $H \to \gamma\gamma$ branching fraction**: the probability of the Higgs boson decaying to two photons is required to calculate the SM expected cross section, but this branching fraction is not known exactly. The uncertainty is currently estimated to be 2% [19]. Since this uncertainty affects the overall normalization of the signal processes, it is included in the signal strength and coupling modifier measurements, and is considered as an uncertainty in the SM predictions for cross section measurements.

6.5.2 Experimental systematics

The experimental systematic uncertainties are divided in shape and normalization uncertainties.

Shape experimental systematics

- **Residual photon energy scale and smearing:** Corrections for the residual differences between the photon energy scales and resolution in the data and simulation are discussed in 4.2.1. The sources of systematic uncertainty that directly affect them are the variations in R_9, electron identification, preselection E_T threshold, and non-closure. Both scale and smearing systematics uncertainties are binned depending on

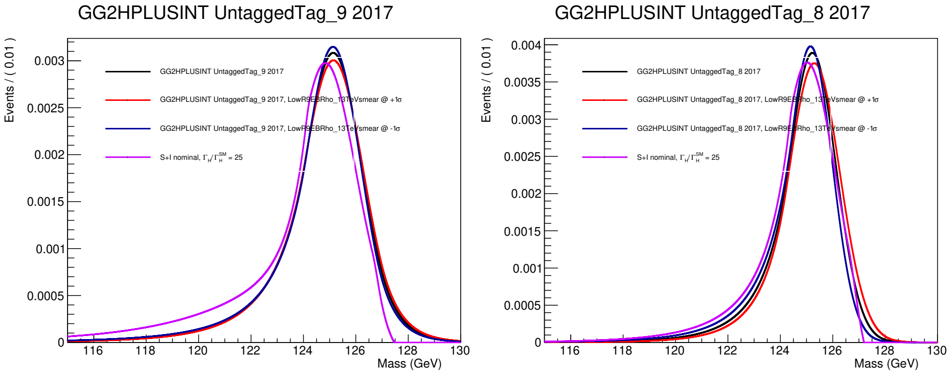

Figure 6.11: Comparison between the nominal signal model, the one for $\Gamma_H/\Gamma_H^{SM} = 25$, and the models when the nuisance corresponding to the smearing in the bin $R_9 < 0.965$ in the barrel in 2017 is equal to $\pm 1\sigma$, for the best (left) and worst (right) category (in terms of impact of this systematic uncertainty).

the photon $|\eta|$ (barrel or endcap) and R_9 value, in two bins per variables for a total of 4. For the scale corrections, all the related uncertainties are added in quadrature for each systematic uncertainty, to assign combined errors for the scale corrections for each bin of η and R_9. The value is estimated by comparing the $H \to \gamma\gamma$ and $Z \to e^+e^-$ scale responses, after appliying the scale corrections, using simulated samples. On the other hand, for the smearing corrections, the assigned systematics is a flat uncertainty of 0.05% regardless of the $R_9/|\eta|$ bin, except for 2016 and 2017, where an inflated uncertainty of 0.8% is applied for all photons with $p_T > 50$ GeV, to cover observed non-closures. These inflated uncertainties are, effectively, the most impactful uncertainties in the analysis. The smearing systematic uncertainties alter the signal shape in a way comparable to the variations due to the interference, therefore a comparison between the nominal signal model, the one for $\Gamma_H/\Gamma_H^{SM} = 25$, and the alternative models for one of the most impactful uncertainties, i.e. the error on the smearing in the bin $R_9 < 0.965$ in the barrel in 2017. The results are shown in Figure 6.11, for $\mathcal{S}_{ggH} + I$ in the best and worst case, and demonstrate that in the most sensitive categories the effect of smearing uncertainties and interference can be disentangled.

- **Non-linearity of energy scale:** In order account for the fact that the scale is calibrated on $Z \to e^+e^-$ and applied to $H \to \gamma\gamma$, where the average p_T of ECAL superclusters is higher, a 0.2% scale uncertainty addressing ECAL non-linearity, derived by comparing E/p for electrons with $E_T \sim 90/2 = 45$ GeV and high-energy electrons (i.e. with $E_T > 50$ GeV), as shown in Figure 6.12, is applied. This uncertainty, because of how large it is compared to the other ones, is applied directly in the DCB mean value parameters, without fitting alternative templates.

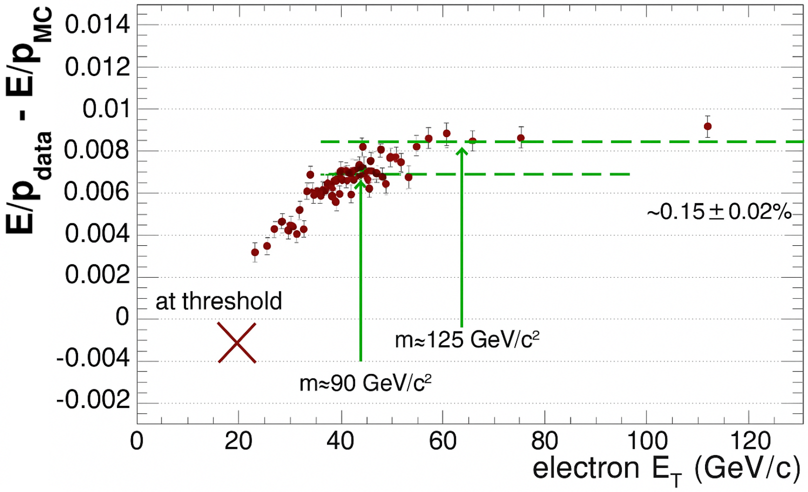

Figure 6.12: Data/MC differences in electrons E/p as a function of electron p_T, showing the difference between ECAL Z and Higgs boson typical energy scales in Run 1. The value of 0.2% was derived in Run 2 with the same technique, yielding a larger coefficient, due to increased integrated radiation doses.

- **Shower Shape Corrections:** Systematic uncertainties related to shower shape variables are due to the fact that the energy regression and scale corrections both employ the shower shape variables, which are different for electrons (used to calibrate) and photons. The impact is ealuated by comparing the simulated energy response of $Z \to e^+e^-$ events, after reweighting the shower shapes variables from the $Z \to e^+e^-$ to the $H \to \gamma\gamma$ case. The uncertainty in the energy scale due to this effect varies from 0.01% to 0.15%, depending on the photon $|\eta|$ and R_9 values.

- **Non-uniformity of light collection:** The systematic uncertainties on the FNUF scale corrections, described in Section 4.2.1, are evaluated by combining the uncertainties from the laser transparency corrections, and the light collection model. The FNUF corrections are only implemented for high R_9 photons, because low R_9 (converted) photons interact with ECAL in the same way as electrons. The largest uncertainty are for the 2018 data-taking, because of the increasing integrated ionizing dose on ECAL crystals, and amounts to about 0.12% in the barrel and 0.3% in the endcap.

- **Modeling of material upstream of ECAL:** The amount of material through which particles traverse before reaching the ECAL, mainly due to tracker sensors and support structures, influences the behavior of electromagnetic showers and may not be accurately represented in simulation, causing differences in the reconstructed scale of photons and electrons. Dedicated MC samples with ±10% variations in the material budget upstream of ECAL are used to estimate the impact on the photon energy scale, by means of a double ratio between the reconstructed over generated energy for ggH($\gamma\gamma$) at $m_H = 90$ GeV (to decouple differences due to non-linearity) and $Z \rightarrow ee$, with respect to the same double ratio with the nominal material budget. This procedure is performed in 6 categories, by dividing the sample two R_9 bins with a boundary at 0.965, and in 3 $|\eta|$ regions. The double ratio as a function of the material budget variation is shown in Figure 6.13 for the category with the largest variations (endcap high R_9).

The uncertainties that modify only the event yield include:

- **Integrated Luminosity:** Uncertainties of 2.5%, 2.3%, and 2.5% are determined by the CMS luminosity monitoring for the 2016, 2017, and 2018 data sets, respectively [58, 76, 77]. These uncertainties are partially correlated across different data sets to account for common sources of uncertainty in the luminosity measurement schemes.

- **Photon Identification BDT Score:** The uncertainty associated with the photon identification BDT score is estimated by varying the

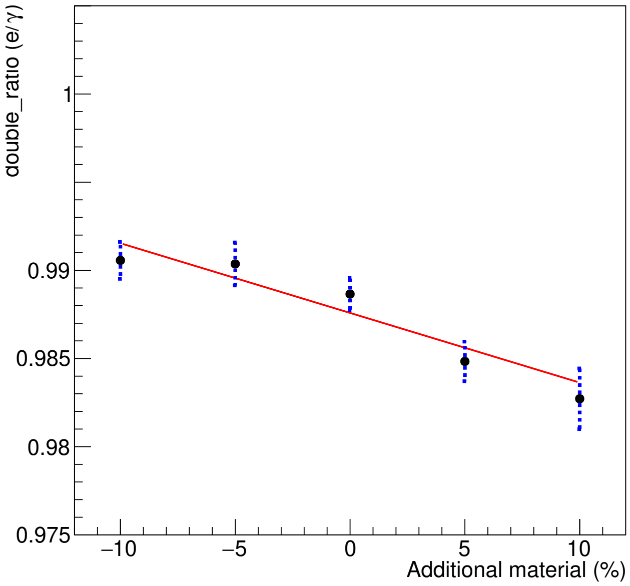

Figure 6.13: Double ratio between E_{reco}/E_{gen} of $ggH(\gamma\gamma)$ at $m_H = 90$ GeV and $Z \to ee$, with respect to the same ratio with the nominal material budget.

set of events used to train the quantile regression corrections, which captures the residual discrepancies between data and simulation. The uncertainty in signal yields is estimated by propagating this uncertainty through the complete analysis chain.

- **Jet Energy Scale and Smearing Corrections:** The jet energy scale is measured using the p_T balance of jets with Z bosons and photons in $Z \to e^+e^-$, $Z \to \mu^+\mu^-$, and $\gamma +$jets events, along with the p_T balance between jets in dijet and multijet events [95]. The uncertainty in the jet energy scale is a few percent and is dependent on p_T and η. The impact of jet energy scale uncertainties on event yields is evaluated by varying the jet energy corrections within their uncertainties and propagating the effect to the final result. This analysis makes use of only 1 uncertainty for jet scale and 1 for resolution, and they are uncorrelated across the years.

- **Per-photon energy resolution estimate**: This uncertainty is designed to cover all differences between data and simulation in the distribution of the mass resolution, predicted from the per-photon resolution, which is an output of the energy regression. The predicted mass resolution is an input of the diphoton BDT, therefore the uncertainties propagate to the event yield. The maximum yield variation due to this systematic uncertainty in the analysis categories is around 5%, however for most categories the impact is below the percent level.

- **Trigger Efficiency:** The efficiency of the trigger selection is evaluated with $Z \rightarrow e^+e^-$ events using the tag-and-probe technique. The magnitude of this uncertainty is less than 1%. An additional uncertainty accounts for a gradual shift in the timing of the ECAL L1 trigger inputs in the region where $|\eta| > 2.0$, which caused specific trigger inefficiencies during the data taking in 2016 and 2017. Both photons and, to a greater extent, jets can be affected by this inefficiency, with the maximum impact on yields for categories targeting VBF production reaching 1.4%.

- **Photon Preselection:** The uncertainty in the preselection efficiency is calculated from the error on the ratio of efficiencies measured in data and simulation, i.e. the so-called scale factors, and its uncertainty is less than 1%.

Chapter 7

Results

The results of this analysis are first derived in the expected case, i.e. without using data in the signal region, and in the observed case, where the final Γ_H constraints are derived. All the results are extracted by performing a binned fit, with 0.25 GeV width, to the data combining all categories and data taking years, using the CMS combine framework. As explained in Section 6.1, the physical parameters μ_F, μ_V, m_H are free in the fit, but are constrained with data, thanks to the scaling of the signal components with μ_F or μ_V and to the dependence of the signal template functions on m_H. Moreover, as explained in Section 6.1, the binning needs to be smaller than the mass resolution, in order to have sensitivity to the shape variations due to the interference, and to reduce the bias due to the background steep derivative.

The fit is performed by minimizing the negative log-likelihood, starting from the likelihood definition in Eq. (6.2):

$$Z = -2\Delta \ln \mathcal{L} = -2 \ln \frac{\mathcal{L}(data|\Gamma_H/\Gamma_H^{SM}, \vec{\theta})}{\mathcal{L}(data|\hat{\Gamma_H}/\Gamma_H^{SM}, \hat{\vec{\theta}})} \tag{7.1}$$

where $\hat{\Gamma_H}/\Gamma_H^{SM}$ and $\hat{\vec{\theta}}$ are the best-fit values of the Higgs boson width modifiers and nuisance parameters. The 68% and 95% confidence intervals are identified when Z = 1.00, and 3.84 respectively. Due to the mathematical boundary on Γ_H in the physics model, it is necessary, in some steps of the limit extraction, to re-parametrize $\sqrt{\mu_F \Gamma_H/\Gamma_H^{SM}}$ as $\sqrt{\mu_F}\lambda$ with $\lambda = \sqrt{\Gamma_H/\Gamma_H^{SM}}$, so that λ can reach negative values allowing the fit to work in a well-behaved fashion, specifically with mathematical continuity both in the likelihood and in its derivatives.

7.1 Expected results

Expected results for the parameter of interest (POI) of the physics model ($\frac{\Gamma_H}{\Gamma_{SM}}$) are obtained in three ways:

1. Via 1D likelihood scans in Γ_H, by comparing, for each Γ_H^{test} value, the Asimov dataset, i.e. a set of histograms (one for each category) interpolating the $S + B$ (signal and background) models in 0.25 GeV bins for $\Gamma_H = 1$, with the same histograms evaluated using the signal models at $\Gamma_H = \Gamma_H^{test}$. In summary, the likelihood is evaluated by using the Asimov datasets as data points for each category;
2. Via frequentist limits calculations at different confidence levels, in the asymptotic approximation, dubbed as *AsymptoticLimits* method;
3. Via the Feldman-Cousins approach, which evaluates the frequentist limits at different confidence levels, by using toy datasets, i.e. several $O(1000)$ sets of histograms, generated from the Asimov histograms using Poisson fluctuations in all bins.

The so-called asymptotic approximation used in the limit calculation relies on a set of statistical assumptions that ensure the validity of the analytic results derived from Wilks' and Wald's theorems [96]. These assumptions are generally satisfied in the large-sample limit and allow the profile likelihood ratio test statistic to follow a simple and universal form. The likelihood function $\mathcal{L}(\lambda, \boldsymbol{\theta})$, in particular, is assumed to be smooth and well-behaved in the vicinity of its maximum. In this regime, the log-likelihood function can be approximated by a quadratic expansion around the global maximum likelihood estimator (MLE), such that

$$-2 \ln \frac{\mathcal{L}(\lambda, \boldsymbol{\theta})}{\mathcal{L}(\hat{\lambda}, \hat{\boldsymbol{\theta}})} \approx \frac{(\lambda - \hat{\lambda})^2}{\sigma^2}, \tag{7.2}$$

where $\hat{\lambda}$ denotes the MLE of the parameter of interest, and σ represents its standard deviation estimated from the curvature of the log-likelihood at the maximum. Under these conditions, the estimator $\hat{\lambda}$ is approximately Gaussian distributed around the true value λ_{true} with variance σ^2.

The *AsymptoticLimits* method is based on an asymptotic approximation of the distributions of the LHC test statistic, which is defined as a profile likelihood ratio, under two hypotheses, usually representing signal and background-only. In this analysis, the parameter of interest is the interference strength modifier, λ, which scales the amplitude of the interference term between the Higgs boson production and the continuum diphoton background. The two hypotheses considered are therefore

- the **background-only hypothesis**, corresponding to $\lambda = 0$, where no interference contribution is present, and
- the **signal-plus-background hypothesis**, corresponding to $\lambda = 1$, which represents the Standard Model prediction for the interference strength.

The test statistic used to test a given value of λ is defined as:

$$q_\lambda = \begin{cases} -2\ln\dfrac{\mathcal{L}(\lambda, \hat{\hat{\boldsymbol{\theta}}}_\lambda)}{\mathcal{L}(\hat{\lambda}, \hat{\boldsymbol{\theta}})} & 0 \leq \hat{\lambda} \leq \lambda, \\ 0 & \hat{\lambda} > \lambda, \end{cases} \tag{7.3}$$

where $\mathcal{L}(\lambda, \boldsymbol{\theta})$ is the likelihood function of the data given the interference strength λ and the set of nuisance parameters $\boldsymbol{\theta}$. The notation $\hat{\lambda}$ and $\hat{\boldsymbol{\theta}}$ denote the unconditional maximum likelihood estimators of these parameters, while $\hat{\hat{\boldsymbol{\theta}}}_\lambda$ are the conditional estimators obtained by maximizing the likelihood for a fixed value of λ. The constraint $\lambda \geq 0$ is enforced in the fit, such that if the unconstrained best-fit value $\hat{\lambda}$ is negative, it is set to zero when evaluating Eq. (7.3). This ensures a one-sided test statistic, appropriate for setting upper limits on λ.

Under the asymptotic approximation [96], the sampling distribution of q_λ is known analytically both under the signal-plus-background hypothesis ($\lambda = 1$) and under the background-only hypothesis ($\lambda = 0$). This allows the computation of two p-values without the need to generate pseudo-experiments:

$$p_\lambda = P(q_\lambda \geq q_\lambda^{\text{obs}} \,|\, \lambda = 1), \tag{7.4}$$

$$p_b = P(q_\lambda \geq q_\lambda^{\text{obs}} \,|\, \lambda = 0). \tag{7.5}$$

The quantity p_λ quantifies the compatibility of the observed data with the signal-plus-background hypothesis (Standard Model interference), while p_b

quantifies the compatibility with the background-only hypothesis (no interference).

The modified frequentist confidence level is defined as

$$CL_s = \frac{p_\lambda}{1 - p_b}, \tag{7.6}$$

which protects against excluding the signal hypothesis in the presence of downward statistical fluctuations of the background. The upper limit on λ at a confidence level α (typically 95%) is determined by finding the largest value of λ for which $CL_s(\lambda) \leq 1 - \alpha$.

Before checking the validity of the asymptotic approximation for this analysis, an Asimov-based estimation of the expected sensitivity is performed. The result of the Negative log-likelihood (NLL) scan on the Asimov dataset corresponding to full Run2 data and MC is shown in Figure 7.1, together with a breakdown of uncertainties in different groups, performed by gradually freezing groups of systematic uncertainties.

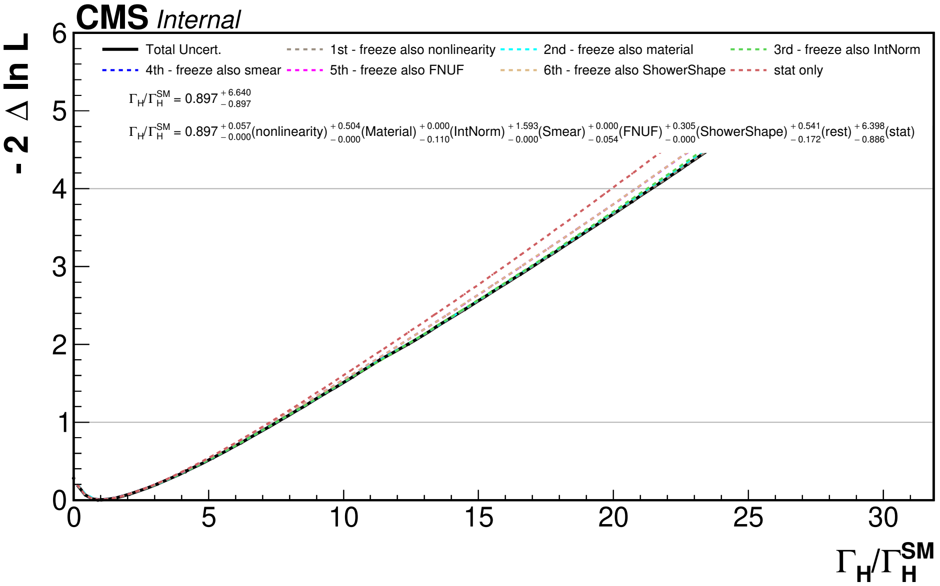

Figure 7.1: Negative log-likelihood scan on the Asimov dataset for full Run2, together with a breakdown of uncertainties in different groups.

The resulting expected limits for full Run2 at 68(95)% CL are, respectively, $\Gamma_H/\Gamma_H^{SM} < 6.6(21.2)$, translating in absolute limits of 27 (87) MeV.

The impact of systematic uncertainties is smaller than the statistical one, amounting to a in quadrature (linear) difference of 27% (4%) between the nominal and statistics only 95% CL limits.

For each nuisance parameter, the impact is evaluated by fixing its value to $\pm 1\sigma$, re-doing the fit and subtracting by the obtained Γ_H values the nominal ones. The impact of the most important profiled parameters is shown in Figure 7.2, and it contains both systematic uncertaintied and profiled quantities, such as μ, M_H and background shape parameters. The full impacts are available in Figure A.41 in the Appendix. Some nuisance parameters describing energy scale and smearing errors, have one-sided impacts because of the non-trivial interplay between scale and resolution and the lineshape fit, when fixing one nuisance parameter at the time.

A successful closure test was performed to validate the signal parametrization, by generating template datasets using directly the MC samples for signal and interference, instead of the functional models, and the best p.d.f. for the background, as shown in Figure 7.3.

The NLL scan on the Asimov dataset with λ as POI is available in Figure 7.4, showing compatibility with the scan in Γ_H/Γ_H^{SM}. The fact that the minimum using an Asimov-based scan is not exactly at $\Gamma_H = 1$ is an effect of the binning. It was verified that decreasing the bin size to 0.125 GeV, the minimum moved to, $\lambda = +0.998^{+1.809}_{-1.824}$ (68% CL). Given the increased run time of the fit with double binning, the small discrepancy with respect to the nominal case, and the successful closure test with MC datasets, it was decided to keep the 0.25 GeV binning.

In order to check for the presence of biases in the statistical procedure, around 500 toys at nominal Γ_H have been generated and fitted with the nominal likelihood, and the histogram of the best-fit values divided by the error (formula in the x-axis) has been fitted. The distribution of the pulls for toy pseudo-experiments is available in Figure 7.5, showing very small bias, measured using the average, but also a variance different from the expected value of 1, meaning that the asymptotic approximation cannot be applied. This behaviour reflects the non-trivial interplay between the shape modification due to the interference and the background, and also the correlation between the floated parameters $(\Gamma_H, m_H, \mu_F, \mu_V)$, and is also coherent with the non-parabolic shape of the $\sqrt{-2\Delta \ln L}(\Gamma_H)$ function, shown in Figure 7.4. Indeed, in the asymptotic approximation, which in this analysis fails, the pull variance is close to 1 and the $\sqrt{-2\Delta \ln L}(\Gamma_H)$ function is parabolic around the minimum, allowing to extract 68% and 95% CL constraints directly from

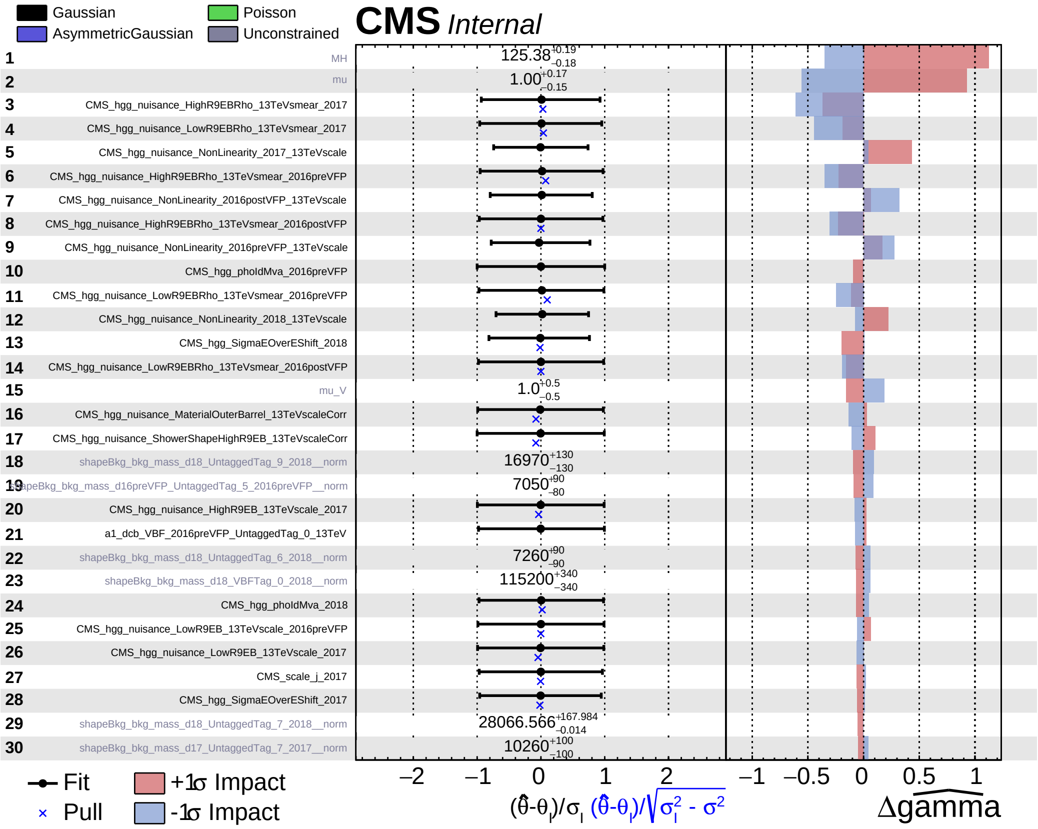

Figure 7.2: Expected impacts of systematic uncertainties. For the unconstrained parameters (i.e. M_H, μ_F, μ_V and background functions normalizations) the reported errors correspond to their 68% error in the nominal fit.

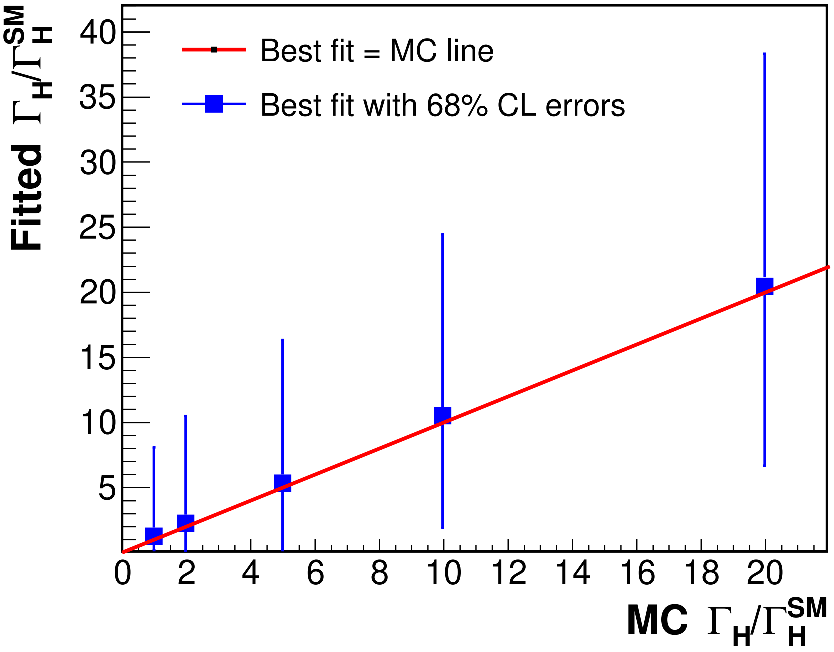

Figure 7.3: Closure test of the Γ_H/Γ_H^{SM} fit, by using histograms from MC, summed with background PDF, to generate the Asimov datasets, instead of the signal models.

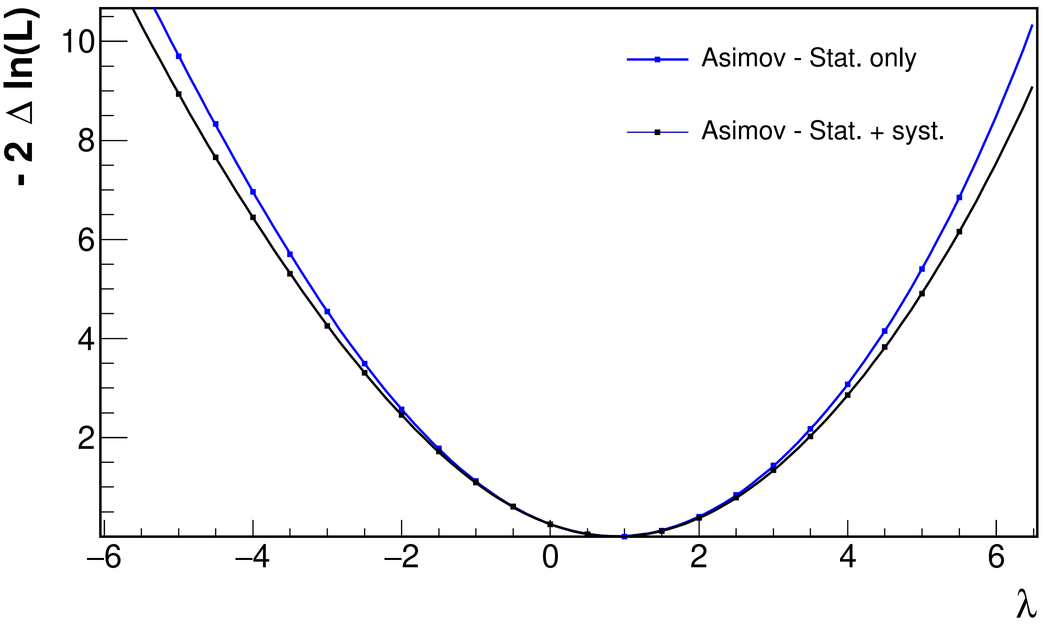

Figure 7.4: NLL scan of the Asimov dataset, extracted from background and signal PDFs, for the parameter λ.

the $\sqrt{-2\Delta \ln L}(\Gamma_H)$ function.

Given the failure of the asymptotic approximation, a toy-based limit extraction procedure must then be applied. Indeed, to extract the upper limit on λ, in this analysis the Feldman-Cousins approach, chosen because of its generality, is employed. The Feldman-Cousins results are also compared, as a control, with the results drawn from a asymptotic frequentist approach. Note also, that both the Feldman-Cousins and the asymptotic approach take into account the physical boundary $\lambda > 0$, to derive the limits, by normalizing the likelihood at its value at $\lambda = 0$ instead that at the minimum, in case where the best fit value is $\lambda < 0$. The results, in terms of upper limit at 95% CL (UL), of the *AsymptoticLimits* method, in the expected (blinded) case are shown in Table 7.1, together with the value obtained in the Asimov-based scan.

Feldman–Cousins confidence intervals The Feldman–Cousins method provides a unified approach for constructing confidence intervals that handles both upper and lower bounds in a statistically consistent way, avoid-

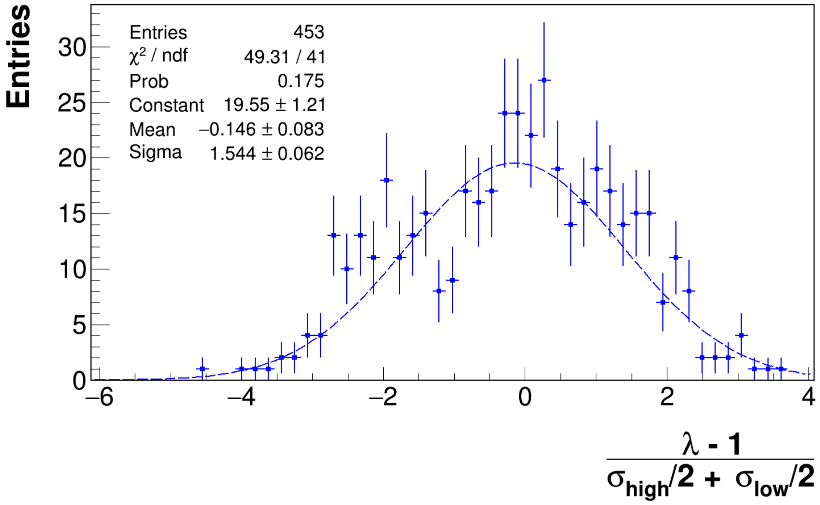

Figure 7.5: Pulls of the best-fit values, divided by the average error, for a set of toys.

UL on λ at 95% CL	**UL on $\Gamma/\Gamma_{\text{SM}}$**	**Exp. quantile**	
1.9	3.4	0.025	-2σ
2.5	6.0	0.160	-1σ
3.3	11.5	0.500	Median
4.6	21.0	*N/A*	*Asimov*
4.7	22.5	0.840	+1σ
6.4	41.3	0.975	+2σ

Table 7.1: Results of the *AsymptoticLimits* method, in the expected (blinded) case.

ing classical flip-flopping between one-sided and two-sided intervals [44].The Feldman-Cousins (FC) procedure for computing confidence intervals, using the test statistic $q(\lambda)$, require, for each point λ to:

- compute the test statistic $q_{\mathrm{obs/quant}}(\lambda)$ at the desired quantile (in the expected case), for example by taking the median, or the $\pm N\sigma$ quantile value, or in the observed case;
- compute the expected distribution of $q(\lambda)$ under the hypothesis of λ as the true value;
- evaluate the confidence level $p_\lambda = P\left[q(\lambda) > q_{\mathrm{obs/quant}}(\lambda) \mid \lambda\right]$

At the end of the procedure, a p_λ function is obtained, and any λ point belongs to a certain confidence region CL_α if p_λ is larger than α (e.g. 0.3173 for a 1σ region, as $1-\alpha = 0.6827$). The p_λ functions at different expected quantiles (median, $\pm 1\sigma$, $\pm 2\sigma$) are shown in Figure 7.6. In particular, in the expected median case, the limit at 95% CL is $\Gamma_H < 140$ MeV.

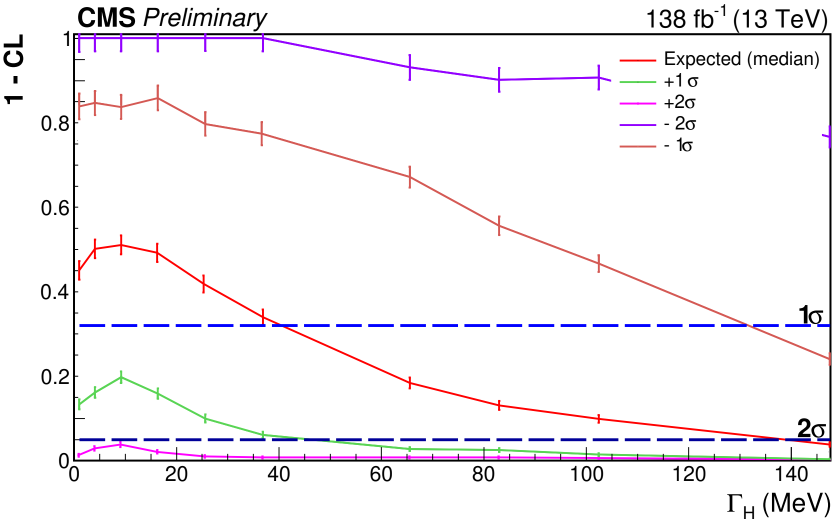

Figure 7.6: P-values p_λ as a function of λ at different expected quantiles (median, $\pm 1\sigma$, $\pm 2\sigma$).

7.2 Pre-unblinding tests

Goodness of fit tests have been performed on toys and data with both the Kolmogorov-Smirnov and Anderson-Darling algorithms, as shown in Figure 7.7, with good results, while blinded observed impacts are shown in Figure 7.8. Only one observed pull in the impacts is above 2.1 σ, and no pull above 3 σ is present.

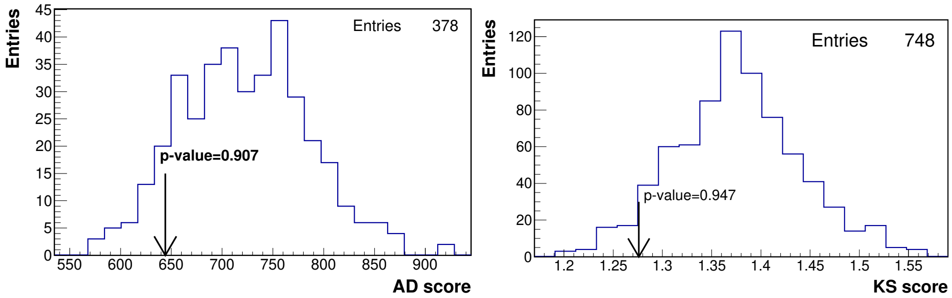

Figure 7.7: Goodness of fit tests on toys, together with the value obtained on data, with both the Kolmogorov-Smirnov (left) and Anderson-Darling (right) algorithms.

Background-only fits, where the signal region 115-135 GeV is not employed, have been performed, and the aggregated plot, summing the event for all categories, is shown in Figure 7.9.

Fits to single categories are available in Figures A.43, A.44, A.45 and A.46 in the Appendix.

7.3 Observed results

Finally, an observed log-likelihood scan in λ has been performed, as shown in Figure 7.10.

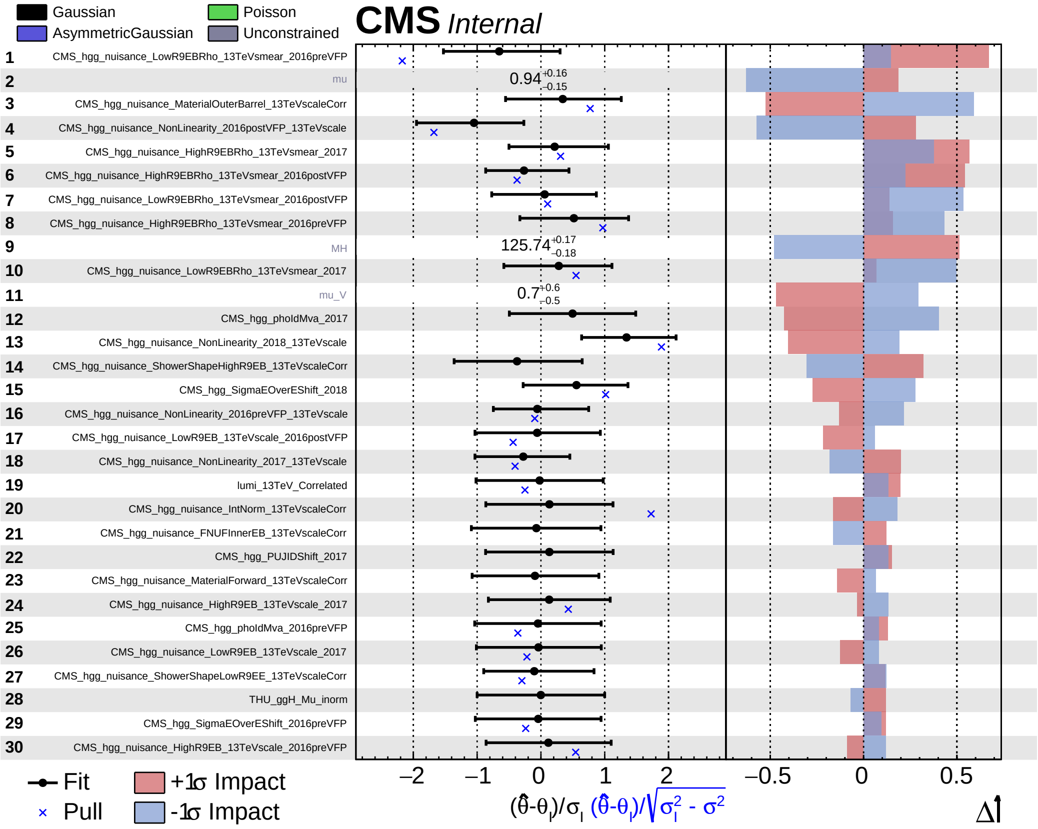

Figure 7.8: Blinded observed impacts.

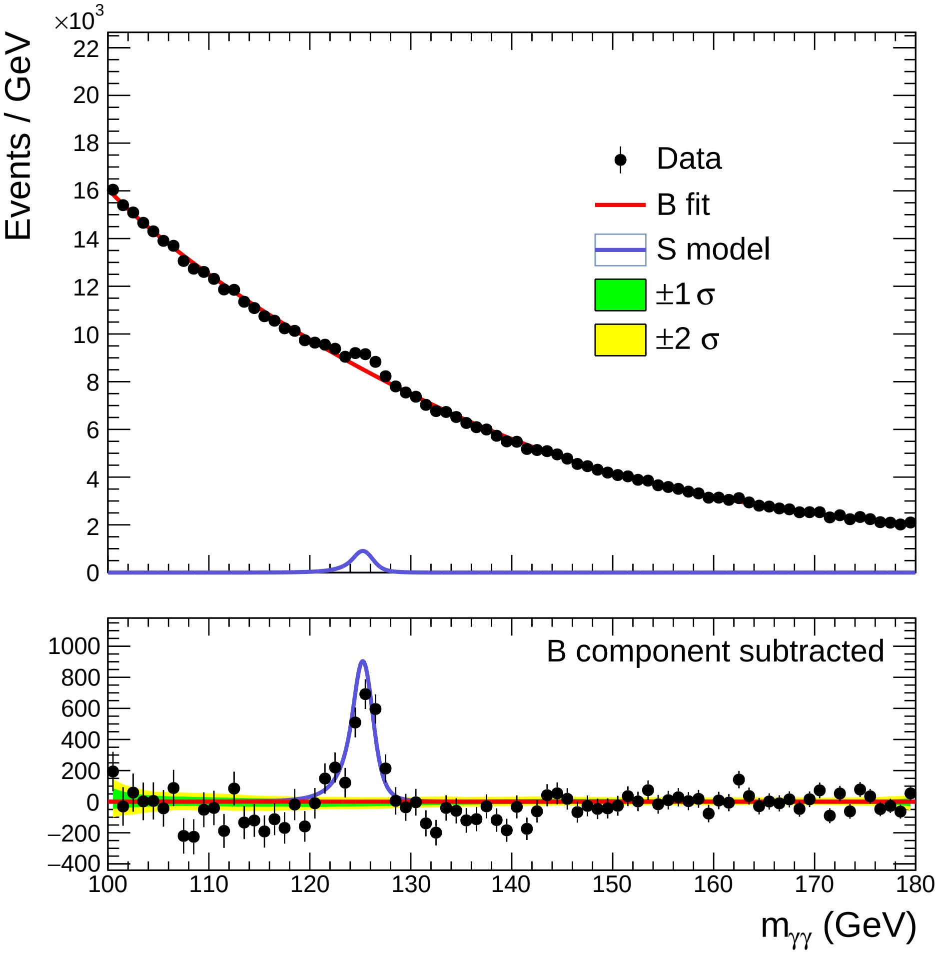

Figure 7.9: Sum of events for all categories, with background-only fit performed in the data sidebands, together with the nominal signal model ($\mu_F = \mu_V = \lambda = 1$).

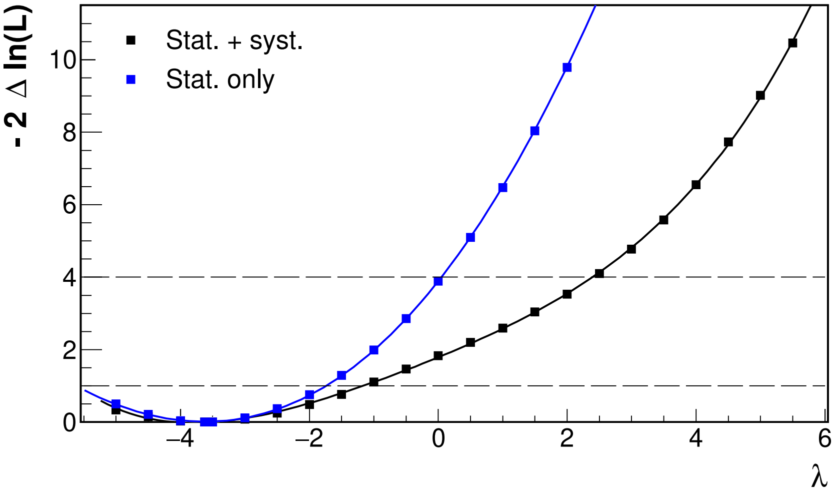

Figure 7.10: Observed log-likelihood scan of λ.

The best fit value is $\lambda = -3.6$, and the impact of systematics is larger than in the expected case, due most likely to the fit converging in a point very far from the nominal one. From the background-only fits, it can be seen that the bin at 123 GeV is under-fluctuating, driving the signal shape to be narrower with respect to the nominal signal model, up to the level where the fit prefers that the tails are suppressed on the left, which is the opposite case as the nominal interference. This effect, which is a kind of under-fluctuation, motivates the negative best fit value. Physically, a negative λ means that the interference yield (the coefficient in the physics model) is negative (i.e. constructive interference instead of destructive as the nominal case).

Fits to the full invariant mass region, without blinding the signal window, and leaving the physical parameters free, i.e. "signal+background" fits are shown in Figure 7.11 after summing events of all categories, without and with a per-category S/B reweighting. Fits to single categories are available in Section A.7 in the Appendix.

A "signal+background fit" is shown in Figure 7.12, summing events of all categories, with a per-category S/B reweighting and background subtraction, together with the signal+background template at different values of λ, where the fit behaviour can be understood more clearly.

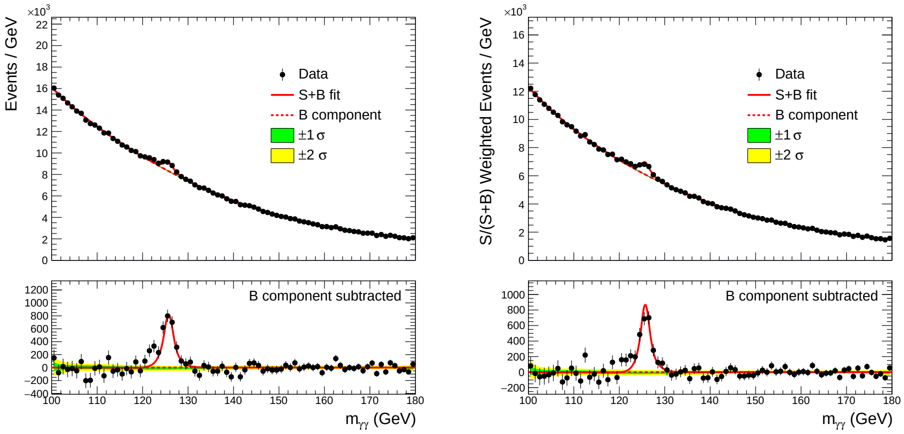

Figure 7.11: Signal+background fits, after summing events of all categories, without (left) and with a S/B reweighting of each category(right).

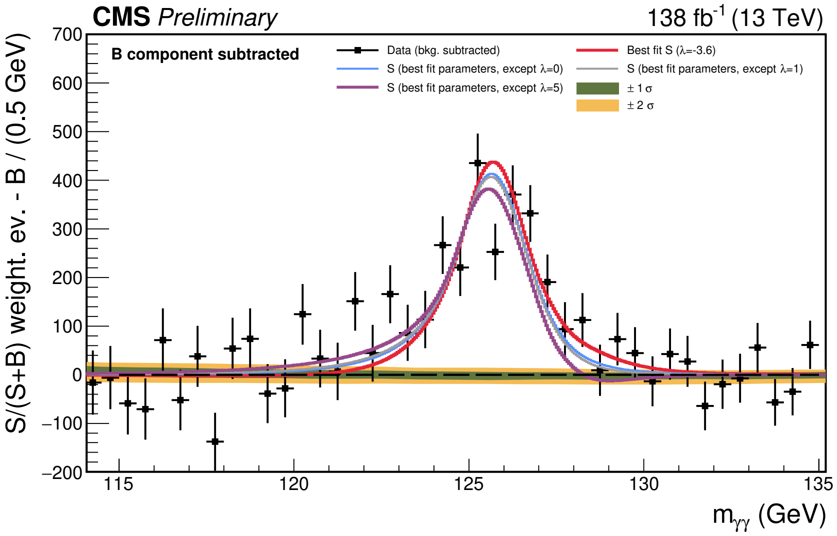

Figure 7.12: Signal+background fit, after summing events of all categories with a per-category S/B reweighting, and background subtraction, together with the signal template at different values of λ.

Observed limits are obtained both in the asymptotic approximation (using the Combine *AsymptoticLimits* routines) and with the Feldman-Cousins approach, and in both cases the physical boundaries $\lambda > 0$ has been applied by subtracting the log-likelihood at $\lambda = 0$ to the log-likelihood values for positive λ. The observed limit at 95% CL obtained with AsymptoticLimits is $\lambda < 3.5$, which is lower than the value expected from the Asimov dataset ($\lambda < 4.6$), and this reflects the best fit value underfluctuation.

Observed results using the Feldman-Cousins are available in Figure 7.13, and summarized in Table 5.2, comparing them with the other methods. As expected, the Feldman-Cousin based limit is the most conservative.

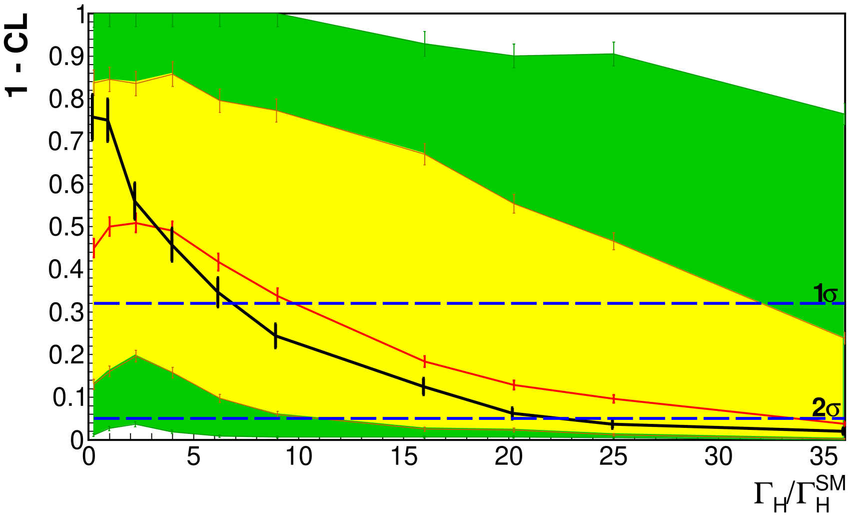

Figure 7.13: Plot of the statistical coverage provided by the Feldman-Cousins approach (with 1000 toys for each point), in terms of 1-CL vs. Γ_H/Γ_H^{SM}, in the median expected scenario (red line), observed (black line), together with 1 and 2σ yellow and green bands.

The final limits are the ones given by the Feldman-Cousins approach: $\Gamma_H < 91.8(138.2)$ MeV in the observed(expected) case. This limit is more than 3 times better than the one put in the $H \to ZZ^* \to 4l$ channel in the on-shell case ($\Gamma_H < 300$ MeV) [11].

	λ	$\Gamma/\Gamma_{\mathrm{SM}} = \lambda^2$	Γ (MeV)
Best fit w/o phys. boundary	-3.6	N/A	N/A
Best fit with phys. boundary	0	0	0
Expected UL (68% CL) in asympt. approx.	2.7	7.4	30.3
Observed UL (68% CL) in asympt. approx.	1.3	1.6	6.4
Expected UL (95% CL) in asympt. approx.	4.6	21.0	86.0
Observed UL (95% CL) in asympt. approx.	3.5	12.3	50.2
Expected UL (68% CL) with Feldman-Cousins	3.1	9.8	40.2
Observed UL (68% CL) with Feldman-Cousins	2.6	6.9	28.1
Expected UL (95% CL) with Feldman-Cousins	5.8	33.7	138.2
Observed UL (95% CL) with Feldman-Cousins	4.7	22.4	91.8

Table 7.2: Summary of the best fits and limits at 68% and 95% CL in different aproaches. The expected limits are provided with Asimov datasets in the asymptotic case and with the median expected quantile for Feldman-Cousins.

Summary

The CMS Higgs boson physics program has yielded a wide range of measurements and constraints of Higgs boson properties, one of which is presented in this dissertation. This work focuses on the determination of the Higgs boson total width in the diphoton decay channel, using CMS data collected during LHC Run 2 at a center-of-mass energy of 13 TeV, corresponding to an integrated luminosity of 137 fb^{-1}. The analysis relies on the interference between signal and continuum background processes, which provides indirect sensitivity to the Higgs boson width. Indeed, in the $\gamma\gamma$ final state, the interference between Higgs boson production and continuum background processes induces both a shift in the apparent resonance peak position and a change in the shape of the diphoton mass spectrum. These effects are modelled by performing a shape fit to the $m_{\gamma\gamma}$ distribution. To optimise the sensitivity, events are categorised according to the transverse momentum of the diphotons, and to a multi-variate score that enhances signal-background discrimination.

The analysis includes the dominant interference contributions from gg and qg initial states. The resulting constraints on the Higgs boson width are obtained leaving also the overall couplings of the Higgs boson to top quark and vector bosons, and the Higgs boson mass, as free parameters in the fit. The observed (expected) upper limit is $\Gamma_H < 92(138)$ MeV at 95% confidence level, where the Standard Model prediction is $\Gamma_H^{\mathrm{SM}} = 4.1$ MeV. The main challenge of the constraint arises from the small amplitude of the interference effects relative to the overall diphoton mass spectrum. This analysis provides the first experimental on-shell constraint on Γ_H that employs the interference-based approach, and improves both the current limit on Γ_H in the diphoton channel and the on-shell limit in the ZZ^* on-shell case ($\Gamma_H < 330$ MeV).

Acknowledgments

Several people supported my effort on this thesis. First and foremost, I would like to express my deepest gratitude to my advisor *Emanuele Di Marco* for his invaluable guidance, unique expertise in this field, huge patience, and continuous support throughout this work. His advice and insights have been fundamental for the completion of this thesis.

I am also grateful to the *CMS Rome group*, especially *Francesca Cavallari*, *Riccardo Paramatti*, and *Chiara Rovelli*, for their support.

I would like to thank the *analysis team* — *Badder*, *Rajdeep*, *Filippo*, *Amrutha*, *Mattia*, and *Toyoko* — for all the contributions and discussions that shaped the analysis presented in this thesis.

Special thanks go to the *Higgs L2 convener Meng Xiao*, and to the *L3 conveners Alessandro Tarabini* and *Fabrice Couderc*, for their helpful feedback and suggestions during the review process.

I am indebted to the *Analysis Review Committee (ARC)* members — *Andrea Benaglia*, *Jonathon Langford*, *Christophe Ochando (chair)*, and *Nicolas Chanon* — for their careful review, constructive comments, and valuable advice that improved the quality of this analysis.

I would also like to thank the members of the CMS *ECAL* and *EGamma* groups for their work and for providing essential tools, corrections, and recommendations.

A sincere acknowledgment goes to all the people who contributed to the *CMS Run 2 data-taking*, to the *LHC technical teams*, and to all those who computed corrections, efficiencies, and ingredients necessary for the Run 2 CMS $\mathrm{H} \to \gamma\gamma$ analyses. Their collective effort made this work possible.

I am grateful also to *Francesco Pandolfi*, for his guidance, and for bringing me to CMS together with *Ivano Sarra*, who taught me that particle physics is made of real detectors.

Finally, I would like to express my deepest gratitude to my *family* for their support throughout my studies.

The most important thank you goes to *Elisa* for her constant love and encouragement, moral support, and unique and pragmatic suggestions on analysis procedures and during the thesis writing and review.

Appendix A

Appendix

A.1 Signal model fits

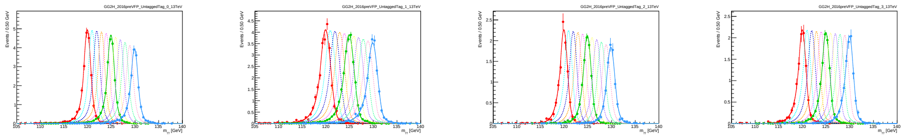

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

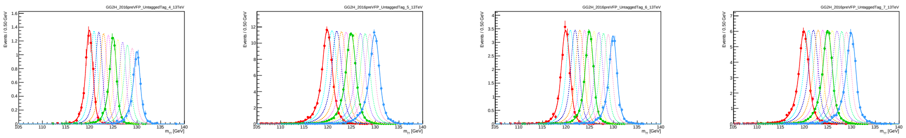

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

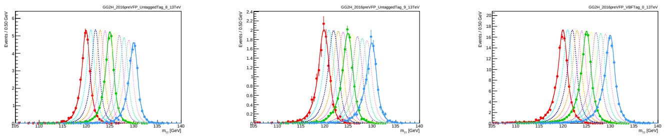

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.1: Signal model fits for GGH process and 2016preVFP categories.

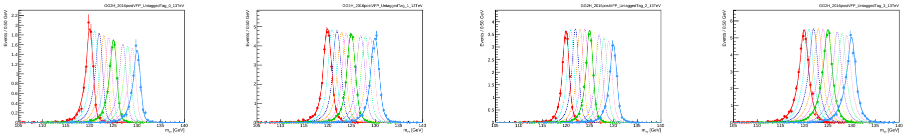

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

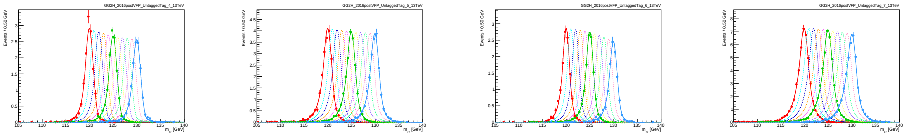

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

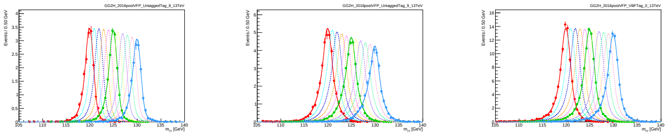

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.2: Signal model fits for GGH process and 2016postVFP categories.

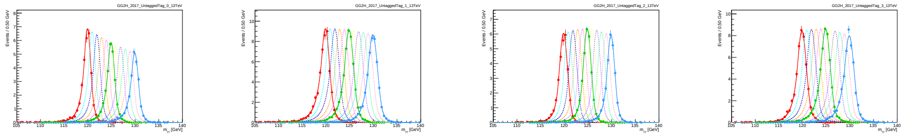

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

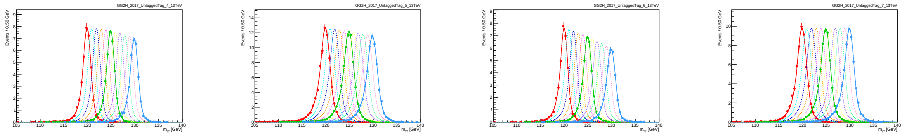

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

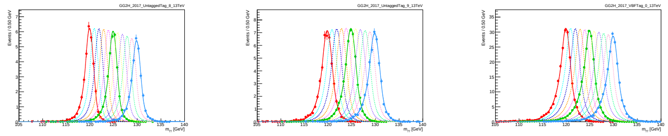

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.3: Signal model fits for GGH process and 2017 categories.

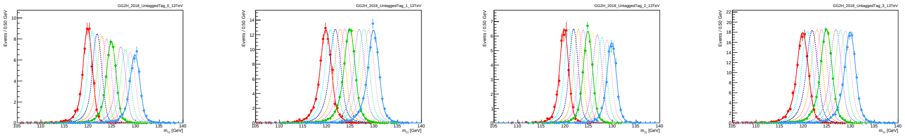

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

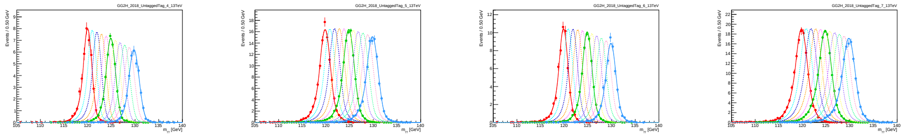

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

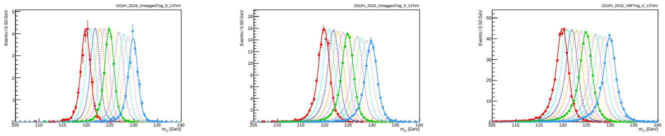

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.4: Signal model fits for GGH process and 2018 categories.

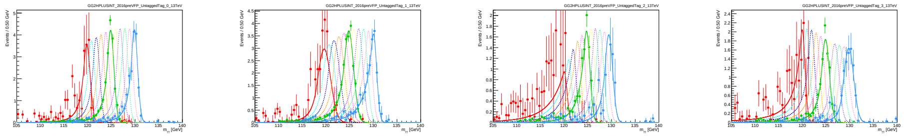

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

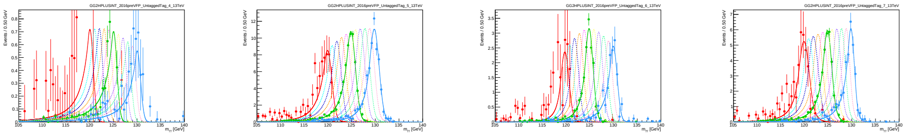

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

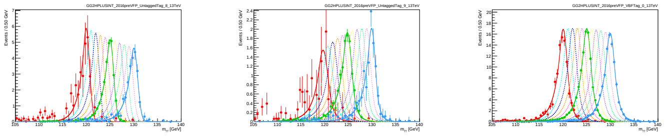

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.5: Signal model fits for GGH+INT. process and 2016preVFP categories.

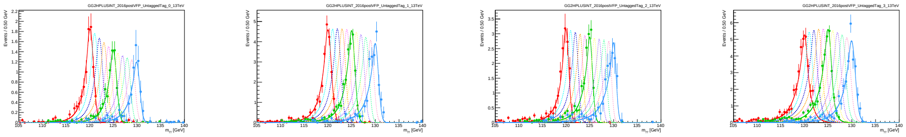

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

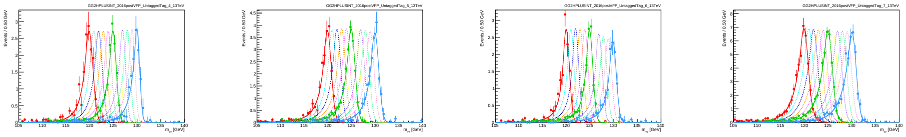

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

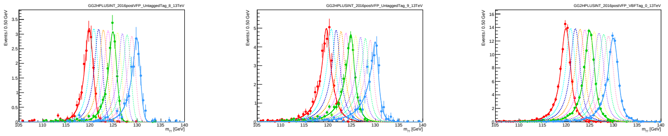

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.6: Signal model fits for GGH+INT. process and 2016postVFP categories.

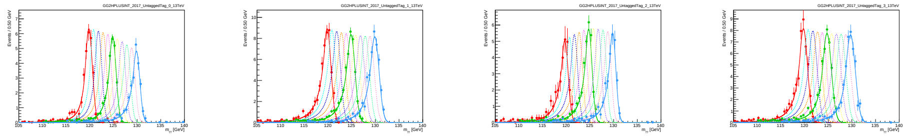

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

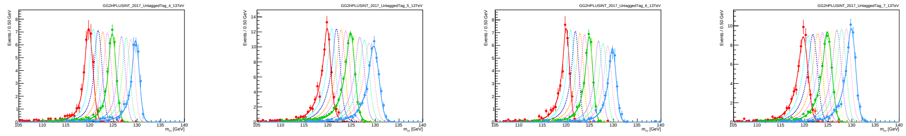

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

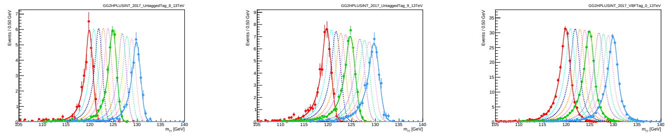

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.7: Signal model fits for GGH+INT. process and 2017 categories.

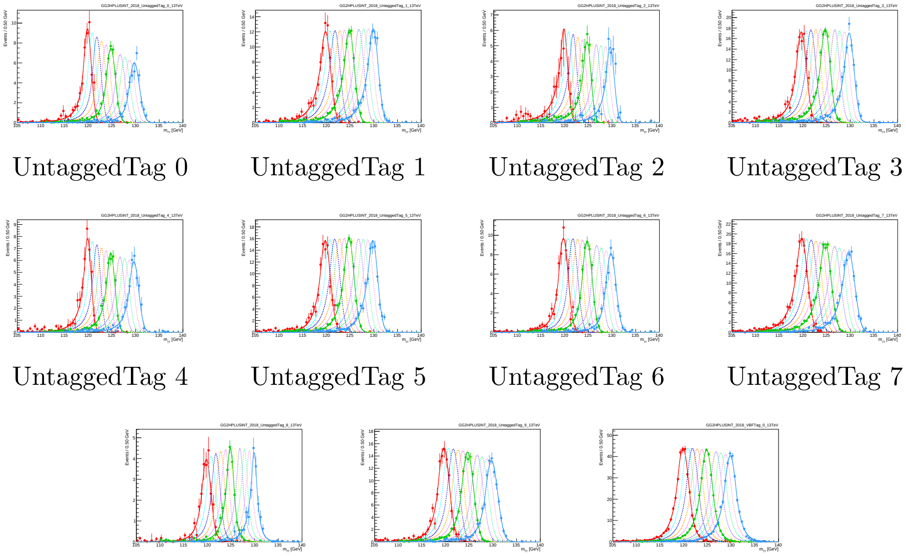

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.8: Signal model fits for GGH+INT. process and 2018 categories.

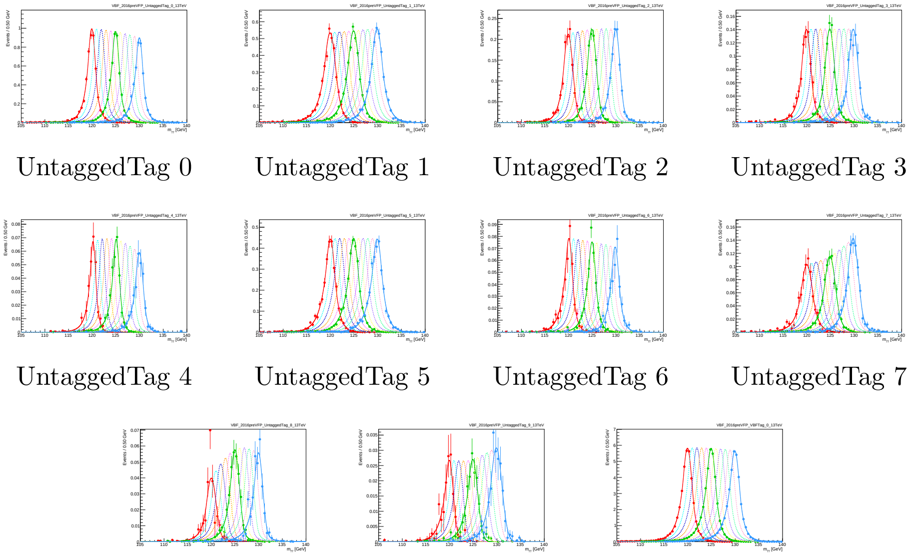

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.9: Signal model fits for VBF process and 2016preVFP categories.

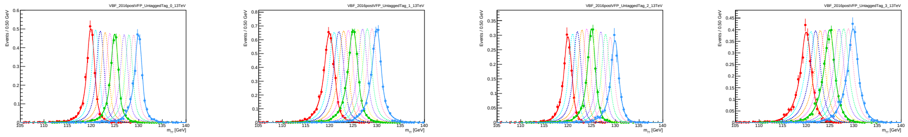

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

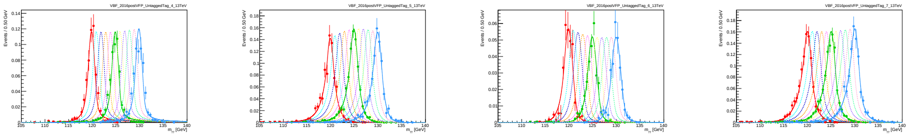

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

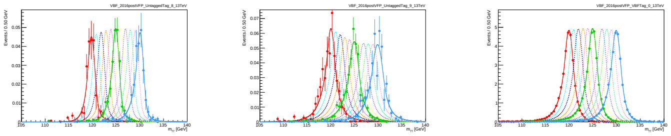

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.10: Signal model fits for VBF process and 2016postVFP categories.

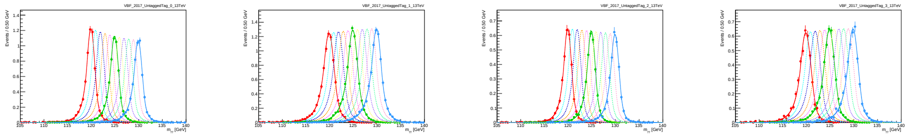

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

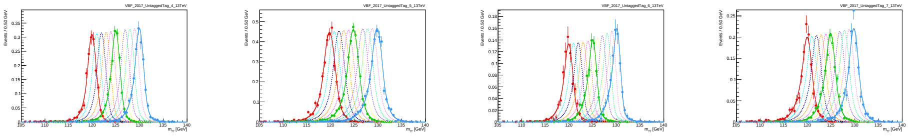

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

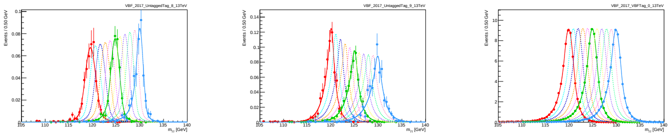

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.11: Signal model fits for VBF process and 2017 categories.

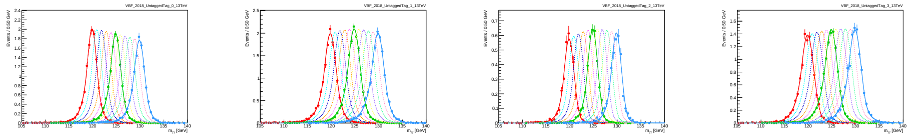

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

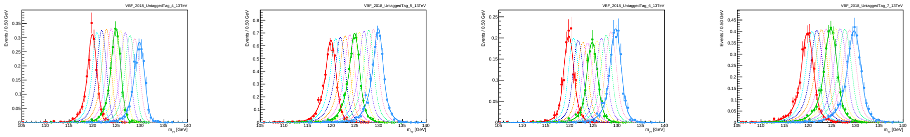

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

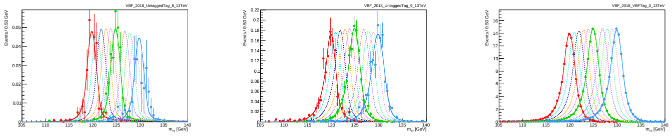

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.12: Signal model fits for VBF process and 2018 categories.

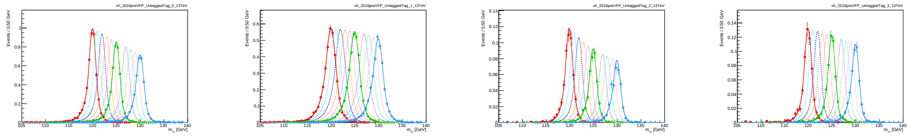

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

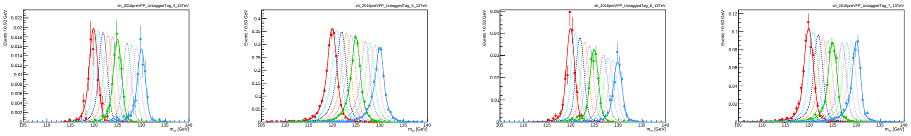

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

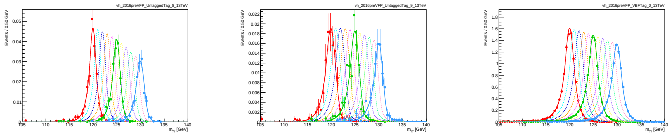

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.13: Signal model fits for VH process and 2016preVFP categories.

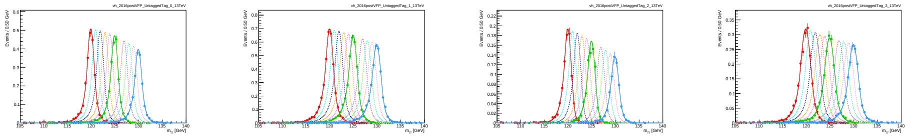

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

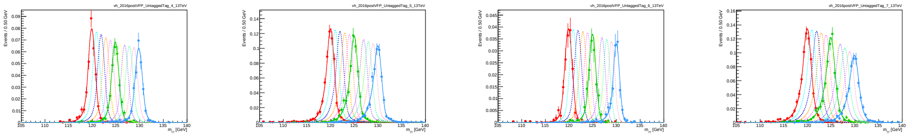

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

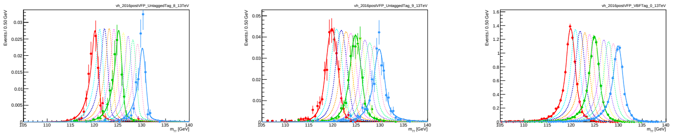

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.14: Signal model fits for VH process and 2016postVFP categories.

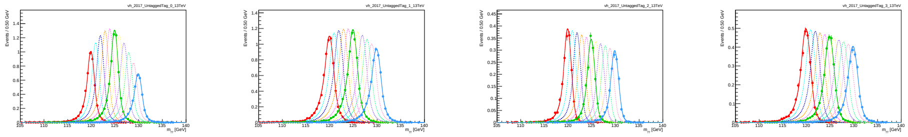

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

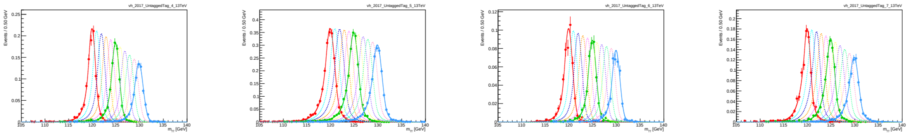

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

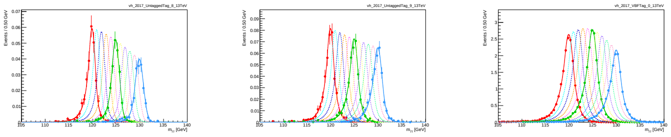

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.15: Signal model fits for VH process and 2017 categories.

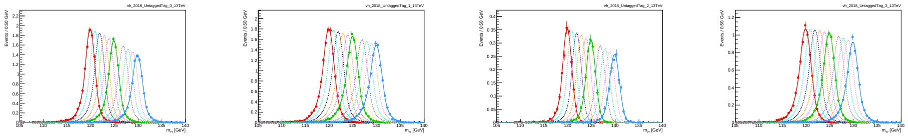

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

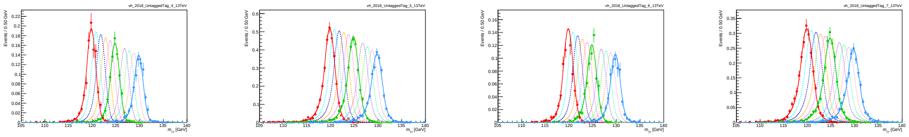

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

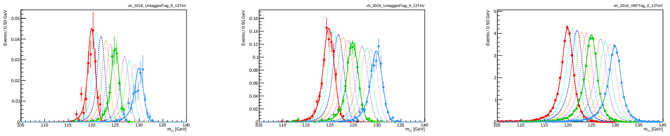

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.16: Signal model fits for VH process and 2018 categories.

A.2 Signal model interpolation functions

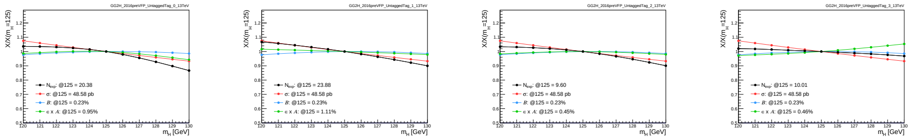

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

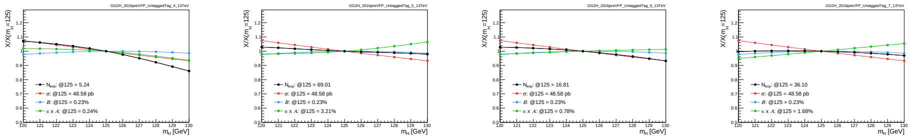

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

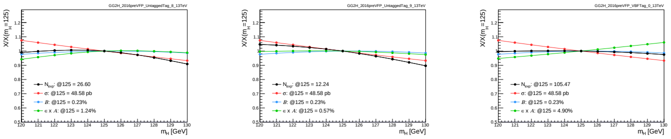

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.17: Signal model interpolation functions for GGH process and 2016preVFP categories.

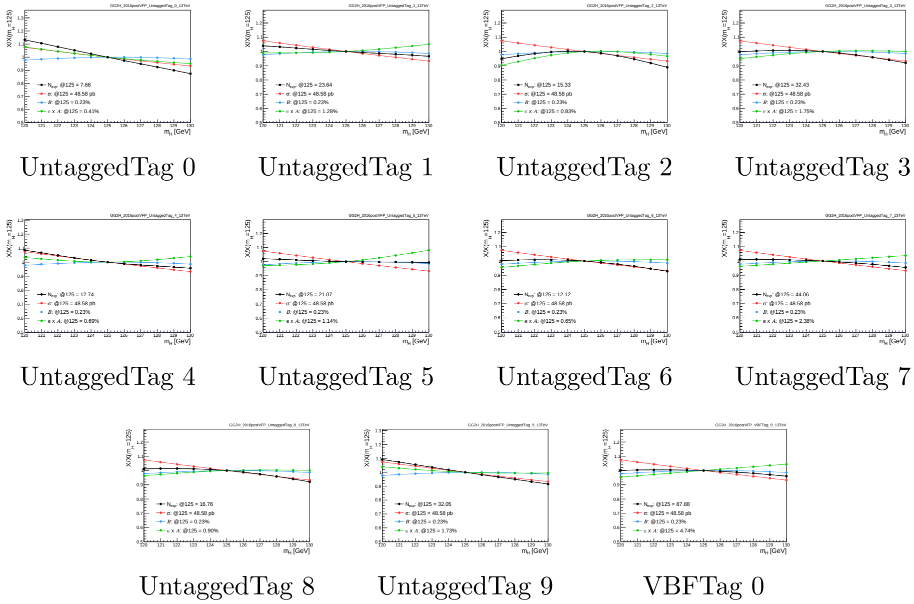

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.18: Signal model interpolation functions for GGH process and 2016postVFP categories.

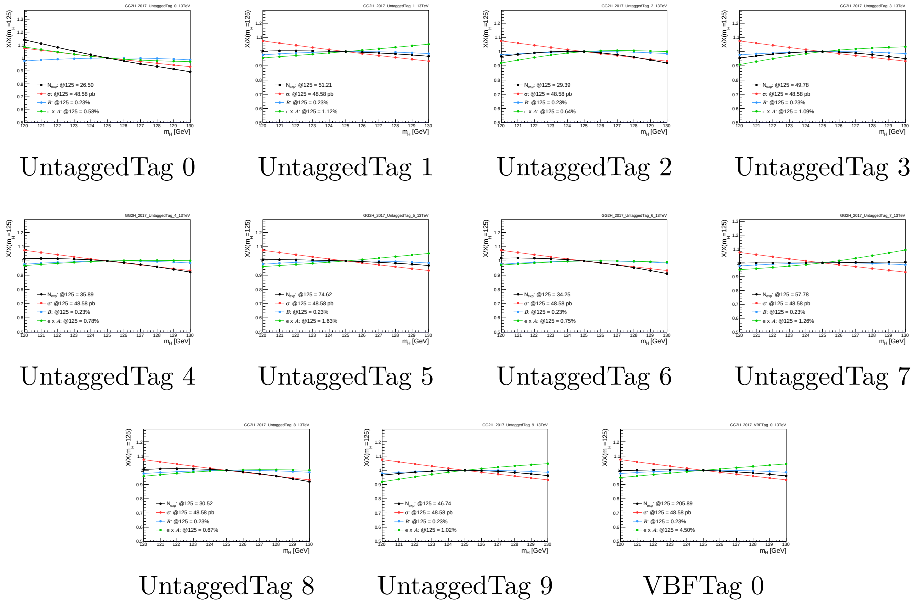

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.19: Signal model interpolation functions for GGH process and 2017 categories.

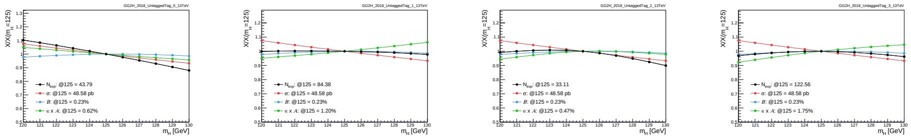

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

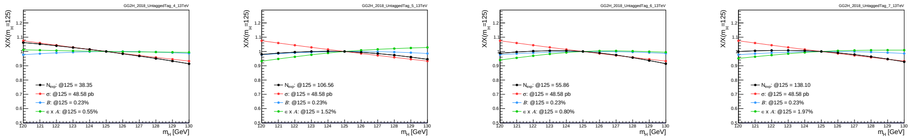

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

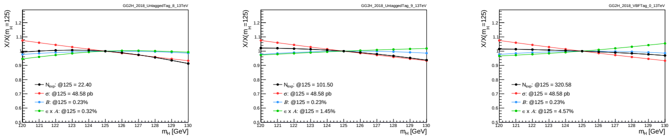

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.20: Signal model interpolation functions for GGH process and 2018 categories.

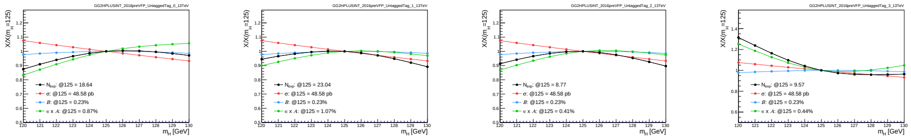

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

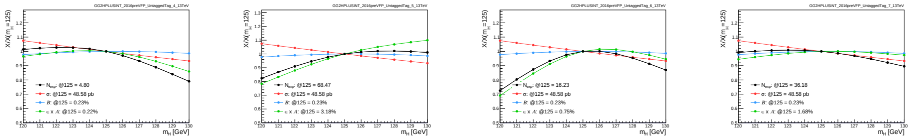

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

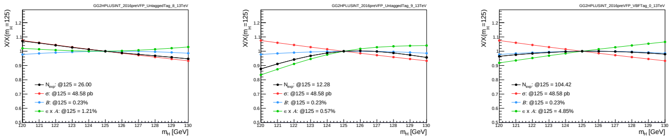

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.21: Signal model interpolation functions for GGH+INT. process and 2016preVFP categories.

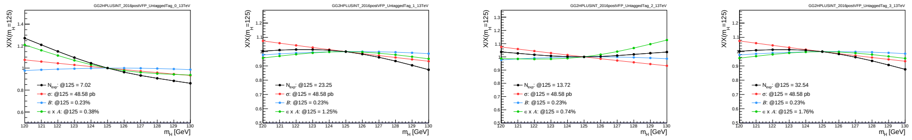

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

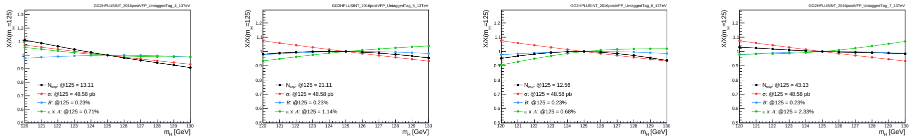

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

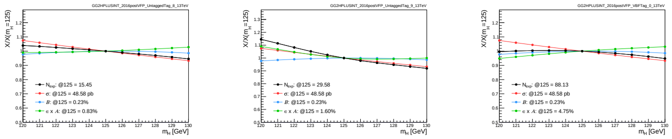

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.22: Signal model interpolation functions for GGH+INT. process and 2016postVFP categories.

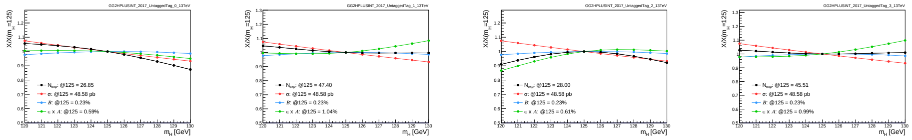

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

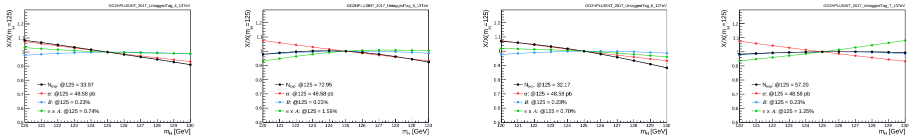

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

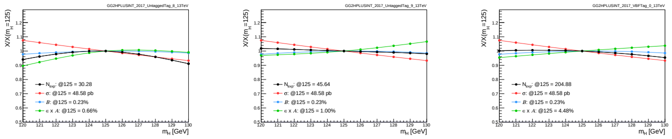

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.23: Signal model interpolation functions for GGH+INT. process and 2017 categories.

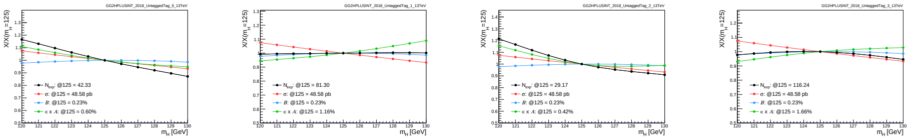

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

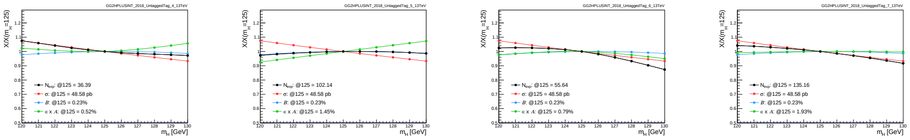

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

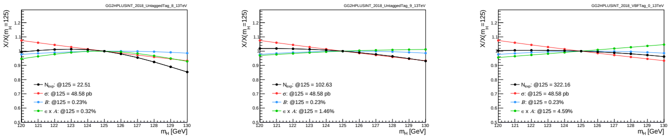

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.24: Signal model interpolation functions for GGH+INT. process and 2018 categories.

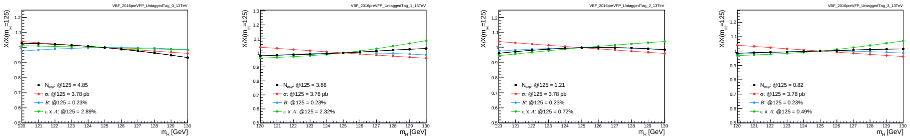

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

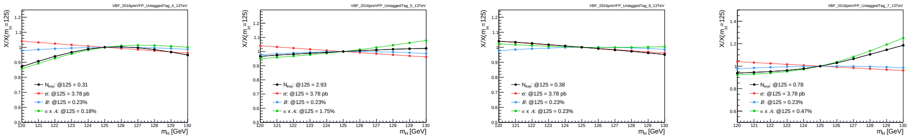

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

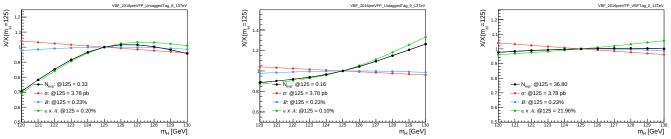

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.25: Signal model interpolation functions for VBF process and 2016preVFP categories.

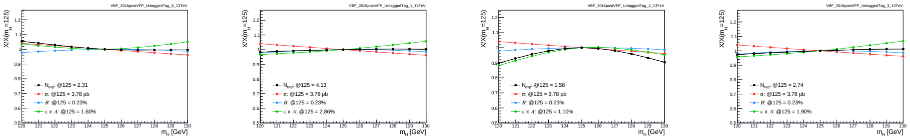

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

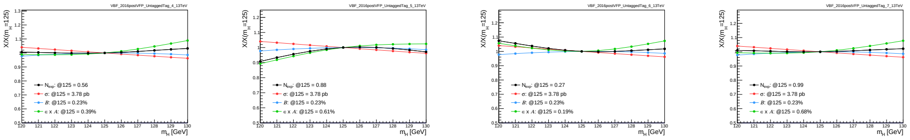

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

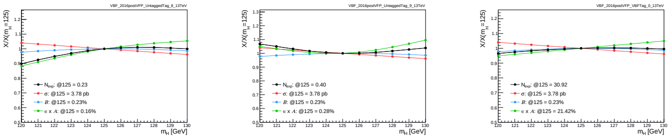

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.26: Signal model interpolation functions for VBF process and 2016postVFP categories.

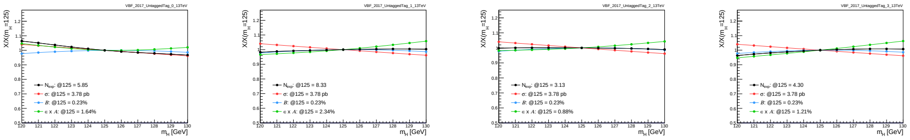

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

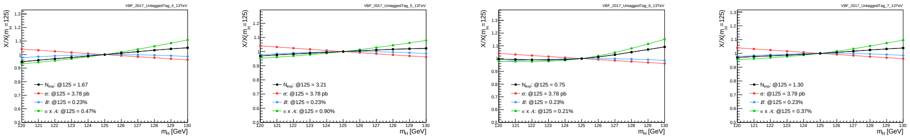

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

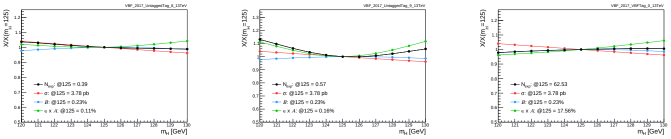

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.27: Signal model interpolation functions for VBF process and 2017 categories.

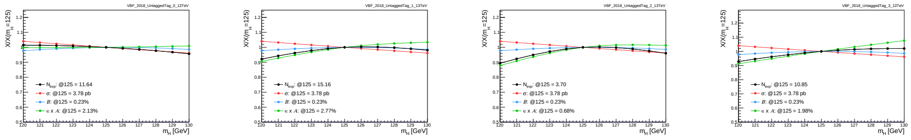

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

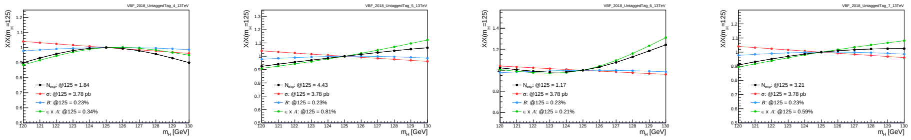

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

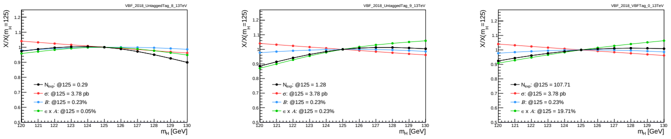

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.28: Signal model interpolation functions for VBF process and 2018 categories.

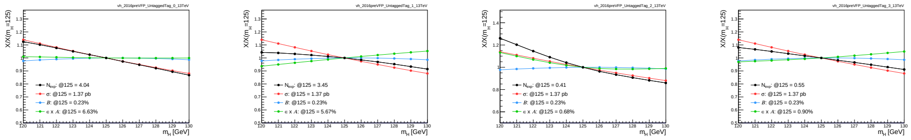

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

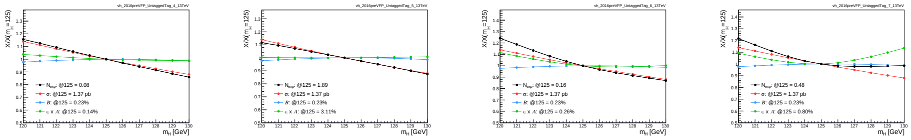

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

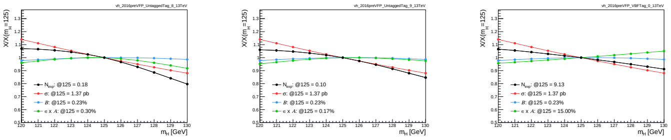

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.29: Signal model interpolation functions for VH process and 2016preVFP categories.

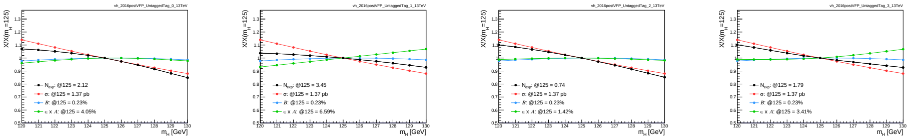

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

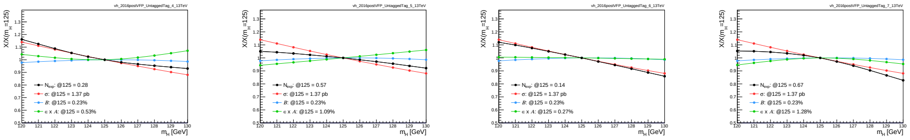

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

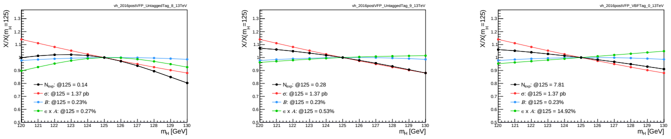

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.30: Signal model interpolation functions for VH process and 2016postVFP categories.

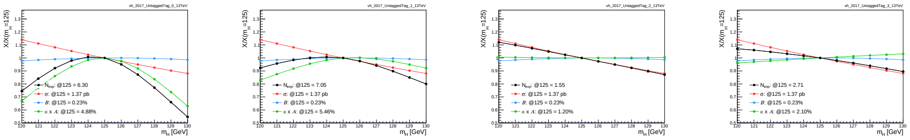

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

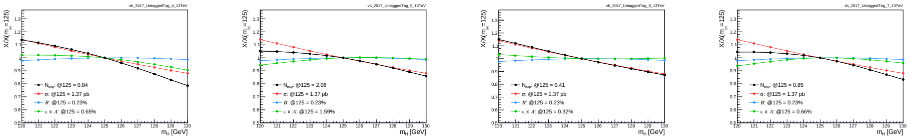

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

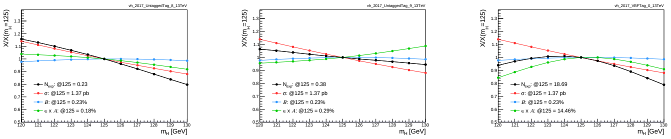

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.31: Signal model interpolation functions for VH process and 2017 categories.

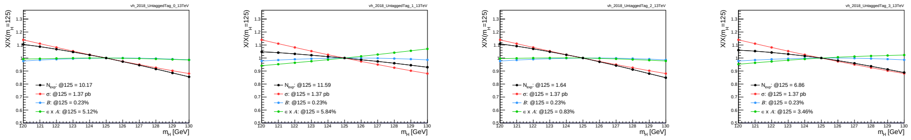

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

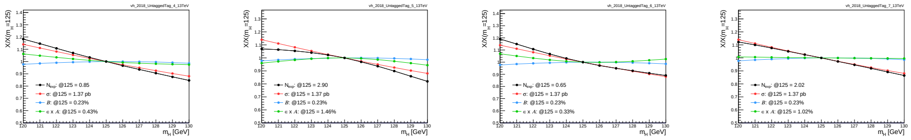

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

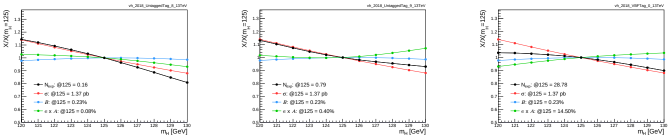

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.32: Signal model interpolation functions for VH process and 2018 categories.

A.3 Signal model summary plots

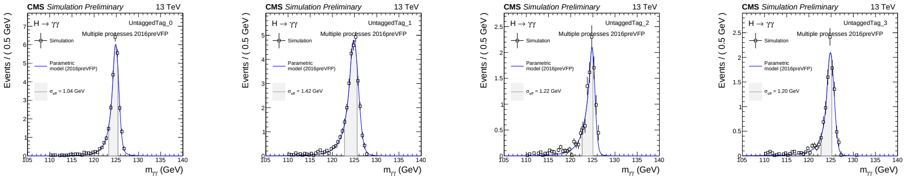

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

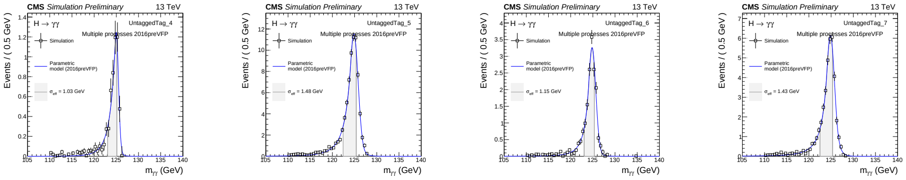

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

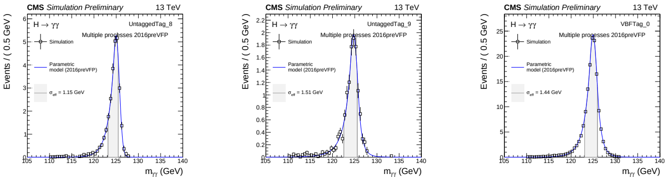

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.33: Signal models including all processes for 2016preVFP categories.

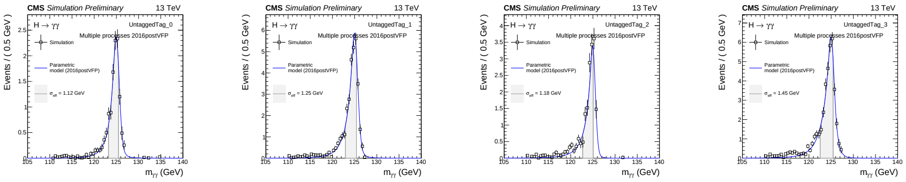

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

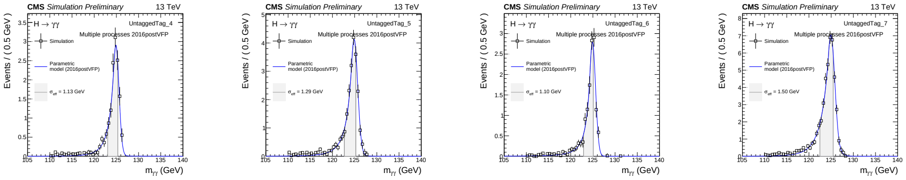

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

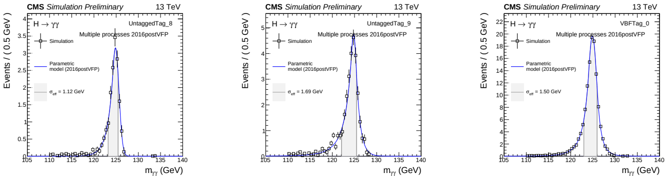

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.34: Signal models including all processes for 2016postVFP categories.

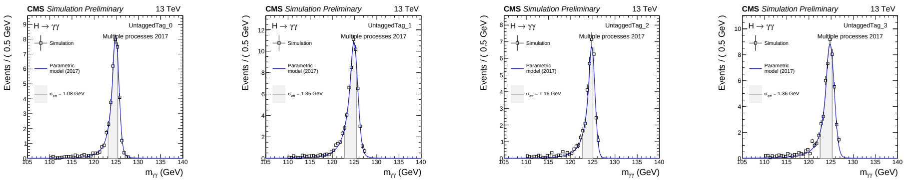

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

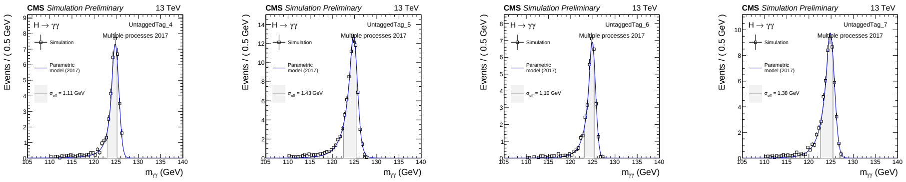

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

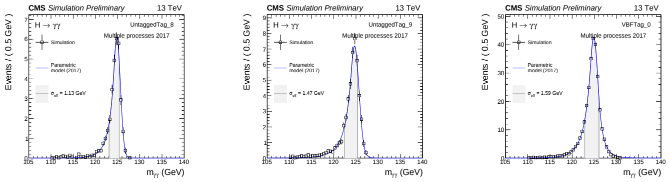

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.35: Signal models including all processes for 2017 categories.

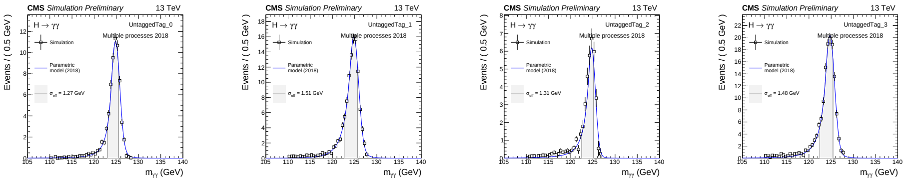

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

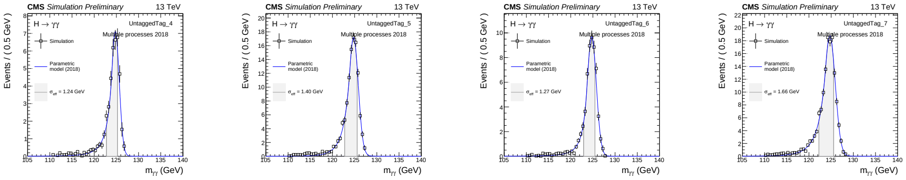

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

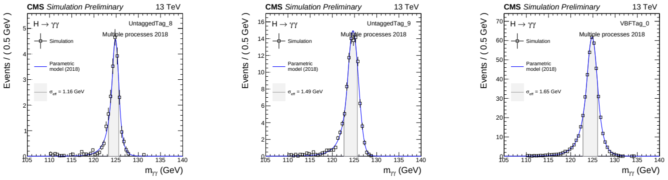

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.36: Signal models including all processes for 2018 categories.

A.4 Background model fits

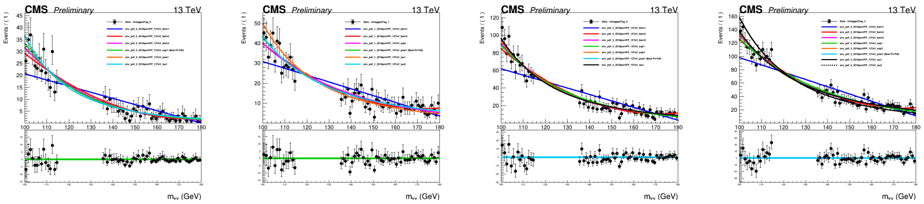

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

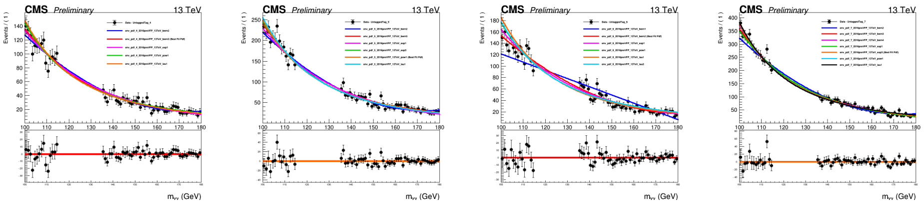

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

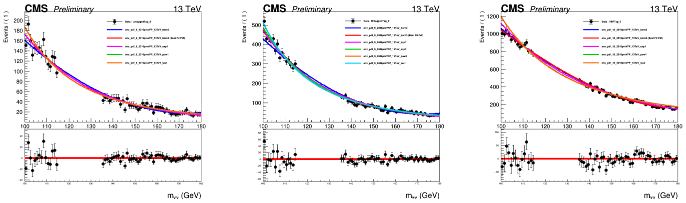

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.37: Fit to the blinded data to determine the collection of functions used to fit the background for 2016preVFP categories.

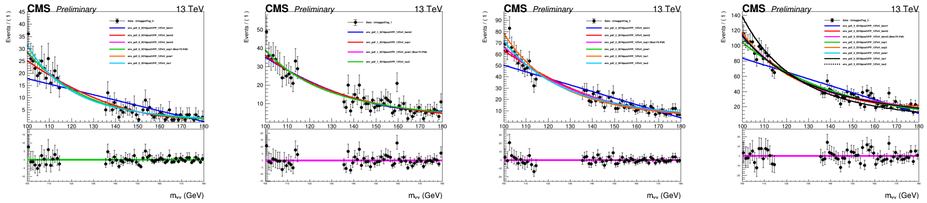

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

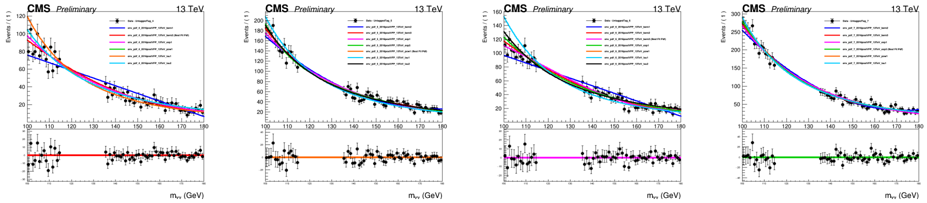

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

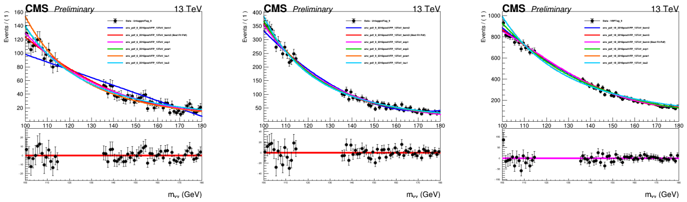

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.38: Fit to the blinded data to determine the collection of functions used to fit the background for 2016postVFP categories.

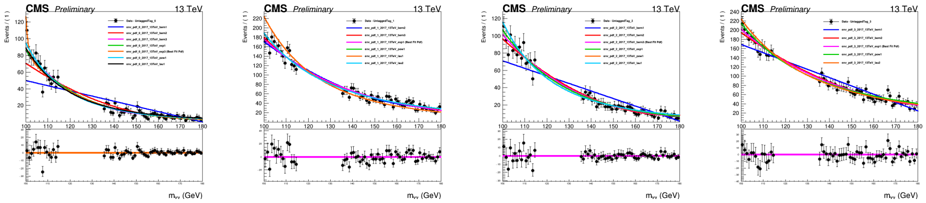

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

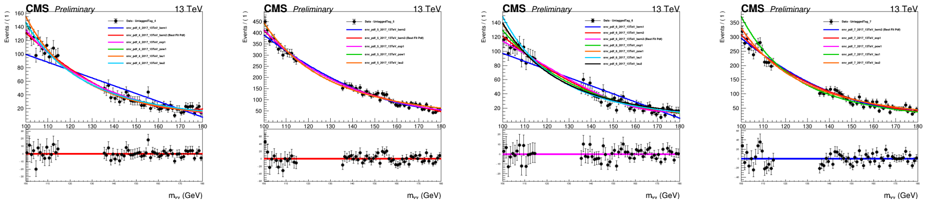

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

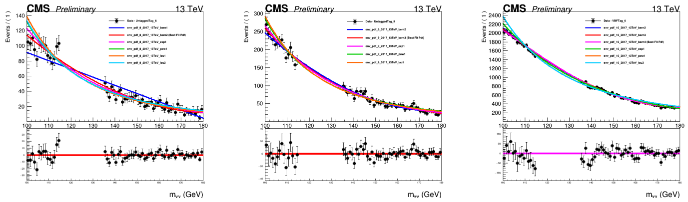

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.39: Fit to the blinded data to determine the collection of functions used to fit the background for 2017 categories.

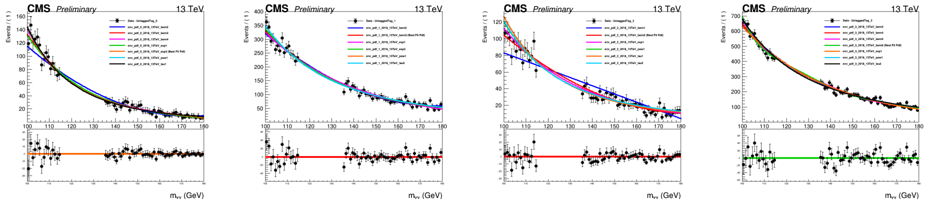

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

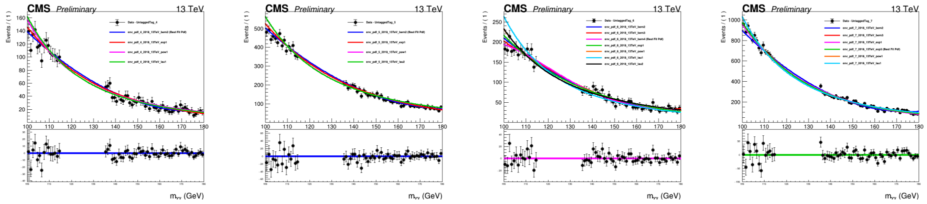

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

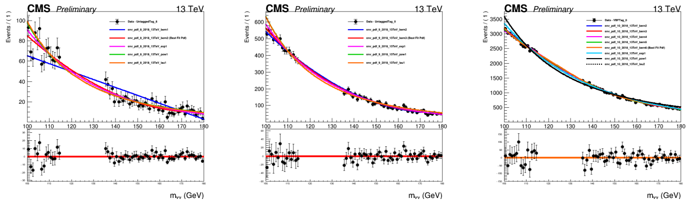

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.40: Fit to the blinded data to determine the collection of functions used to fit the background for 2018 categories.

A.5 Impacts

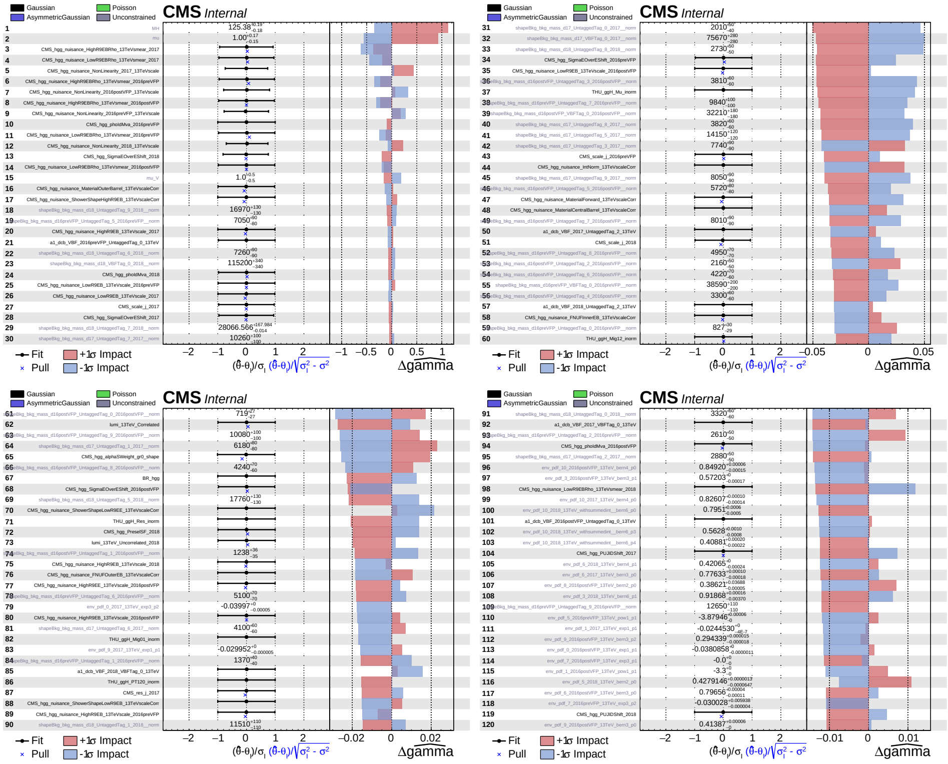

Figure A.41: Full expected impacts for Γ_H.

Expected impacts on Γ_H

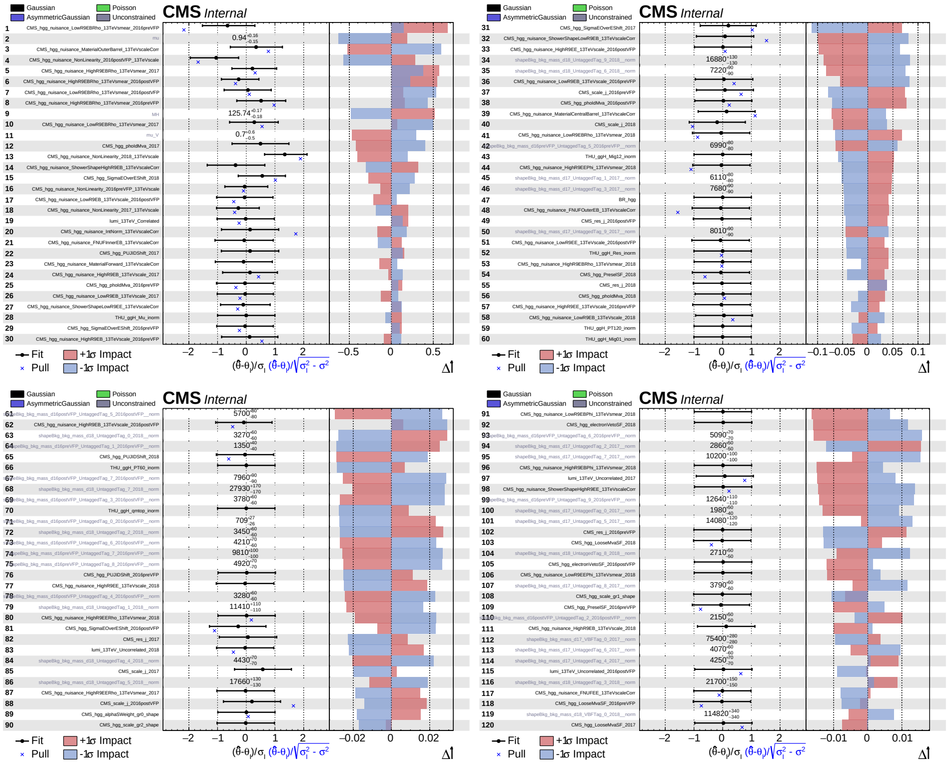

Figure A.42: Full observed impacts for λ.

Observed impacts on λ

A.6 Background-only fits

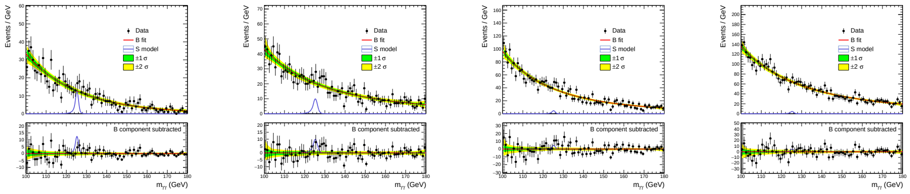

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

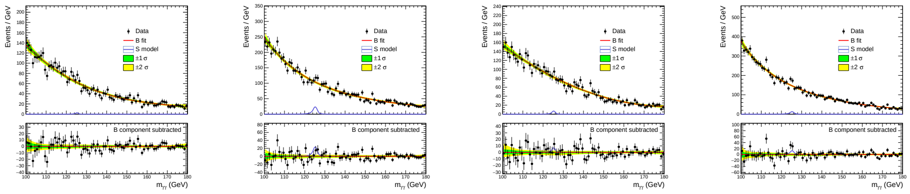

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

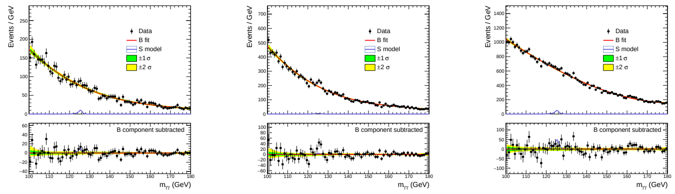

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.43: Background-only fits for 2016preVFP categories.

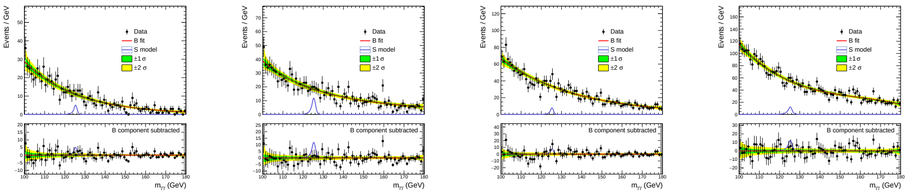

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

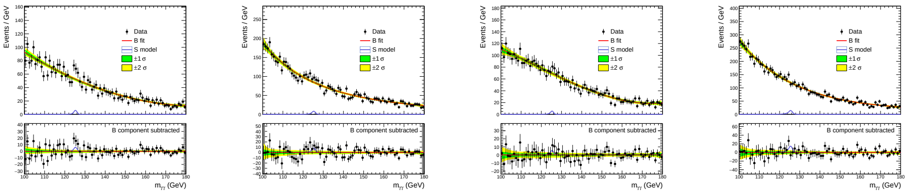

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

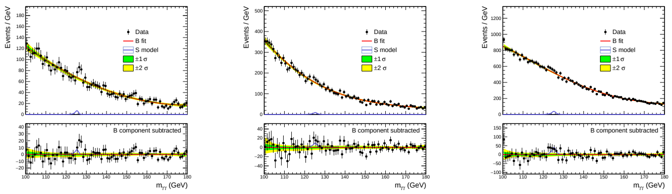

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.44: Background-only fits for 2016postVFP categories.

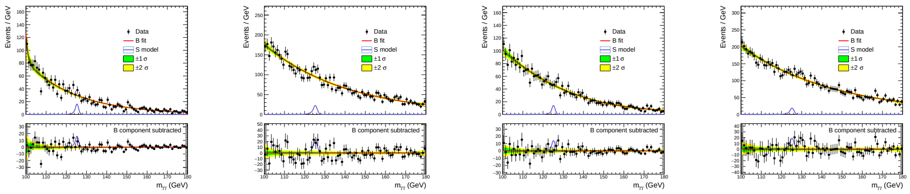

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

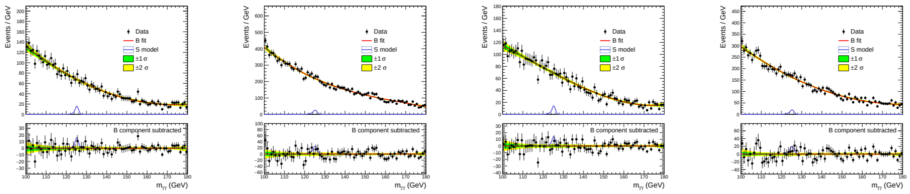

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

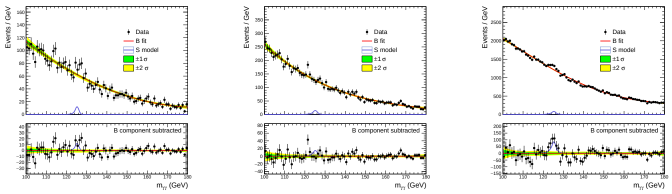

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.45: Background-only fits for 2017 categories.

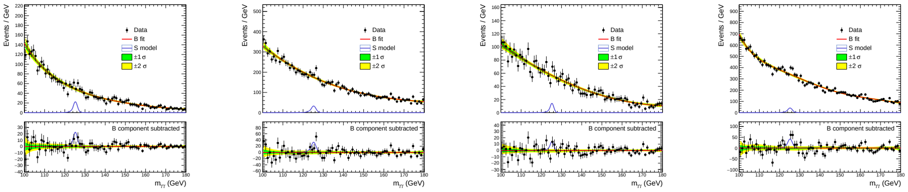

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

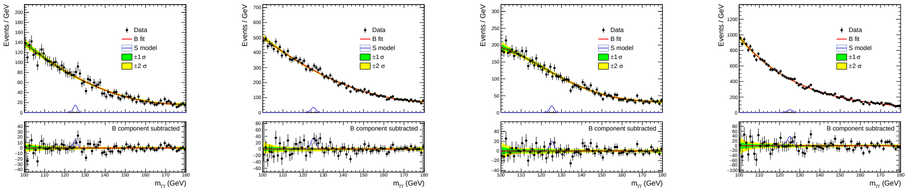

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

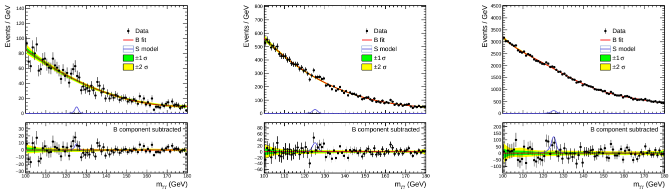

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.46: Background-only fits for 2018 categories.

A.7 Signal+background fits

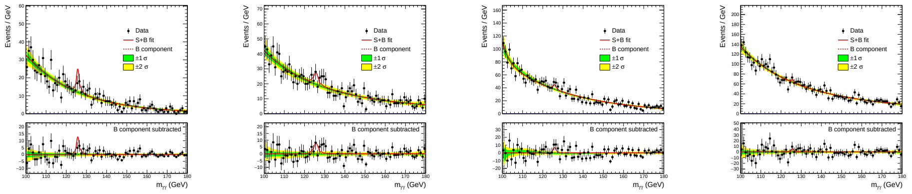

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

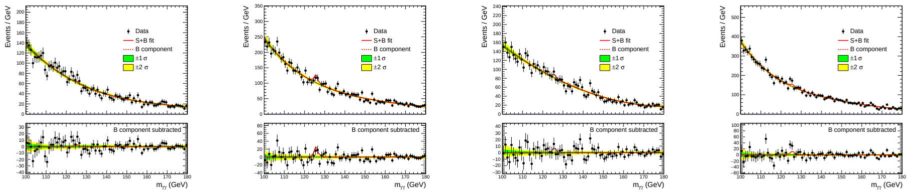

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

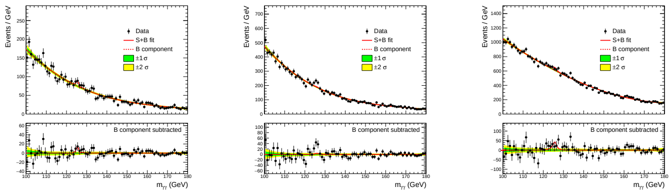

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.47: Signal + background fits for 2016preVFP categories.

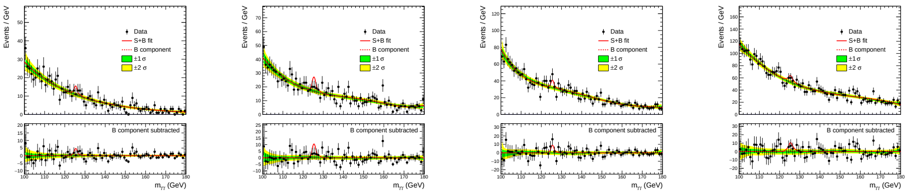

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

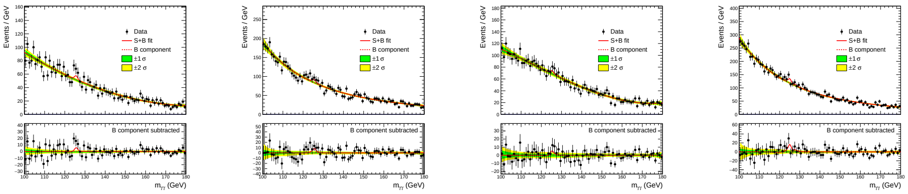

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

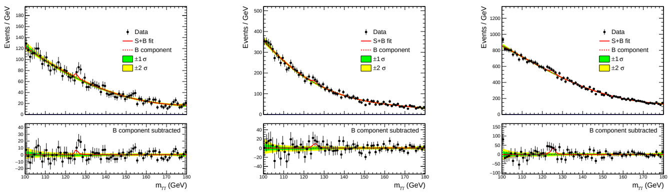

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.48: Signal + background fits for 2016postVFP categories.

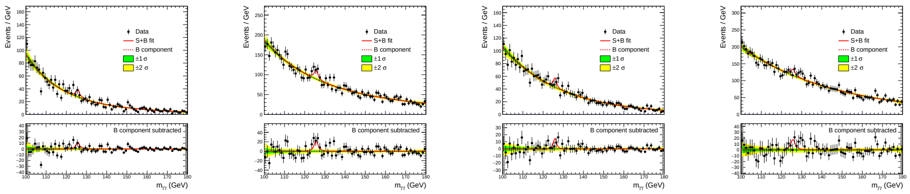

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

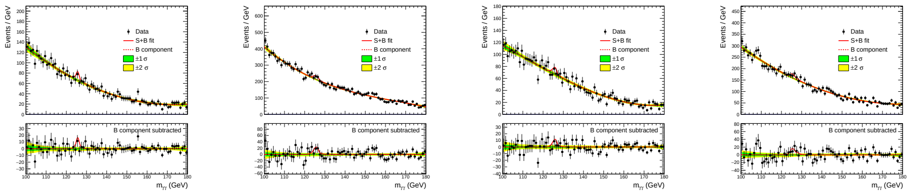

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

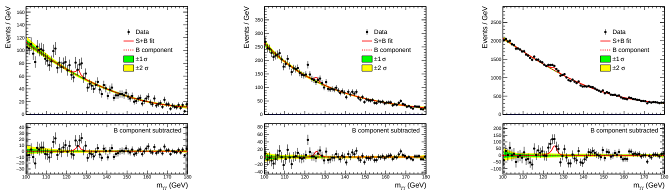

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.49: Signal + background fits for 2017 categories.

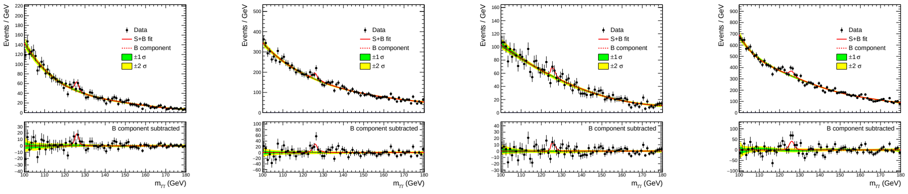

UntaggedTag 0 UntaggedTag 1 UntaggedTag 2 UntaggedTag 3

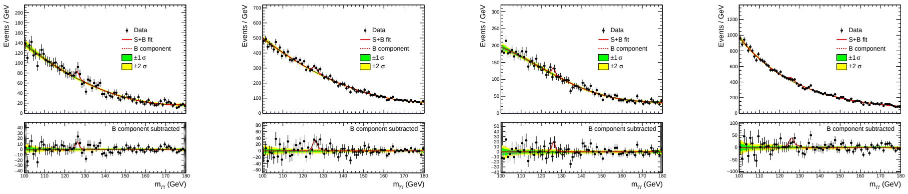

UntaggedTag 4 UntaggedTag 5 UntaggedTag 6 UntaggedTag 7

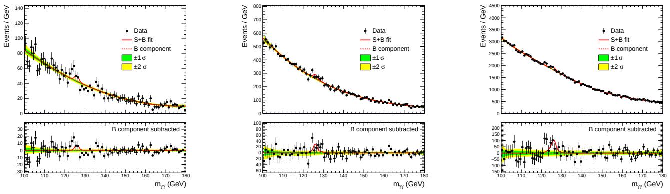

UntaggedTag 8 UntaggedTag 9 VBFTag 0

Figure A.50: Signal + background fits for 2018 categories.

Bibliography

[1] Constraints on the Higgs boson total decay width using signal-background interference in the diphoton final state with proton-proton collisions at sqrts = 13 TeV. Technical Report CMS-PAS-HIG-25-004, CERN, Geneva, 2025. URL: `https://cds.cern.ch/record/2941115`.

[2] Measurement of the inclusive 3-jet production differential cross section in proton–proton collisions at 7 TeV and determination of the strong coupling constant in the TeV range. *The European Physical Journal C*, 75(5). URL: `http://dx.doi.org/10.1140/epjc/s10052-015-3376-y`, `doi:10.1140/epjc/s10052-015-3376-y`.

[3] Andreas Frey and Nguyen Kien. The Higgs Mechanism, 2009. Lecture notes, LMU Munich, July 2009. URL: `https://www.theorie.physik.uni-muenchen.de/lsfrey/teaching/archiv/sose_09/rng/higgs_mechanism.pdf`.

[4] R. L. Workman and Others. Review of Particle Physics. *PTEP*, 2022:083C01, 2022. `doi:10.1093/ptep/ptac097`.

[5] A. Barczyk, J. Kirkby, L. Malgeri, J. Berdugo, J. Casaus, C. Mañá, J. Marín, G. Martínez, E. Sánchez, C. Willmott, C. Casella, M. Pohl, K. Deiters, P. Dick, and C. Petitjean. Measurement of the Fermi constant by FAST. *Physics Letters B*, 663(3):172–180, May 2008. URL: `http://dx.doi.org/10.1016/j.physletb.2008.04.006`, `doi:10.1016/j.physletb.2008.04.006`.

[6] ATLAS collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1–29, 2012. `doi:https://doi.org/10.1016/j.physletb.2012.08.020`.

[7] CMS collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, 2012. URL: `https://www.sciencedirect.com/science/article/pii/S0370269312008581`.

[8] A portrait of the Higgs boson by the CMS experiment ten years after the discovery. *Nature*, 607(7917):60–68, Jul 2022. `doi:10.1038/s41586-022-04892-x`.

[9] Characterising the Higgs boson with ATLAS data from the LHC Run-2. *Physics Reports*, 1116:4–56, 2025. Breaking boundaries — ATLAS physics highlights and milestones from the LHC Run 2. URL: `https://www.sciencedirect.com/science/article/pii/S037015732400382X`, `doi:https://doi.org/10.1016/j.physrep.2024.11.001`.

[10] Combined measurement of the higgs boson mass from the and $h \to \gamma\gamma$ and $h \to zz^* \to 4l$ decay channels with the atlas detector using $\sqrt{s}$ =7, 8, and 13 tev pp collision data. *Physical Review Letters*, 131(25), December 2023. URL: `http://dx.doi.org/10.1103/PhysRevLett.131.251802`, `doi:10.1103/physrevlett.131.251802`.

[11] Measurement of the Higgs boson mass and width using the four leptons final state. Technical Report CMS-PAS-HIG-21-019, CERN, Geneva, 2023. URL: `https://cds.cern.ch/record/2871702`.

[12] A measurement of the Higgs boson mass in the diphoton decay channel. *Physics Letters B*, 805:135425, 2020. `arXiv:2002.06398`, `doi:10.1016/j.physletb.2020.135425`.

[13] Evidence for Higgs boson decay to a pair of muons. *Journal of High Energy Physics*, 2021(1). URL: `http://dx.doi.org/10.1007/JHEP01(2021)148`, `doi:10.1007/jhep01(2021)148`.

[14] ATLAS Collaboration. Evidence for the dimuon decay of the Higgs boson in pp collisions with the ATLAS detector, 2025. URL: `https://arxiv.org/abs/2507.03595`, `arXiv:2507.03595`.

[15] Search for the Higgs boson decay to a pair of electrons in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Lett. B*, 846(CMS-HIG-21-015, CERN-EP-2022-131, CMS-HIG-21-015-003):137783, 2023. URL:

`https://cds.cern.ch/record/2823465`, `arXiv:2208.00265`, `doi:10.1016/j.physletb.2023.137783`.

[16] A. Sopczak. Prospects for higgs boson research at the lhc. *Acta Physica Polonica B*, 56(9):1, September 2025. URL: `http://dx.doi.org/10.5506/APhysPolB.56.9-A4`, `doi:10.5506/aphyspolb.56.9-a4`.

[17] Michael E. Peskin. What is the hierarchy problem?, 2025. URL: `https://arxiv.org/abs/2505.00694`, `arXiv:2505.00694`.

[18] Bernhard Mistlberger. Higgs boson production at hadron colliders at n3lo in qcd. *Journal of High Energy Physics*, 2018(5), May 2018. URL: `http://dx.doi.org/10.1007/JHEP05(2018)028`, `doi:10.1007/jhep05(2018)028`.

[19] CERN. CERN Yellow Reports: Monographs, Vol 2 (2017): Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector, 2017. `doi:10.23731/CYRM-2017-002`.

[20] Search for the standard model Higgs boson produced through vector boson fusion and decaying to $b\bar{b}$. *Phys. Rev. D*, 92:032008, Aug 2015. URL: `https://link.aps.org/doi/10.1103/PhysRevD.92.032008`, `doi:10.1103/PhysRevD.92.032008`.

[21] CMS collaboration. Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D*, 105:092007, May 2022. URL: `https://link.aps.org/doi/10.1103/PhysRevD.105.092007`, `doi:10.1103/PhysRevD.105.092007`.

[22] CMS collaboration. Observation of Higgs boson decay to bottom quarks. *Phys. Rev. Lett.*, 121:121801, Sep 2018. `doi:10.1103/PhysRevLett.121.121801`.

[23] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, and R. Tanaka (Eds.). Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables. *CERN-2011-002*, CERN, Geneva, 2011. `arXiv:1101.0593`, `doi:10.5170/CERN-2011-002`.

[24] Measurement of the Higgs boson inclusive and differential fiducial production cross sections in the diphoton decay channel with pp collisions at $\sqrt{s} = 13$ TeV. *Journal of High Energy Physics*, 2023(7). URL: `http://dx.doi.org/10.1007/JHEP07(2023)091`, `doi:10.1007/jhep07(2023)091`.

[25] Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC. *Journal of Instrumentation*, 16(05). URL: `http://dx.doi.org/10.1088/1748-0221/16/05/P05014`, `doi:10.1088/1748-0221/16/05/p05014`.

[26] Electron and photon energy calibration with the ATLAS detector using LHC Run 2 data. *Journal of Instrumentation*, 19(02), 2024. URL: `http://dx.doi.org/10.1088/1748-0221/19/02/P02009`, `doi:10.1088/1748-0221/19/02/p02009`.

[27] Measurements of differential cross-sections in four-lepton events in 13 tev proton-proton collisions with the atlas detector. *Journal of High Energy Physics*, 2021(7), 2021. URL: `http://dx.doi.org/10.1007/JHEP07(2021)005`, `doi:10.1007/jhep07(2021)005`.

[28] CMS Collaboration. Measurement of the higgs boson mass and width using the four-lepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D*, 111:092014, May 2025. `doi:10.1103/PhysRevD.111.092014`.

[29] CMS collaboration. Measurements of the Higgs boson production cross section and couplings in the W boson pair decay channel in proton-proton collisions at $\sqrt{s} = 13$ tev. *Eur. Phys. J. C*, 83:667, 2023. `doi:10.1140/epjc/s10052-023-11632-6`.

[30] CMS collaboration. Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector. *Physics Letters B*, 779:283–316, 2018. `doi:https://doi.org/10.1016/j.physletb.2018.02.004`.

[31] Andrey Elagin, Pavel Murat, Aliaksandr Pranko, and Alexei Safonov. A New Mass Reconstruction Technique for Resonances Decaying to di-tau. *Nucl. Instrum. Meth. A*, 654:481–489, 2011. `arXiv:1012.4686`, `doi:10.1016/j.nima.2011.07.009`.

[32] S. Chatrchyan et al. Evidence for the 125 GeV Higgs boson decaying to a pair of τ leptons. *JHEP*, 05:104, 2014. `arXiv:1401.5041`, `doi:10.1007/JHEP05(2014)104`.

[33] L. Bianchini, J.S. Conway, A. Fowlie, and J. Sjoelin. Reconstruction of the Higgs mass in events with Higgs bosons decaying into a pair of τ leptons using matrix element techniques. *J. Phys. Conf. Ser.*, 513:022035, 2014. `doi:10.1088/1742-6596/513/2/022035`.

[34] ATLAS collaboration. Test of CP invariance in vector-boson fusion production of the Higgs boson in the $h \to \tau\tau$ channel in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Physics Letters B*, 805:135426, 2020. `doi:https://doi.org/10.1016/j.physletb.2020.135426`.

[35] CMS Collaboration. Analysis of the CP structure of the yukawa coupling between the Higgs boson and τ leptons in proton-proton collisions at $\sqrt{s} = 13$ tev. *J. High Energ. Phys.*, 2022(6):12, 2022. `doi:10.1007/JHEP06(2022)012`.

[36] Observation of $h \to bb$ decays and vh production with the atlas detector.

[37] Fe Taylor, Aram Apyan, Richard Barbieri, Austin Baty, Katharina Bierwagen, Stephanie Brandt, Wit Busza, Ivan Cali, Leonardo Matteo, Guillelmo Gomez-Ceballos, Maxim Goncharov, Doga Gulhan, Gian Innocenti, Dmytro Kovalskyi, Yue Shi Lai, Yen-Jie Lee, Andrew Levin, Luckey Jr, and Victoria Zhukova. Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments. 03 2015.

[38] The ATLAS collaboration. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 tev. *Journal of High Energy Physics*, 2016(8):45, Aug 2016. `doi:10.1007/JHEP08(2016)045`.

[39] The ATLAS Collaboration. A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery. *Nature*, 607(7917):52–59, Jul 2022. `doi:10.1038/s41586-022-04893-w`.

[40] S Heinemeyer, C Mariotti, G Passarino, R Tanaka, J R Andersen, P Artoisenet, E A Bagnaschi, A Banfi, T Becher, F U Bernlochner, S Bolognesi, P Bolzoni, R Boughezal, D Buarque, J Campbell, F Caola, M Carena, F Cascioli, N Chanon, T Cheng, S Y Choi, A David, P de Aquino, G Degrassi, D Del Re, A Denner, H van Deurzen, S Diglio, B Di Micco, R Di Nardo, S Dittmaier, M Dührssen, R K Ellis, G Ferrera, N Fidanza, M Flechl, D de Florian, S Forte, R Frederix, S Frixione, S Gangal, Y Gao, M V Garzelli, D Gillberg, P Govoni, M Grazzini, N Greiner, J Griffiths, A V Gritsan, C Grojean, D C Hall, C Hays, R Harlander, R Hernandez-Pinto, S Höche, J Huston, T Jubb, M Kadastik, S Kallweit, A Kardos, L Kashif, N Kauer, H Kim, R Klees, M Krämer, F Krauss, A Laureys, S Laurila, S Lehti, Q Li, S Liebler, X Liu, E Logan, G Luisoni, M Malberti, F Maltoni, K Mawatari, F Maierhoefer, H Mantler, S Martin, P Mastrolia, O Mattelaer, J Mazzitelli, B Mellado, K Melnikov, P Meridiani, D J Miller, E Mirabella, S O Moch, P Monni, N Moretti, A Mück, M Mühlleitner, P Musella, P Nason, C Neu, M Neubert, C Oleari, J Olsen, G Ossola, T Peraro, K Peters, F Petriello, G Piacquadio, C T Potter, S Pozzorini, K Prokofiev, I Puljak, M Rauch, D Rebuzzi, L Reina, R Rietkerk, A Rizzi, Y Rotstein-Habarnau, G P Salam, G Sborlini, F Schissler, M Schönherr, M Schulze, M Schumacher, F Siegert, P Slavich, J M Smillie, O Stål, J F von Soden-Fraunhofen, M Spira, I W Stewart, F J Tackmann, P T E Taylor, D Tommasini, J Thompson, R S Thorne, P Torrielli, F Tramontano, N V Tran, Z Trócsányi, M Ubiali, M Vazquez Acosta, T Vickey, A Vicini, W J Waalewijn, D Wackeroth, C Wagner, J R Walsh, J Wang, G Weiglein, A Whitbeck, C Williams, J Yu, G Zanderighi, M Zanetti, M Zaro, P M Zerwas, C Zhang, T J E Zirke, and S Zuberi. *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties.* CERN Yellow Reports: Monographs. 2013. Comments: 404 pages, 139 figures, to be submitted to CERN Report. Working Group web page: https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CrossSections. URL: `https://cds.cern.ch/record/1559921`, `doi:10.5170/CERN-2013-004`.

[41] Aleksandr Azatov, Adam Falkowski, Andrei V. Gritsan, Christophe Grojean, Nikolas Kauer, Ennio Salvioni, Marion Thomas, Eleni Vryonidou, Lucas Kang, Jorge de Blas, and Ulascan Sarica. Off-shell Higgs In-

terpretations Task Force : Models and Effective Field Theories Subgroup Report. Technical Report LHCHWG-2022-001, CERN, Geneva, 2022. URL: `http://cds.cern.ch/record/2801789`, `doi:10.17181/LHCHWG-2022-001`.

[42] CMS collaboration. Combined measurements and interpretations of Higgs boson production and decay at $\sqrt{s} = 13$ TeV. Technical Report CMS-PAS-HIG-21-018, CERN, Geneva, 2025. URL: `https://cds.cern.ch/record/2929999`.

[43] M. Oreglia et al. Study of the reaction $\psi' \to \gamma\gamma\frac{J}{\psi}$. *Phys. Rev. D*, 25:2259, May 1982. `doi:10.1103/PhysRevD.25.2259`.

[44] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998. URL: `https://link.aps.org/doi/10.1103/PhysRevD.57.3873`, `doi:10.1103/PhysRevD.57.3873`.

[45] CMS collaboration. Limits on the Higgs boson lifetime and width from its decay to four charged leptons. *Phys. Rev. D*, 92:072010, Oct 2015. URL: `https://link.aps.org/doi/10.1103/PhysRevD.92.072010`, `doi:10.1103/PhysRevD.92.072010`.

[46] Lance J. Dixon and Ye Li. Bounding the Higgs boson width through interferometry. *Physical Review Letters*, 111(11), September 2013. `arXiv:1305.3854`, `doi:10.1103/physrevlett.111.111802`.

[47] Fabrizio Caola and Kirill Melnikov. Constraining the Higgs boson width with zz production at the LHC. *Phys. Rev. D*, 88:054024, Sep 2013. URL: `https://link.aps.org/doi/10.1103/PhysRevD.88.054024`, `doi:10.1103/PhysRevD.88.054024`.

[48] Measurement of off-shell Higgs boson production in the $H^* \to ZZ \to 4\ell$ decay channel using a neural simulation-based inference technique in 13 TeV pp collisions with the ATLAS detector. *Reports on Progress in Physics*, 88(5):057803, may 2025. `arXiv:2412.01548`, `doi:10.1088/1361-6633/adcd9a`.

[49] Ashley Ellen Mcdougall. *Determining the Higgs boson width in the diphoton decay channel with the ATLAS experiment.* Phd thesis, Ams-

terdam University, Nov 27 2024. CERN-THESIS-2024-230, CERN, May 15 2024. URL: `https://repository.cern/records/qssqj-9sh31`.

[50] Stephen P. Martin. Shift in the lhc higgs diphoton mass peak from interference with background. *Physical Review D*, 86(7), October 2012. URL: `http://dx.doi.org/10.1103/PhysRevD.86.073016`, `doi:10.1103/physrevd.86.073016`.

[51] Piotr Bargieła, Federico Buccioni, Fabrizio Caola, Federica Devoto, Andreas von Manteuffel, and Lorenzo Tancredi. Signal-background interference effects in Higgs-mediated diphoton production beyond NLO. *The European Physical Journal C*, 83(2):174, Feb 2023. `doi:10.1140/epjc/s10052-023-11337-w`.

[52] F. Coradeschi, D. de Florian, L.J. Dixon, N. Fidanza, S. Hoeche, H. Ita, Y. Li, and J. Mazzitelli. Interference effects in the $\gamma\gamma$ + 2jets channel at the LHC. *Physical Review D*, 92(1), July 2015. URL: `http://dx.doi.org/10.1103/PhysRevD.92.013004`, `doi:10.1103/physrevd.92.013004`.

[53] Michelangelo L Mangano. Introduction to QCD. (CERN-OPEN-2000-255), 2000. URL: `https://cds.cern.ch/record/454171`, `doi:10.5170/CERN-1999-004.53`.

[54] Torbjorn Sjostrand, Stephen Mrenna, and Peter Skands. A brief introduction to pythia 8.1. *Computer Physics Communications*, 178(11):852, June 2008. `arXiv:0710.3820`, `doi:10.1016/j.cpc.2008.01.036`.

[55] Gavin Bewick, Silvia Ferrario Ravasio, Stefan Gieseke, Stefan Kiebacher, Mohammad R. Masouminia, Andreas Papaefstathiou, Simon Plätzer, Peter Richardson, Daniel Samitz, Michael H. Seymour, Andrzej Siódmok, and James Whitehead. Herwig 7.3 release note, 2024. URL: `https://arxiv.org/abs/2312.05175`, `arXiv:2312.05175`.

[56] T Gleisberg, S Hoche, F Krauss, M Schonherr, S Schumann, F Siegert, and J Winter. Event generation with SHERPA 1.1. *Journal of High Energy Physics*, 2009(02):007, February 2009. URL: `http://dx.doi.org/10.1088/1126-6708/2009/02/007`, `doi:10.1088/1126-6708/2009/02/007`.

[57] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, aug 2008. `doi:10.1088/1748-0221/3/08/S08001`.

[58] Albert M Sirunyan et al. Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS. *Eur. Phys. J. C*, 81(CMS-LUM-17-003, CERN-EP-2021-033):800, 2021. `arXiv:2104.01927`, `doi:10.1140/epjc/s10052-021-09538-2`.

[59] CMS Collaboration. Cumulative luminosity per year. `https://cmslumi.web.cern.ch/publicplots/multiYear/int_lumi_cumulative_pp_3.pdf`. [Accessed 22-12-2025].

[60] David Barney. CMS Detector Slice. CMS Collection., 2016. URL: `https://cds.cern.ch/record/2120661`.

[61] CMS Collaboration.

[62] CMS Collaboration. Material budget for tracker phase 1 in units of radiation length. `https://twiki.cern.ch/twiki/pub/CMSPublic/TrackerMaterialBudgetplots/MaterialBudgetTrackerPhase1_FractionRadiationLength.png`, 2020. [Accessed 22-12-2025].

[63] CMS Collaboration. Material budget for tracker phase 1 in units of hadron interaction lengths. `https://twiki.cern.ch/twiki/pub/CMSPublic/TrackerMaterialBudgetplots/MaterialBudgetTrackerPhase1_HadronicInteractionLengthFraction.png`, 2020. [Accessed 22-12-2025].

[64] The CMS experiment at the CERN LHC. *JINST*, 3:S08004, 2008. `doi:10.1088/1748-0221/3/08/S08004`.

[65] Francesca Cavallari. Performance of calorimeters at the LHC. *Journal of Physics: Conference Series*, 293(1):012001, apr 2011. `doi:10.1088/1742-6596/293/1/012001`.

[66] Wolfgang Adam, R Frühwirth, Are Strandlie, and T Todor. Reconstruction of Electrons with the Gaussian-Sum Filter in the CMS Tracker at the LHC. Technical Report CMS-NOTE-2005-001, CERN, Geneva, 2005. URL: `https://cds.cern.ch/record/815410`.

[67] Giuseppe Cerati, Peter Elmer, Brian Gravelle, Matti Kortelainen, Vyacheslav Krutelyov, Steven Lantz, Matthieu Lefebvre, Mario Masciovecchio, Kevin McDermott, Boyana Norris, Allison Reinsvold Hall, Daniel Riley, Matevž Tadel, Peter Wittich, Frank Würthwein, and Avi Yagil. Parallelized and vectorized tracking using Kalman filters with CMS detector geometry and events. *EPJ Web of Conferences*, 214:02002, 2019. URL: `http://dx.doi.org/10.1051/epjconf/201921402002`, `doi:10.1051/epjconf/201921402002`.

[68] Reconstruction of signal amplitudes in the CMS electromagnetic calorimeter in the presence of overlapping proton-proton interactions. *Journal of Instrumentation*, 15(10):P10002, oct 2020. `doi:10.1088/1748-0221/15/10/P10002`.

[69] M. Anfreville, D. Bailleux, J.P. Bard, A. Bornheim, C. Bouchand, E. Bougamont, M. Boyer, R. Chipaux, V. Daponte-Puill, M. Dejardin, J.L. Faure, P. Gras, P. Jarry, C. Jeanney, A. Joudon, J.P. Pansart, Y. Penichot, J. Rander, J. Rolquin, J.M. Reymond, J. Tartas, P. Venault, P. Verrecchia, L. Zhang, K. Zhu, and R.Y. Zhu. Laser monitoring system for the CMS lead tungstate crystal calorimeter. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 594(2):292–320, 2008. URL: `https://www.sciencedirect.com/science/article/pii/S0168900208001599`, `doi:https://doi.org/10.1016/j.nima.2008.01.104`.

[70] Common tools for analyses of Higgs boson in the diphoton decay channel with Run 2 ultra-legacy data. CMS Analysis Note CMS AN-2021/025, CMS Collaboration, 2021. Internal CMS document. URL: `https://cms.cern.ch/iCMS/analysisadmin/cadi?ancode=AN-2021/025`.

[71] Measurement of the mass of the Higgs boson in the diphoton decay channel in pp collisions at 13 TeV with the legacy Run 2 dataset. CMS Analysis Note CMS AN-2020/217, CMS Collaboration, 2020. Internal CMS document. URL: `https://cms.cern.ch/iCMS/analysisadmin/cadi?ancode=AN-2020/217`.

[72] Alexandre Zabi. Triggering on electrons and photons with the CMS experiment at the LHC. *Journal of Physics: Conference Series*, 404(1):012056, dec 2012. `doi:10.1088/1742-6596/404/1/012056`.

[73] The FLUKA code: Overview and new developments. *EPJ - Nuclear Sciences & Technologies*, 10, 2024. `doi:https://doi.org/10.1051/epjn/2024015`.

[74] François-Xavier Gentit. Litrani: a general purpose Monte-Carlo program simulating light propagation in isotropic or anisotropic media. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 486(1):35–39, 2002. Proceedings of the 6th International Conference on Inorganic Scintillators and their Use in Scientific and Industrial Applications. `doi:https://doi.org/10.1016/S0168-9002(02)00671-X`.

[75] Joshua Lorne Bendavid. *Evidence for a narrow Higgs-like diphoton resonance with a mass of 125 GeV in pp collisions at* $\sqrt{s} = 7$–8 *TeV*. Ph.d. thesis, Massachusetts Institute of Technology, Department of Physics, Cambridge, MA, 2013. Includes bibliographical references (pp. 153–158). Cataloged from PDF version of thesis. URL: `https://dspace.mit.edu/handle/1721.1/79539`.

[76] CMS Collaboration. CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV. CMS Physics Analysis Summary CMS-PAS-LUM-17-004, 2018. URL: `https://cds.cern.ch/record/2621960/`.

[77] CMS Collaboration. CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV. CMS Physics Analysis Summary CMS-PAS-LUM-18-002, 2019. URL: `https://cds.cern.ch/record/2676164/`.

[78] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *Journal of High Energy Physics*, 2014(7), July 2014. `arXiv:1405.0301`, `doi:10.1007/JHEP07(2014)079`.

[79] Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements. *The European Physical Journal C*, 80(1), January 2020. `arXiv:1903.12179`, `doi:10.1140/epjc/s10052-019-7499-4`.

[80] Richard D. Ball, Valerio Bertone, Stefano Carrazza, Luigi Del Debbio, Stefano Forte, Patrick Groth-Merrild, Alberto Guffanti, Nathan P. Hartland, Zahari Kassabov, Jose I. Latorre, Emanuele R. Nocera, Juan Rojo, Luca Rottoli, Emma Slade, and Maria Ubiali. Parton distributions from high-precision collider data: NNPDF Collaboration. *The European Physical Journal C*, 77(10), October 2017. `arXiv:1706.00428`, `doi:10.1140/epjc/s10052-017-5199-5`.

[81] Pier Francesco Monni, Emanuele Re, and Marius Wiesemann. MiNNLO$_{\mathrm{PS}}$: optimizing $2 \to 1$ hadronic processes. *Eur. Phys. J. C*, 80(CERN-TH-2020-084, LAPTH-027/20, MPP-2020-75):1075, 2020. `arXiv:2006.04133`, `doi:10.1140/epjc/s10052-020-08658-5`.

[82] Enrico Bothmann, Gurpreet Singh Chahal, Stefan Hoeche, Johannes Krause, Frank Krauss, Silvan Kuttimalai, Sebastian Liebschner, Davide Napoletano, Marek Schonherr, Holger Schulz, Steffen Schumann, and Frank Siegert. Event generation with Sherpa 2.2. *SciPost Phys.*, 7:034, 2019. URL: `https://scipost.org/10.21468/SciPostPhys.7.3.034`, `doi:10.21468/SciPostPhys.7.3.034`.

[83] Piotr Bargieła, Federico Buccioni, Fabrizio Caola, Federica Devoto, Andreas von Manteuffel, and Lorenzo Tancredi. Signal-background interference effects in higgs-mediated diphoton production beyond nlo. *The European Physical Journal C*, 83(2), February 2023. URL: `http://dx.doi.org/10.1140/epjc/s10052-023-11337-w`, `doi:10.1140/epjc/s10052-023-11337-w`.

[84] GEANT4 Collaboration. Geant4 - a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250, 2003. `doi:10.1016/S0168-9002(03)01368-8`.

[85] Transverse-momentum resummation for the signal-background interference in the $h \to \gamma\gamma$ channel at the LHC.

[86] Stefan Höche and Stefan Prestel. The midpoint between dipole and parton showers. *The European Physical Journal C*, 75(9), September 2015. URL: `http://dx.doi.org/10.1140/epjc/s10052-015-3684-2`, `doi:10.1140/epjc/s10052-015-3684-2`.

[87] Mrinal Dasgupta, Frédéric A. Dreyer, Keith Hamilton, Pier Francesco Monni, and Gavin P. Salam. Logarithmic accuracy of parton showers: a fixed-order study. *Journal of High Energy Physics*, 2018(9), September 2018. URL: `http://dx.doi.org/10.1007/JHEP09(2018)033`, `doi:10.1007/jhep09(2018)033`.

[88] Performance of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 13$ tev. *Journal of Instrumentation*, 19(09):P09004, sep 2024. URL: `10.1088/1748-0221/19/09/P09004`, `doi:10.1088/1748-0221/19/09/P09004`.

[89] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. `doi:10.1109/4235.996017`.

[90] Roger Barlow. Extended maximum likelihood. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 297(3):496, 1990. `doi:10.1016/0168-9002(90)91334-8`.

[91] R. A. Fisher. On the interpretation of χ^2 from contingency tables, and the calculation of p. *Journal of the Royal Statistical Society*, 85(1):87, 1922. `doi:10.2307/2340521`.

[92] P. Dauncey, M. Kenzie, N. Wardle, and G. Davies. Handling uncertainties in background shapes: the discrete profiling method. *Journal of Instrumentation*, 10(04):P04015, April 2015. `arXiv:1408.6865`, `doi:10.1088/1748-0221/10/04/p04015`.

[93] T. W. Anderson and D. A. Darling. Asymptotic theory of certain "goodness of fit" criteria based on stochastic processes. *Annals of Mathematical Statistics*, 23:193–212, 1952. `doi:10.1214/aoms/1177729437`.

[94] Felipe L. Gewers, Gustavo R. Ferreira, Henrique F. De Arruda, Filipi N. Silva, Cesar H. Comin, Diego R. Amancio, and Luciano Da F. Costa. Principal Component Analysis: A natural approach to data exploration. *ACM Comput. Surv.*, 54(4), May 2021. `doi:10.1145/3447755`.

[95] Jet algorithms performance in 13 TeV data. Technical Report CMS-PAS-JME-16-003, CERN, Geneva, 2017. URL: `https://cds.cern.ch/record/2256875`.

[96] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2), February 2011. URL: `http://dx.doi.org/10.1140/epjc/s10052-011-1554-0`, `doi:10.1140/epjc/s10052-011-1554-0`.