Measurements of Higgs boson properties in the four-lepton final state at sqrt(s) = 13 TeV with the CMS experiment at the LHC.

Simon Regnard

NNT : 2016SACLX081

Thèse de doctorat de l'Université Paris-Saclay
préparée à l'École polytechnique

École doctorale n°576
Particules, Hadrons, Énergie, Noyau, Instrumentation, Imagerie, Cosmos et Simulation (PHENIICS)

Spécialité de doctorat : Physique des particules

par

Simon Regnard

Measurements of Higgs boson properties in the four-lepton final state at $\sqrt{s} = 13\,\text{TeV}$ with the CMS experiment at the LHC

Mesure des propriétés du boson de Higgs dans l'état final à quatre leptons à $\sqrt{s} = 13\,\text{TeV}$ avec l'expérience CMS au LHC

Thèse présentée et soutenue à l'École polytechnique le 7 novembre 2016
devant le jury composé de :

M. Ivica Puljak	FESB, Split	Président du jury
Mme Isabelle Wingerter-Seez	LAPP, Annecy-le-Vieux	Rapporteur
Mme Lydia Roos	LPNHE, Paris	Rapporteur
M. Tiziano Camporesi	CERN, Meyrin	Examinateur
M. Abdelhak Djouadi	LPT, Orsay	Examinateur
M. Yves Sirois	LLR, Palaiseau	Directeur de thèse
M. Roberto Salerno	LLR, Palaiseau	Co-directeur de thèse

Abstract

This thesis reports a study of Higgs boson production in proton-proton (pp) collisions at $\sqrt{s} = 13\,\text{TeV}$ recorded with the CMS detector at the CERN Large Hadron Collider (LHC), exploiting the decay channel into a pair of Z bosons that in turn decay into pairs of electrons or muons ($\text{H} \to \text{ZZ} \to 4\ell$, $\ell = \text{e}, \mu$).

This work is carried out in the context of the beginning of Run II of the LHC, a new data-taking period that started in 2015, following a two-year shutdown. This restart is marked by an increase of the centre-of-mass energy from 8 TeV to 13 TeV and an upgrade of collider parameters that increases the luminosity, setting new constraints on the triggering, reconstruction and analysis of pp collision events. Therefore, considerable effort is devoted to the improvement and reoptimization of the CMS trigger system for Run II, focusing on the reconstruction and selection of electrons and on the preparation of multilepton trigger paths that preserve a maximal efficiency for the $\text{H} \to \text{ZZ} \to 4\ell$ channel.

Secondly, the offline algorithms for electron and muon selection are optimized and their efficiencies are measured in data, while the selection logic of four-lepton candidates is improved. In order to extract rare production modes of the Higgs boson such as vector boson fusion, VH associated production, and $\text{t}\bar{\text{t}}\text{H}$ associated production, a new classification of selected events into exclusive categories is introduced, using discriminants based on matrix-element calculations and jet flavour tagging.

Results of the analysis of first 13 TeV data are presented for two data sets recorded in 2015 and early 2016, corresponding to integrated luminosities of $2.8\,\text{fb}^{-1}$ and $12.9\,\text{fb}^{-1}$, respectively. A standalone rediscovery of the Higgs boson in the four-lepton channel is achieved at the new energy. The signal strength relative to the standard model prediction, the mass and decay width of the boson, and a set of parameters describing the contributions of its main predicted production modes are measured. All results are in good agreement with standard model expectations for a 125 GeV Higgs boson within the uncertainties, which are dominated by their statistical component with the current data set. Finally, a search for an additional high-mass resonance decaying to four leptons is performed, and no significant excess is observed.

Résumé

Cette thèse présente une étude de la production de boson de Higgs dans les collisions proton-proton (pp) à $\sqrt{s} = 13\,\text{TeV}$ enregistrées avec le détecteur CMS au Grand collisionneur de hadrons (LHC) du CERN. Les résultats sont obtenus par une analyse du canal à quatre leptons, où le boson de Higgs se désintègre en une paire de bosons Z qui se désintègrent à leur tour en paires d'électrons ou de muons ($\text{H} \to \text{ZZ} \to 4\ell$, $\ell = \text{e}, \mu$). En dépit de son faible rapport de branchement, ce canal bénéficie d'une grande sensibilité : son grand rapport signal sur bruit, la possibilité de reconstruire complètement l'état final et l'excellente résolution de CMS pour l'impulsion des leptons facilitent l'extraction du signal et permettent de mesurer un large éventail de propriétés du boson.

Ce travail s'inscrit dans le contexte du début du Run II du LHC, une nouvelle période de prise de données qui a commencé en 2015 après une interruption de deux ans. Ce redémarrage est marqué par une augmentation de l'énergie dans le centre de masse de 8 TeV à 13 TeV, et par une mise à niveau des paramètres du collisionneur augmentant la luminosité, notamment le resserrement de l'espacement entre paquets de protons de 50 à 25 ns. En augmentant le taux d'occurrence des processus physiques d'intérêt et en accentuant le phénomène d'empilement, ces nouveaux paramètres placent des contraintes inédites sur les algorithmes de déclenchement, de reconstruction et d'analyse des données de collisions pp.

Un effort important est consacré à l'amélioration et à la réoptimisation du système de déclenchement de haut niveau de CMS pour le Run II, en mettant l'accent sur les algorithmes permettant la sélection d'événements à quatre leptons. La reconstruction en ligne des électrons bénéficie d'abord d'une remise à niveau des algorithmes de reconstruction des agrégats dans le calorimètre électromagnétique et des traces dans le trajectographe, permettant ainsi la synchronisation avec les progrès correspondants des algorithmes hors ligne. La sélection des électrons est ensuite améliorée : une nouvelle stratégie d'isolation est notamment développée, qui combine les idées de l'isolation détecteur utilisé en ligne au Run I et de l'isolation Particle-flow utilisée hors ligne, et permet de retrouver les performances accrues de cette dernière tout en maîtrisant le temps de calcul. Une procédure de correction des effets de l'empilement sur l'isolation est également implémentée, de façon à maintenir une efficacité de déclenchement des électrons stable dans les conditions de prise de données du Run II. Tous ces développements sont finalement utilisés pour concevoir un ensemble de voies de déclenchement sélectionnant le signal $\text{H} \to \text{ZZ} \to 4\ell$ avec une efficacité maximale tout en respectant les contraintes de taux sortant d'événements. La stratégie consiste à combiner des voies di-leptons et tri-leptons, respectivement avec et sans coupures d'isolation, et utilisant un point de fonctionnement commun pour l'ensemble des coupures de sélection des électrons.

L'analyse hors ligne du canal $\text{H} \to \text{ZZ} \to 4\ell$ est ensuite mise en place, en commençant par les algorithmes de sélection des électrons, des muons, des photons de rayonnement

d'état final et des jets hadroniques. Les critères d'isolation des leptons et le classificateur multivarié permettant l'identification des électrons sont notamment réoptimisés pour les nouvelles conditions de prise de données. L'échelle et la résolution de l'impulsion des leptons sont étalonnées grâce aux résonances dileptoniques connues. L'efficacité de la reconstruction et de la sélection des leptons est mesurée par la méthode du Tag-and-Probe dans les données et dans la simulation, ce qui permet de repondérer cette dernière par des facteurs correctifs.

Une fois définie la sélection des particules et objets physiques, la sélection des événements s'articule autour de l'assemblage et du choix d'un candidat à quatre leptons, selon une procédure réoptimisée. Une nouvelle répartition des événements ainsi sélectionnés en six catégories exclusives est introduite, dans le but d'extraire des modes de production rares du boson de Higgs tels que la fusion de bosons vecteurs, la production associée VH et la production associée $\mathrm{t\bar{t}H}$. Cette catégorisation exploite la présence de jets et de leptons additionnels, l'étiquetage des jets de b, ainsi que quatre discriminants de production qui utilisent le calcul d'éléments de matrice et les techniques de discrimination entre jets de quarks et de gluons. Les ajustements multidimensionnels permettant d'extraire le signal dans chaque catégorie exploitent ensuite deux observables de désintégration : la masse invariante à quatre leptons et un discriminant par éléments de matrice sensible à la configuration cinématique du système de quatre leptons. Les nouvelles catégories sont intégrées dans la procédure de modélisation du signal, dans l'estimation des bruits de fond irréductibles, qui repose sur la simulation corrigée par des facteurs rendant compte des corrections d'ordre supérieur, et dans celle des bruits de fond réductibles, fondée sur deux méthodes indépendantes exploitant des régions de contrôle dans les données.

L'analyse des premières collisions pp à 13 TeV utilise dans un premier temps les données enregistrées par CMS en 2015, qui correspondent à une luminosité intégrée de $2.8\,\mathrm{fb}^{-1}$. Cet échantillon modeste est exploité pour exercer toutes les méthodes d'analyse et produire un premier ensemble de résultats publics. Le boson de Higgs découvert au Run I est remis en évidence à la nouvelle énergie, avec une signification statistique observée (attendue) de 3.0 (3.1) déviations standard au minimum de la valeur p, et de 2.5 (3.4) déviations standard à la masse de $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$ mesurée au Run I. L'intensité du signal μ, normalisée à l'attendu pour le boson de Higgs du modèle standard, vaut $\mu = 0.82^{+0.57}_{-0.43}$ à cette masse. Une catégorisation simplifiée des événements permet de mesurer deux paramètres associés aux modes de production par couplage aux fermions et aux bosons vecteurs, pour des valeurs respectives de $\mu_{\mathrm{ggH,t\bar{t}H}} = 0.95^{+0.64}_{-0.49}$ et $\mu_{\mathrm{VBF,VH}} = 0.0^{+2.5}_{-0.0}$. Ces résultats sont en accord avec les prévisions du modèle standard aux incertitudes de mesure près.

Dans un second temps, les données enregistrées par CMS pendant la première partie de l'année 2016, correspondant à une luminosité intégrée plus importante de $12.9\,\mathrm{fb}^{-1}$, sont exploitées pour produire un ensemble plus vaste de résultats publics. Avec une signification statistique observée (attendue) de 6.2 (6.5) déviations standard à la masse de $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$, le boson de Higgs est cette fois redécouvert de façon indépendante à $\sqrt{s} = 13\,\mathrm{TeV}$ dans l'état final à quatre leptons. À la même masse, l'intensité du signal relative à la prédiction du modèle standard est mesurée à une valeur de $\mu = 0.99^{+0.33}_{-0.26}$. Le déploiement de la nouvelle stratégie de catégorisation profite à l'extraction des modes de production du boson : d'une part, les paramètres de couplage associés aux fermions et aux bosons vecteurs sont respectivement $\mu_{\mathrm{ggH,t\bar{t}H}} = 1.00^{+0.39}_{-0.32}$ et $\mu_{\mathrm{VBF,VH}} = 0.91^{+1.56}_{-0.91}$; d'autre part, quatre nouveaux paramètres associés aux quatre principaux modes de

production du boson de Higgs sont mesurés. Cela permet de sonder indépendamment pour la première fois les processus VH et $t\bar{t}H$ dans l'état final à quatre leptons, en prélude à de futures combinaisons avec les autres états finaux. Par une procédure utilisant les incertitudes de masse par événement en plus des observables déjà citées, la masse du boson de Higgs est mesurée à une valeur de $m_{\mathrm{H}} = 124.50^{+0.48}_{-0.46}\,\mathrm{GeV}$. La largeur de désintégration Γ_{H} du boson est mesurée en exploitant à la fois la région du pic et la région de haute masse, où le boson est produit hors couche de masse ; cela conduit à une contrainte $\Gamma_{\mathrm{H}} < 41\,\mathrm{MeV}$ à 95% de niveau de confiance. Tous ces résultats sont en bon accord avec les prévisions du modèle standard pour un boson de Higgs d'environ 125 GeV, aux incertitudes de mesure près, ces dernières étant dominées par la composante statistique avec l'échantillon de données actuel. Enfin, une recherche de nouvelle résonance se désintégrant en quatre leptons est effectuée dans un intervalle de masse atteignant 2.5 TeV, et pour plusieurs hypothèses de largeur de désintégration. Aucun excès significatif n'est observé, ce qui conduit à définir des limites supérieures d'exclusion à 95% de niveau de confiance sur la section efficace de production de la résonance hypothétique.

Acknowledgements

As this eventful and rewarding three-year adventure is coming to an end, I would like to express my gratitude to the many people who have accompanied and helped me at various stages of this thesis work.

This begins with my two outstanding supervisors: Roberto with his patience, optimism, and availability, and Yves with his experience, perspective, and motivation. Thank you both for mentoring me, for giving me the amount of autonomy I needed, and for pushing me to the forefront whenever you could.

I have been fortunate enough to find an extremely favourable working environment in the CMS group of Laboratoire Leprince-Ringuet at École polytechnique. To Giacomo, Florian, Philipp, Toni, Christophe, Claude, Jean-Baptiste, Olivier, Alexandre, Luca, Stéphanie, Thomas, Iurii, Philippe, and all the others: thank you for all the discussions and encouragements, for the shared knowledge and skills, and simply for your enthusiasm. And thanks to the IT and administration teams of LLR for their willingness to make my life easier.

Many thanks to the CMS Collaboration at CERN: I know that this thesis builds on the work of hundreds of physicists who designed, constructed, and now run this experiment. My own work is essentially a piece of a vast collective analysis, and I have easily found help from competent people whenever I needed it. I am particularly indebted to Matteo, Hugues, and other TSG participants for guiding me in the demanding world of the trigger; to Nicola for transmitting his great expertise with many analysis technicalities; to my CJLST colleagues Andrei, Ulascan, Meng, Bilal, Candice, Heshy, Pedro, Ivica, Piotr, Stefano, Roberto for their help and advice throughout months of shared analysis work; to David for sharing the editorship effort in my first CMS analysis approval process; to Pedja, Giovanni, Tongguang, Nate, Hengne, Roko, and the whole HZZ team: the analysis and results I am presenting here directly rely on much work of you all, and it has been a pleasure working alongside you.

Beyond everyday work, I will keep good memories of meeting other young physicists from various horizons, at CERN or at workshops and conferences. To Nadir, Xavier, Yacine, Luca, Alice, Márton, Michaël, Christophe, Tutanon, Léo, plus already mentioned people and many others, thank you for the interesting conversations we had, be it to discuss particle physics or to take our minds off it.

Finally, without naming them all, I wish to thank my family and my close friends: your steady support for my joining the LHC adventure has come in diverse, sometimes unexpected forms, and has helped me more than you think through these three years.

Contents

Introduction

Elucidating the mechanism of electroweak symmetry breaking has long been one of the main challenges of the study of the standard model (SM) of particle physics, i.e. the quantum field theory that describes the electromagnetic, weak, and strong interactions. A leading hypothesis has been provided by the Brout–Englert–Higgs (BEH) mechanism, which builds on the Goldstone theorem and on a doublet of complex scalar fields to generate masses for the otherwise massless weak vector bosons. Moreover, it naturally provides mass terms for the fermionic matter fields via their Yukawa interactions with the Higgs doublet. The main experimental implication of the minimal BEH mechanism is the existence of a new physical scalar boson, the Higgs boson (H), the mass m_{H} of which is a free parameter of the model. While the BEH mechanism was first proposed in 1964, expectations became increasingly insistent over the last decades as the discoveries of the $\mathrm{W}^{\pm}$/ Z^{0} bosons and of the top quark and various precision measurements solidly established the SM as the working theory of fundamental interactions, leaving the hypothetical Higgs boson as a missing ingredient of the model. The allowed range for m_{H} is restricted from theoretical arguments and got progressively constrained by experimental searches.

After large intervals of the m_{H} phase space were excluded at the Large Electron Positron collider and at the Tevatron in the 1990s and early 2000s, the existence of the Higgs boson became one of the inescapable questions of the early physics programme of the new energy frontier machine, the CERN Large Hadron Collider (LHC). The LHC was intended to provide proton-proton (pp) collisions with a nominal centre-of-mass energy of 14 TeV, and came to life accompanied by two multi-purpose detectors, the Compact Muon Solenoid (CMS) and A Toroidal LHC Apparatus (ATLAS). These were primarily designed to search for the long-awaited boson in a variety of production and decay channels, and to explore physics at the TeV scale. After a two-decade-long adventure of development, construction, and eventful commissioning, the LHC began delivering pp collisions in 2009, reached conditions suitable for physics in 2010, and finally delivered two high-luminosity data samples in 2011 and 2012, at centre-of-mass energies of 7 TeV and 8 TeV, respectively. This initial data taking period, referred to as Run I, was intended to quickly answer the Higgs boson riddle with intermediate energies, postponing a restart at design energy to 2015 after a two-year shutdown.

After new regions of the phase space were excluded in late 2011, salvation came from the last non-excluded interval in July 2012, when the ATLAS and CMS collaborations reported the discovery of a new boson with mass near 125 GeV, with properties consistent with the standard model Higgs boson, marking an historic milestone in physics. While the discovery relied on a combination of studies in five different decay channels, two of them provided most of the sensitivity and a measurement of the mass: the decay to a pair of photons ($\mathrm{H}\to\gamma\gamma$) and the decay to a pair of Z bosons that both decay into pairs of electrons or muons ($\mathrm{H}\to\mathrm{ZZ}\to 4\ell$, $\ell = \mathrm{e}, \mu$). The new boson was soon shown to

have spin-parity 0^+, and all subsequent property measurements from both collaborations confirmed it as very consistent so far with expectations for a minimal scalar sector with one physical H boson as incorporated into the standard model.

The $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ process is the focus of the present thesis work. Taking advantage of its large signal-to-background ratio and of the complete reconstruction of the final state with excellent resolution, CMS has maximally exploited the handful of Run I Higgs boson events collected in this channel to produce a large set of property measurements, such as the mass, spin-parity, width, lifetime, and fiducial cross section of the new boson, and constraints on anomalous couplings at the HZZ vertex. Moreover, $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ is a rich channel in the sense that requires inputs from all subdetectors and expertise areas of CMS. From the early days of the experiment, it has catalysed the development of many analysis techniques such as electron reconstruction and identification down to low transverse momentum, precise efficiency measurements by the Tag-and-Probe technique, and the use of matrix-element based discriminants to optimally extract angular information from candidate events.

With the collective focus shifting from the exploitation of Run I to the prospect of a restart at $\sqrt{s} = 13\,\mathrm{TeV}$ in 2015, exploring the four-lepton channel at this new unexplored energy has remained very promising in many respects. On the one hand, thanks to the increase of Higgs boson production cross sections and the prospects of large data samples, the aforementioned properties of the boson can be measured with ever-growing precision, hoping for possible deviations from the SM expectation to provide a portal to physics beyond the SM. In this work, the emphasis is put on the couplings, which are explored by searching for rare production modes of the Higgs boson such as vector boson fusion, VH associated production and $\mathrm{t\bar{t}H}$ associated production. On the other hand, the four-lepton channel is relevant to direct searches of hypothetical new scalar resonances decaying to ZZ, since very little background is expected at high mass, and the good four-lepton mass resolution is an important asset for resonance searches. A very similar event selection can be used to both study the boson discovered near 125 GeV and search for a new one.

This thesis work started at the end of Summer 2013, in the context of the so-called Long Shutdown 1 (LS1), with the prospect of working for about a year and a half without fresh data. This allowed me to organize my thesis work from the bottom up, starting from the very base of the future $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis, i.e. trigger level. I was indeed initially involved in a collective long-term technical work to re-optimize the CMS trigger system so as to cope with the challenging event rates and pileup conditions that were expected for Run II. My task consisted in optimizing the high-level trigger (HLT) algorithms related to electron selection, more specifically isolation. I replaced the detector-based isolation that was used at trigger level in Run I by a new algorithm that is consistent with the overall Particle-flow paradigm used throughout CMS, and added pileup subtraction to it. When these studies ended in early 2014, my focus shifted from optimizing such object-level algorithms to actually building the trigger paths that would have to select the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ events in Run II. Here, the goal was to tune various selection thresholds in a way that preserves a signal efficiency relative to final selected events as close to 100% as possible, as it had been in Run I, while fitting in the rate budget. A first CMS-wide proposal for a Run II trigger menu was delivered in the summer of 2014, in two different versions corresponding to different levels of expected instantaneous luminosity. As the data taking approached, some final technical ingredients such as HCAL local reconstruction and ECAL calibrations came

in late 2014 and early 2015, requiring to spend a substantial amount of time refining the trigger paths and menus.

From late 2014 on, I moved from trigger to analysis level and contributed to several important aspects of the new $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis. An important point was the tuning of the selection algorithms for offline electrons and muons, such as identification and isolation, and a rationalization of the selection of four-leptons candidates. In parallel, I designed an extended event categorization scheme, with the goal of extracting all main Higgs boson production mechanisms from the rich event topologies of the four-lepton channel. Even though reaching a standalone observation of all these processes was known to require much more data than would be available by mid-2016, this served a larger, twofold purpose: paving the way for the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis that will gradually exploit the full Run II data set until 2018, and contributing to the upcoming CMS combinations of Higgs physics results from all decay channels, which will set more stringent constraints on the couplings.

Meanwhile, in the spring of 2015, the LHC delivered its first pp collisions at $\sqrt{s} = 13\,\mathrm{TeV}$. Many weeks of data taking were used for the LHC operators to fully tame what had practically become a new machine after the upgrade work of LS1. The first pp collisions were delivered with a bunch spacing of 50 ns, and the CMS detector began recording data using a modified version of the Run I Level-1 trigger, allowing the introduction of the first version of the Run II HLT menu. As the LHC moved to 25 ns bunches in July 2015, CMS switched to a new 'Stage 1' Level-1 trigger, which was used until the end of 2015, still with modest luminosity and pileup levels, to complete a commissioning-oriented year. The period was spiced by some problems with the cooling system of the superconducting solenoid, so that CMS found itself recording a fraction of the collision data with the magnet off. At the end of 2015, what was available for mainstream physics studies such as $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ was a data sample corresponding to an integrated luminosity of $2.8\,\mathrm{fb}^{-1}$ recorded with the magnet on. The decision was taken to prepare a first set of preliminary results to be presented at the *51st Rencontres de Moriond on Electroweak interactions and unified theories* in La Thuile, in March 2016. I was chosen as one of the two co-editors of the Physics Analysis Summary (PAS) that CMS released on that occasion.

After the adventures of the 2015 data taking, the year 2016 was considered as the first real Run II production year for the LHC, with initial hopes for about $30\,\mathrm{fb}^{-1}$ of data for the full year. CMS began recording data with a novel 'Stage 2' Level-1 trigger setup, which had been commissioned in the late 2015 runs. On top of this, high-luminosity versions of the HLT menu were progressively deployed as the LHC was breaking luminosity records. From early on, the CMS Higgs physics analysis group decided that the first 13 TeV-based $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ paper would still have to wait for the complete 2016 data sample to be delivered. But in the meantime, an extended, intermediary set of preliminary results and a new PAS were prepared for the main summer conference, namely the *38th International Conference on High Energy Physics* (ICHEP) in Chicago. For this so-called 2016 analysis, a data sample corresponding to an integrated luminosity of $12.9\,\mathrm{fb}^{-1}$ recorded by the CMS detector by mid-July 2016 was used, and I presented the results on behalf of the CMS Collaboration at ICHEP on 4th August 2016. For this new milestone, the analysis was substantially extended and improved, in particular by deploying the complete event categorization aimed at hunting for all the main production mechanisms of the Higgs boson.

This thesis is structured as follows. A theory-oriented introduction to $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$

studies is given in Chapter 1, while an experiment-oriented introduction to the LHC and the CMS detector is made in Chapter 2. Chapter 3 reports the preparation and commissioning of the triggering of four-lepton events in Run II. Chapter 4 focuses on objects at offline level, presenting the successive selection requirements that are applied to all reconstructed particles involved in the study of the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ channel. These pieces assemble in Chapter 5 to build four-lepton candidates. There lies the core of the analysis, where event-level discriminating variables are chosen to best extract the Higgs boson signal, backgrounds are estimated, uncertainties are assessed and measurements are performed. This leads to two result-oriented chapters that correspond to the two successive rounds of public results. Chapter 6 first reports the outcome of the analysis based on the data sample recorded in 2015. Then, Chapter 7 presents the extended results of the 2016 analysis, in which all benefits of this thesis work are collected.

Chapter 1

The Higgs boson in the standard model

Experimental observations of subnuclear phenomena are nowadays best explained by the standard model of particle physics (SM), a quantum field theory that describes the fundamental constituents of matter and the fundamental interactions that affect them, more precisely the electroweak and strong interaction. In its current formulation, the SM incorporates a minimal scalar sector which explains the nature of electroweak symmetry breaking and the origin of the masses of fermions. This extension of the SM has recently received an experimental foundation with the discovery of a particle consistent with the physical Higgs boson that it predicts. Measuring the properties of the new observed particle is the topic of the present thesis work.

This first chapter is intended as an introduction to both objects: the Higgs boson as predicted by the theory, and the Higgs-boson-like particle that was discovered at the Large Hadron Collider. The theoretical formulation of the minimal scalar sector as currently included in the SM is first recalled, and Higgs boson phenomenology at a pp collider is discussed, introducing in particular the production mechanisms that are studied throughout this thesis. A brief history of experimental Higgs boson searches and discovery is then given. Finally, the distinctive features of the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ decay channel ($\ell = \mathrm{e}, \mu$) are presented.

1.1 Standard model of particle physics

1.1.1 Gauge groups and fields of the standard model

The standard model of particle physics is a renormalizable quantum field theory that describes the electromagnetic, weak, and strong interactions of elementary particles with the gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$. It does not encompass the gravitational interaction, which is negligible at subnuclear scales.

In the SM, matter is described by spin-$\frac{1}{2}$ fermion fields, while interactions are mediated by spin-1 boson fields. The fields live in representations of symmetry groups that act on some internal degrees of freedom of the particles. The strong interaction is mediated by *gluons*, and described by *quantum chromodynamics* (QCD), a theory based on the colour group $SU(3)_C$. Electromagnetism and the weak interaction are respectively mediated by the *photon* and the $\mathrm{W}^{\pm}$ and Z^0 *weak bosons*, and are unified in the *electroweak* (EW) *theory* (or Glashow–Salam–Weinberg model) [1, 2, 3], which relies on the $SU(2)_L \times$

$U(1)_Y$ group associated to weak isospin and weak hypercharge. As will be explained in Section 1.1.2, this group reduces to the $U(1)_{em}$ gauge group of electromagnetism after spontaneous breaking of the electroweak symmetry.

Twelve fundamental fermion fields are known: six *quarks* and six *leptons*. They are arranged in three generations, each of which is made of two quarks of electrical charge $+\frac{2}{3}$ and $-\frac{1}{3}$, and two leptons of charge -1 and 0. The first generation builds up ordinary matter, while the other two are replicas of it. Each fermion particle has an associated antiparticle with opposite quantum numbers.

- **Quarks** are subject to both strong and electroweak interactions. The three successive doublets are the *up* (u) and *down* (d), which are stable constituents of the atomic nuclei and have masses of a few MeV; the *charm* (c) and *strange* (s), with masses of about 1.3 GeV and 0.1 GeV; and the *top* (t) and *bottom* (b), with masses of about 172 GeV and 4.2 GeV. Owing to the QCD property of confinement, quarks do not exist as free states: they can only be observed in *hadrons*, either *mesons* (bound states of one quark and one anti-quark) or *(anti)baryons* (bound states of three (anti)quarks), except for the short-lived top quark. While quark flavour is conserved by the strong and electromagnetic interactions, it is not by the weak interaction. The mass eigenstates of down-type quarks mix with their weak eigenstates via the Cabibbo–Kobayashi–Maskawa (CKM) matrix.

- **Leptons** are only sensitive to the weak and electromagnetic forces. The three negatively charged leptons are the *electron* (e), the *muon* (μ), and the *tau* (τ), with masses of about 0.511 MeV, 106 MeV, and 1.8 GeV. The electron is stable, the muon has a lifetime long enough to be considered stable in collider experiments, whereas taus are only observed via their decay products. The three corresponding uncharged leptons are the *neutrinos*: ν_e, ν_μ, and ν_τ, which only interact with other particles via the nuclear weak force, and are not directly detectable at collider experiments. The observation of neutrino oscillations prove that these have masses, albeit tiny; the mixing of weak and mass eigenstates is described by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix.

The fundamental fermion fields are listed in Table 1.1 in terms of the $SU(2)_L$ representation, together with their main quantum numbers for the EW theory.

The SM is a combination of two *gauge theories*, i.e. quantum field theories where the Lagrangian density is invariant under a *local* symmetry, in which the transformation parameters can be function of the space-time position. Imposing such an invariance naturally implies the very existence of interactions. The formalism will first be introduced here for the case of QCD, and then applied to the EW interaction.

Strong interaction

QCD relies on the symmetry group $SU(3)_C$, a non-abelian group with 8 generators, which can be chosen as the 8 Gell-Mann matrices $\frac{\lambda_a}{2}$. The group is associated with a quantum number called *colour*, which can take three values. Quarks live in the simplest non-trivial representations of the group: **3** and $\overline{\mathbf{3}}$.

The quark fields q_f transform linearly under the gauge symmetry:

$$q_f(x) \longrightarrow e^{-i\,\alpha_a(x)\,\frac{\lambda_a}{2}}\, q_f(x)\ , \tag{1.1}$$

Table 1.1: Classification of the three generations of fundamental fermions in terms of the $SU(2)_L$ representation, with their weak isospin (I_3), weak hypercharge (Y) and electric charge (Q) quantum numbers.

	1st gen.	2nd gen.	3rd gen.	I_3	Y	Q
Quarks	$\begin{pmatrix} u_L \\ d_L \end{pmatrix}$	$\begin{pmatrix} c_L \\ s_L \end{pmatrix}$	$\begin{pmatrix} t_L \\ b_L \end{pmatrix}$	$\begin{pmatrix} \frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$\begin{pmatrix} \frac{1}{3} \\ \frac{1}{3} \end{pmatrix}$	$\begin{pmatrix} \frac{2}{3} \\ -\frac{1}{3} \end{pmatrix}$
	u_R	c_R	t_R	0	$\frac{4}{3}$	$\frac{2}{3}$
	d_R	s_R	b_R	0	$-\frac{2}{3}$	$-\frac{1}{3}$
Leptons	$\begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}$	$\begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}$	$\begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}$	$\begin{pmatrix} \frac{1}{2} \\ -\frac{1}{2} \end{pmatrix}$	$\begin{pmatrix} -1 \\ -1 \end{pmatrix}$	$\begin{pmatrix} 0 \\ -1 \end{pmatrix}$
	e_R	μ_R	τ_R	0	-2	1

but the space-time derivatives of the fields do not. For the Dirac Lagrangian $\bar{q}_f\,(i\gamma^\mu\partial_\mu - m)\,q_f$ to be gauge invariant, ∂_μ has to be replaced by a *covariant derivative*:

$$D_\mu = \partial_\mu + i\,g_s\,\frac{\lambda_a}{2}\,G^a_\mu\ , \tag{1.2}$$

where a coupling constant g_s and 8 gauge vector fields G^a_μ corresponding to the 8 gluons have been introduced. These have to transform under $SU(3)_C$ as

$$G^a_\mu \ \longrightarrow\ G^a_\mu + \alpha^b(x)\,f^{abc}\,G^c_\mu + \frac{1}{g_s}\,\partial_\mu\alpha^a(x)\ , \tag{1.3}$$

where the f^{abc} are the structure constants of the group, defined in the commutation rules $[\frac{\lambda_a}{2}, \frac{\lambda_b}{2}] = i f_{abc}\frac{\lambda_c}{2}$ of the generators. There are always as many gauge fields as generators, and they live in the adjoint representation, here called **8**.

The full Lagrangian of QCD, like any gauge theory, can be written as:

$$\mathcal{L}_{\mathrm{QCD}} = \bar{q}_f\,(i\gamma^\mu\partial_\mu - m)\,q_f - g_s\,\bar{q}_f\,\gamma^\mu\,\frac{\lambda_a}{2}\,q_f\,G^a_\mu - \frac{1}{4}G^{\mu\nu}_a\,G^a_{\mu\nu}\ , \tag{1.4}$$

where a summation over all quark fields is implied. The first term is the free Dirac Lagrangian and the second one describes the interactions between matter and gauge fields. The last one is a gauge-invariant kinetic term that involves the field strength tensor of the gauge bosons:

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - g_s\,f^{abc}\,G^b_\mu\,G^c_\nu\ . \tag{1.5}$$

This introduces trilinear and quadrilinear terms in the Lagrangian, corresponding to interactions among the gauge bosons (here gluons), which is a consequence of non-abelianity. An important remark is that the gauge symmetry forbids the presence of mass terms for the gauge bosons in the Lagrangian.

Electroweak interaction

The unbroken electroweak interaction arises from gauge invariance under the $SU(2)_L \times U(1)_Y$ symmetry group. In the EW theory, fermions of opposite chirality are considered

as different fields and have different interactions. This forbids the presence of mass terms for fermions in the Lagrangian.

$SU(2)_L$ is a non-abelian group associated to weak isospin (I_3), with three generators that can be chosen as the Pauli matrices $\frac{\sigma_i}{2}$. The three gauge fields are denoted as W_μ^i ($i = 1, 2, 3$). As hinted in Table 1.1, fermions of left-handed chirality are in $SU(2)_L$ doublets (i.e. they live in the representation **2**), while the corresponding right-handed fermions are $SU(2)_L$ singlets and do not interact with the W_μ^i gauge fields.

$U(1)_Y$ is an abelian group associated to weak hypercharge (Y), analogous to the well-known $U(1)_{em}$ group of electromagnetism, but distinct from it. It acts on both chiral components of fermions, and its single gauge field is called B_μ. The weak hypercharge is linked to the electric charge by the relation $Q = I_3 + \frac{Y}{2}$.

Denoting the $SU(2)_L$ doublets as $L = \begin{pmatrix}\psi_L \\ \psi'_L\end{pmatrix}$, and the singlets as ψ_R and ψ'_R, a gauge transformation can be expressed as

$$L \longrightarrow e^{-i\,\alpha_i(x)\,\frac{\sigma_i}{2} - i\,\beta(x)\,\frac{Y}{2}} L\ , \qquad \psi_R^{(\prime)} \longrightarrow e^{-i\,\beta(x)\,\frac{Y}{2}} \psi_R^{(\prime)}\ , \tag{1.6}$$

which calls for two different versions of the covariant derivative:

$$\begin{array}{lrcl} L: & D_\mu & = & \partial_\mu + i\,g_w\,\frac{\sigma_i}{2}\,W_\mu^i + i\,g\,\frac{Y}{2}\,B_\mu\ , \\ \psi_R, \psi'_R: & D_\mu & = & \partial_\mu + i\,g\,\frac{Y}{2}\,B_\mu\ . \end{array} \tag{1.7}$$

Two coupling constants g_w and g have been introduced, as well as the 3+1 gauge fields, which transform as:

$$\begin{array}{rcl} W_\mu^i & \longrightarrow & W_\mu^i + \alpha^j(x)\,\epsilon^{ijk}\,W_\mu^k + \frac{1}{g_w}\,\partial_\mu \alpha^i(x)\ , \\ B_\mu & \longrightarrow & B_\mu + \frac{1}{g}\,\partial_\mu \beta(x)\ , \end{array} \tag{1.8}$$

where the ϵ^{ijk} originate from the commutation rules $[\frac{\sigma_i}{2}, \frac{\sigma_j}{2}] = i\epsilon_{ijk}\frac{\sigma_k}{2}$.

The physical weak boson fields $W_\mu^\pm$ and Z_μ and the photon field A_μ are linear combinations of the W_μ^i and B_μ gauge fields:

$$\begin{array}{rcl} W_\mu^\pm & = & \frac{1}{\sqrt{2}}\,(W_\mu^1 \mp i\,W_\mu^2)\ , \\ Z_\mu & = & W_\mu^3\,\cos\theta_w - B_\mu\,\sin\theta_w\ , \\ A_\mu & = & W_\mu^3\,\sin\theta_w + B_\mu\,\cos\theta_w\ , \end{array} \tag{1.9}$$

where θ_w is the weak mixing angle, defined as:

$$\cos\theta_w = \frac{g_w}{\sqrt{g_w^2 + g^2}}\ , \qquad \sin\theta_w = \frac{g}{\sqrt{g_w^2 + g^2}}\ . \tag{1.10}$$

The electroweak unification then simply materializes as $e = g\cos\theta_w = g_w \sin\theta_w$ where e is the electric charge.

The full electroweak Lagrangian can be expressed as:

$$\mathcal{L}_{\mathrm{EW}} = \bar{L}\,i\,\gamma^\mu D_\mu\,L + \bar{\psi^{(\prime)}}_R\,i\,\gamma^\mu D_\mu\,\psi_R^{(\prime)} - \frac{1}{4} W_i^{\mu\nu}\,W_{\mu\nu}^i - \frac{1}{4} B^{\mu\nu}\,B_{\mu\nu}\ , \tag{1.11}$$

where a summation over all $SU(2)_L$ doublets L and singlets $\psi_R^{(\prime)}$ is implied, and the field strength tensors are:

$$\begin{array}{rcl} W_{\mu\nu}^i & = & \partial_\mu W_\nu^i - \partial_\nu W_\mu^i - g_w\,\epsilon^{ijk}\,W_\mu^j\,W_\nu^k\ , \\ B_{\mu\nu} & = & \partial_\mu B_\nu - \partial_\nu B_\mu\ . \end{array} \tag{1.12}$$

Developing the two first terms of $\mathcal{L}_{\text{EW}}$ provides a massless Dirac Lagrangian, and interaction terms that lead to Feynman rules for the interaction of fermions with the γ, $\text{W}^{\pm}$, and Z^0 mediators of the electromagnetic and weak interactions. The two last terms of $\mathcal{L}_{\text{EW}}$ describe interactions among gauge bosons, which are trilinear (ZWW, γWW) and quadrilinear (ZZWW, $\gamma\gamma$WW, γZWW, WWWW).

1.1.2 The Brout–Englert–Higgs mechanism

A fundamental shortcoming of the above description of the electroweak theory is that it requires the weak vector bosons and the fermions to be massless, because introducing ad hoc mass terms such as $-M^2 W^\mu W_\mu$ and $-m\,(\bar{\psi}_R\psi_L + \bar{\psi}_L\psi_R)$ would violate the gauge invariance of the Lagrangian. However, the masses of $\text{W}^{\pm}/\text{Z}^0$ bosons can be generated by a procedure known as the Brout–Englert–Higgs (BEH) mechanism, which was first proposed in 1964 in three independent papers from Englert & Brout [4], Higgs [5], and Guralnik, Hagen & Kibble [6].

The BEH mechanism requires to postulate the existence of a new field which is symmetric under the gauge transformations, and acquires a non-zero expectation value in the vacuum state, breaking the electroweak symmetry. The simplest field that can be introduced is a $SU(2)_L$ doublet of complex scalar fields:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi^1 + i\,\phi^2 \\ \phi^3 + i\,\phi^4 \end{pmatrix} \ , \tag{1.13}$$

which is introduced in the Lagrangian density of the SM via the term

$$\mathcal{L}_{\text{BEH}} = (D^\mu\phi)^\dagger (D_\mu\phi) + V(\phi^\dagger\phi) \ , \tag{1.14}$$

where $D_\mu = \partial_\mu + i\,g_w\,\frac{\sigma_i}{2}\,W^i_\mu + i\,g\,\frac{Y}{2}\,B_\mu$ is the covariant derivative.

The potential is chosen as:

$$V(\phi^\dagger\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2 \ , \tag{1.15}$$

whereby $\lambda > 0$ is required for the vacuum to be stable, and $\mu^2 < 0$ is chosen so as to induce spontaneous symmetry breaking. The ground state is given by:

$$\phi^\dagger\phi = -\frac{\mu^2}{2\lambda} \ , \tag{1.16}$$

which corresponds to a manifold of non-zero values.
If one chooses to fix this vacuum state on the ϕ_0 axis:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \ , \ v^2 \equiv -\frac{\mu^2}{\lambda} \ , \tag{1.17}$$

it can be shown that this state is only invariant under $T_3 + \frac{Y}{2}$, or in other words, under the $U(1)_{em}$ symmetry. The field can be re-parameterized as:

$$\phi(x) = \frac{1}{\sqrt{2}} e^{i\,\theta^i(x)\,\sigma_i} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix} \ , \tag{1.18}$$

where four real scalar fields have been introduced: a massive field h, and three massless fields θ^1, θ^2, and θ^3. The apparition of the θ^i fields is actually predicted by the Goldstone theorem [7], which states that the spontaneous breaking of a continuous symmetry

conjures as many massless scalar fields (*Goldstone bosons*) as there are broken generators. Here, the $SU(2)_L \times U(1)_Y$ symmetry into the $U(1)_{em}$ one indeed implies the existence of $4 - 1 = 3$ Goldstone bosons. However, none of these have been observed.

At this point, the BEH mechanism consists in taking advantage of gauge invariance, by performing the local transformation:

$$\phi(x) \longrightarrow e^{-i\,\alpha_i(x)\,\frac{\sigma_i}{2}}\,\phi(x) \tag{1.19}$$

with $\alpha_i(x) = 2\theta^i(x)$, which eliminates the θ^i fields, actually turning their three degrees of freedom into the longitudinal degrees of freedom of the now massive $\mathrm{W}^{\pm}$ and Z^0 vector bosons. This leaves the vacuum state containing a single scalar field h, the quanta of which are associated with a physical massive particle called the *Higgs boson* (H).

Injecting the expression of $\phi(x)$ and of the covariant derivative into Eq. 1.14 and neglecting the constant term leads to:

$$\begin{aligned}\mathcal{L}_{\mathrm{BEH}} \;=\; & \tfrac{1}{2}\,\partial_\mu h\,\partial^\mu h + \mu^2 h^2 \\ & + \tfrac{g_w^2\,v^2}{4}\,W_\mu^-\,W^{+\mu} + \tfrac{g_w^2\,v^2}{8\cos^2\theta_w}\,Z_\mu\,Z^\mu \\ & + \tfrac{g_w^2\,v}{2}\,h\,W_\mu^-\,W^{+\mu} + \tfrac{g_w^2}{4}\,h^2\,W_\mu^-\,W^{+\mu} + \tfrac{g_w^2\,v}{4\cos^2\theta_w}\,h\,Z_\mu\,Z^\mu + \tfrac{g_w^2}{8\cos^2\theta_w}\,h^2\,Z_\mu\,Z^\mu \\ & + \tfrac{\mu^2}{v}\,h^3 + \tfrac{\mu^2}{4v^2}\,h^4\ .\end{aligned} \tag{1.20}$$

The first line provides the mass of the Higgs boson: $m_{\mathrm{H}} = \sqrt{2}|\mu|$, which is a free parameter of the theory.
Using the second line, the $W_\mu^{\pm}$ fields are identified as charged under $U(1)_{em}$, with masses $m_{\mathrm{W}^+} = m_{\mathrm{W}^-} = \frac{1}{2}\,g_w\,v$, while the neutral Z^0 boson acquires a mass $m_{\mathrm{Z}^0} = \frac{g_w\,v}{2\cos\theta_w}$. The absence of an $A_\mu A^\mu$ term confirms that the photon remains massless.
The third and fourth line contain trilinear and quadrilinear terms that describe the couplings of the Higgs boson to weak vector bosons (HWW, HZZ, HHWW, and HHZZ) and its self-couplings (H^3 and H^4).
The scale of electroweak symmetry breaking can be determined noticing that the W boson mass links the v parameter to the Fermi constant G_F:

$$v = 2\frac{m_{\mathrm{W}}}{g_w} = (\sqrt{2}\,G_F)^{-1/2} \simeq 246\,\mathrm{GeV}\ . \tag{1.21}$$

Besides providing a description of electroweak symmetry breaking, the field ϕ is also used to extend the SM Lagrangian with gauge-invariant *Yukawa terms* that generate masses for the fermions. Such terms have to be proportional to $\bar{L}\,\phi\,\psi'_R + \bar{\psi'}_R\,\phi\,L$ for down-type fermions (here overlooking fermion mixing for simplicity), and to $i\,\bar{L}\,\sigma_2\,\phi^*\,\psi_R + i\,\bar{\psi}_R\,\sigma_2\,\phi^*\,L$ for up-type fermions. After electroweak symmetry breaking, the additional terms in the Lagrangian take the form:

$$\mathcal{L}_{\mathrm{Yukawa}} = \sum_f -m_f\,\bar{\psi}\,\psi\,\left(1 + \frac{h}{v}\right) + \sum_{f'} -m_{f'}\,\bar{\psi'}\,\psi'\,\left(1 + \frac{h}{v}\right)\ , \tag{1.22}$$

where the first and second sum runs over up-type and down-type fermions, respectively. Thus, besides providing the needed fermion mass terms, this procedure implies that the Higgs boson field interacts with fermions, with couplings proportional to their masses. It should however be noted that these masses are not predicted by the theory, and that the overall symmetry of the theory gets reduced here since the three generations of matter no longer appear as identical.

Finally, the introduction of the Higgs boson also has the virtue of ensuring the calculability of the SM. In particular, the perturbative unitarity of the scattering matrix is preserved at high energies, as the longitudinal W/Z boson scattering amplitude no longer grows as the centre-of-mass energy increases.

1.2 Higgs boson phenomenology at a pp collider

1.2.1 Production mechanisms

At the Large Hadron Collider (LHC), the SM Higgs boson is produced in a variety of processes [8], five of which are studied in detail in this thesis work. As shown in the leading order Feynman diagrams of Fig. 1.1, their final states can involve other particles than just the Higgs boson, hence leading to different experimental signatures that help extracting these processes.

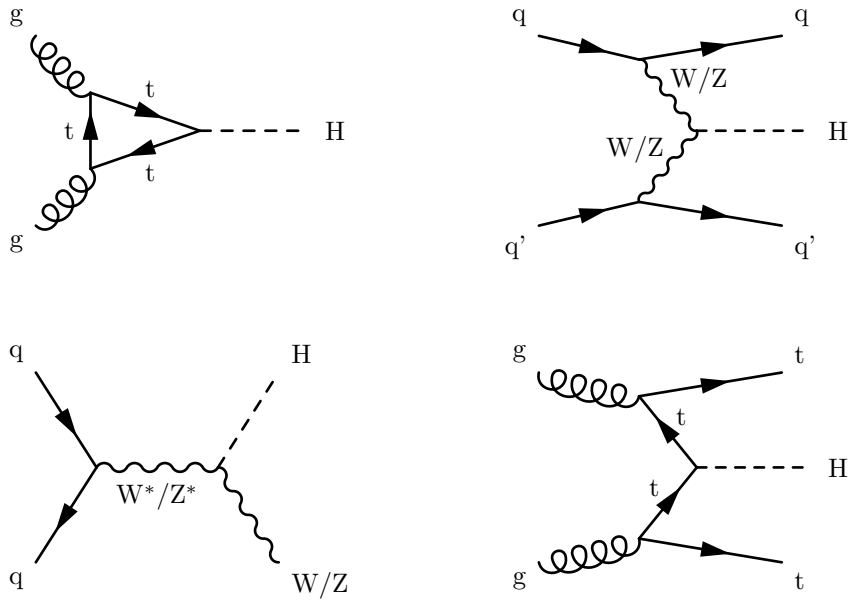

Figure 1.1: Leading order Feynman diagrams for the main production modes of the SM Higgs boson at the LHC: gluon fusion (top left), vector boson fusion (top right), VH associated production (bottom left) and $t\bar{t}H$ associated production (bottom right).

Gluon fusion

The gluon fusion process (hereafter denoted as ggH) is induced by a pair of gluons that fuse into a Higgs boson via an intermediate quark loop, as illustrated in Fig. 1.1 (top left). As the most massive quark, the top quark gives the largest contribution to the loop, but heavy undiscovered particles could also contribute to it. Since the gluon luminosity is very large in pp collisions at the high centre-of-mass energies provided by the LHC, ggH is the leading Higgs boson production mechanism and dominates all the other ones by more than one order of magnitude.
Higher-order QCD corrections converge slowly and have been found to increase the LO

ggH cross section by as much as a factor 2–2.5, which makes it crucial to use state-of-the-art computations. Recent N3LO QCD + NLO EW computations [9, 10] are used throughout this thesis work. As a consequence of the magnitude of these QCD corrections, ggH final states with one or more jets are frequent, and have to be treated as backgrounds in searches for the hereafter described subdominant modes.

Vector boson fusion

Vector boson fusion (VBF) is the second production mode at the LHC, its cross section being an order of magnitude smaller than that of ggH.
At leading order, it is a $qq' \to qq'H$ process where two vector bosons are radiated off quarks and merge into a Higgs boson (Fig. 1.1, top right). This results in a very clean experimental signature featuring two forward and backward energetic jets with a large dijet invariant mass, while the Higgs boson decay system is boosted and ends up in a more central region of the detector. Moreover, since no colour is exchanged, central hadronic activity is also suppressed. This characteristic topology helps rejecting backgrounds from SM processes and ggH production in association with two jets. W boson fusion and Z boson fusion cannot be distinguished experimentally.
The VBF cross section is known to a good accuracy, with higher-order QCD corrections being smaller than for ggH. It becomes larger than the ggH one for very high Higgs boson mass hypotheses, in a mass range above 1 TeV, relevant to the searches for additional non-standard high-mass scalar resonances [10].

Associated production with a vector boson

Associated production with a vector boson, which is often referred to as *VH associated production* or *Higgsstrahlung*, and further split into WH and ZH production, is the third most prominent mechanism at the LHC (about twice less frequent than VBF).
In this $q\bar{q}$-induced process, the Higgs boson is radiated off a W or Z boson (Fig. 1.1, bottom left), leading to an experimental signature where the Higgs boson decay products are boosted, and accompanied by the products of the associated W or Z. When this vector boson decays hadronically, a pair of nearby boosted jets with invariant mass close to the nominal $m_{W^\pm}$ or m_{Z^0} can be sought for, whereas leptonic decays either provide one lepton and missing transverse energy (for WH), or either a pair of leptons or missing transverse energy (for ZH).
The high-order QCD corrections are quite large for this process; like for VBF, the cross sections are currently computed to NNLO QCD and NLO EW accuracy [10].

Associated production with a top quark pair

Associated production with a top quark pair ($t\bar{t}H$) is about 100 times rarer than ggH, but it is the first production mode that allows direct probing of the Yukawa coupling between the top quark and the Higgs boson, which is key to SM precision tests. Indeed, although the $Ht\bar{t}$ vertex is already involved in gluon fusion through the top quark loop, one cannot exclude that an unknown heavy fundamental particle contributes to the loop.
In this mainly gluon-induced $t\bar{t}H$ process, the Higgs boson is accompanied by a $t\bar{t}$ pair in the final state, as illustrated in Fig. 1.1 (bottom right). Each associated top quark then decays into a bottom quark and a $W^\pm$ boson which can in turn decay leptonically or hadronically, leading to several possible experimental signatures and extraction strategies.

Another production mode, associated production with a bottom quark pair ($b\bar{b}H$), has a cross section of a similar order of magnitude to $t\bar{t}H$ at $\sqrt{s} = 13\,\text{TeV}$, but this elusive process is not taken into account in the remainder of this work, because no Monte Carlo simulated samples were available yet to study it. Other Higgs boson production mechanisms have too low cross sections to be even considered here.

Figure 1.2 presents the centre-of-mass-energy dependency of the cross sections of the main production modes at the LHC [10]. Table 1.2 numerically shows by how much the cross sections of the five processes studied in this thesis increase from 8 TeV to 13 TeV, i.e. from Run I to Run II of the LHC. Most of them gain a factor of about 2, except $t\bar{t}H$ which benefits from a factor 4 because of the large mass of the involved objects.

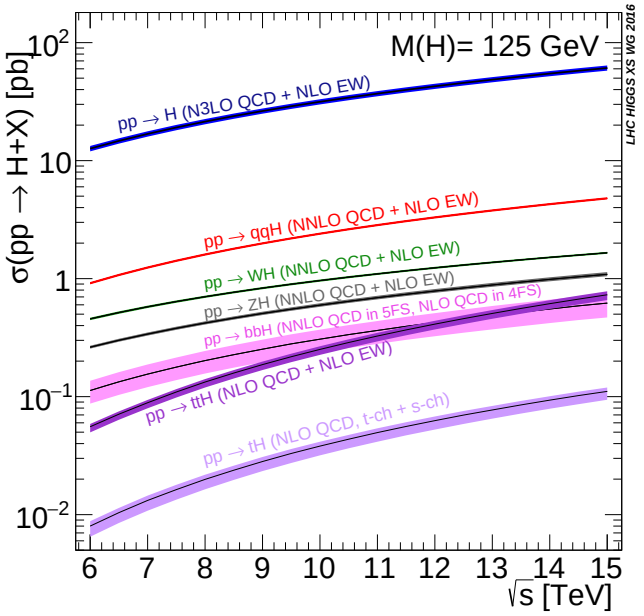

Figure 1.2: Total production cross section for a 125 GeV SM Higgs boson at the LHC as a function of the centre-of-mass energy.

Table 1.2: Variations of the cross sections of the main Higgs boson production modes at the LHC for $m_H = 125\,\text{GeV}$, moving from late Run I (8 TeV) to Run II (13 TeV) [10].

Process	Computation accuracy	$\sigma_{8\,\text{TeV}}$	$\sigma_{13\,\text{TeV}}$	$\sigma_{8\,\text{TeV}}/\sigma_{13\,\text{TeV}}$
ggH	N3LO QCD, NLO EW	21.42 pb	48.58 pb	2.27
VBF	NNLO QCD, NLO EW	1.601 pb	3.782 pb	2.36
WH	NNLO QCD, NLO EW	0.7026 pb	1.373 pb	1.95
ZH	NNLO QCD, NLO EW	0.4208 pb	0.8839 pb	2.10
$t\bar{t}H$	NLO QCD, NLO EW	0.1330 pb	0.5071 pb	3.81

1.2.2 Decay modes

Since the SM Higgs boson directly couples to all massive particles of the standard model and can also couple to massless particles via intermediate loops, it can decay in a variety of channels. The total decay width of the Higgs boson and the relative branching fractions of its decay channels are fully determined by the value m_{H} of its mass. Figure 1.3 presents the values of the branching fractions as a function of the hypothesized m_{H} [11], thus illustrating the strategic issues that existed before the discovery. Five main decay channels were studied at that time, and are still being exploited these days. Their respective relevancy to the discovery and to the property measurements does not only depend on their branching fraction, but also on the experimental capability of extracting the corresponding signals while rejecting their backgrounds.

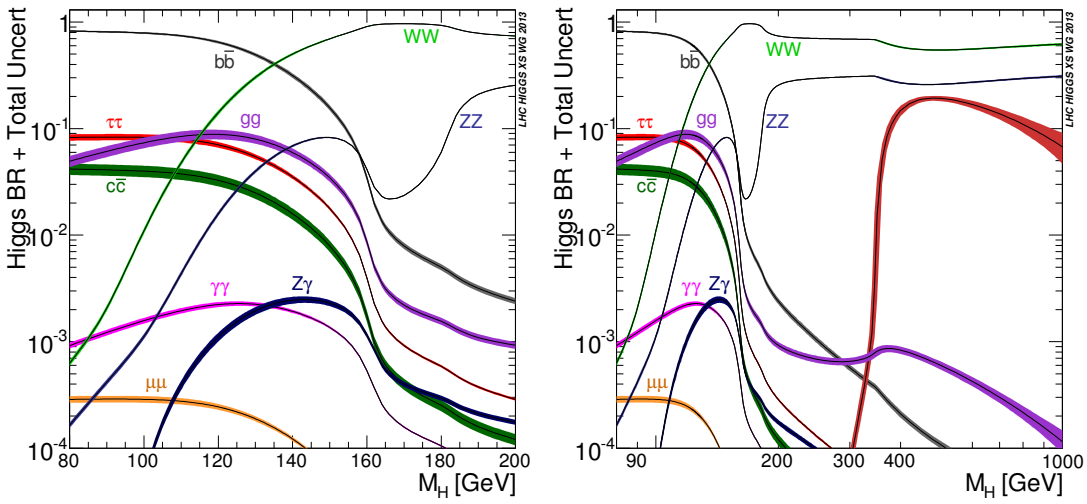

Figure 1.3: Decay branching ratios with uncertainties for the SM Higgs boson as a function of m_{H}, in a low-mass range (left) and in a range extending to high mass (right).

For values of m_{H} up to about 135 GeV, the Higgs boson mainly decays into a $\mathrm{b\bar{b}}$ pair, but the inclusive signal is overwhelmed by the QCD production of bottom quarks. Therefore, this channel is exploited in boosted regimes, mainly in the VH production mode with the associated W or Z boson decaying leptonically.

The sensitivity at masses below 120 GeV is actually dominated by a channel with low branching fraction, namely the decay to two photons[1]. Indeed, if a good experimental resolution can be achieved for the diphoton invariant mass, the Higgs boson signal appears as a clear peak on top of the backgrounds from QCD production of two photons or jet fragments misidentified as photons. This has made $\mathrm{H} \to \gamma\gamma$ one of the two so-called golden channels for the discovery.

As the hypothesized m_{H} increases, decays to pairs of massive gauge bosons ($\mathrm{H} \to \mathrm{WW}^{(*)}$ and $\mathrm{H} \to \mathrm{ZZ}^{(*)}$) open up and their branching ratios grow. The WW channel has the largest branching ratio of all as of 135 GeV, and it particularly dominates the ZZ channel around the WW mass threshold. Accurate lepton identification and missing transverse

[1]Since photons are massless, the Higgs boson does not couple directly to them, but the $\mathrm{H} \to \gamma\gamma$ decay happens through a fermion loop or W boson loop.

energy reconstruction have made the $WW \to \ell\nu\ell\nu$ decay a sensitive channel for Higgs boson searches at intermediate masses, although the m_H resolution is poor because of the escape of neutrinos. By contrast, the ZZ channel can offer a complete reconstruction of the final state, with excellent mass resolution. The privileged $H \to ZZ \to 4\ell$ channel, where both Z bosons decay to pairs of either electrons or muons, is the focus of this thesis. It constitutes the other golden channel for the discovery and property measurements, and its characteristics will be discussed in detail from Section 1.4 onwards.

Three other channels contribute significantly to the decay width at 125 GeV and up to masses near 150 GeV. The decay to a $\tau^+\tau^-$ pair provides intermediate sensitivity and is probed with a variety of experimental strategies, depending on the decays of the tau leptons. The decays to a pair of gluons or to a pair of charm quarks are generally not exploited, because their experimental signatures cannot be distinguished from the overwhelming QCD dijet production. Finally, the decay to a $t\bar{t}$ pair opens up around $2m_t$, and is hence relevant to searches of new high-mass resonances.

1.3 Previous Higgs boson searches and discovery

To guide Higgs boson searches at colliders, the allowed range for the Higgs boson mass m_H has been constrained by theoretical arguments [12, 8], in particular by imposing the energy scale Λ up to which the SM is valid. An upper limit called *triviality* is obtained by requiring that the running quartic coupling λ of the potential V remains finite up to the scale Λ. A lower limit comes from *vacuum stability*, requiring that λ remains positive after including radiative corrections, at least up to Λ, which implies that the minimum of the potential is absolute. A looser *metastability* constraint is found by requiring the minimum to just be local. If Λ were of the order of the Planck scale ($\approx 10^{19}$ GeV), which is now disproved, then m_H would be constrained between 130 and 170 GeV. For $\Lambda \simeq 1$ TeV, the allowed range goes up to 700 GeV. It should be noted that another long-known upper bound of about 710 GeV arises from the requirement of the unitarity of the scattering matrix.

Direct Higgs boson searches were first carried out at the Large Electron Positron (LEP) collider exploiting the $e^+e^- \to HZ$ Higgsstrahlung process, which led to a lower bound of $m_H > 114.4$ GeV at a 95% confidence level (CL) [13]. At the Tevatron proton–antiproton collider, early measurements excluded the mass range 162–166 GeV, exploiting the WW decay channel [14]. Independently from direct searches, precision electroweak measurements using data from different colliders provided an indirect constraint on the SM Higgs boson mass, with an upper limit of 158 GeV at a 95% CL [15].

At the LHC, the first direct Higgs boson searches were based on data from pp collisions collected in 2011 at a centre-of-mass energy of $\sqrt{s} = 7$ TeV, which amounted to an integrated luminosity of 5.1 fb^{-1}. Using the five aforementioned decay channels, the CMS Collaboration excluded a range of masses from 127 to 600 GeV at a 95% CL [16], while the ATLAS Collaboration excluded the ranges 111.4–116.6 GeV, 119.4–122.1 GeV, and 129.2–541 GeV at a 95% CL [17]. Within the remaining allowed mass region, both experiments reported an excess of events between 2 and 3 standard deviations near 125 GeV.

In 2012, the pp centre-of-mass energy of the LHC was raised to $\sqrt{s} = 8$ TeV, and an additional data sample of around 5.3 fb^{-1} was collected by the end of June by each of the two experiments, allowing elucidation of the last non-excluded mass region. On 4th July

2012, the ATLAS and CMS collaborations concomitantly reported the observation of a new boson with mass near 125 GeV, compatible with the SM Higgs boson [18, 19, 20]. That same month, the CDF and D0 experiments published an analysis of the full Tevatron data sample and reported an excess of events of about 3 standard deviations in the range 120–135 GeV [21, 22], consistent with the LHC observations.

The remainder of the 2012 data taking increased the integrated luminosity of 8 TeV data to approximately $20\,\mathrm{fb}^{-1}$ per experiment. The full Run I LHC data sample, both at $\sqrt{s} = 7\,\mathrm{TeV}$ and 8 TeV, was thoroughly analysed in various decay channels and production modes by both collaborations, shedding more light on the newly discovered boson. Improved calibrations allowed for a precise measurement of its mass using the two high-resolution channels [23, 24, 25, 26], and ultimately combining data from both experiments [27], leading to a value of $m_{\mathrm{H}} = 125.09 \pm 0.21(\mathrm{stat.}) \pm 0.11(\mathrm{syst.})\,\mathrm{GeV}$. More importantly, the new boson was soon shown to have spin-parity $J^P = 0^+$ [28], and its production and decay rates and its coupling strengths to SM particles turned out to be consistent with expectations for the SM Higgs boson [26, 29]. The combined analyses of ATLAS and CMS data strengthened the fact that the measured properties are consistent across channels and experiments, and with the SM predictions [30].

1.4 The $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ channel

The present thesis work exploits the decay of the Higgs boson to two Z bosons that in turn decay to pairs of electrons or muons, a channel which will be either denoted as $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ or as $\mathrm{H} \to 4\ell$. Background sources to this process include an irreducible four-lepton contribution from the production of ZZ or $\mathrm{Z}\gamma^*$ pairs via $\mathrm{q\bar{q}}$ annihilation and gluon fusion, and a reducible contribution from processes such as $\mathrm{Z} + \mathrm{b\bar{b}}$ and $\mathrm{t\bar{t}}$ production that involve two prompt leptons and two secondary leptons from b-quark jets. The 4ℓ channel involves three possible final states, hereafter referred to as 4e, 4μ, and $2\mathrm{e}2\mu$; since these have different reducible background rates and mass resolutions, they are analysed separately.

Among the five main Higgs boson decay channels that have been exploited in Run I of the LHC, $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ has garnered a particular interest. In spite of its low branching ratio (0.01240% for a 125.0 GeV SM Higgs boson), it has been the most sensitive channel for the discovery, thanks to three assets:

- the complete reconstruction of its final state, which provides the very discriminating four-lepton invariant mass variable (denoted as $m_{4\ell}$) as well as powerful angular information,
- the good momentum resolution of electrons and muons in the LHC detectors, which makes the resonance peak in the $m_{4\ell}$ distribution narrow,
- its very large signal-to-background ratio, which is typically of the order of 2 to 1 in a 5 GeV window around the peak.

Building on this, $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ has been a key contributor to many property measurements of the Higgs boson. Using the full Run I data sample, the CMS and ATLAS collaborations have measured the following parameters:

- the mass of the boson [23, 24, 31],

- a signal strength parameter, as well as two coupling parameters controlling production modes related to fermions [23, 31],
- its spin-parity quantum numbers, establishing the boson as a scalar particle [28, 32, 23, 33, 34],
- its decay width, with a first loose upper bound from a direct method [23, 24], followed by a much more stringent one from a combined analysis of the on-shell and off-shell regions [35, 36] (also phrased by ATLAS as a constraint on the off-shell signal strength), and finally a lower bound from a study of the flight distance within the CMS detector [37],
- its fiducial cross section at $\sqrt{s} = 8\,\text{TeV}$ [38, 39] and at $\sqrt{s} = 7\,\text{TeV}$ [39], both integrated and differential as a function of a few parameters (which differ between ATLAS and CMS),
- anomalous contributions to the tensor structure of the interaction between the new boson and the SM vector bosons [33, 34, 37].

Moreover, the exploitation of the four-lepton decay has gone further than the study of the new boson. Searches for an additional heavy Higgs boson at high mass were performed with the full LHC Run I data sample, and were combined with other $H \to ZZ$ channels in the ATLAS analysis [40], plus $H \to WW$ channels in the CMS analysis [41].

This thesis takes over from part of the above CMS studies, in the sense that it aims at exploiting pp collision data of Run II of the LHC and the $H \to ZZ \to 4\ell$ channel to perform the first measurements of properties of the new boson at the larger centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$. New measurements of the mass, signal strength and decay width will be presented, and particular emphasis will be put on a study that prepares future measurements of the Higgs boson couplings, namely the search for the rare production modes mentioned in Section 1.2.1, i.e. vector boson fusion, VH associated production and $t\bar{t}H$ associated production. A first high-mass resonance search at $\sqrt{s} = 13\,\text{TeV}$ will also be presented. To initiate the description of the long process that leads to such results, Chapter 2 will introduce the experimental apparatus.

Chapter 2

Experimental apparatus

The European Organization for Nuclear Research (CERN) was created in 1954 by twelve countries of Western Europe, as a laboratory devoted to the study of atomic nuclei. Notwithstanding its name, it soon moved to catalysing the studies of interactions between subnuclear particles, developing the field that came to be referred to as particle physics or high-energy physics. Over its six-decade history, CERN has hosted a number of particles accelerators and experiments, playing a major role in the construction of the standard model of particle physics. The major milestones of this road toward higher energies include the discovery of neutral currents in the Gargamelle bubble chamber in 1973, the discovery of the W and Z bosons in the UA1 and UA2 experiments in 1983, and the discovery of a Higgs boson by the ATLAS and CMS experiments in 2012.

This thesis is carried out within the CMS experiment, exploiting data from proton-proton (pp) collisions delivered by the CERN Large Hadron Collider (LHC). The present chapter provides an introduction to the LHC and CMS apparatus, and to the baseline algorithms for particle and jet reconstruction that are used centrally in CMS.

2.1 Large Hadron Collider

The Large Hadron Collider [42, 43, 44] is a two-ring superconducting hadron accelerator and collider. It was designed to collide proton beams with a nominal centre-of-mass energy of $\sqrt{s} = 14\,\text{TeV}$ (i.e. 7 TeV per beam) and an instantaneous luminosity of $10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$, and heavy ion beams (lead nuclei) with a nominal energy of 2.76 TeV per nucleon and a peak luminosity of $10^{27}\,\text{cm}^{-2}\,\text{s}^{-1}$, making it the most powerful machine ever built for particle physics research. Figure 2.1 gives an overview of the unprecedented event rates of known SM processes that are expected in pp collisions at the LHC, as compared to the Tevatron. The magnitude of the nominal 14 TeV energy was not only chosen to probe the scalar sector over the whole Higgs boson mass range allowed by the unitarity constraint, but also to probe new physics in the hypothesis of an absence of any Higgs boson, a case where unitarity in WW scattering would have required new physics to appear at a scale $\Lambda \lesssim 1.2\,\text{TeV}$.

2.1.1 Design and detectors

The LHC is installed near Geneva in a 26.7 km-long subcircular tunnel passing through Switzerland and France, which was constructed between 1984 and 1989 for an earlier CERN collider, the Large Electron Positron (LEP). The LHC project was given a first approval by the CERN Council in 1994, and updated two years later to its definitive

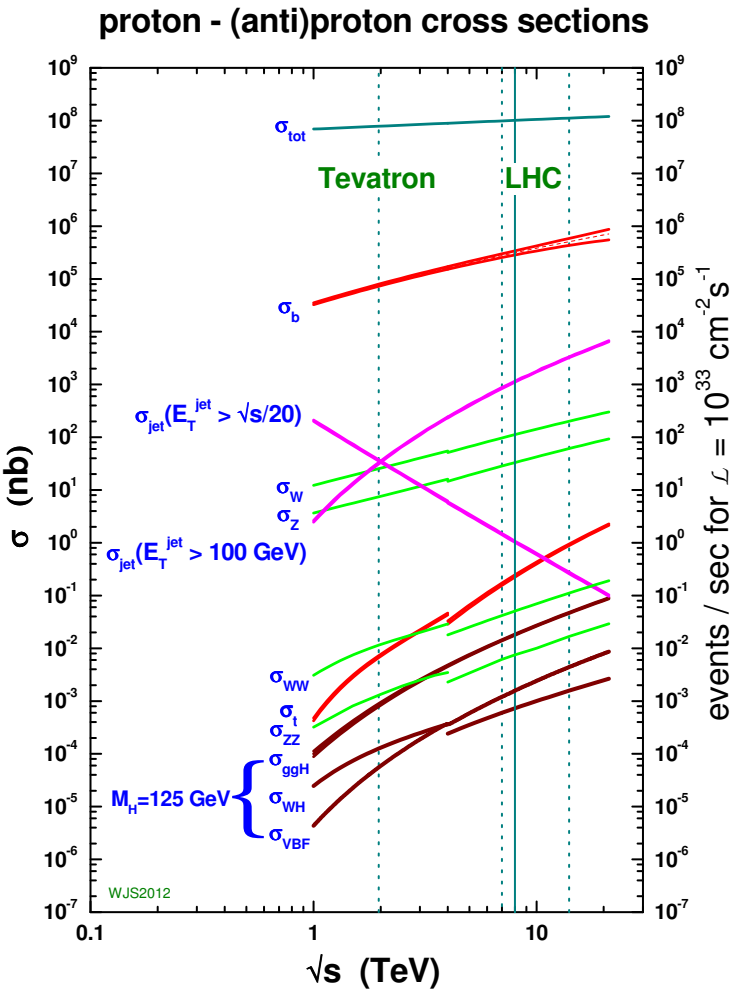

Figure 2.1: Expected production cross sections and event rates for signal and background processes at hadron colliders, as a function of the centre-of-mass energy [45]. Discontinuities are due to the Tevatron being a proton-antiproton collider while the LHC is a proton-proton collider.

form of a 14 TeV machine. The closedown of LEP in 2000 then liberated the tunnel for the LHC construction. As a particle-particle collider, the LHC uses two separate parallel rings with counter-rotating beams. Owing to limited space in the 3.7-metre-large arc section of the LEP tunnel, these two rings are included in one single twin-bore magnet system. They intersect at four interaction points. The beams are bent by 1232 15-metre-long dipole magnets made of copper-clad niobium-titanium. They are kept focused by 392 quadrupole magnets, each 5–7 metres long, while some stronger quadrupole electromagnets squeeze them further close to the intersection points to maximize the probability of interaction. Superfluid helium-4 is used to cool the magnets and maintain them at their operating temperature of 1.9 K (−271.25°C).

Prior to being injected into the LHC, proton beams are prepared by a chain of pre-accelerators that increase the energy in steps. These systems are visible in the diagram

of the CERN accelerator complex shown in Fig. 2.2. Protons are first accelerated to an energy of 50 MeV in the Linear Accelerator (LINAC2), which feeds the Proton Synchrotron Booster (PSB), where they are accelerated to 1.4 GeV. They then reach 26 GeV in the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS) further increases their energy to 450 GeV before they are injected into the LHC ring at Point 2 and Point 8. In the LHC, the acceleration process and shaping of proton bunches is done by 16 radiofrequency cavities, where the electromagnetic field oscillates at 400 MHz.

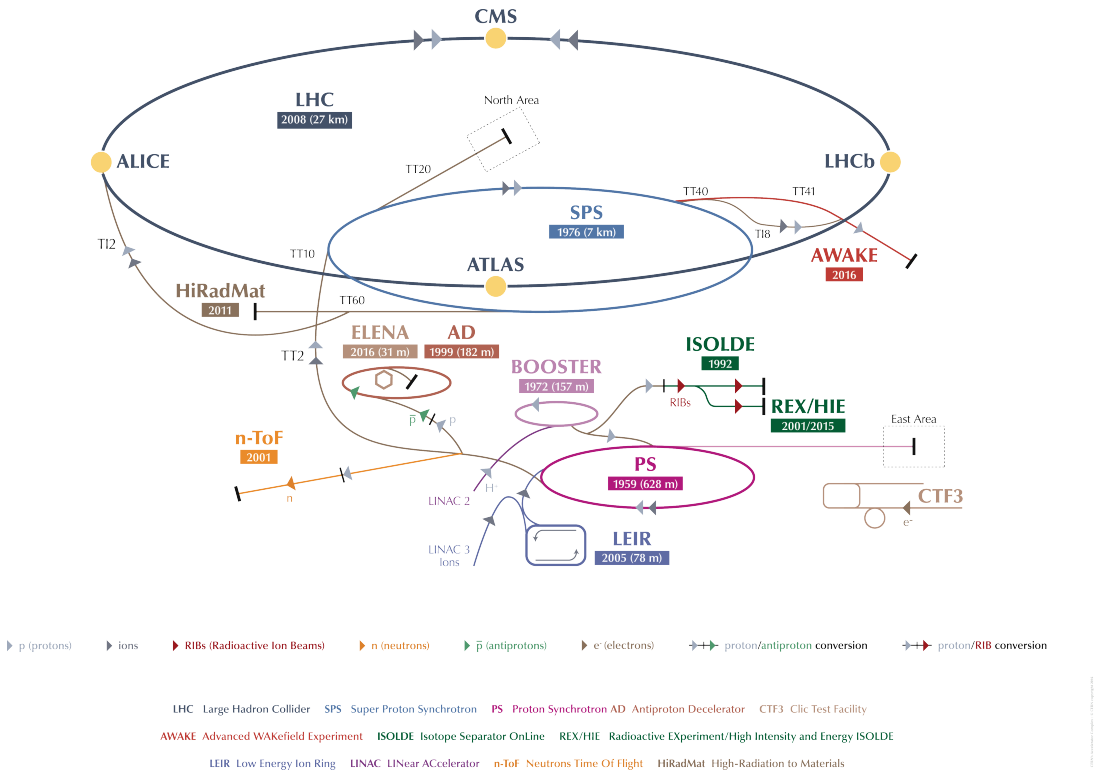

Figure 2.2: The CERN accelerator complex [46]. The proton injection chain for the LHC starts from the LINAC2 and proceeds through the Booster, PS, and SPS.

The machine instantaneous luminosity only depends on the beam parameters, and can be written as:

$$L = \frac{f_{\text{rev}}\, N_b^2\, n_b\, \gamma_r}{4\pi\, \epsilon_n\, \beta^*} F \ , \tag{2.1}$$

where N_b is the number of particles per bunch, n_b is the number of bunches per beam, $f_{\text{rev}} = 11\,\text{kHz}$ is the revolution frequency, ϵ_n the normalized transverse beam emittance, β^* is the beta function at the collision point, which measures the beam focalization and is corrected by the relativistic gamma factor γ_r, and F is a geometric luminosity reduction factor that accounts for the crossing angle at the interaction point [44]. The nominal value of the beta function is $\beta^* = 0.55\,\text{m}$, and the nominal luminosity is reached with $n_b = 2808$ bunches per beam, and $N_b = 1.15 \times 10^{11}$ protons per bunch. This choice corresponds to a spacing of 25 ns (or 7.5 m) between bunches, i.e. a bunch crossing rate of 40 MHz.

The record-breaking collision parameters of the LHC machine have important consequences on the design of detectors. Under nominal conditions, the number of inelastic

collision events is of the order of 10^9 per second, with approximately 20 collisions per bunch crossing. The occurrence of such overlapping pp interactions is called *pileup*, and calls for a high granularity to distinguish particles from different interactions. Moreover, fast response and good time resolution are needed in order to distinguish events from consecutive crossings, i.e. avoiding the phenomenon of overlap of consecutive signals called *out-of-time pileup*.

Four main particle detectors are installed in underground caverns at the four beam intersection points. The two largest ones, *A Toroidal LHC ApparatuS* (ATLAS) and *Compact Muon Solenoid* (CMS) are located at the symmetrically opposite Point 1 and Point 5 of the LHC ring, respectively, where the provided instantaneous luminosities are expected to be highest. ATLAS and CMS are multi-purpose experiments, designed to cover a wide physics programme in the scalar, electroweak, and strong sectors, with optimized sensitivity for Higgs boson searches and possible physics beyond the standard model (BSM) at the TeV scale. The two other detectors have more specific goals: *Large Hadron Collider beauty* (LHCb, located at Point 8) is aimed at studying CP violation in B-hadron interactions using lower peak luminosities ($\approx 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$), while *A Large Ion Collider Experiment* (ALICE, Point 2) is dedicated to heavy-ion collisions, studying quark-gluon plasma.

2.1.2 Operations

The first injections of proton beams into the LHC were carried out on 10^{th} September 2008, but the initial testing was interrupted on the 19^{th} of that month by a major magnet quench incident due to a faulty electrical connection between two magnets, which caused extensive damage to over 50 superconducting magnets and delayed operations for more than one year. The injections started over in late November 2009, and were shortly followed by the first high-energy collisions. Beam energies were progressively increased, reaching 3.5 TeV in March 2010, when the first physics runs started. It was decided to not immediately aim at the design LHC beam parameters yet, and to only operate the collider with a 50 ns bunch spacing (1404 bunches) and intermediate centre-of-mass energies until 2012, hoping to discover the long-sought Higgs boson in this first data taking era referred to as *Run I*. Design energies would have to wait for *Run II* in 2015, after a two-year upgrade period called *Long Shutdown 1* (LS1).

Following a first data sample of $47\,\mathrm{pb}^{-1}$ in 2010 at $\sqrt{s} = 7\,\mathrm{TeV}$, the LHC delivered a high-luminosity data set of about $6\,\mathrm{fb}^{-1}$ during the 2011 runs, a large fraction of which was collected by ATLAS and CMS, allowing the exclusion of most of the allowed range for the Higgs boson mass (see Section 1.3). Meanwhile, instantaneous luminosities had been steadily growing, reaching above $10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. This can be seen in Fig. 2.3 which shows the cumulative integrated luminosities and peak instantaneous luminosities delivered to the CMS experiment in every year of data taking. In early 2012, it was decided to increase the centre-of-mass energy to 4 TeV per beam for the year 2012. By the end of that year, the LHC had delivered a data sample corresponding to an integrated luminosity of $23\,\mathrm{fb}^{-1}$ at $\sqrt{s} = 8\,\mathrm{TeV}$, whereby the Higgs boson discovery was already accomplished using the first $\approx 5\,\mathrm{fb}^{-1}$. The exploitation of the large collected data samples went on during LS1, providing a quantity of precision measurements in the electroweak and strong sectors.

LS1 brought considerable upgrade and consolidation to the LHC, with the goal of starting Run II in 2015 at a centre-of-mass energy of $\sqrt{s} = 13\,\mathrm{TeV}$, and of soon reaching beyond-design instantaneous luminosities of $1.4 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Although not reaching

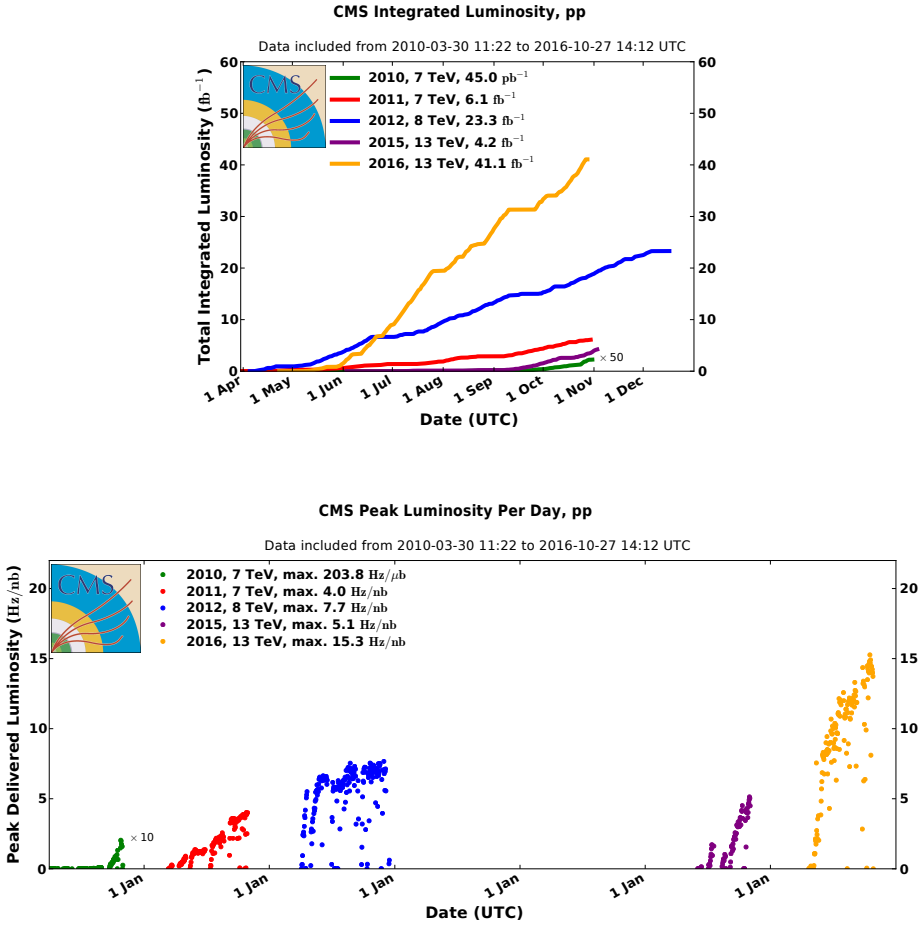

Figure 2.3: (top) Cumulative integrated luminosity versus day and (bottom) peak instantaneous luminosity versus day, delivered to CMS during stable beams for pp collisions, for the 2010 (green), 2011 (red), 2012 (blue), 2015 (purple), and 2016 (orange) data taking.

14 TeV, the energy increase substantially extends the physics reach of the experiments with respect to 2012, as evidenced by the ratios of parton luminosities illustrated in Fig. 2.4. After several months of training of the dipole magnets, the first Run II beam was injected on 5$^{\text{th}}$ April 2015, and the first 6.5 TeV beam was obtained on 10$^{\text{th}}$ April. The year 2015 was intended as a commissioning year and started with 50 ns collisions, before moving to the nominal bunch spacing of 25 ns over the summer. Figure 2.3 (bottom) illustrates the progressive luminosity ramp-up; a record number of bunches was reached on 24$^{\text{th}}$ October, with an instantaneous luminosity of $5.2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$. In total, a data sample corresponding to more than $4\,\text{fb}^{-1}$ was delivered in 2015 to ATLAS and CMS, as illustrated in Fig. 2.5 (left) for the case of CMS.

2016 started as a production year, still at $\sqrt{s} = 13\,\text{TeV}$ and with a bunch spacing of 25 ns, but the experience gained in 2015 led to now choose a relatively bold set of operational parameters. Even though some problems with the SPS beam dump made it necessary to restrain the number of bunches to about 2100, the nominal LHC luminosity of $10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ was reached in June 2016, notably thanks to a reduction of the β^* parameter from 80 cm to 40 cm. Instantaneous luminosities further grew over the

summer, reaching more than $1.4 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The 2016 pp data taking ended in late October, with a total delivered integrated luminosity of about $41\,\mathrm{fb}^{-1}$, as illustrated in Fig. 2.5 (right).

This thesis work exploits both the 2015 data sample and the part of the 2016 one that was delivered up to mid-July 2016. Characteristics of these samples relevant to the $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ analysis will be discussed in Sections 6.1 and 7.1, respectively.

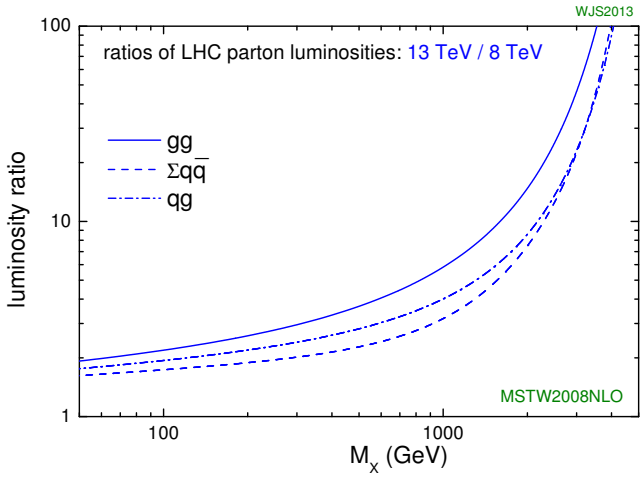

Figure 2.4: Ratios between parton luminosities at the LHC using 13 TeV and 8 TeV centre-of-mass energies, as a function of the mass of the final-state produced particles [45].

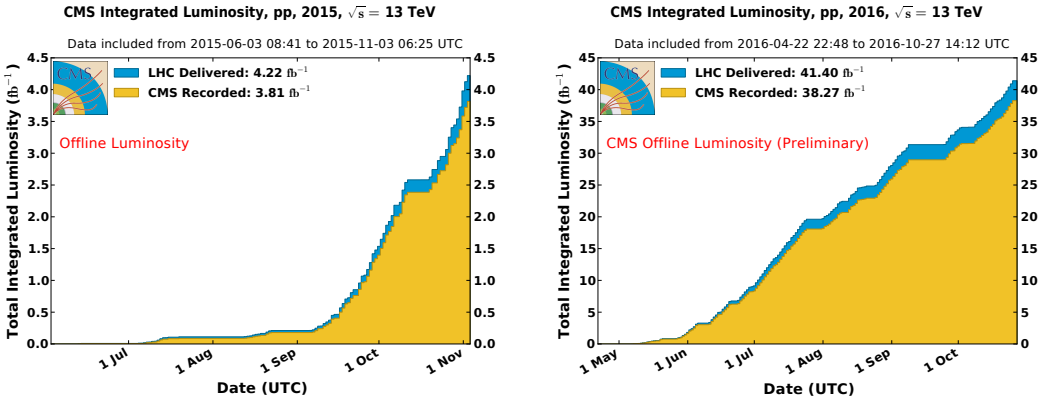

Figure 2.5: Cumulative offline luminosity versus day delivered to CMS (blue), and recorded by CMS (yellow) during stable beams for pp collisions at a centre-of-mass energy of $\sqrt{s} = 13\,\mathrm{TeV}$, in 2015 (left) and 2016 (right). The delivered luminosity accounts for the luminosity delivered from the start of stable beams until the LHC requests CMS to turn off the sensitive detectors to allow a beam dump or beam studies. The luminosity is given as determined from counting rates measured by the luminosity detectors after offline validation.

2.2 Compact Muon Solenoid

As one of the two multipurpose experiments of the LHC, the CMS detector was designed to cover a wide range of measurements in particle physics and heavy ion physics. Since one of the prime goals was to elucidate the nature of electroweak symmetry breaking, the detector was optimized to search for the then hypothetical Higgs boson in a variety of decay channels. It was also meant to explore various aspects of particle physics at the TeV energy scale, one example being the search for particles predicted by supersymmetric models.

2.2.1 Detector layout

The design of CMS devoted particular care to the measurement of muons, electrons, and photons. A good resolution was necessary over a large range of momenta, going all the way from e.g. dileptonic decays of low-mass hadronic resonances to possible new resonances at the TeV scale. Moreover, the unprecedented collision parameters of the LHC set a strong threefold technical constraint on the design:

- *Fast response*: The necessity of handling one LHC bunch crossing calls for a typical response time of 20–50 ns, thus requiring high-performance readout electronics.

- *High granularity*: To minimize the probability that particles from pileup interactions be in the same detector element as particles from the main pp interaction, LHC detectors need a high number of electronic channels.

- *Radiation resistance*: The high flux of particles from pp collisions damages detector components over the long term, especially inner tracking and forward calorimetry.

The name *Compact Muon Solenoid* summarizes the design choices that were made to meet these challenges. The CMS detector layout is organized around a superconducting *solenoid* magnet of 6 m internal diameter and 12.5 m length, providing a large magnetic field of 3.8 T. The overall apparatus is rather compact: it materializes as a 21.6-metre-long, 14.6-metre-wide, 14 000-tonne cylinder around the LHC beam axis. The magnet coil is large enough to accommodate three major subsystems within its volume: a silicon pixel and strip *tracker* which measures the trajectories of charged particles, a lead tungstate crystal *electromagnetic calorimeter* (ECAL) that mainly collects the energies of electrons and photons, and a brass and scintillator *hadron calorimeter* (HCAL) which stops the more penetrating hadrons. Some *forward calorimeters* further improve hermeticity. The measurement of muons relies on a combination of inner tracking and information from the *muon chambers*, which are gas-ionization detectors embedded in the steel flux-return yoke outside the solenoid. The overall layout of the detector is illustrated in Fig. 2.6.

CMS is located in an underground cavern at Point 5 of the LHC, about 100 m below the village of Cessy, in France. The aforementioned major detector subsystems will be described in the following sections, whereas the trigger system that applies online selection to the 40 MHz of LHC collisions will be described in Chapter 3. A comprehensive description of the whole CMS apparatus can be found in Reference [47].

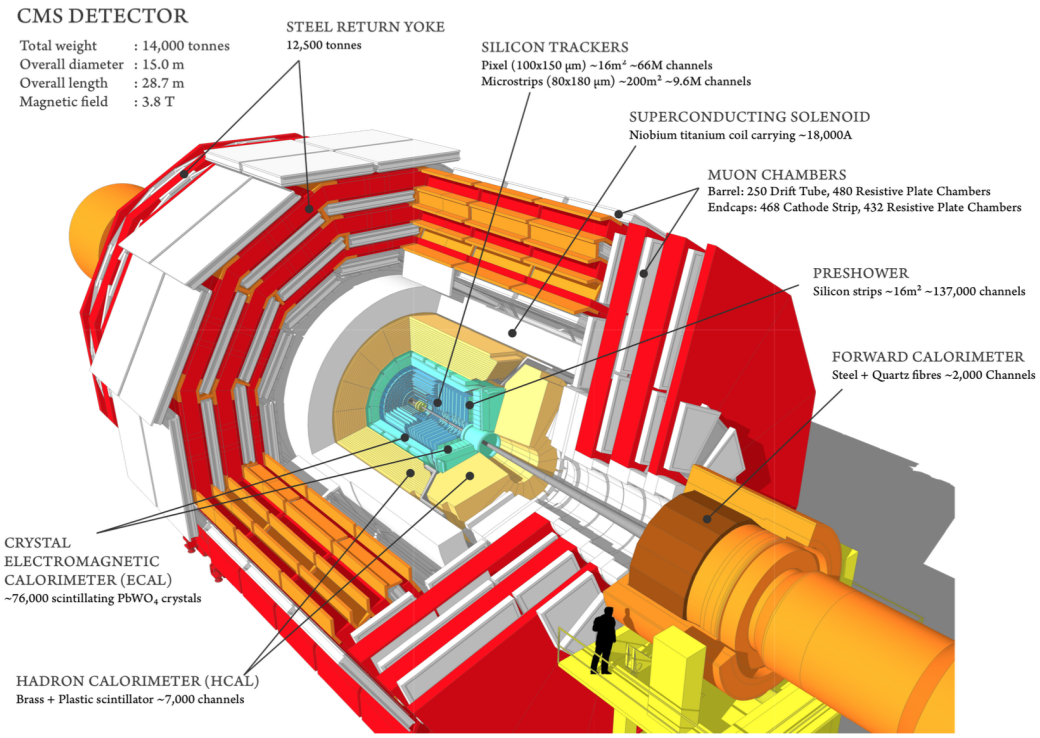

Figure 2.6: A perspective view of the CMS detector, illustrating its major subsystems.

Coordinate system

A conventional coordinate system has been adopted to describe the CMS detector, and will be used throughout this thesis. Its origin is centred at the nominal collision point, and its z axis coincides with the proton beam direction and points toward the Jura mountains from LHC Point 5. The y axis points vertically upwards, while the x axis points radially inwards, toward the centre of the LHC ring. The azimuthal angle ϕ is measured from the x axis in the $x-y$ plane and takes values in $[-\pi, \pi]$. The radial coordinate in this plane is denoted by r. The polar angle θ is measured from the z axis and takes values in $[0, \pi]$, and the pseudorapidity is defined as $\eta \equiv -\ln[\tan(\theta/2)]$.

The particle momentum and energy transverse to the beam direction, respectively denoted as p_{T} and E_{T}, are computed from the x and y components of the energy and momentum, which e.g. implies that $E_{\mathrm{T}} = E\sin(\theta) = E/\cosh(\eta)$. The imbalance of the total transverse energy measurement in a collision is referred to as the *missing transverse energy* and denoted as $E_{\mathrm{T}}^{\mathrm{miss}}$. The angular distance $\Delta\mathrm{R}$ between two particles i and j is defined as:

$$\Delta\mathrm{R}(i,j) = \sqrt{(\eta^i - \eta^j)^2 + (\phi^i - \phi^j)^2}\ . \tag{2.2}$$

Based on the η coordinate, the detector is divided into a central part called the *barrel*, and two opposite forward parts called the *endcaps*. The exact boundary depends on the subsystem.

Inner tracking system

The innermost subdetector of CMS is the silicon tracker, which occupies a cylindrical volume of 5.8 m in length and 2.5 m in diameter, and is immersed in the homogeneous magnetic field of 3.8 T provided by the solenoid. The tracker is designed to reconstruct the trajectories of charged particles emerging from the LHC collisions, thus allowing the determination of their charge and momentum. The high granularity allows the reconstruction of several vertices in every LHC bunch crossing, which is key to identifying particles of the main hard pp collision among those from pileup interactions. Secondary vertices can also be reconstructed, indicating late decays of particles such as B hadrons.

Since several hundreds of particles go through the tracker during each bunch crossing, high granularity and fast response are especially important to its design. This implies a high power density of on-detector electronics, which in turn requires efficient cooling. Nevertheless, the amount of material has to be kept to the minimum, in order to limit phenomena such as multiple scattering, bremsstrahlung, photon conversion, and nuclear interactions, all of which complicate particle reconstruction. The constraint of radiation resistance is also particularly strong, since the tracker directly surrounds the interaction point.

All of this has led to a design entirely based on silicon detector technology. Charged particles traversing a sensor produce electron-hole pairs which drift under an applied electric field, giving rise to a current pulse. The overall tracker layout is illustrated in Fig. 2.7; it involves two classes of sensors.

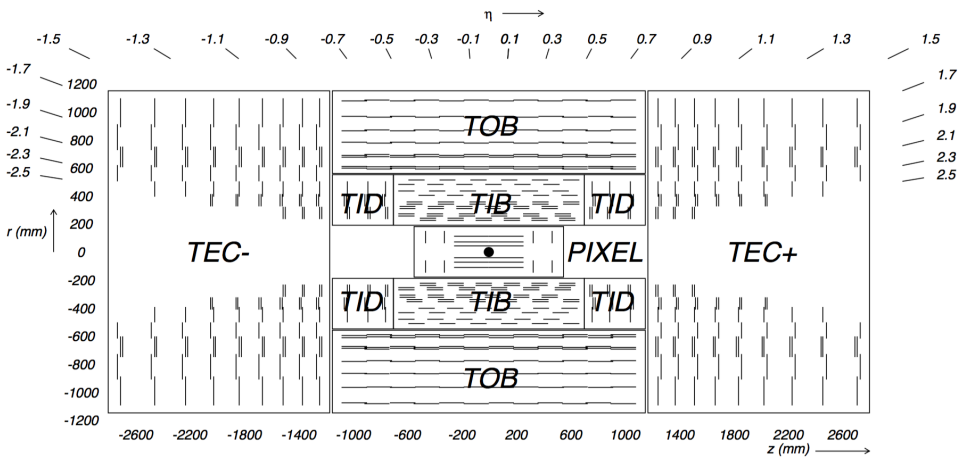

Figure 2.7: Longitudinal sectional view of the CMS tracker, showing the inner pixel detector with its barrel and endcap modules, and the strip detector with two collections of barrel modules, the tracker inner barrel (TIB) and the tracker outer barrel (TOB), and two collections of endcap modules, the tracker inner discs (TID) and the tracker endcaps (TEC) [47].

- The innermost part is made of silicon *pixels* of size $100 \times 150\,\mu\text{m}^2$, which are arranged in three cylindrical barrel layers at radii between 4.4 cm and 10.2 cm, and two endcap disk layers. This pixel detector achieves a spatial resolution of 15–20 μm, allowing for a three-dimensional vertex reconstruction.
- The rest of the tracker consists of silicon *strips*. They are arranged parallelly to the beam axis in the $4 + 6$ barrel layers, which extend outwards to $r = 1.1$ m,

and radially in the $3 + 9$ endcap layers, which provide a pseudorapidity coverage of $|\eta| < 2.5$. The distance between neighbouring strips varies from 80 to 205 μm depending on the location in the detector.

In total, this represents about 200 m^2 of active silicon area and 75 million readout channels. Cooling is ensured by liquid perfluorohexane, maintaining the sensor temperature at around –10°C.

Electromagnetic calorimeter

The electromagnetic calorimeter, or ECAL, measures the energies of electrons and photons. Its design was driven by the prospect of detecting Higgs boson decays to pair of photons, which called for an excellent energy and position resolution. The ECAL is also of primary importance to the $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ channel, since it allows the reconstruction of electrons, together with information from the inner tracker.

The ECAL is a hermetic, homogeneous, high-granularity calorimeter made of scintillating *crystals* of lead tungstate ($PbWO_4$). Thanks to the short radiation length (0.89 cm) and high density of this material, electromagnetic showers can be absorbed within relatively short crystals, while the small Molière radius (2.2 cm) allows for a good shower separation. 80% of the scintillation light is emitted in 25 ns, which is fast enough to cope with the LHC bunch spacing. The barrel part of the ECAL (EB) is made of 61 200 crystals of length 230 mm (25.8 X_0) and frontal cross section $22 \times 22\,\mathrm{mm}^2$, covering the pseudorapidity range $|\eta| < 1.479$. The endcap parts (EE) involve 7324 crystals each, with length 220 mm (24.7 X_0) and frontal cross section $28.62 \times 28.62\,\mathrm{mm}^2$, covering the $1.479 < |\eta| < 3.0$ range, as shown in Fig. 2.8.

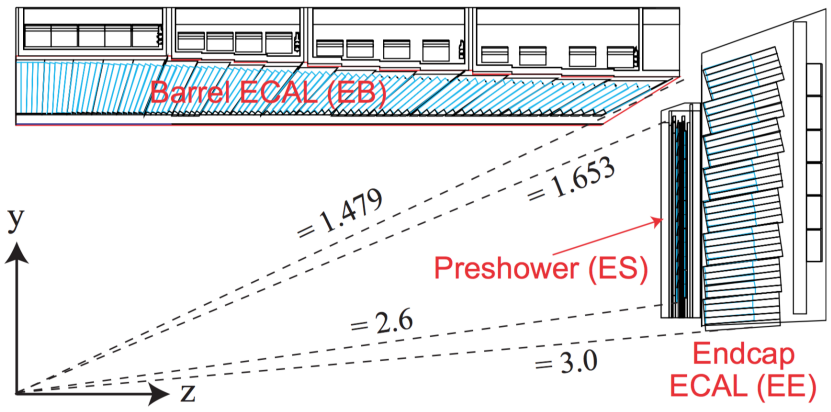

Figure 2.8: Longitudinal sectional view of a quarter of the CMS detector, showing the ECAL barrel and endcap, with the preshower in front.

The structure of the EB relies on 36 *supermodules* that cover half of the barrel length and 20° in ϕ. Each supermodule is made of four *modules* that each contain 400 or 500 crystals in an alveolar structure. Each EE is made of two semi-circular *dees* containing 3662 crystals. This general structure is illustrated in Fig. 2.9. The crystals are mounted in a quasi-projective geometry, so that their axes make a 3° angle with respect to the direction of the nominal interaction point in both the η and ϕ projections, thus avoiding to align inter-crystal gaps with particle trajectories. Still, some gaps (referred to as

cracks) remain between modules and complicate the energy reconstruction. Larger cracks are present at $\eta = 0$ and at the EB–EE transition.

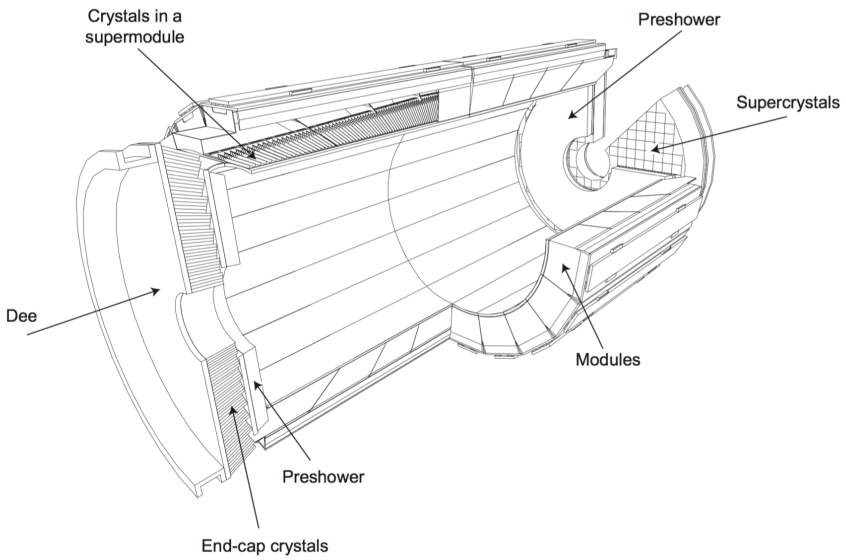

Figure 2.9: Layout of the CMS electromagnetic calorimeter showing the arrangement of crystal modules, supermodules and endcaps, with the preshower in front [47].

Two preshower detectors (ES) are installed at each end of the tracker, in front of the EE, covering the $1.653 < |\eta| < 2.6$ region. These sampling calorimeters are made of a lead radiator layer that initiates electromagnetic showers from incoming particles, followed by silicon strip sensors that measure the deposited energy and transverse shower profiles. The ES help distinguish $\pi^0 \to \gamma\gamma$ decays from single photons, and to identify electrons against minimum ionizing particles.

Scintillation light from ECAL crystals is read by fast, radiation-tolerant *photodetectors*, which need to amplify the small light yield and to be insensitive to particles traversing them. Owing to different magnetic field configurations and expected radiation levels, different photodetector technologies are used in the EB and EE, namely avalanche photodiodes and vacuum phototriodes, respectively. Since the crystal response is temperature-dependent, a cooling system stabilizes the temperature of both crystals and photodetectors to 18°C with a ±0.05°C accuracy in order to preserve energy resolution.

Although lead tungstate crystals are radiation resistant, they are known to undergo a limited but rapid loss of optical transmission under irradiation. This phenomenon depends on the luminosity and crystal pseudorapidity, and is partly balanced by an annealing effect. This results in a cyclic transparency behaviour between LHC collision runs and machine refills. To measure this effect, laser pulses are injected into the crystals via optical fibres, and the response is normalized by the laser pulse magnitude measured using silicon photodiodes. Time-dependent corrections are then applied to the measured particle energies.

Another technical challenge for ECAL operation is calibration, which consists of a global component that gives the absolute energy scale, and of a channel-to-channel relative

component, referred to as intercalibration. The latter accounts for the differences in crystal scintillation light yields, which can vary by ≈15% and ≈25% among EB and EE crystals, respectively. This was first mitigated via laboratory measurements of crystal light yield and photodetector response, and cosmic-ray-based studies. Direct usage of physics events now provides a percent-level precision on intercalibration constants, thus allowing accurate energy measurements for electrons and photons.

The ECAL energy resolution is composed of a stochastic, a noise, and a constant contribution. For instance, a typical parameterization of the resolution obtained from incident electrons in beam tests [48] was:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{2.8\%}{\sqrt{E}}\right)^2 + \left(\frac{0.12}{E}\right)^2 + (0.30\%)^2 \,, \tag{2.3}$$

where E is given in GeV.

Hadronic calorimeter

The hadron calorimeter, or HCAL, collects the energy of hadrons, therefore playing a major role in the reconstruction of jets. Moreover, its wide extension in pseudorapidity captures a large fraction of particles emerging from the interaction point, allowing for a reliable measurement of missing transverse energy which is a signature of otherwise undetected particles such as neutrinos.

The structure of the CMS HCAL is illustrated in Fig. 2.10. The HCAL barrel (HB) is located between the outer extent of the EB ($r = 1.77\,\text{m}$) and the inner extent of the magnet coil ($r = 2.95\,\text{m}$). Since it is not enough to fully absorb hadronic showers, an outer hadron calorimeter (HO) is placed outside of the solenoid. The HB covers the pseudorapidity range $|\eta| < 1.3$, which the HCAL endcaps (HE) extend to $|\eta| < 3.0$. Beyond this, two forward hadron calorimeters (HF) are located 11.2 m away from the interaction point and reach $|\eta| = 5.2$, thus ensuring good hermeticity.

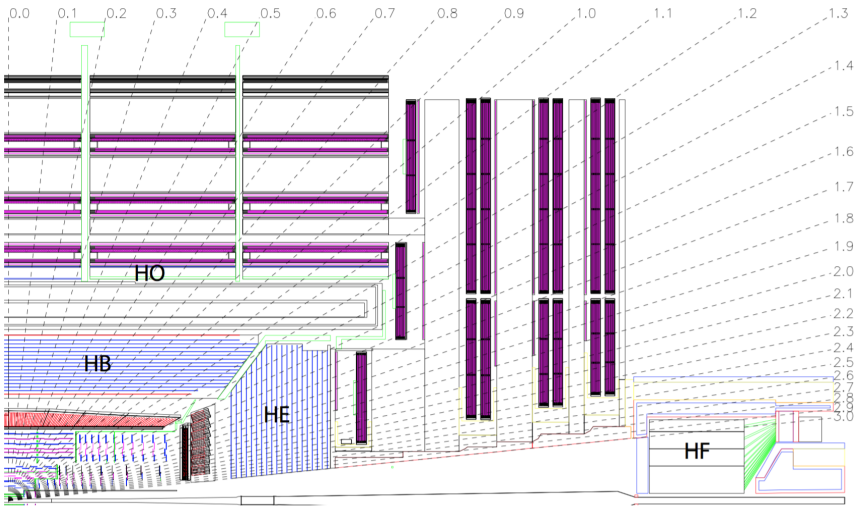

Figure 2.10: Longitudinal sectional view of a quarter of the CMS detector, showing the locations of the hadron barrel (HB), endcap (HE), outer (HO), and forward (HF) calorimeters [47].

The HCAL is a sampling calorimeter. In the HB and HE, brass layers are used as the passive, absorber material, and are interspersed with plastic scintillator tiles which serve as the active material. The calorimeter cells are grouped in projective towers of granularity $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ in the HB and $\approx 0.17 \times 0.17$ in the HE. The scintillation light is converted by wavelength-shifting fibres embedded in the scintillator tiles, channelled to photodetectors via clear fibres, and detected by hybrid photodiodes. The HFs, on the other hand, have to sustain a harsher radiation environment, which led to choose quartz fibres as the active material. The fibres are placed between steel absorber plates and emit Cherenkov light, which is detected by photomultipliers.

Muon chambers

Muons play a major role in many physics analyses, and particularly in the study of the $H \to ZZ \to 4\ell$ channel. Thanks to other particles being absorbed by calorimeters, muons are quite easy to identify and to distinguish from their backgrounds. The outer muon system of CMS has been designed to achieve high-precision measurement of muon momenta and charge, even without the help of the inner tracker. As shown in Fig. 2.11, this subsystem is made of *muon chambers* embedded in the iron return yoke of the CMS magnet, which are in principle not reached by other detectable particles. It is divided into a cylindrical barrel section and two planar endcap regions. Three types of gas-ionization chambers are used, for a total of 25 000 m^2 of detection planes.

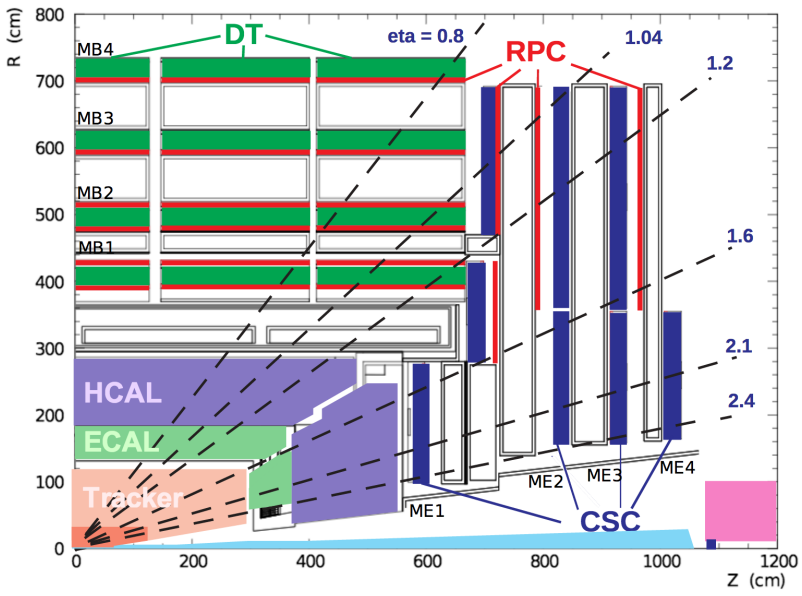

Figure 2.11: Longitudinal sectional view of a quarter of the CMS detector, showing the four DT stations in the barrel (MB1–MB4, green), the four CSC stations in the endcap (ME1–ME4, blue), and the RPC stations (red) [49].

The *Drift Tube Chambers* (DTs) are located in the barrel region ($|\eta| < 1.2$), where the neutron-induced background is small, the muon rate is low, and the magnetic field is quite uniform. The DTs are organized into four stations interspersed among the layers of the flux return plates. Their basic constituents are rectangular drift cells. These are bounded by two parallel aluminium planes, and aluminium beams that serve as

cathodes, while the anodes are 50 μm stainless steel wires located in the centre of the cells. A muon passing through a cell ionizes the gas mixture that fills the cell volume. The drift time of the resulting electrons is then used to measure the distance between the muon track and the wire. The drift cells of each chamber are offset by a half-cell width with respect to their neighbour, in order to eliminate dead spots in the efficiency. Each chamber has a resolution of 100 μm in the $r-\phi$ plane. The number and orientation of chambers in each station were chosen in a way that helps accurately link muon hits from different stations into a single muon track.

The *Cathode Strip Chambers* (CSCs) are used in the endcaps ($0.9 < |\eta| < 2.4$), where the muon rates and background levels are high and the magnetic field is large and non-uniform. The CSCs are multiwire proportional chambers, made of 6 anode wire planes interleaved among 7 cathode panels, with the wires running approximately perpendicular to the strips. A muon passing through a chamber generates an avalanche, inducing a charge on several cathode strips. The ensuing interpolation allows for a fine spatial resolution, of the order of 50 μm. There are four stations of CSCs in each endcap; the chambers are positioned perpendicular to the beam line and interspersed between the flux return plates.

The *Resistive Plate Chambers* (RPCs) are found both in the barrel and in the endcaps ($|\eta| < 1.6$), and are only used at trigger level. While both DTs and CSCs can already trigger muons independently with good efficiency and background rejection, the RPCs provide a complementary trigger system with moderate spatial resolution but excellent time resolution (of the order of 1 ns), thus helping measure the correct beam-crossing time, even at the largest LHC luminosities. The RPCs are double-gap chambers, operated in avalanche mode to ensure reliable operation at high rates; they are arranged in six layers in the barrel and three layers in the endcaps.

2.2.2 Particle and jet reconstruction

All collision events recorded by the CMS detector are centrally processed through *event reconstruction*, i.e. a common set of reconstruction algorithms that take raw detector data as input, and provide as output a collection of detected particles (electrons, muons, taus, photons, and hadrons) with associated properties such as momentum and charge, as well as higher-level physics objects such as jets and missing transverse energy. The present section is intended as an introduction to the subset of reconstruction algorithms that are relevant to this thesis.

Electron and muon reconstruction will first be described in some detail, since they are the most important ingredients of the analysis of the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ channel. The *Particle-flow* (PF) paradigm will then be introduced; it considers event reconstruction in a global approach, with one unified algorithm coherently reconstructing every particle (including electrons and muons) via a combination of information from all subsystems. The resulting PF particle candidates are then used as inputs to the jet reconstruction algorithm. Figure 2.12 illustrates the signatures of the main types of particles passing through the CMS detector.

Electron reconstruction

Not only do electrons deposit most of their energy in the ECAL, but since they are charged particles, they also leave hits in the inner tracker. Electron reconstruction in CMS uses an elaborate algorithm that combines the inputs from both subsystems: it

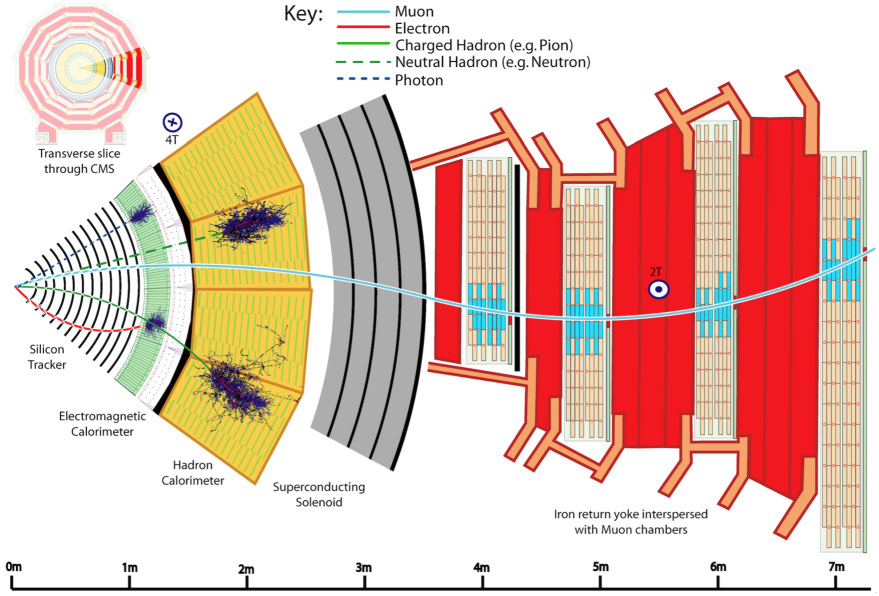

Figure 2.12: Schematic view of a transverse slice of the CMS detector, illustrating the specific signatures of different types of detected particles.

associates a reconstructed track with a cluster of energy in the ECAL, and ultimately exploits both sides of the information to estimate the electron momentum. The complete reconstruction procedure is described in Reference [50], for the version used in Run I for pp collisions at $\sqrt{s} = 8\,\text{TeV}$. Since then, it has been revisited during LS1 to be made more consistent with the global Particle-flow event reconstruction.

One of the main challenges of electron reconstruction in CMS is the amount of tracker material located between the collision point and the ECAL, as illustrated in Fig. 2.13. This indeed causes significant bremsstrahlung along the electron trajectory, with the resulting photons possibly converting to electron pairs. On average, 33% of the electron energy is radiated this way around $\eta^{\text{e}} = 0$ where the material is minimal, and about 86% for $\eta^{\text{e}} \approx 1.4$ where the material is largest, and this energy can spread over a large volume, mostly in the ϕ direction. Dedicated techniques have long been developed to accurately account for this effect, and will be summarized in what follows.

The electron energy usually spreads out over several crystals of the ECAL, and the first step of reconstruction consists in clustering these energy deposits. The two algorithms that were used in Run I in the EB and EE have now been replaced by a general algorithm called *PF clustering*, which was developed for PF reconstruction and consists of three steps. First, *cluster seeds* are identified as local crystal energy maxima above a given threshold[1]. Second, *topological clusters* are grown from the seeds by aggregating crystals with at least one side in common with a crystal already in the cluster, and with an energy exceeding another threshold. A topological cluster gives rise to as many *PF clusters* as it has seeds. Third, the energy of each cell is shared among all PF clusters according to the cell-cluster distance, with an iterative determination of cluster energies and positions.

[1]The threshold represents two standard deviations of the electronic noise, $E > 230\,\text{MeV}$ in the EB, and $E > 600\,\text{MeV}$ or $E_{\text{T}} > 150\,\text{MeV}$ in the EE.

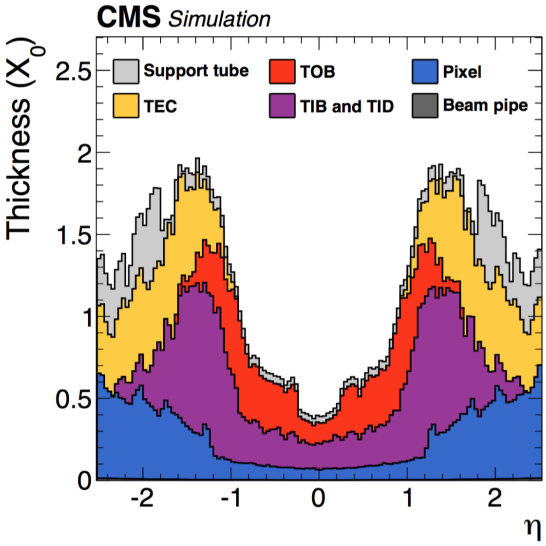

Figure 2.13: Total thickness of tracker material traversed by a particle produced at the centre of the CMS detector, expressed in units of X_0, as a function of particle pseudorapidity η [50]. Contributions of the tracker elements presented in Fig. 2.7 are shown separately, as well as contributions from the beam pipe and the support tube that surrounds the tracker.

PF clusters are then assembled into *PF superclusters*, starting from a *seed cluster* and gathering the presumptive clusters of bremsstrahlung photons and conversions products. In the endcaps, clusters reconstructed in the preshower detector are also collected with those of the EE. The superclustering procedure has been improved for Run II: it now relies on an advanced geometrical shape which exploits the η–ϕ correlation between the position of the electron and of bremsstrahlung clusters, and its E_{T} dependency. This improves the energy collection and the rejection of pileup clusters, particularly for low-p_{T} and large-η electrons.

Electron tracks are not reconstructed with the standard *KF tracking* algorithm used for other charged particles, because the large radiative losses in the tracker cause changes in curvature, which reduces the efficiency of collecting hits and hampers the estimation of track parameters. A dedicated *GSF tracking* procedure is thus used for electrons. As it can be very time consuming, it is preceded by a *seeding* procedure.

Seeding consists in finding and selecting the two or three first hits in the tracker from which the track can be initiated. Two complementary seeding algorithms are used:

- *ECAL-driven seeding* starts from the supercluster energy and position, which is used to extrapolate the electron trajectory toward the innermost layers of the tracker (both for positive and negative charge hypotheses) to predict the seeding hits. Actual seeds are formed by combining pairs or triplets of hits with the vertices obtained from pixel tracks, and they are selected using geometrical windows around the predicted ones, which are optimized to ensure a good compromise between efficiency and CPU time.

- *Tracker-driven seeding* relies on tracks reconstructed by the general algorithm for charged particles, which are extrapolated toward the ECAL and matched to a supercluster. If such tracks have a small number of hits or bad quality, they

are refitted using the procedure described in the next paragraph. A multivariate technique is then used to determine whether the seed can be used as an electron seed.

The selected seeds are then used to initiate electron track reconstruction. Track building proceeds iteratively from the track parameters provided in each layer, modelling the electron energy loss with a Bethe–Heitler function. To maintain good efficiency in the presence of bremsstrahlung, compatibility requirements between the predicted and the found hits in each layer are quite loose. If several hits are compatible with the predicted one, then several trajectory candidates are created and developed, with a limit of five candidate trajectories for each layer. At most one missing hit is allowed per trajectory. Once all hits are collected, track parameters are estimated by a fit that uses a Gaussian Sum Filter (GSF) [51, 52], instead of the Kalman Filter (KF) [53] used for non-electron tracks. The energy loss in each layer is approximated by a mixture of Gaussian distributions, each of which is attributed a weight that describes its associated probability. The parameters of electron tracks are estimated from the most probable value (the mode) of all the track components[2].

At this point, a *refining* procedure is applied to the PF superclusters. Based on the GSF track, the tracker-driven KF track and the seed cluster of the PF supercluster, a search is performed for additional tracks and clusters corresponding to (possibly converted) bremsstrahlung photons, exploiting the track tangents and energy–momentum compatibility between candidate conversion tracks and corresponding clusters. Partner tracks are also searched for whenever a conversion track is identified. The resulting *refined supercluster* is only used in the cases where the electron is tracker-seeded.

Some loose requirements are applied to variables characterizing the geometrical association between the track and the supercluster, using different approaches depending whether seeding was ECAL-driven or tracker-driven. Moreover, a disambiguation procedure is applied to cases where two nearby GSF tracks share the same superclusters, which can happen e.g. when a bremsstrahlung photon carries a significant fraction of the initial electron energy.

Charge estimation is then performed. While a natural way would consist in taking the sign of the GSF track curvature, this can be complicated by cases of early bremsstrahlung followed by a photon conversion, where the contributions from conversion can be mistakenly included in track fitting. Two other charge estimates are thus used: that from the possible KF track associated to the GSF track, and the sign of the difference in ϕ between the vector joining the beam spot to the supercluster position and the one joining it to the first hit of the electron GSF track. The electron charge is chosen as the sign obtained by at least two of the three estimates.

The final step is the estimation of electron momentum, which relies on a combination of the energy of the supercluster and the momentum estimate of the GSF track. This proceeds as follows:

- To improve the estimate, electrons are first classified into five categories: three of them ('golden', 'big brem' and 'showering') are associated to different bremsstrahlung patterns, another one ('crack') contains electrons located near cracks of the ECAL, and the last one ('bad track') handles electrons with a poorly fitted track.

[2]This was preferred to a weighted mean of all components, which was found to give a slight bias in the momentum value and less good resolution.

- The accuracy of supercluster energy obtained by summing all crystal energies (plus preshower energies in the endcaps) is harmed by several effects such as energy leakage out of the supercluster, in the ECAL cracks or into the HCAL, energy loss in the tracker material, and energy from pileup. The supercluster energy is thus corrected via a multivariate regression that exploits variables related to all of these effects. This was shown to improve the resolution by 20% on average, while reducing the bias in the peak position. A similar regression calculates the uncertainty on the corrected energy. Calibrations are further applied to correct for e.g. the imperfect description of the tracker material in simulation and the evolution of crystal transparency and ECAL noise during data taking.

- The final estimate of the electron momentum combines the corrected and calibrated supercluster energy E^{SC} with the GSF track momentum p. The latter is expected to be the more precise one below 15 GeV, or for electrons near cracks. The combination again uses a multivariate regression, whereby the main input variables are E^{SC}, p, their relative uncertainties, the electron category, and its position in the barrel or endcaps. This regression generally reduces the bias in the electron momentum while much improving the resolution, e.g. by $\approx 25\%$ for electrons with $p_{\mathrm{T}}^{\mathrm{e}} \approx 15\,\mathrm{GeV}$ in the barrel.

The typical relative transverse momentum resolution for electrons produced in Z boson decays eventually ranges from 1.7% for single-cluster electrons in the barrel region to 4.5% for poorly measured or multi-cluster electrons in the endcaps.

Muon reconstruction

In the standard CMS reconstruction, muon tracks are first reconstructed independently in the inner tracker and in the outer muon system, and are called *tracker tracks* and *standalone-muon tracks*, respectively. The muon reconstruction algorithm then combines the information from both subsystems using two approaches [49].

- *Global muons* are formed by propagating standalone-muon tracks inwards to the inner tracker. If a matching tracker track is found, a global-muon track is fitted combining the hits from the tracker and standalone-muon track, using the Kalman Filter technique [54].

- *Tracker muons* are formed by extrapolating tracker tracks outwards to the outer muon system, requiring that at least one muon track segment made of DT or CSC hits matches the extrapolated track. The possibility for tracker muons to have one single matched segment in the muon system makes this algorithm more efficient than global-muon reconstruction at low momentum ($p^{\mu} \lesssim 5\,\mathrm{GeV}$).

About 99% of muons produced in pp collisions with high enough momentum are reconstructed by either algorithm. Muon candidates found by both methods (sharing the same tracker track) are merged into one single candidate. Since muons reconstructed only as standalone-muon tracks have worse momentum resolution and are more contaminated by cosmic-ray muons, they are not used in analyses.

The charge and momentum of muons with $p_{\mathrm{T}}^{\mu} < 200\,\mathrm{GeV}$ are taken from the tracker track, because the precision of the muon system measurement at low momentum is limited by multiple scattering. Above 200 GeV, the additional lever arm provided by the outer muon detectors becomes a significant advantage, and the charge and momentum

are extracted from the combined trajectory fit. The relative transverse momentum resolution for muons with $20 < p_{\mathrm{T}}^{\mu} < 100\,\mathrm{GeV}$ ranges from 1.3–2.0% in the barrel to $\approx 6\%$ in the endcaps.

Global event reconstruction

Event reconstruction in CMS relies of the Particle-flow algorithm [55, 56], which aims at reconstructing and identifying each stable particle in the event (namely electrons, muons, photons, charged hadrons, and neutral hadrons) via a thorough and coherent combination of all CMS subdetectors, allowing for an optimal determination of their direction, energy and type. This task is all the more delicate as most stable particles produced in pp collisions have a rather low p_{T}: for instance, in a quark or gluon jet with a total $p_{\mathrm{T}}^{\mathrm{jet}}$ below 100 GeV, the average p_{T} carried by the stable constituent particles is of the order of a few GeV, and accurately reconstructing as many of them as possible is key to analyses that rely on jets.

The CMS detector is very well suited for a Particle-flow approach. Indeed, its silicon inner tracker immersed in the uniform 3.8 T magnetic field allows reconstruction of charged-particle tracks with large efficiency and acceptance, its electromagnetic calorimeter is very granular and provides very good energy resolution, and limitations in the performance of its hadronic calorimeter can be compensated by combinations with ECAL and tracker information, e.g. allowing neutral hadrons to be detected as an energy excess on top of the energy deposited by charged hadrons pointing to the same calorimeter cells. While the PF algorithm covers all species of stable particles, it primarily improves performance for jets, tau leptons and missing transverse energy.

It should first be stressed that the use of PF reconstruction has motivated the development of specific advanced algorithms for clustering and tracking.

- **Clustering**. The algorithms that cluster calorimeter energy deposits serve several purposes:
  - detecting and measuring the energy and direction of stable neutral particles such as neutral hadrons and photons,
  - separating these neutral particles from the energy deposits of charged hadrons,
  - reconstructing and identifying electrons and their associated bremsstrahlung photons,
  - helping the energy measurement of charged hadrons for which the track parameters were not determined accurately.

  The clustering that was developed for PF reconstruction is the *PF clustering* that has already been described for electrons. It is actually also performed in the HCAL subsystems (except for the HF), involving cells instead of crystals. Compared to other possible clustering algorithms, PF clustering allows higher detection efficiency down to low-energy particles, and separation of close energy deposits.

- **tracking**. The CMS tracker provides precise measurements of the momentum of charged hadrons and of the charged-particle direction at the production vertex. Since charged hadrons detected solely by calorimeters are reconstructed with reduced efficiency, degraded energy resolution and biased direction, tracking efficiency has to be as close to 100% as possible. Nevertheless, maintaining a low

tracking fake rate is essential, because fake tracks with randomly distributed momenta may cause potentially large energy excesses. To satisfy these contradictory constraints, an *iterative tracking* strategy [57] was adopted, consisting in a sequence of six steps. Tracks are first seeded and reconstructed with very tight criteria, ensuring a negligibly small fake rate at the cost of efficiency. The next steps proceed by removing hits that are unambiguously assigned to the already reconstructed tracks (thus keeping a low fake rate thanks to reduced combinatorial effects), and by progressively loosening track seeding criteria (thus increasing the efficiency). In the last iterations, vertex constraints are relaxed, so as to reconstruct secondary charged particles from e.g. photon conversions or decays of long-lived hadrons.

At its earliest stage, the PF algorithm used in CMS relies on some algorithmic bricks, which can be of several types:

- ECAL and HCAL PF clusters,
- Charged-particle tracks, as obtained from iterative tracking,
- Reconstructed muon candidates, as already described,
- PF superclusters, which are not only used in the reconstruction of electrons, but also in that of photons, which can also involve conversions in the tracker material,
- GSF tracks, as defined in electron reconstruction,
- Converted bremsstrahlung photon candidates, which are reconstructed by a dedicated algorithm.

Since a single particle can in general produce several of these bricks, they need to be connected to one another by a *link algorithm*, so that each particle can be fully reconstructed while avoiding possible double counting from different detectors. This algorithm defines a distance between every pair of elements in the event, quantifying the quality of their hypothetical link. It then decides whether these elements are linked, based on a set criteria that depends on the types of both elements. This creates *PF blocks* made of several elements connected directly or indirectly.

The refining of PF superclusters, which was mentioned in the description of electron reconstruction but also applies to photons, is performed at this stage of the algorithm. This allows for the completion of electron and photon reconstruction, and all the clusters and tracks that the electron and photon candidates finally involve can later be removed from their blocks to avoid energy double counting.

The core part of the Particle-flow algorithm can then be performed, using the PF blocks as inputs. For each block, the algorithm proceeds through the following sequence.

1. First, muon candidates are applied a set of identification requirements that have high efficiency in selecting genuine muons in jets, while rejecting e.g. charged hadrons misidentified as muons. The criteria will be described in Section 4.1.2 when analysis-specific lepton identification is discussed[3]. Muons candidates that pass this selection are called *PF muons*, and their tracks are removed from the considered PF block. The small amount of energy deposited by these muons in the ECAL and HCAL is estimated for later use.

[3]It should nevertheless be noted that the analysis of the $H \to ZZ \to 4\ell$ channel also uses muons that do not pass the PF identification requirements, as will be explained in Chapter 4.

2. The algorithm then addresses the electron candidates. To be accepted as *PF electrons*, these also have to pass an identification criterion, defined by an algorithm that exploits tracking and calorimetric variables[4]. The building bricks of PF electrons (tracks and ECAL PF clusters) are removed from further processing of the block.

3. Tracks for which the transverse momentum is measured with an uncertainty larger than the expected calorimeter resolution are then removed from the block, thus rejecting many fake tracks while retaining a set of high-quality tracks.

4. Before examining the remaining tracks and the ECAL and HCAL PF clusters, the energies of said clusters have to undergo a calibration procedure, correcting for the non-linear calorimeter response and threshold effects. At this point, the aforementioned expected calorimetric energy deposits of PF muons are subtracted from the energies of matching PF clusters.

5. Detecting neutral particles (photons and neutral hadrons) in a block involves a comparison between track momenta and ECAL and HCAL cluster energies. However, a track can be linked to a number of ECAL or HCAL clusters, while a cluster may be linked to more than one track. An involved disambiguation procedure is thus applied, whereby links are either kept or discarded based on criteria related to the distance between the involved tracks and clusters.

6. If the total calibrated calorimetric energy is smaller than the total track momentum by more than three standard deviations, a relaxed search for muons and fake tracks is performed. This concerns less than 0.03% of all tracks.

7. Each of the remaining tracks in the block gives rise to a *PF charged hadron*, with momentum either equal to the track momentum if the calorimetric energy is not consistent with the former, or computed from a fit to the track and cluster energy otherwise.

8. If the total calibrated cluster energy is significantly larger than the total associated track momentum, additional neutral particles are defined. Specifically, if the total calorimetric energy excess is larger than the total ECAL energy, a *PF photon* is created with this ECAL energy and a *PF neutral hadron* is created with the remaining part of the excess of calibrated ECAL and HCAL energy. Otherwise, the uncalibrated excess only gives rise to a PF photon. Similarly to PF electrons, photons have to satisfy some identification criteria to be accepted as PF photons.

9. The remaining unlinked ECAL and HCAL clusters give rise to PF photons and PF neutral hadrons, respectively.

At this point, a final collection of individual PF particles (also called *PF candidates*) of five possible types is obtained for every event. They are then used to perform a variety of higher-level tasks, such as:

- reconstructing jets, as described in the next paragraph,

[4]The identification requirement of PF electrons is not used in the $H \to ZZ \to 4\ell$ analysis. Chapter 4 will explain that all electron candidates are considered at first, and specific identification criteria are then defined.

- determining the missing transverse energy, which provides an estimate of the direction and energy of the particles that escape detection, such as neutrinos or particles predicted by BSM theories,
- reconstructing and identifying tau leptons from their decay products,
- building isolation variables, as will be explained in Sections 3.2.2 and 4.1.3.

Jet reconstruction

While quarks and gluons are abundantly produced in high-energy processes such as hard scattering of partons in pp collisions, they cannot be observed directly as they immediately fragment and hadronize, giving rise to collimated showers of particles referred to as *jets*. Obtaining information on the original partons thus requires to reconstruct jets, i.e. to accurately combine the reconstructed particles and to determine the total jet momentum. It is also crucial to reach a good understanding of the jet energy scale and resolution, since they are usually an important component of systematic uncertainties.

Jets considered in this thesis are reconstructed from the aforementioned collection of PF particles. They are grouped using a clustering algorithm that consists in defining some measure of distance between pairs of particles, and iteratively combining the closest ones. More precisely, CMS uses the anti-k_{T} algorithm [58], as implemented in the FastJet package [59], and here operated with a size parameter $R = 0.4$. This algorithm tends to cluster jets around the hardest particles, resulting in conic jet shapes. It is known to be infrared and collinear safe, which means that the number of jets is unaffected by soft collinear gluon emission or parton splitting.

The momentum of a reconstructed jet is computed as the vectorial sum of the momenta of all the particles it involves, and it typically differs by 5 to 10% from the true hadron-level momentum. The difference is caused by a number of effects such as the non-linear response of calorimeters, detector noise, and additional energy deposits from pileup interactions. Thus, a sequence of *jet energy corrections* is applied, both in simulation and in data [60]. First, the contribution from pileup and detector noise is subtracted from every jet, using a calculation of the mean energy density in the event [61, 62]. Then, the jet energy response is made uniform with respect to $p_{\mathrm{T}}^{\mathrm{jet}}$ and η^{jet}. The corrections are derived from simulation, and are confirmed with in situ measurements of the energy balance in dijet and photon + jet events [60]. The final jet energy resolution typically amounts to 15% at 10 GeV, 8% at 100 GeV, and 4% at 1 TeV, which can be compared to the values of 40%, 12%, and 5% obtained with another jet clustering algorithm that only uses calorimetric information.

2.3 Conclusion

This thesis is carried out in a context where the LHC is breaking new records, both in terms of energy, reaching $\sqrt{s} = 13\,\mathrm{TeV}$ in Run II, and in terms of luminosity. Together with the versatility of the CMS detector, this provides good prospects for the exploration of the recently discovered scalar sector.

Having now described the experimental setup and the most general algorithmic objects, the description of the analysis of the $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ channel will unfold in a logical order, from the recording of collision events to the measurements of physics parameters. This starts in Chapter 3 with the presentation of the work done at the level of online event selection.

Chapter 3

Optimizing the CMS trigger system

The first step of the analysis workflow of a collider experiment such as CMS consists in recording collision data. The LHC bunch crossings occur at a rate of 40 MHz, thus delivering orders of magnitude more collisions events than can be stored for offline analysis. The CERN processing capacities can only absorb an input rate of the order of the kilohertz. Nevertheless, only a small fraction of the bunch crossings contain hard scattering events of interest to the physics programme of the LHC experiments, and these events thus have to be accurately selected.

The system that takes care of this selection and reduces the event rate to a manageable amount is called the *trigger*, and is a fundamental ingredient of the performance of a high-energy physics detector. The trigger selects in real time the small fraction of collision events that are most relevant to physics analyses, and while these *triggered* events are stored on tape for later analysis, the non-triggered events are lost for good. Efficient triggering of electrons and photons down to low p_{T} values has been a cornerstone of the most important CMS achievement of Run I, namely the discovery of the Higgs boson in the four-lepton and diphoton channels.

By focusing on the CMS trigger system, this chapter presents a low-level but crucial aspect of the selection of $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ events. It summarizes the long-term preparatory work that ensured efficient triggering of four-lepton events in the new data taking conditions of Run II. My main contribution includes the improvement of trigger-level selection algorithms for electrons, the definition of the trigger paths specific to $\mathrm{H}\to 4\ell$, and the optimization of the operating points they use.

3.1 The CMS trigger

The CMS trigger system is described in detail in Reference [63]. Unlike many high-energy physics experiments that rely on a three-level trigger, CMS performs online event selection via two successive layers:

- The first level, called the *Level-1 trigger* (L1), is implemented on custom-designed hardware, and is adjusted to bring the event rate down to about 100 kHz, which is the upper limit imposed by the CMS readout electronics.
- The second level, called the *high-level trigger* (HLT), is implemented in software and runs a streamlined version of the CMS reconstruction algorithm to select an average rate of 1 kHz. It then transmits the data to the CMS Tier-0 computing centre for storage and offline processing.

3.1.1 Level-1 trigger

The hardware-based Level-1 trigger of CMS performs a fast readout of the detector with a limited granularity, selecting events that contain such distinctive detector signals as ionization deposits consistent with a muon, or energy clusters consistent with an electron, photon, tau lepton, or jet. The L1 selection relies on a programmable *menu* made of 128 algorithms or *seeds*, each of which selects a particular type of objects and passes them to the HLT for subsequent processing. Every seed is assigned an adjustable *prescale* value n, meaning that it accept a fraction $1/n$ of events that pass its specific selection criteria. Thresholds and prescales are adjusted to the LHC instantaneous luminosity during data taking, so as to restrict the output rate to the 100 kHz upper limit.

The Level-1 trigger has a fixed latency: it has 4 μs to decide to accept or reject an event, using information from the calorimeters and muon detectors. Trigger primitives are computed from energy deposits in the trigger towers of the ECAL and HCAL calorimeters in the one hand, and from track segments and hit patterns in the DT, CSC and RPC muon chambers on the other hand. As illustrated in Fig. 3.1, the information is processed through two separate flows: the *calorimeter trigger* builds the EG (electron or photon), jet, and tau candidates, as well as variables such as $E_{\mathrm{T}}^{\mathrm{miss}}$ and H_{T} (the scalar transverse energy sum of jets above a certain threshold), while the *muon trigger* builds the muon candidates. The combined event information is finally evaluated in the *global trigger*, which makes the final decision based on the menu.

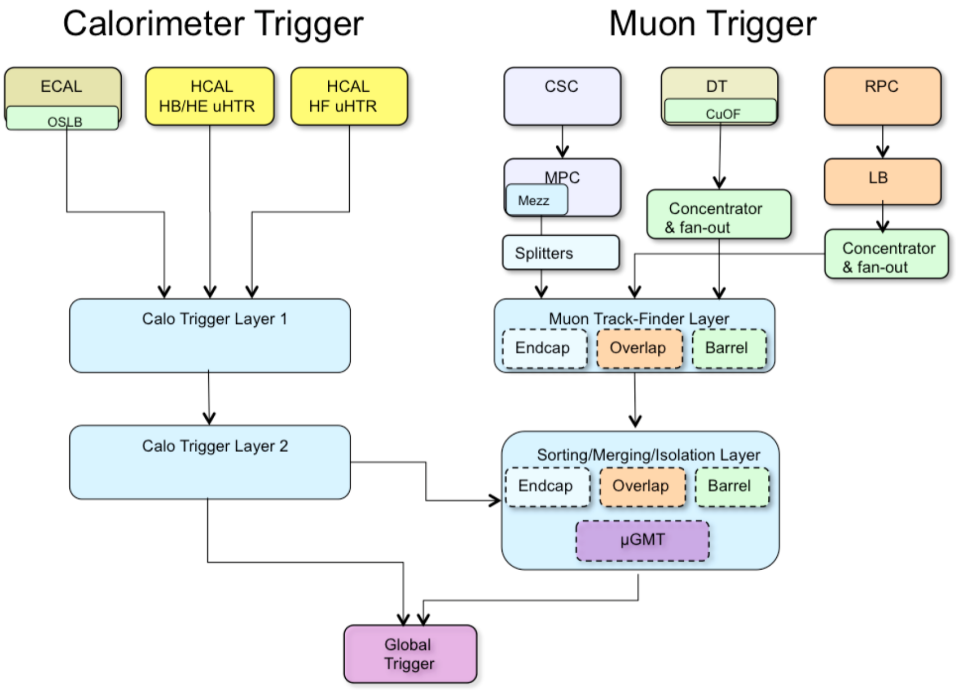

Figure 3.1: Architecture of the CMS L1 trigger system in 2016.

For Run II of the LHC, with prospects of peak instantaneous luminosities of $1.6 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ or more, and of pileup levels of 50 simultaneous inelastic collisions per crossing, the Level-1 trigger was known not to be able to keep low trigger thresholds while staying within the 100 kHz bandwidth. Hence, two major updates have been successively introduced [64] to preserve good performance as the LHC ramps up.

- A first upgrade, referred to as the *Stage 1*, went online in the course of 2015. In the calorimeter trigger, data communication from the ECAL was improved with new optical links, and an earlier subsystem called *global calorimeter trigger* was replaced with a new *MP7* data processing card with better data throughput and computational power, allowing the execution of improved algorithms and the inclusion of event-by-event pileup subtraction. Minor improvements were also introduced in the muon trigger.

  I contributed to the commissioning of the updated calorimeter trigger, as illustrated by the turn-on curves for electron L1 triggering efficiency shown in Fig. 3.2.

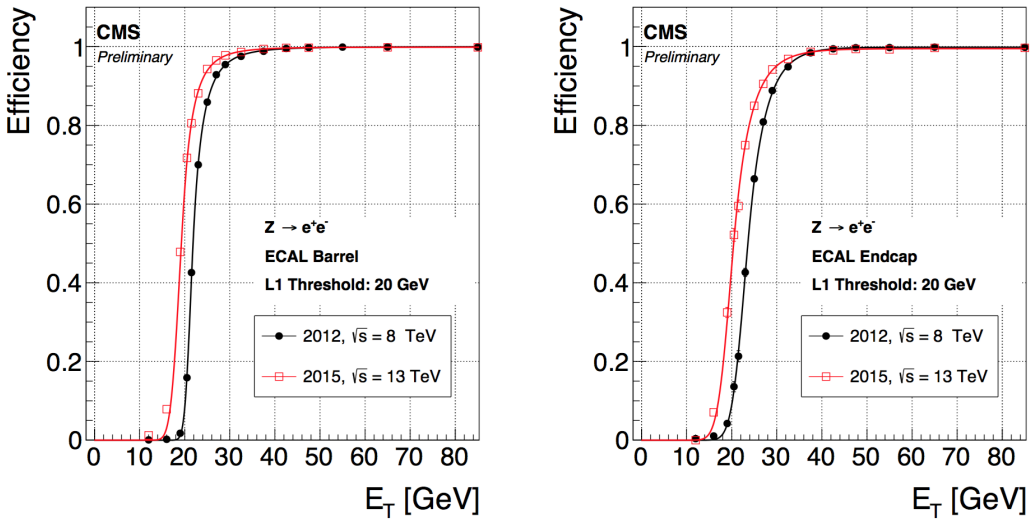

Figure 3.2: Electron trigger efficiency at Level 1 (EG E_T threshold: 20 GeV) as a function of the offline reconstructed electron E_T for electrons in the ECAL barrel (left) and in the ECAL endcaps (right), for 2012 (black dots) and 2015 data (red squares). Efficiencies are measured with the Tag-and-Probe technique. Besides a new ECAL response correction strategy, new ECAL calibrations with higher p_T granularity have been applied at Level 1 during the 2015 data taking, which explains the overall shift of turn-on curves.

- A *Stage 2* upgrade was then introduced, running in parallel to the Stage 1 for commissioning in late 2015, and going online as of the beginning of the 2016 data taking.

  In the calorimeter trigger, events are now read by two new layers of data processors, with a new time-multiplexed architecture. The 18 cards of *Layer 1* perform pre-processing steps that only require a regional view of the detector. *Layer 2* then relies on nine MP7 cards, each of which has access to a whole event, now with full trigger tower granularity, and benefits from extra latency thanks to time multiplexing. All of this allows for more complex algorithms for clustering and

isolation. The position and energy resolution of jet, EG, and tau candidates is improved, which provides the additional background rejection required to cope with the increased instantaneous luminosity and pileup.

In parallel, the muon trigger has also been upgraded, now combining all three muon systems to perform an integrated track finding, allowing for a more elaborate p_{T} measurement. The L1 trigger architecture after the Stage 2 upgrade is the one illustrated in Fig. 3.1.

3.1.2 High-level trigger

The high-level trigger is implemented in software and performs a full readout of the CMS detector. Events are reconstructed with the same software as used for offline processing, but in a very optimized configuration which is two orders of magnitude faster. Based on the raw data from events accepted by the L1 trigger, all main classes of physics objects can be reconstructed at HLT, such as electrons, muons, photons, taus, missing transverse energy, and jets, including some more advanced techniques like b tagging or jet substructure study. Specific selection criteria are applied to these objects so as to both keep the rate under control and retain the subset of events most relevant to subsequent data analysis. As opposed to the L1, this HLT filtering procedure is able to exploit the full precision of the data from the CMS detector, with offline-quality algorithms.

The HLT runs on a single dedicated farm of commercial computers. This so-called *event filter farm* consists of builder units that collect and assemble individual event fragments from the detector, and of filter units that unpack the raw data into detector-specific data structures, and perform event reconstruction and trigger filtering. In total, the farm currently comprises approximately 16 000 CPU cores. Its computing power has been increased by about 50% since Run I to cope with pileup[1] and code complexity.

Data processing at the HLT is structured around the concept of a *path*, which consists in a predefined sequence of algorithmic processing steps of increasing complexity, that both reconstructs a certain type of physics objects and applies a selection to it. Like L1 seeds, HLT paths can be prescaled, and the set of all paths used at a given time is also called a menu. All paths are run in parallel and independently of each other, but the common modules and sequences are shared among different paths. The successive reconstruction modules and selection filters are organized in such a way that the fastest selections such as those relying on information from the calorimeters and muon detectors are run first. This helps reducing the event rate as soon as possible, before considering CPU-expensive steps such as track reconstruction.

Finally, all events that are selected by at least one path are directed to one of various *data streams*, which are illustrated in Fig. 3.3. The main physics data stream transmits events as full raw detector data, for prompt offline event reconstruction and permanent mass storage. Its maximum average rate has been increased from 400 Hz in Run I to about 1 kHz in Run II. Other physics streams include *data parking*, i.e. the storage of full event content from special loose HLT paths for a delayed offline processing during LHC long shutdowns, and *data scouting*, i.e. the storage of reduced, non-reprocessable event content from very loose HLT paths. Some special streams are dedicated to data quality monitoring (online and offline), and to detector alignment and calibration workflows.

[1] While the HLT output rate is made robust against pileup, the CPU usage still tends to increase linearly with it.

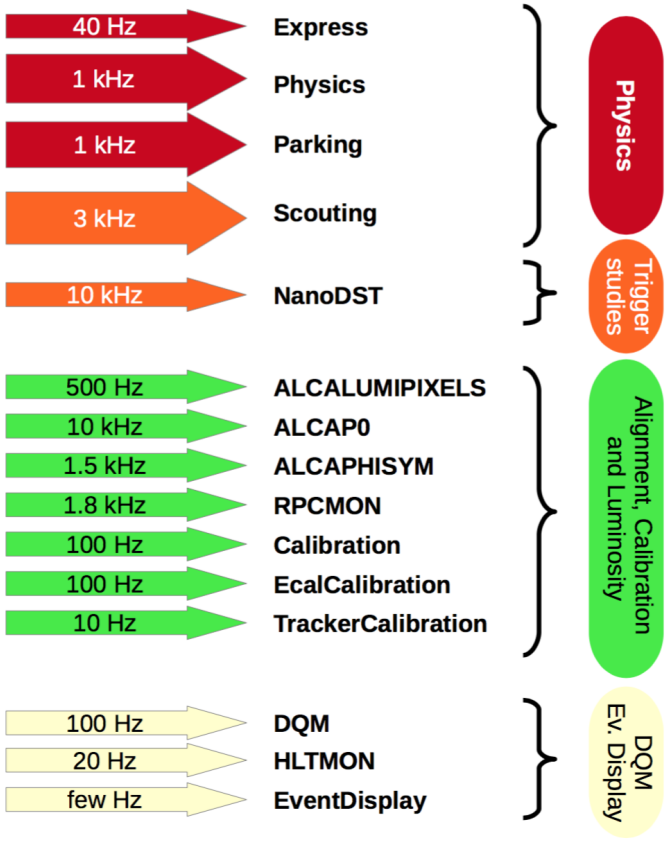

Figure 3.3: Schematic representation of HLT data streams in early Run II, grouped by purpose.

3.2 HLT Run II optimization

High-level trigger selection represents a unique optimization challenge, which relies on three metrics.

- **Efficiency**. Since any non-triggered event is lost, the HLT algorithm must apply as loose a selection as affordable, in order to maximize the efficiency of selecting all interesting physics signals, but also to build all control regions needed by analyses. Moreover, strategies for the measurement of trigger efficiencies in collision data have to be designed from the start.

- **Rate**. The output event rate is constrained by the downstream computing resources: for all HLT-selected events to be processed with the offline event reconstruction, the maximal allowed HLT rate in Run II is of the order of 1 kHz for the main physics data stream.

- **Timing**. Every event has to be processed in limited time, with an average budget of about 160 ms per event. This entails that the HLT software is thoroughly optimized, and any update of the algorithm needs to cope with the timing constraint.

The two years of LS1 were used by the CMS HLT group to perform a major upgrade of online reconstruction and selection algorithms in view of Run II of the LHC, and to

redesign and optimize the menus from scratch. The task of containing the HLT rate within the available bandwidth while maintaining the physics reach of the experiment became harder than in Run I, for two reasons. First, moving from $\sqrt{s} = 8\,\text{TeV}$ to 13 TeV makes the cross sections of the physics processes of interest gain an average factor of 2. Second, as a result of the new beam parameters, in particular the narrowing of the bunch spacing from 50 ns to 25 ns, luminosities were expected to reach $1.4 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ in 2015, which is about twice more than the 2012 peak luminosity. This implies that the HLT rate was expected to increase by a factor ≈ 4, while the increase of offline storage and processing only allowed it to double, thus calling for a reduction of the effective rate by a factor ≈ 2.

The default solution to reduce the rate consists in tightening the selection cuts, for example by raising the minimal p_T thresholds on various types of particles. In order to avoid this as much as possible, the HLT upgrade strategy of LS1 mainly relied on porting online some advanced algorithms that were already in use in the offline reconstruction in Run I. This included a wider use of tracking and Particle-flow-based techniques, and allowed HLT thresholds to be put closer to the offline cuts without losing much efficiency.

Another problem was the increase of pileup, which was expected to reach about 40 interactions per bunch crossing at peak luminosity, and to be further aggravated by *out-of-time* pileup from different crossings, as a consequence of narrower bunch spacing. While pileup only moderately impacts the event rate, it hampers some of the reconstruction and selection algorithms at the expense of efficiency, calling for the implementation of mitigation strategies.

The remainder of this section first describes what was done in the particular case of electron reconstruction and electron selection, and then shows how the new algorithms were exploited to build and optimize trigger paths specific to the $\text{H} \to \text{ZZ} \to 4\ell$ analysis.

3.2.1 Online electron reconstruction

Electron reconstruction in CMS has been described in Section 2.2.2 in its general, offline version. The HLT version of this reconstruction is slightly simplified, and it is implemented in an algorithmic unit called an *electron sequence*, which consists in a series of modules that performs the successive steps of both reconstruction and selection. Table 3.1 presents a typical electron sequence of the Run II HLT: any path that uses electrons has to include a certain variant of this sequence.

The different steps of Table 3.1 will be progressively discussed in this subsection and in the next one, but three general characteristics of HLT reconstruction are already visible here:

- Steps 2, 11, and 20 are described as 'local' or 'regional', referring to the fact that the HLT algorithm only unpacks and processes regions of the detector that it strictly needs. For example, most electron sequences only initiate reconstruction around the EG candidates provided by the Level-1 trigger.

- To optimize timing, as soon as the reconstruction provides enough elements to build a selection variable, it is built and the cut is applied, rejecting some background right away. For example, instead of a cut on the transverse momentum p_T^e, the electron sequence contains a cut on the supercluster transverse energy E_T^SC (step 5), because the supercluster is reconstructed before the track.

Table 3.1: Summary of a typical electron sequence in the Run II HLT. Besides the description of the successive steps, the 'type' column indicates if the steps contributes to reconstructing the electron itself, to preparing variables that will be used for selection, or to performing the selection.

	step of the electron sequence	type
1	unpacking of ECAL information	electron reconstruction
2	local ECAL PF clustering	electron reconstruction
3	building of the PF supercluster	electron reconstruction
4	matching to L1 seeds	electron reconstruction
5	cut on $E_{\mathrm{T}}^{\mathrm{SC}}$	electron selection cut (main threshold)
6	cut on $\sigma_{i\eta i\eta}$	electron selection cut (identification)
7	unpacking of HCAL information	preparation of electron selection
8	computation of the energy density ρ	preparation of electron selection
9	cut on ECAL PF-cluster isolation	electron selection cut (isolation)
10	cut on H/E	electron selection cut (identification)
11	local HCAL PF clustering	preparation of electron selection
12	cut on HCAL PF-cluster isolation	electron selection cut (isolation)
13	building of pixel seeds	electron reconstruction
14	matching to pixel seeds	electron reconstruction
15	electron track reconstruction	electron reconstruction
16	cut on $\frac{1}{E} - \frac{1}{p}$	electron selection cut (identification)
17	cut on $\Delta\eta$	electron selection cut (identification)
18	cut on $\Delta\phi$	electron selection cut (identification)
19	pixel tracking with vertex constraint	preparation of electron selection
20	regional iterative tracking for isolation	preparation of electron selection
21	cut on track isolation	electron selection cut (isolation)

- Priority is given to objects and variables that are least time-consuming and carry the largest background rejection power. This drives the ordering of modules in the electron sequence, as well as the ordering of sequences in a path: for example, in paths using both electrons and muons, the latter are treated first.

The four first steps of the electron sequence are related to the *clustering* of the electron energy in the ECAL, i.e. the algorithm that collects the energy which is usually spread in several crystals of the ECAL. The HLT upgrade of LS1 included the replacement of the Run I clustering algorithm by the *PF clustering* one, for both HLT electrons and HLT photons, thus ensuring synchronization with the similar change made in offline electron reconstruction. As mentioned in Section 2.2.2, PF clustering brings the possibility of sharing cluster energies among crystals, which allows for a finer reconstruction of individual showers, improving performance for low-energy clusters and robustness against pileup. Like in the offline algorithm, the resulting PF clusters are then used as inputs to a superclustering step, whereby the possible additional bremsstrahlung photons and the electrons arising from photon conversions are recovered and included in the PF superclusters. The implementation of this algorithm at HLT was achieved within the time budget thanks to the fact that it is processed regionally, around Level-1 candidates. Figure 3.4 (top) shows that the efficiency of PF clustering and superclustering at HLT

remains very similar to that the previous algorithm[2]. Most importantly, the energy resolution of PF superclusters benefits from new regression-based energy corrections, as illustrated in Fig. 3.4 (bottom).

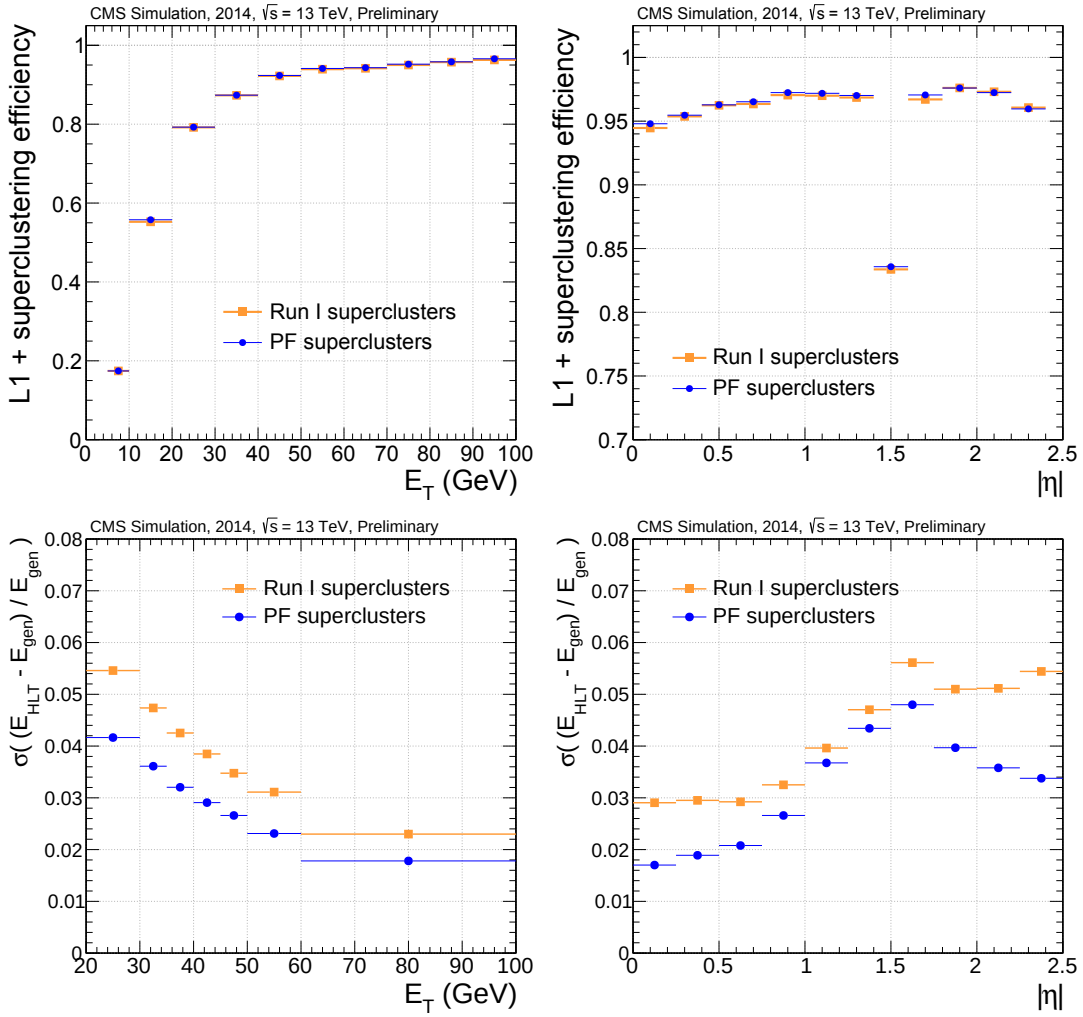

Figure 3.4: Efficiency of the HLT superclustering algorithm on electrons of $p_{\mathrm{T}}^{\mathrm{e}} > 5\,\mathrm{GeV}$ (top, including L1 efficiency) and energy resolution of HLT electron superclusters (bottom), as a function of the energy (left) and pseudorapidity (right) of the supercluster. Particle-flow superclustering is compared to the algorithm used in Run I. Electrons from $\mathrm{Z}\to\mathrm{e}^{+}\mathrm{e}^{-}$ simulated events at $\sqrt{s} = 13\,\mathrm{TeV}$, with an average number of 40 pileup interactions per event, and a bunch spacing of 25 ns. Superclusters are required to match ($\Delta\mathrm{R} < 0.1$) a generated electron from the Z decay.

Steps 13 to 15 of the electron sequence materialize the second part of electron reconstruction, namely the reconstruction of the electron track, which was also upgraded during LS1. As mentioned in Section 2.2.2, electrons have unique tracking challenges in CMS, because they often radiate bremsstrahlung photons in the silicon tracker before reaching the ECAL. In Run I, the tracking algorithm for HLT electrons was similar to the one used for other charged particles, and thus relied on the Kalman Filter, which

[2]This efficiency is close to 100%, whereas the plots exhibit lower values because they also include the Level-1 efficiency.

is known not to be optimal in this context because the energy loss of electrons cannot be accurately modelled with a single Gaussian distribution. Thus, HLT electron reconstruction has now been equipped with the electron-dedicated *GSF tracking* procedure which was already in use for offline electrons. The parameters of the algorithm are different at HLT, but track hits are still collected with a relaxed χ^2, in order to better accomodate deviations due to bremsstrahlung. Figure 3.5 (top left) confirms that more hits are collected this way. Then, the final fit also uses a Gaussian Sum Filter, whereby the energy loss in each layer is approximated as a sum of Gaussian functions with weights that described their associated probabilities. The estimation of the electron momentum is improved (Fig. 3.5, top right), as well as that of its direction (Fig. 3.5, bottom). Moreover, tests performed on 8 TeV data showed that GSF tracking at HLT provides a 25% rate reduction with respect to KF tracking, for a similar efficiency.

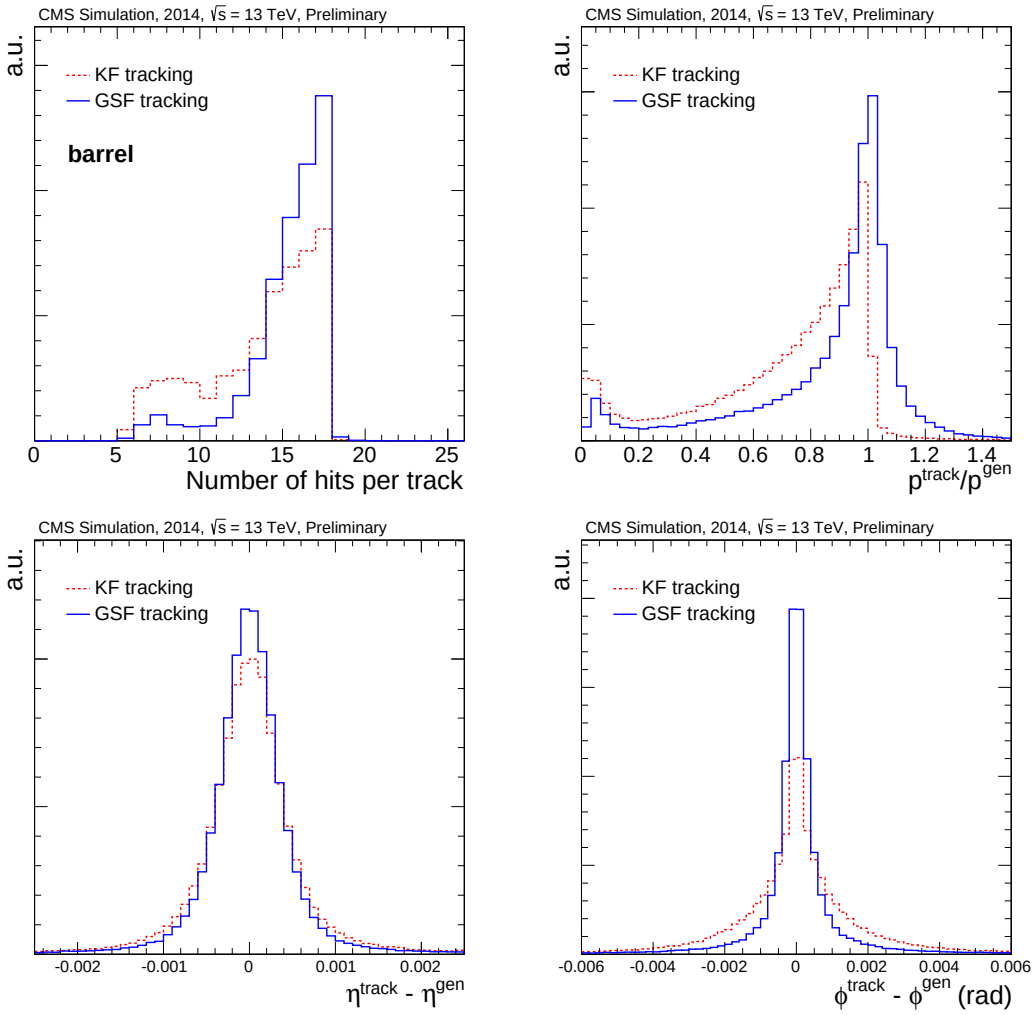

Figure 3.5: Number of reconstructed hits per track in the barrel (top left), ratio between the reconstructed track momentum and the generated one (top right), and difference between the reconstructed track direction and the generated one in the η (bottom left) and ϕ (bottom right) direction, for electrons tracks at HLT. Distributions are shown for tracks reconstructed with the Gaussian Sum Filter and the Kalman Filter. The selection of electrons is similar as in Fig. 3.4.

3.2.2 Online electron selection

The concept of electron selection at HLT covers three main types of requirements, which are applied via 9 different cuts in the standard HLT electron sequence of Table 3.1.

First, electrons are required to pass a lower threshold on the transverse energy $E_{\mathrm{T}}^{\mathrm{SC}}$ of their supercluster. Such a cut is the most important one in any HLT path, because it critically determines the rate of the path and its acceptance for the signal it targets. Low energy thresholds can only be afforded when several particles are combined in a path, and/or if other selection requirements are tight enough. It should be noted that the supercluster energy is used here, rather than the final electron momentum computed from the combination of the supercluster energy and the track momentum. Thanks to this choice, the cut can be applied early in the sequence (step 5) for the benefit of timing.

Second, 5 of the 9 electron selection variables pertain to electron *identification*, which refers to the quality requirements that select genuine electrons among all objects that can be reconstructed as electrons. These variables characterize the electron supercluster and its consistency with its associated electron track[3]:

- $\sigma_{i\eta i\eta}$ characterizes the width of the ECAL cluster in the η direction,
- H/E is the ratio of the HCAL energy collected behind the supercluster to the supercluster energy,
- $\frac{1}{E} - \frac{1}{p}$ measures the deviation of the supercluster energy from the electron momentum measured from the electron track,
- $\Delta\eta$ and $\Delta\phi$ describe the distance between the energy-weighted position of the supercluster and the track position of closest approach to the supercluster.

The ordering of the 5 cuts is dictated by timing constraints: $\sigma_{i\eta i\eta}$ (step 6) is exploited as soon as the supercluster is formed, H/E (step 10) requires to additionally retrieve information from the HCAL, and the three other variables (steps 16 to 18) have to wait for the electron track to be reconstructed.

Third, three variables pertain to the electron *isolation*, which is the main selection item that was revisited during LS1, and will now be described in detail.[4] Isolation refers to a set of powerful cuts that force the electron candidate to be located in a region devoid of other high-energy particles such as those which arise in jets. Thus, it both rejects the objects that are misidentified as electrons and the backgrounds from non-prompt electrons arising from hadronic processes, which are typically surrounded by large hadronic activity. Isolation is not relevant to all analyses, but it is of great importance to the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis, which relies on low-p_{T} isolated leptons.

In Run I, CMS has been using two different isolation paradigms.

- *Detector-based isolation*, which was used at HLT, is defined as the sum of calorimeter deposits and track transverse momenta in a geometrical cone of radius $\Delta\mathrm{R} = 0.3$ around the electron direction at the interaction vertex. This variable is made of three components, which are handled as three separate cuts at HLT:

[3] Compared to this HLT sequence, offline electron identification involves many more variables, as will be described in Section 4.1.2.

[4] The described optimization of isolation is used both for the selection of electrons and photons at HLT. Photons are reconstructed from the ECAL in a similar way as electron superclusters, and are known to behave quite analogously to electrons as far as isolation is concerned.

– *ECAL isolation* is defined as the sum of ECAL energy deposits in the isolation cone, divided by the transverse energy of the electron supercluster:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL}} \equiv \frac{1}{E^{\mathrm{SC}}_{\mathrm{T}}} \sum_{\mathrm{ECAL}} E_{\mathrm{T}} \ , \tag{3.1}$$

– *HCAL isolation* is the similar quantity for the HCAL:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL}} \equiv \frac{1}{E^{\mathrm{SC}}_{\mathrm{T}}} \sum_{\mathrm{HCAL}} E_{\mathrm{T}} \ , \tag{3.2}$$

– *tracker isolation* is the sum of transverse momenta of tracks in the cone (whereby said tracks are reconstructed via a dedicated regional tracking around the electron track), divided by the transverse momentum of the electron:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{tr.}} \equiv \frac{1}{p^{\mathrm{e}}_{\mathrm{T}}} \sum_{\mathrm{tracks}} p_{\mathrm{T}} \ . \tag{3.3}$$

- *Particle-flow isolation* is a more recent approach that has been used offline in most physics analyses that involve electrons. In this algorithm, the CMS PF algorithm is run locally, and the isolation variable is built from the transverse momenta of individual PF candidates, divided by the transverse momentum of the electron. There are again three components, because three types of PF candidates are considered:

  – *photon isolation* uses the sum of the transverse momenta of photons in the isolation cone:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{pho.}} \equiv \frac{1}{p^{\mathrm{e}}_{\mathrm{T}}} \sum_{\mathrm{photons}} p_{\mathrm{T}} \ , \tag{3.4}$$

  – *neutral hadron isolation* uses neutral hadrons instead:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{n.h.}} \equiv \frac{1}{p^{\mathrm{e}}_{\mathrm{T}}} \sum_{\substack{\mathrm{neutral} \\ \mathrm{hadrons}}} p_{\mathrm{T}} \ , \tag{3.5}$$

  – *charged hadron isolation* uses charged hadrons originating from the primary vertex:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{c.h.}} \equiv \frac{1}{p^{\mathrm{e}}_{\mathrm{T}}} \sum_{\substack{\mathrm{charged} \\ \mathrm{hadrons}}} p_{\mathrm{T}} \ . \tag{3.6}$$

Figure 3.6 shows two diagrams that illustrate the detector-based and Particle-flow isolation schemes. The three components of the former cannot be compared one by one to the those of the latter, but there is a strong correlation between ECAL isolation and photon isolation on the one hand, and between tracker isolation and charged hadron isolation on the other hand. The discrimination power of HCAL isolation is shared between neutral hadron and charged hadron isolation. Several benefits are expected from the PF approach: it avoids energy double counting, it better handles detector noise thanks to its reliance on PF clustering, and it circumvents a problem of the detector-based approach, namely the necessity of manually removing the energy of the electron from its own isolation variables.

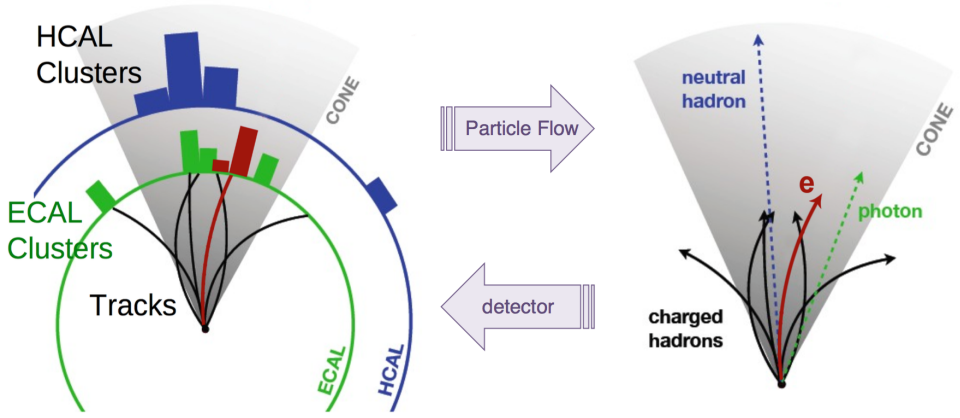

Figure 3.6: Illustration of two schemes of isolation: the detector-based approach used at HLT in Run I (left) and the Particle-flow approach used offline (right).

Besides these two strategies, I implemented a third, intermediate algorithm called *PF-cluster isolation* for the first time. As opposed to fully-fledged PF isolation, the step of assembling PF inputs into PF candidates, which is the most time-consuming one, is not run at all in this method. Instead, PF clusters and tracks are directly used to build three isolation components, in a similar way as in detector-based isolation:

– *ECAL PF-cluster isolation* is defined as the sum of transverse energies of ECAL PF clusters in the isolation cone, divided by the transverse energy of the electron supercluster:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}} \equiv \frac{1}{E^{\mathrm{SC}}_{\mathrm{T}}} \sum_{\substack{\mathrm{ECAL} \\ \mathrm{PF\ clusters}}} E_{\mathrm{T}} \ , \tag{3.7}$$

– *HCAL PF-cluster isolation* is the similar quantity for the HCAL:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}} \equiv \frac{1}{E^{\mathrm{SC}}_{\mathrm{T}}} \sum_{\substack{\mathrm{HCAL} \\ \mathrm{PF\ clusters}}} E_{\mathrm{T}} \ , \tag{3.8}$$

– *iterative track isolation* is similar to tracker isolation, except that the regional tracking that provides the tracks in the isolation cone is based on the new iterative tracking algorithm used in PF reconstruction, instead of the Run I tracking algorithm:

$$\mathcal{I}^{\mathrm{e}}_{\mathrm{it.tr.}} \equiv \frac{1}{p^{\mathrm{e}}_{\mathrm{T}}} \sum_{\substack{\mathrm{iter.} \\ \mathrm{tracks}}} p_{\mathrm{T}} \ . \tag{3.9}$$

The idea behind this intermediate algorithm is to decouple the effects of energy double counting from those of the new clustering and tracking used in PF reconstruction. Like in PF isolation, the electron footprint can be easily removed from the isolation sums.

The performance of the three available techniques was studied in simulated samples at $\sqrt{s} = 13\,\mathrm{TeV}$, with an average number of 40 pileup interactions and a bunch spacing of 25 ns, and selecting HLT reconstructed electrons that pass a $E^{\mathrm{SC}}_{\mathrm{T}} > 27\,\mathrm{GeV}$ requirement. True electrons were taken from a $\mathrm{Z} \to \mathrm{e}^{+}\mathrm{e}^{-}$ sample and geometrically matched

($\Delta R < 0.1$) to generator-level electrons from the Z boson decay, whereas fake electrons were retrieved from a QCD sample containing jets in a similar energy range. ROC curves were used to compare the three methods, as illustrated in the plots of Fig. 3.7, for the three components of each methods and for their sum. As already mentioned, comparing the individual components can be deceptive in the case of PF isolation; thus, the sum of components is used as the main estimator of overall performance, even though three separate cuts are applied. Results show that both PF-related approaches outperform the detector-based one, and that PF-cluster isolation does a little better than PF isolation.

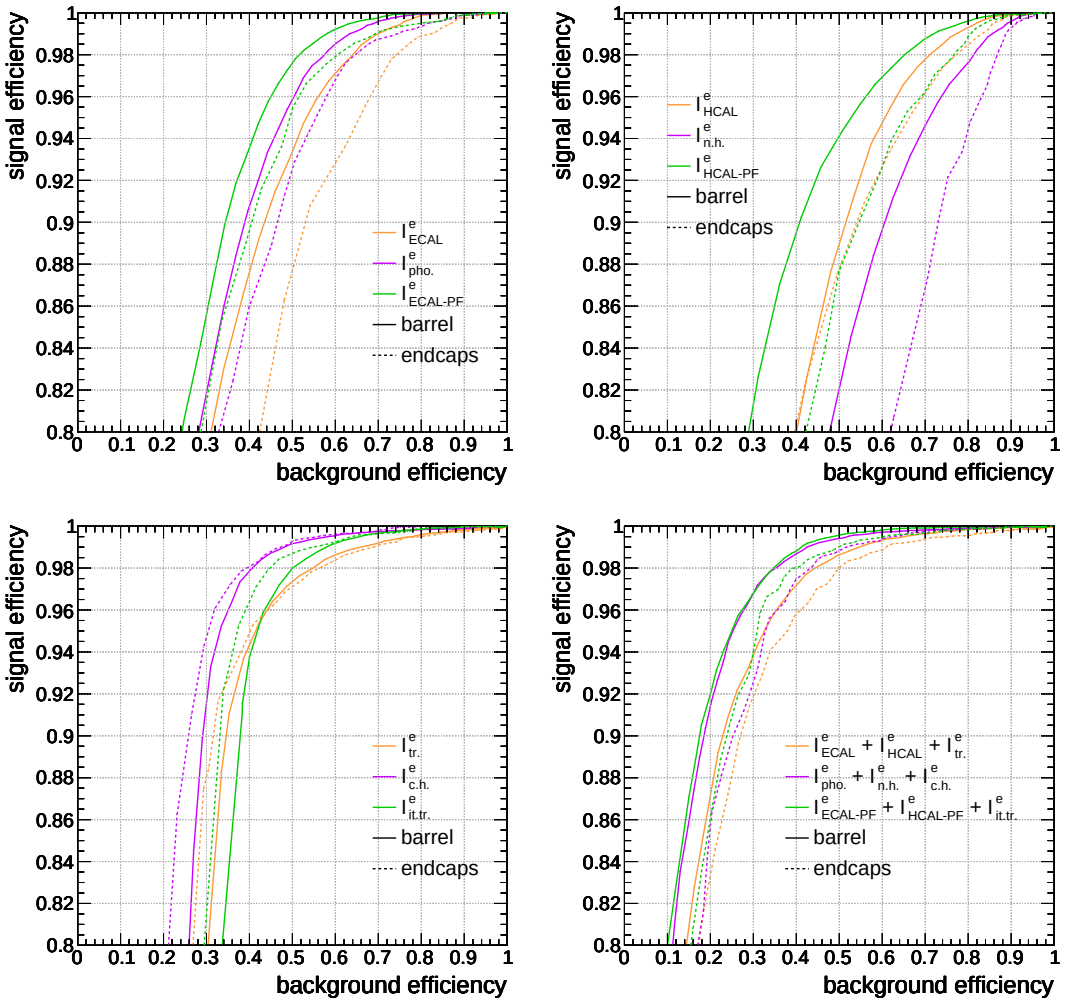

Figure 3.7: ROC curves for electron isolation at HLT, comparing the detector-based isolation setup used in Run I (orange curves) to PF isolation (violet curves) and PF-cluster isolation (green curves). Comparisons are first shown individually for ECAL-related (top left), HCAL-related (top right) and tracking-related (bottom left) components, and for a combined isolation variable that sums the three components (bottom right).

While ROC curves describe the compromise between signals and backgrounds (which translates to efficiencies and event rates in the case of the HLT), the final choice of an isolation strategy also depends on the third metric of HLT optimization, i.e. timing. In this respect, PF-cluster isolation is algorithmically similar to detector-based isolation,

because the cuts on $\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}}$, $\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}}$, and $\mathcal{I}^{\mathrm{e}}_{\mathrm{it.tr.}}$ (steps 9, 12, and 21 in Table 3.1) are applied at the same place of the electron sequence as the cuts on $\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL}}$, $\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL}}$, and $\mathcal{I}^{\mathrm{e}}_{\mathrm{tr.}}$. PF isolation, on the other hand, requires to assemble PF inputs into PF candidates, which not only is time-consuming, but also implies that $\mathcal{I}^{\mathrm{e}}_{\mathrm{pho.}}$, $\mathcal{I}^{\mathrm{e}}_{\mathrm{n.h.}}$, and $\mathcal{I}^{\mathrm{e}}_{\mathrm{c.h.}}$ can only be defined after the regional iterative tracking (step 20), which thus has to be processed around more fake electrons. Numerical studies confirmed the expectations: for instance, a measurement of the average processing time of a single-electron path with the aforementioned background sample provided values of 8.5 ms, 10.3 ms, and 8.1 ms for detector-based, PF and PF-cluster isolation, respectively. Together with the outcome of ROC curves, this led to choose PF-cluster isolation as the isolation technique for HLT electrons in Run II.

Besides these studies, electron isolation was also assigned a pileup subtraction strategy, for the first time at the HLT. Indeed, pileup interactions (both from in-time and out-of-time crossings) result in undesirable energy deposits in the subdetectors, which shifts the values of isolation variables upwards, no matter which of the three isolation schemes is chosen. As a result, if nothing is done, the isolation efficiency exhibits a decreasing trend as a function of the number n_{v} of vertices per event. Two changes have been made to mitigate this effect. First, tracker-related components such as $\mathcal{I}^{\mathrm{e}}_{\mathrm{it.tr.}}$ have been protected by only considering tracks that originate from the same primary vertex as the electron track, thus discarding pileup tracks. Second, calorimeter-related components such as $\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}}$ and $\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}}$ have been applied the FastJet technique [61, 62, 59], which is also used for offline electrons (see Section 4.1.3). This consists in making the transformation:

$$\sum_{\text{PF clusters}} E_{\mathrm{T}} \quad \longrightarrow \quad \sum_{\text{PF clusters}} E_{\mathrm{T}} \; - \; \rho \times A_{\mathrm{eff}} \; , \tag{3.10}$$

where ρ is defined as the median of the energy-density distribution of neutral particles within the area of any jet in the event, and the effective area A_{eff} is a scale factor proportional to the geometric area of the isolation cone. Since both the average isolation sum and the average ρ are empirically found to be affine functions of n_{v}, defining A_{eff} as the ratio of the slopes of these affine functions is expected to make the average of the transformed isolation (Eq. 3.10) independent on n_{v}. Figure 3.8 illustrates the impact of such a correction on the signal efficiency versus n_{v}, which indeed becomes flatter. Since the correction works on average over a large sample, it also slightly degrades the signal-to-background compromise. Moreover, values of A_{eff} need to be determined beforehand, separately for every isolation component and for different data taking conditions, and they are η-dependent. I thus tuned some appropriate A_{eff} values for $\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}}$ and $\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}}$, in the barrel and the endcaps, and separately for the triggers targeting 50 ns and 25 ns runs.

3.2.3 Building trigger paths

Besides the electron-related improvement described in Sections 3.2.1 and 3.2.2, several upgrades were also performed in parts of the HLT algorithm related to other types of particles: for instance, isolation was also implemented for the first time in muon sequences. With these tools in hand, the CMS trigger group proceeded to design new HLT menus for Run II, i.e. the entire sets of HLT paths that would trigger the events relevant to all physics analyses of CMS, and to calibration and monitoring. Instead of just updating the Run I menus, it was decided to restart the menu from a blank

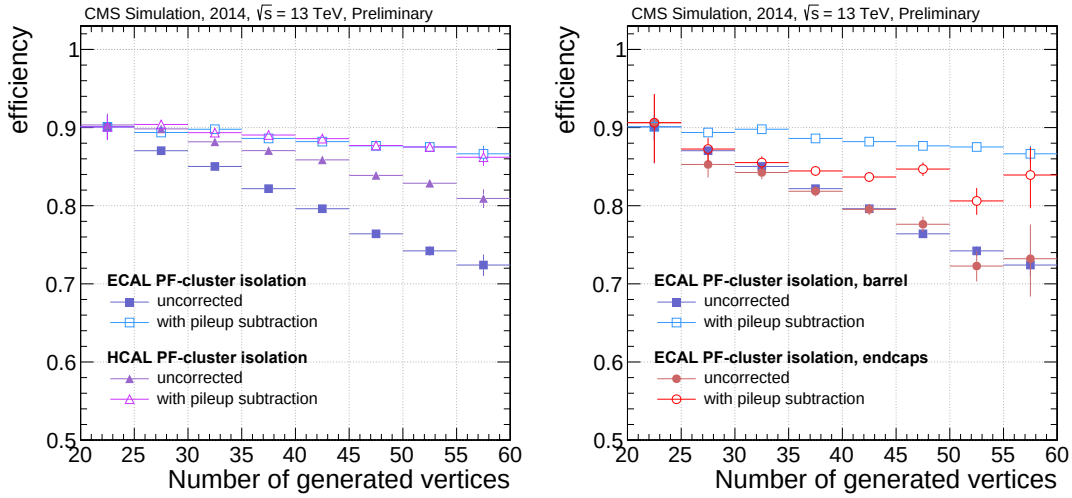

Figure 3.8: Signal efficiency of PF-cluster isolation vs. the number of generated vertices, with and without pileup subtraction. Comparison of (left) ECAL and HCAL PF-cluster isolation in the barrel, and (right) ECAL PF-cluster isolation in the barrel and in the endcaps. The operating points are adapted for illustration purposes.

page, given how different the algorithms and data taking conditions of Run II were. To organize the work, preliminary menus were required to target three main scenarios, each of which was studied with dedicated simulated samples:

- The 1.4e34 scenario targeted what was hoped to be the main production regime of early Run II, with a luminosity expectation of $1.4 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, a bunch spacing of 25 ns, and an average number of 40 pileup interactions per event.
- The 7e33 scenario corresponded to a less optimistic hypothesis of $7\times10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, still for 25 ns runs, and with 20 pileup interactions per event on average.
- The 5e33 scenario was designed for early runs at a Run I-like bunch spacing of 50 ns, for a luminosity hypothesis of $5 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The 5e33 menu was interfaced with a modified version of the Run I Level-1 trigger, unlike the 1.4e34 and 7e33 menus which relied on the new Stage 1 trigger.

This section reports the design of the subset of the menu that is meant to trigger the events that will be used in the $\mathrm{H}\to\mathrm{ZZ}\to4\ell$ analysis. To this end, a collection of HLT paths has to be built out of electron and muon sequences, and the thresholds for momentum, lepton identification, and lepton isolation cuts have to be tuned so that the set of paths fits the rate and timing constraints and selects the signal with as good efficiency as possible in all three final states (4e, 4μ and $2\mathrm{e}2\mu$). The 4ℓ channel can be viewed as privileged trigger-wise, because the presence of four leptons in the final state makes triggering easier for combinatorial reasons. Nevertheless, these electrons and muons generally have low p_T, and do not necessarily have clean enough signatures to be selected with tight identification and selection cuts. The main challenge concerns the electrons, because the background objects that can mimic prompt electrons come at a much larger rate than for muons.

The rate computation procedure here consists in running the HLT on a well-chosen set of simulated samples, covering the processes that are liable to contribute a significant

fraction of the output of the HLT menu (in this case: QCD, $Z \to \ell\ell$, and $W \to \ell\nu$). The timing constraint does not really affect the optimization of $H \to ZZ \to 4\ell$ paths, because it has already been brought under control in particle-level studies such as mentioned in Sections 3.2.1 and 3.2.2. The measurement of efficiency is specific to the signal, and it is achieved here by running the HLT over $H \to ZZ \to 4\ell$ simulated samples for the dominant ggH production mode.[5] Efficiencies are computed with a dedicated kinematic selection on generator-level leptons, which puts them in a similar region of the phase space as defined by the offline kinematic cuts that will be described in Sections 4.1.1 and 5.2.2.[6] To estimate the cost and benefit of a particular path within a larger set of paths, its *pure efficiency* and its *pure rate* are computed, i.e. the amount of additional efficiency and rate that this path brings on top of the other paths of the set.

Table 3.2 presents an example of a possible set of paths for the $H \to ZZ \to 4\ell$ channel, with results of rate and efficiency computations. This illustrates the general strategy, which relies on two classes of HLT paths:

- Most of the efficiency is provided by *dilepton paths*, using asymmetric p_T thresholds and lepton isolation. Dielectron, dimuon, electron+muon, and muon+electron combinations are all considered.
- To cover the parts of the phase space that are missed by dilepton paths, some *trilepton paths* are added, relaxing the p_T thresholds and removing isolation requirements. Every possible combination of three leptons is covered.

All 4ℓ-dedicated dilepton and trilepton paths are collectively referred to as *multilepton paths*.

Negotiation within the CMS trigger community led to assign a rate budget of about 85 Hz to the set of all multilepton paths in the 1.4e34 scenario. These paths were built on top of Level-1 seeds that were in turn part of a 1.4e34 Level-1 menu. The E_T thresholds of these seeds were defined upstream according to rate constraints, and their values influenced the choice of E_T and p_T thresholds of the corresponding multilepton HLT paths. The operating points for the 8 cuts of electron identification and isolation were optimized in a way that eases bookkeeping and monitoring: all multilepton paths that use electrons have the same identification thresholds, and all dileptons paths that use electrons have the same isolation thresholds. Of the several sets of operating points that were tested, the final chosen one is presented in Table 3.3, with its two different versions: one with isolation for the dilepton paths (left), and one without isolation for the trilepton paths (right). The distributions of the 8 variables are illustrated for signal electrons in Fig. 3.9, together with the final thresholds. It was not deemed necessary to cut on the $\frac{1}{E} - \frac{1}{p}$ variable to satisfy the rate constraint.

Most of the above optimization work was first performed for the 1.4e34 scenario, i.e. the most demanding one. Then, to design multilepton paths for the 7e33 scenario, it was decided to keep the same identification and isolation operating points and to relax some of the E_T and p_T cuts of the dilepton paths, thus accompanying a similar lowering of E_T thresholds in the 7e33 Level-1 menu. Finally, for the 5e33 scenario, the HLT

[5] The other, subdominant modes were not explicitly studied at trigger level, but the kinematics and the reconstruction quality of the leptons that they involve is not expected to be significantly different for the sake of measuring trigger efficiencies.

[6] Since this kinematic-based denominator selection does not take the impact of offline lepton quality into account, the resulting efficiencies are lower than those measured with respect to the event selection of the final analysis. The latter efficiency measurement will be described in Section 5.2.1.

Table 3.2: Example of the computation of rates and efficiencies of a set of multilepton paths designed for the 1.4e34 scenario. The simulated samples have an average number of 40 pileup interactions per event, and the rates are computed assuming an instantaneous luminosity of $1.4 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The efficiencies are computed separately for the three final states of the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ processes, using the denominator selection described in the text.

requirements of the HLT path	efficiency (pure efficiency)			rate (pure rate)
	$\mathrm{H}\to\mathrm{ZZ}\to 4\mathrm{e}$	$\mathrm{H}\to\mathrm{ZZ}\to 4\mu$	$\mathrm{H}\to\mathrm{ZZ}\to 2\mathrm{e}2\mu$	
2 isolated electrons, $E_\mathrm{T} > 23$ and 12 GeV	85.9 % (18.9 %)		39.1 % (0.9 %)	13.7 Hz (13.5 Hz)
3 non-isolated electrons, $E_\mathrm{T} > 16$, 12 and 8 GeV	70.8 % (3.8 %)			4.4 Hz (4.2 Hz)
2 isolated muons, $p_\mathrm{T} > 17$ and 8 GeV		98.6 % (0.0 %)	59.9 % (0.0 %)	24.7 Hz (2.8 Hz)
1 isolated muon, $p_\mathrm{T} > 17$ GeV 1 isolated tracker muon, $p_\mathrm{T} > 8$ GeV		99.6 % (0.5 %)	63.6 % (0.6 %)	34.0 Hz (12.2 Hz)
3 non-isolated muons, $p_\mathrm{T} > 12$, 10 and 5 GeV		97.1 % (0.4 %)		7.0 Hz (5.7 Hz)
1 isolated electron, $E_\mathrm{T} > 23$ GeV 1 isolated muon, $p_\mathrm{T} > 8$ GeV			60.7 % (2.8 %)	20.5 Hz (18.3 Hz)
1 isolated muon, $p_\mathrm{T} > 23$ GeV 1 isolated electron, $E_\mathrm{T} > 12$ GeV			46.8 % (0.9 %)	4.0 Hz (2.3 Hz)
2 non-isolated electrons, $E_\mathrm{T} > 12$ GeV each 1 non-isolated muon, $p_\mathrm{T} > 8$ GeV			48.3 % (0.6 %)	2.6 Hz (2.3 Hz)
2 non-isolated muons, $p_\mathrm{T} > 9$ GeV each 1 non-isolated electron, $E_\mathrm{T} > 9$ GeV			62.9 % (1.0 %)	3.6 Hz (2.7 Hz)
All paths	89.7 %	99.8 %	97.0 %	88.3 Hz

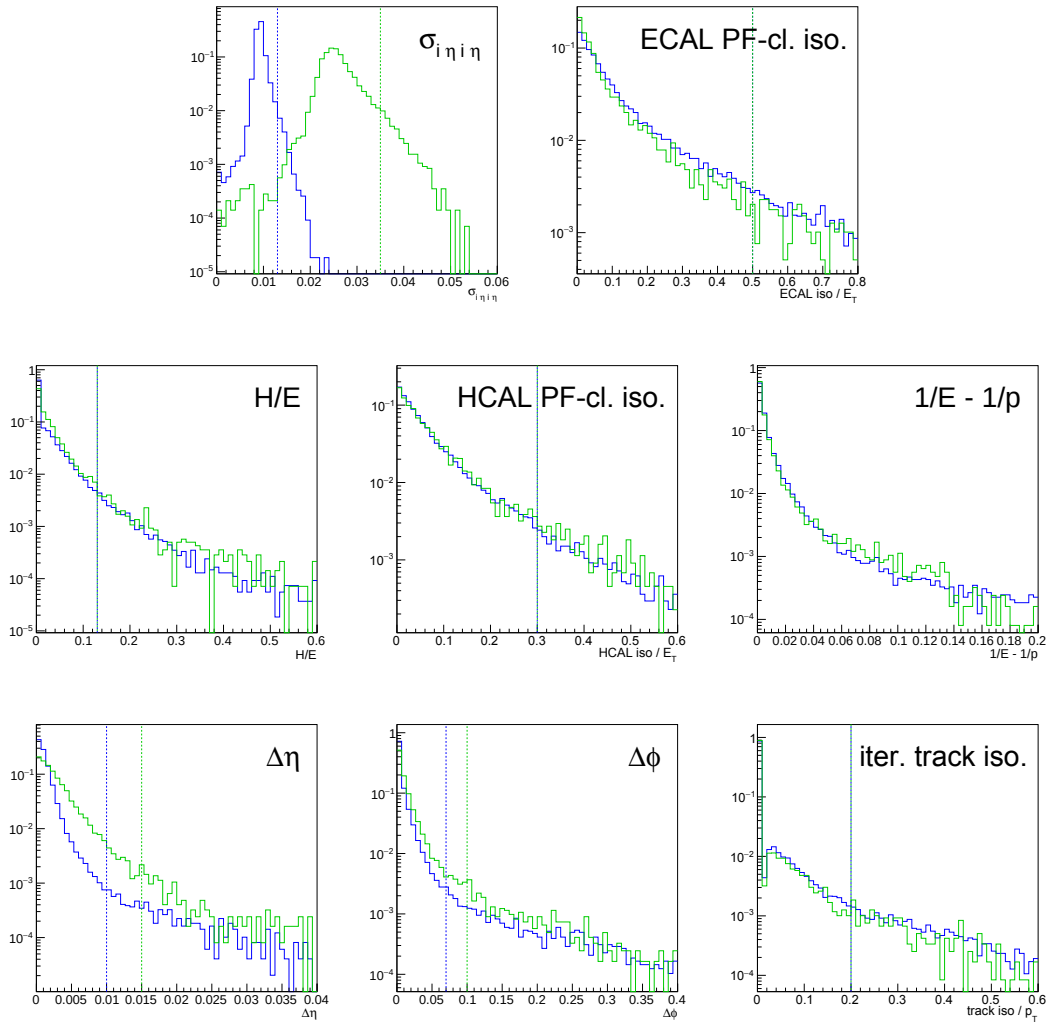

Figure 3.9: Distributions of the 8 HLT electron identification and isolation variables, obtained from running the double electron path of Table 3.2 relaxing all cuts, on a signal $H \to 4\ell$ sample simulated with average pileup 40. True HLT electron candidates are selected with a geometrical matching to generated electrons. The blue and green histograms show the distributions for barrel and endcap electrons, respectively, while the blue and green vertical dashed lines show the corresponding final operating points. The plots are shown in the order in which the cuts are applied.

Table 3.3: Final 2015 operating point for electrons in (left) dilepton and (right) trilepton HLT paths.

filter	threshold	
	barrel	endcaps
$\sigma_{i\eta i\eta}$	0.013	0.035
$\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}}$	0.5	0.5
H/E	0.13	0.13
$\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}}$	0.3	0.3
$\frac{1}{E} - \frac{1}{p}$	-	-
$\Delta\eta$	0.01	0.015
$\Delta\phi$	0.07	0.1
$\mathcal{I}^{\mathrm{e}}_{\mathrm{it.tr.}}$	0.2	0.2

filter	threshold	
	barrel	endcaps
$\sigma_{i\eta i\eta}$	0.013	0.035
$\mathcal{I}^{\mathrm{e}}_{\mathrm{ECAL-PF}}$	-	-
H/E	0.13	0.13
$\mathcal{I}^{\mathrm{e}}_{\mathrm{HCAL-PF}}$	-	-
$\frac{1}{E} - \frac{1}{p}$	-	-
$\Delta\eta$	0.01	0.015
$\Delta\phi$	0.07	0.1
$\mathcal{I}^{\mathrm{e}}_{\mathrm{it.tr.}}$	-	-

menu of the 7e33 scenario was re-interfaced with the modified Run I Level-1 trigger. Since event reconstruction in the first 2015 runs was expected to happen with a still limited control over the alignment of the detector, the cuts on the alignment-sensitive variables $\Delta\eta$ and $\Delta\phi$ were removed from the operating points of Table 3.3. The final list of all multilepton paths is presented in Table 3.4, along with their L1 seeds and with temporary prescale values that indicate whether the paths are activated in the 1.4e34 or in the 7e33 scenario. Two remarks can be made:

- The `_DZ` label in the names of the dielectron and dimuon paths refers to an additional cut, whereby the two triggered leptons are required to have nearby longitudinal positions at the interaction vertex. Such a cut decreases the rate of these paths by about 60%, for an efficiency loss of 0.2 to 0.8%.

- Besides the main signal paths, the menus include some prescaled control paths, to serve as backup plans or for monitoring and efficiency measurements. In the present study, dielectron paths without the vertex constraint were included, as well as single-electron paths corresponding to halves of the dielectron ones.

It should be noted that this is the list of paths and prescales as defined in early 2015, i.e. before the first pp collisions of Run II. A much more complete prescale strategy was then designed to cover all possible levels of instantaneous luminosity encountered during data taking. Moreover, minor changes are steadily being introduced in the HLT software as the data taking progresses, but the overall strategy is not altered. The underlying Stage 1 Level-1 trigger was replaced by the Stage 2 in early 2016, which means that the multilepton HLT paths used as of 2016 are seeded by different algorithms.

Besides the multilepton paths, it was decided to use some of the single-lepton paths that have been developed by other CMS groups for other physics analyses: indeed, even though these paths have larger p_{T} thresholds and tighter identification requirements, they were shown to trigger part of the small fraction of $\mathrm{H} \to 4\ell$ events that escape multilepton paths. The exact list of HLT paths and L1 seeds that have actually been exploited by the $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ analysis in 2015 and 2016 will be quoted in Tables 6.2 and 7.2, in the sections that describe the corresponding data samples.

Table 3.4: List of final multilepton HLT paths, with their L1 seeds in the Stage 1 menu, and their prescales in preliminary menus corresponding to the 7e33 and 1.4e34 scenarios. Paths and seeds are named after the internal nomenclature of CMS: the names include the number and the type of triggered objects (Ele/Mu for electrons/muons), the value of E_{T} or p_{T} thresholds in GeV, and some other labels that characterize the identification and isolation operating points.

HLT path	L1 seed	L1*HLT prescale	
		7e33	1.4e34
Signal paths			
HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL_DZ	L1_DoubleEG_22_10	1	1
HLT_Ele17_Ele12_CaloIdL_TrackIdL_IsoVL_DZ	L1_DoubleEG_15_10	1	0
HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL	L1_TripleEG_14_10_8	1	1
HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ	L1_DoubleMu_12_5 / _10_3p5	1	1
HLT_Mu17_TrkIsoVVL_TkMu8_TrkIsoVVL_DZ	L1_DoubleMu_12_5 / _10_3p5	1	1
HLT_TripleMu_12_10_5	L1_TripleMu_5_5_3	1	1
HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL	L1_Mu5_EG20	1	1
HLT_Mu8_TrkIsoVVL_Ele17_CaloIdL_TrackIdL_IsoVL	L1_Mu5_EG15	1	0
HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL	L1_Mu20_EG10	1	1
HLT_Mu17_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL	L1_Mu12_EG10	1	0
HLT_Mu8_DiEle12_CaloIdL_TrackIdL	L1_Mu6_DoubleEG10	1	1
HLT_DiMu9_Ele9_CaloIdL_TrackIdL	L1_DoubleMu7_EG7	1	1
Control paths			
HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL	L1_DoubleEG_22_10	5	10
HLT_Ele17_Ele12_CaloIdL_TrackIdL_IsoVL	L1_DoubleEG_15_10	10	0
HLT_Ele23_CaloIdL_TrackIdL_IsoVL	L1_SingleEG_20	2000	4000
HLT_Ele12_CaloIdL_TrackIdL_IsoVL	L1_SingleEG_10	10000	20000

3.3 Conclusion

A series of advanced reconstruction and selection techniques was implemented in the CMS HLT in order to maintain high triggering performance for electrons and photons in the challenging luminosity and pileup conditions of the Run II data taking of the LHC. A set of customized trigger paths was then defined to specifically select $\mathrm{H} \to 4\ell$ events with a very large efficiency, while satisfying the harsh rate constraints.

The fraction of events that pass at least one HLT path can then be fully exploited offline for physics analyses, without the constraint of real-time processing. From this point on, *reconstruction* gets decoupled from *selection*:

- Event reconstruction (previously described in Section 2.2.2) is performed centrally for all of CMS as soon as the data are recorded. It comes with an average timing of a few tens of seconds per event, and a rate constraint which is satisfied by construction since this is the one applied at HLT: all triggered events are reconstructed.
- Event selection involves physicists as end users of the sets of reconstructed events. It is of course completely specific to physics analyses, and free from timing or rate constraints.

For clarity, the study of the $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ channel will now be split in two chapters, starting with the selection of individual physics objects in Chapter 4, and moving on to the event-level selection, signal extraction and measurements in Chapter 5.

Chapter 4

Selecting particles and jets

As events that are selected *online* by the HLT enter the *offline* world, they are processed through the full CMS event reconstruction described in Section 2.2.2. For every event, this centrally-managed step results in a large collection of *physics objects*, such as particle candidates, jets and missing transverse energy, all with mild quality requirements. Objects that are specifically sought for in the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis are prompt electrons and muons, photons originating from final-state radiation, and jets.

This chapter describes the analysis-specific selection that all of these particles and jets are required to pass before they are used to build event-level observables. This starts with the selection of offline electrons and muons, which has a certain similarity with the HLT selection described in Chapter 3, but is optimized, performed, and monitored independently from it. The main conceptual difference is that the rate and timing constraints do not apply anymore here: the only metric for optimization is the sensitivity of the physics analysis. In parallel, one has to ensure that the involved selection variables behave coherently in data and simulation, and to correct for possible modelling imperfections.

Most variables discussed in this chapter were developed in the past for the CMS analyses of Run I, but they had to be fully revisited for the Run II restart at a new centre-of-mass energy and in new pileup conditions. I mainly contributed to the optimization of new operating points for electron identification, lepton impact parameter cuts, and lepton isolation, and I was involved in the commissioning of all main lepton selection variables. Moreover, since jets play a crucial part in the extraction of Higgs boson production modes, I also studied b tagging and quark/gluon tagging variables, exploiting these techniques for the first time in the $\mathrm{H}\to 4\ell$ analysis.

4.1 Lepton selection

In the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis, the selection of leptons (hereafter referring to electrons and muons) involves three main types of requirements: impact parameter, identification, and isolation. Optimally exploiting these various tools is all the more important as no fewer than four leptons have to be selected, which implies that the per-lepton efficiencies propagate to event selection with a power four. As a consequence, the present lepton selection is quite loose in comparison with other analyses using leptons.

Throughout this section, all studies of lepton selection optimization rely on a common setup, which consists in comparing two different sources of leptons from simulated samples:

- *Signal leptons* are taken from a $gg \to H(125) \to ZZ \to 4\ell$ sample, and are defined as the reconstructed electrons (muons) that are geometrically matched ($\Delta R < 0.1$) to generated prompt electrons (muons) from the Higgs boson decay. These are exactly the objects that should be selected to form four-lepton candidates for the analysis.

- *Background leptons* (or *fake leptons*) are taken from a Drell-Yan + jets sample, and are defined as the reconstructed electrons (muons) that are geometrically separated ($\Delta R > 0.1$) from both generated prompt electrons (muons) from $Z^{(*)}/\gamma^*$ decay. These objects can for instance be jets misidentified as electrons, or secondary leptons from decays of hadrons or photon conversions, and they are here taken from the Drell-Yan + jets process because it constitutes the largest reducible background contribution to the signal region.

4.1.1 Lepton preselection

Electron and muon reconstruction have been described in Section 2.2.2. The electron and muon candidates that are considered at the beginning of the present selection flow are not the PF electrons and PF muons, but simply the reconstructed electron and muon candidates that the PF algorithm uses as 'bricks'. The identification requirements of the PF muons will be applied later, while those of the PF electrons are not applied at all and will be replaced by a specific selection.

Reconstruction acceptance

Since electrons are only reconstructed within the extent of silicon tracker, they are applied a pseudorapidity cut of $|\eta^e| < 2.5$, while the configuration of the outer muon system dictates a muon acceptance cut of $|\eta^\mu| < 2.4$. To be considered for the analysis, reconstructed electrons are required to have a transverse momentum p_T^e larger than 7 GeV, while muons are used down to $p_T^\mu = 5$ GeV. This momentum acceptance is mostly limited by the increase of fake leptons at very low p_T^ℓ and the difficulty of reliably measuring efficiencies and momentum scales.

Early preselection

Electron candidates are preselected using loose cuts on track-cluster matching observables, so as to preserve the highest possible efficiency while already rejecting part of the jet background.

Muons are preselected as either global muons or tracker muons. Standalone muons tracks that are only reconstructed in the muon system are rejected. Tracker muons that are not global muons are required to be arbitrated, and if two muons share 50% or more of their segments, the lower-quality one is rejected.

Impact parameter cuts

Non-prompt muons that originate from in-flight decays of hadrons and cosmic rays, and electrons from photon conversions can be suppressed using the fact that their track does not point to the *primary collision vertex* of the event. This vertex is defined by first associating tracks to reconstructed primary vertices based on vertex fits and proximity criteria, then clustering said tracks with a jet algorithm, and finally picking the primary

vertex that has the largest sum of squared transverse momenta of the resulting clusters and the per-vertex missing momentum.

Electrons and muons are first required to satisfy the following loose vertex constraint:

$$d_{xy} < 0.5\,\text{cm} \quad \text{and} \quad d_z < 1\,\text{cm}\ , \tag{4.1}$$

where d_{xy} and d_z are the electron or muon impact parameter with respect to the primary collision vertex in the transverse plane and in the longitudinal direction, respectively. These two variables are illustrated in Fig. 4.1 for signal and background electrons. Distributions for muons are similar.

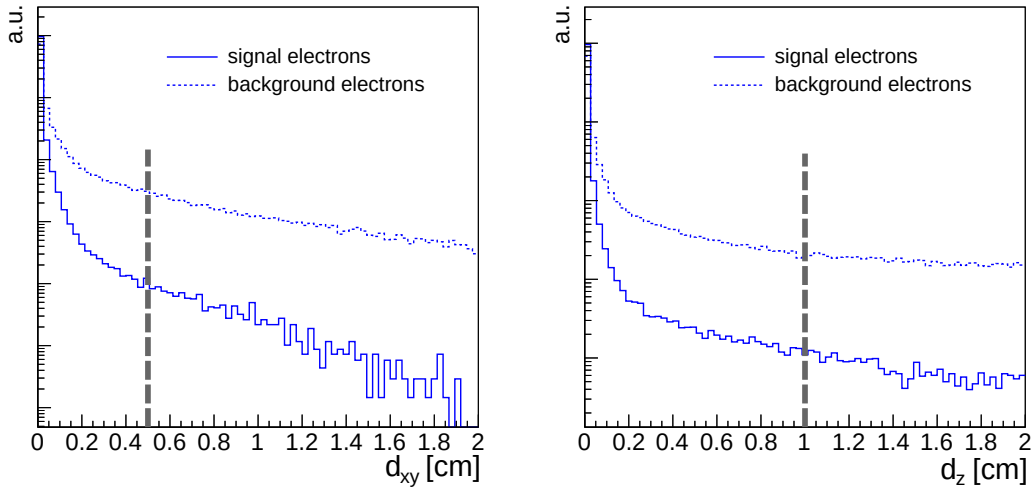

Figure 4.1: Distribution of the d_{xy} (left) and d_z (right) impact parameter variables for signal and background electrons. The two types of electrons are defined in the text, and the acceptance and early preselection cuts are applied. The vertical grey dashed lines denote the operating points.

A second vertex-related cut is applied as

$$\text{SIP}_{3\text{D}} < 4\ , \tag{4.2}$$

whereby the 3D impact parameter significance $\text{SIP}_{3\text{D}}$ is defined as the ratio of the impact parameter of the lepton track ($\text{IP}_{3\text{D}}$) in three dimensions with respect to the primary collision vertex position, and its uncertainty:

$$\text{SIP}_{3\text{D}} \equiv \frac{\text{IP}_{3\text{D}}}{\sigma_{\text{IP}_{3\text{D}}}}\ . \tag{4.3}$$

This variable is illustrated in Fig. 4.2 for signal and background leptons. Moreover, to assess the modelling of $\text{SIP}_{3\text{D}}$, Fig. 4.3 compares its distributions in 2016 data and simulation, for electrons and muons. This plot and similar ones are obtained from a $\text{Z}+1\ell$ control region, which is defined by requiring the presence of a pair of opposite-sign same-flavour leptons that pass the full lepton selection and a $60 < m_{\ell^+\ell^-} < 120\,\text{GeV}$ invariant mass cut, plus exactly one additional lepton passing all the above requirements but the $\text{SIP}_{3\text{D}}$ one. The latter lepton is the one used in the plot. Distributions of $\text{SIP}_{3\text{D}}$ are found to be moderately well modelled in this background-dominated region.

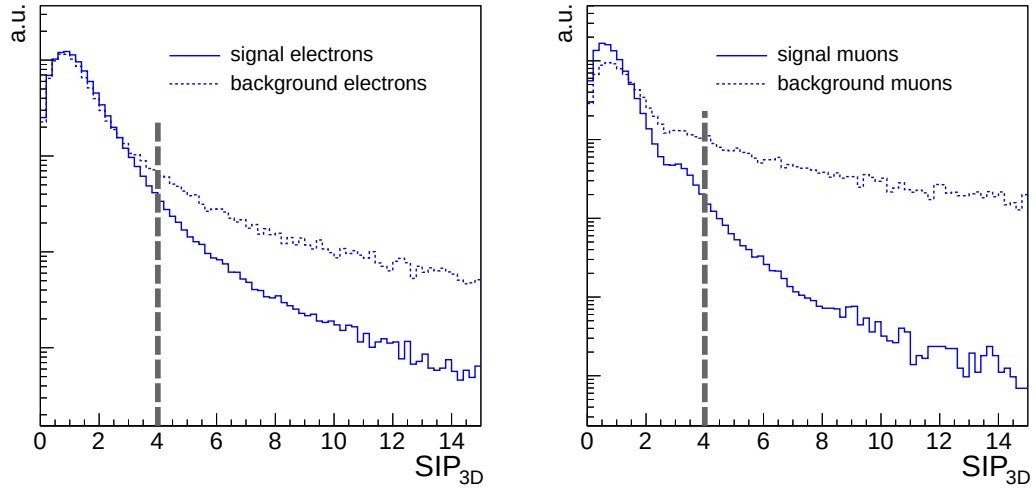

Figure 4.2: Distribution of the 3D impact parameter significance SIP_{3D}, for signal and background electrons (left) and muons (right). The leptons are required to pass all preselection cuts but SIP_{3D} itself. The vertical grey dashed lines denote the operating points.

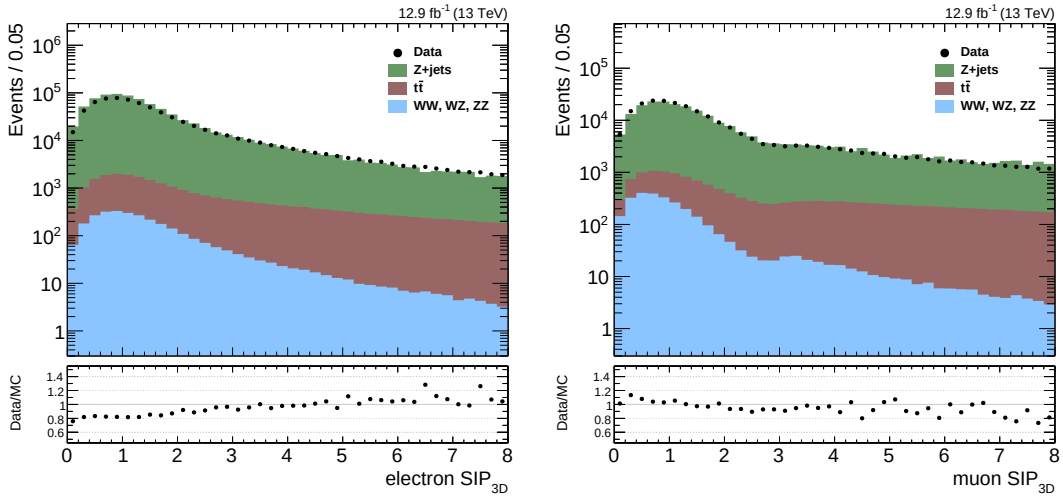

Figure 4.3: Distribution of the 3D impact parameter significance SIP_{3D} in 2016 data and simulation, for electrons (left) or muons (right) selected as the third lepton of a $\mathrm{Z}(\to \ell\ell) + 1\ell$ control region.

Loose leptons

Leptons that satisfy all of the above requirements are called *loose leptons*. This level of selection is not advanced enough for the leptons to be considered for the signal region yet, but it is used in the definition of some control regions, as will be explained in Section 5.5.2.

4.1.2 Lepton identification

The concept of *lepton identification* refers to all the electron-specific or muon-specific quality requirements that aim at selecting genuine electrons and muons among all objects that can be mistakenly reconstructed as electron or muon candidates. As described earlier, the definition of loose leptons already includes a few such requirements, but the main stage of identification is described here.

Muon identification

In this analysis, muon candidates are said to pass identification if they satisfy the requirements of the PF algorithm, i.e. if they are PF muons in the sense described in Section 2.2.2. These identification criteria are originally meant for the PF algorithm to correctly identify muons among all individual particles produced in the collision, in particular within jets, because the unidentified or fake muons can bias the jet or $E_{\mathrm{T}}^{\mathrm{miss}}$ measurements.

PF muon identification proceeds in three steps [65]. First, a strict detector-based isolation criterion (different from the analysis-specific isolation that will be described in Section 4.1.3) is applied to already select muons that have little neighbouring activity and hence do not pose any difficulty for the PF algorithm. Second, non-isolated muon candidates are selected as PF muons if they both include a minimum number of hits in the muon track and satisfy a compatibility criterion of the muon segment and calorimeter deposits, as defined by a template based on simulation. Finally, the remaining muon candidates can also be selected if they pass a looser requirement on the number of hits and a matching requirement of the track to hits in the muon stations.

The reason the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis does not use PF muons from the outset is the need to have a looser muon selection level in order to build control regions, hence the definition of loose muons in Section 4.1.1.

Electron identification

It has been shown in Section 3.2.2 that electron identification at the HLT is achieved via a sequence of cuts on a few well-chosen variables (not including the cuts on $E_{\mathrm{T}}^{\mathrm{SC}}$ and on isolation). The offline-level electron identification of the $\mathrm{H}\to 4\ell$ analysis exploits a much broader range of observables, but they are condensed into a single multivariate classifier, by means of boosted decision trees (BDT). Three main classes of observables are used as inputs: ECAL observables related to the electron supercluster in the ECAL, observables related to electron track reconstruction, and observables that describe the energy-momentum and geometrical matching between the supercluster and the electron track. All are described in Table 4.1.

The BDT classifier is trained on signal electrons from a $\mathrm{gg}\to\mathrm{H}(125)\to\mathrm{ZZ}\to 4\ell$ sample, and background electrons from a Drell-Yan + jets sample, as defined earlier. Signal electrons are geometrically matched to generated electrons from $\mathrm{Z}^{(*)}$ decays, while background electrons are the reconstructed electrons that are either matched to electrons originating from decays within hadronic jets, or not matched to any generated electron. This selection excludes electrons from tau decays, which are not used at all in the training. To improve identification performance of low-energy electrons, the training proceeds in two bins of transverse momentum: $5 < p_{\mathrm{T}}^{\mathrm{e}} < 10\,\mathrm{GeV}$ and $p_{\mathrm{T}}^{\mathrm{e}} > 10\,\mathrm{GeV}$. Since electrons are reconstructed by different subdetectors such as the barrel and endcap

Table 4.1: List of input variables to the electron identification BDT classifier. Variables that were introduced in Run II are marked with (*).

	observable	definition
pure cluster shape observables	$\sigma_{i\eta i\eta}$	energy-weighted standard deviation of single crystal η (in crystal index) within a 5×5 block of crystals centred on the highest-energy crystal of the cluster seed
	$\sigma_{i\phi i\phi}$	similar as $\sigma_{i\eta i\eta}$, but in the ϕ direction
	η width	supercluster width along η
	ϕ width	supercluster width along ϕ
	$1 - E_{5\times1}/E_{5\times5}$	circularity: $E_{5\times5}$ is the energy computed in the 5×5 block of crystals centred on the highest-energy crystal of the cluster seed, and $E_{5\times1}$ is the energy computed in the strip of crystals containing it
	$R_9 = E_{3\times3}/E_{\mathrm{SC}}$	energy sum of the 3×3 block of crystals centred on the highest-energy crystal, divided by the supercluster energy
	H/E	energy collected by the HCAL towers within a cone of $\Delta\mathrm{R} = 0.15$ centred on the supercluster position, divided by the supercluster energy
	$E_{\mathrm{PS}}/E_{\mathrm{raw}}$	energy fraction deposited in the preshower subdetector, divided by the untransformed supercluster energy (endcap bins only)
pure tracking observables	$f_{\mathrm{brem}} = 1 - p_{\mathrm{out}}/p_{\mathrm{in}}$	fractional momentum loss of the GSF track, which measures the amount of bremsstrahlung emission
	N_{KF}	number of hits of the KF track
	N_{GSF}	number of hits of the GSF track (*)
	χ^2_{KF}	goodness of fit of the KF track
	χ^2_{GSF}	goodness of fit of the GSF track
	$N_{\mathrm{miss.hits}}$	number of expected but missing inner hits in the first tracker layers (*)
	$P_{\mathrm{conv.}}$	fit probability for a conversion vertex associated with the electron track (*)
track-cluster matching observables	$E_{\mathrm{SC}}/p_{\mathrm{in}}$	ratio of the supercluster energy to the track momentum at the innermost track position
	$E_{\mathrm{e}}/p_{\mathrm{out}}$	ratio of the energy of the ECAL cluster closest to the electron track extrapolation to ECAL to the track momentum at the outermost track position
	$\frac{1}{E_{\mathrm{SC}}} - \frac{1}{p}$	deviation of the supercluster energy from the electron momentum obtained by combining ECAL and tracker information
	$\Delta\eta_{\mathrm{in}} = \lvert\eta_{\mathrm{SC}} - \eta_{\mathrm{in}}\rvert$	pseudorapidity distance between the energy weighted position of the supercluster and the track position of closest approach to the supercluster, extrapolated from the innermost track position and direction
	$\Delta\phi_{\mathrm{in}} = \lvert\phi_{\mathrm{SC}} - \phi_{\mathrm{in}}\rvert$	similar as $\Delta\eta_{\mathrm{in}}$, but in the ϕ direction
	$\Delta\eta_{\mathrm{seed}} = \lvert\eta_{\mathrm{seed}} - \eta_{\mathrm{out}}\rvert$	pseudorapidity distance between the seed cluster position and the electron track extrapolation to the ECAL

parts of the ECAL and tracker, with more material budget in front of the ECAL as $|\eta^{e}|$ increases, the training bins are further split into three ranges of pseudorapidity: $|\eta^{e}| < 0.8$, $0.8 < |\eta^{e}| < 1.479$, and $|\eta^{e}| > 1.479$, finally defining six training categories. Guiding the training this way was shown to provide better performance than directly including η^{e} as an input variable.

Three of the track-based variables quoted in Table 4.1 had not yet been considered in the Run I version of the classifier: the number of hits of the GSF track, the number of expected but missing hits in the inner tracker layers, and the fit probability for a conversion vertex associated with the electron track. These variables are sensitive to the background of genuine electrons resulting from photon conversions in the material in front of the ECAL. These most often occur in the middle and endcap part of the tracker, resulting in tracks with fewer hits in the first tracker layers. Rejection of such electrons was previously done separately from the multivariate classifier, by requiring reconstructed electrons to have fewer than two missing expected inner hits. By including such variables into the multivariate identification instead, selection performance is improved, as shown in Fig. 4.4, where ROC curves for the new Run II classifier are compared to the combination of a retrained version of the Run I classifier and the missing inner hit cut.

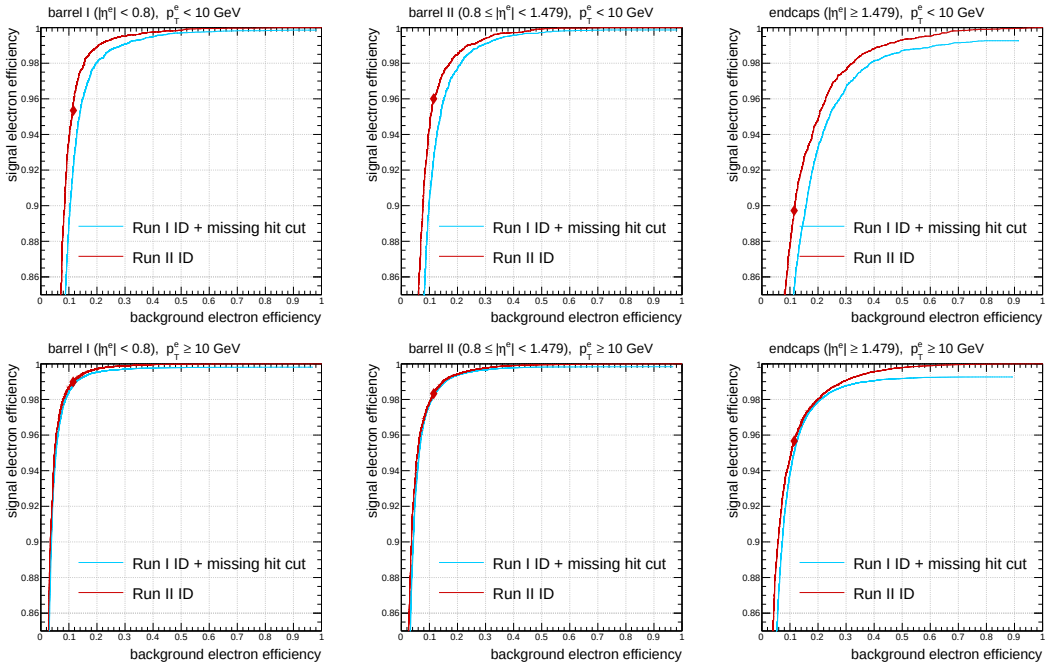

Figure 4.4: ROC curves of the electron identification BDT classifier of the 2015 analysis, for each training category. The chosen operating point is shown by the diamond markers. The red curves correspond to the improved classifier used in this Run II analysis, while the light blue curves show a retraining of the Run I classifier on 13 TeV simulated samples. The gap between the upper right end of the latter curves and the corner of the frame materializes the efficiency of the missing inner hit cut, which used to be applied separately from the Run I BDT classifier.

Table 4.2 presents the operating points that I defined in each of the six training categories, i.e. the minimum cut values on the BDT output score. Also quoted are their efficiencies on signal and background electrons, which are computed from the same

simulated samples and illustrated on the ROC curves of Fig. 4.4. As the separation power of the classifier differs among categories, several sets of six operating points were tested, with the aim of optimizing the compromise between the *global* signal efficiency and background rejection, i.e. the one computed merging all training categories. The most satisfying solution was achieved using the same background electron efficiency in every category, fixing it to a value (11.5%) that preserves the same global signal electron efficiency (97.6%) as was used with the Run I classifier. With this compromise, the global background electron efficiency decreases by 13%. Operating points were again redefined for the 2016 analysis after the BDT was retrained.

Table 4.2: Operating points of the electron identification BDT classifier for the 2015 analysis in the six training categories, and corresponding efficiencies for signal and background electrons, evaluated on simulated samples.

training bin		operating point	signal electron efficiency	background electron efficiency
$5 < p_{\mathrm{T}}^{\mathrm{e}} < 10\,\mathrm{GeV}$	$\lvert\eta^{\mathrm{e}}\rvert < 0.8$	-0.265	0.953	0.115
	$0.8 < \lvert\eta^{\mathrm{e}}\rvert < 1.479$	-0.556	0.960	0.115
	$\lvert\eta^{\mathrm{e}}\rvert > 1.479$	-0.551	0.897	0.115
$p_{\mathrm{T}}^{\mathrm{e}} > 10\,\mathrm{GeV}$	$\lvert\eta^{\mathrm{e}}\rvert < 0.8$	-0.072	0.990	0.115
	$0.8 < \lvert\eta^{\mathrm{e}}\rvert < 1.479$	-0.286	0.983	0.115
	$\lvert\eta^{\mathrm{e}}\rvert > 1.479$	-0.267	0.957	0.115

While this optimization work was based on simulation, the BDT classifier has also been validated in 13 TeV data. Figure 4.5 compares distributions of the BDT output score in 2016 data and simulation in the Z+1ℓ control region. A reasonable level of agreement is observed, except for large values of the variable, but this part of the spectrum is found to be better modelled in signal-dominated regions.

4.1.3 Lepton isolation

As already mentioned in HLT studies, lepton isolation is a powerful handle to suppress background from non-prompt or misidentified leptons. Since it is a very analysis-dependent quality cut, offline isolation is usually kept separate from the more generic lepton identification requirements: in particular, it is not an input to the electron identification BDT. The selection of the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ signal very much relies on offline isolation, and uses its own definition of it. For both electrons and muons, the chosen strategy is Particle-flow isolation, but the three components described in Section 3.2.2 (photon, neutral hadron, and charged hadron isolation) are here gathered into one single variable, which is corrected globally for pileup, and one single cut is made. The variable can be written as:

$$\mathcal{I}_{\mathrm{PF}}^{\ell} \equiv \frac{1}{p_{\mathrm{T}}^{\ell}}\left(\max\left[0, \sum_{\substack{\text{photons}\\ \text{(except FSR)}}} p_{\mathrm{T}} + \sum_{\substack{\text{neutral}\\ \text{hadrons}}} p_{\mathrm{T}} - p_{\mathrm{T}}^{\mathrm{PU}}\right] + \sum_{\substack{\text{charged}\\ \text{hadrons}}} p_{\mathrm{T}}\right), \tag{4.4}$$

and its distributions are shown in Fig. 4.6 for signal and background leptons.

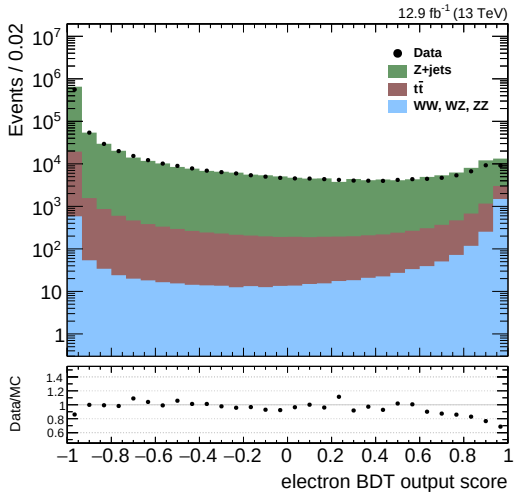

Figure 4.5: Distribution of the electron identification BDT classifier in 2016 data and simulation, for electrons selected as the third lepton of a $\mathrm{Z}(\to\ell\ell)+1\ell$ control region. The six electron categories used for the training are merged in this plot.

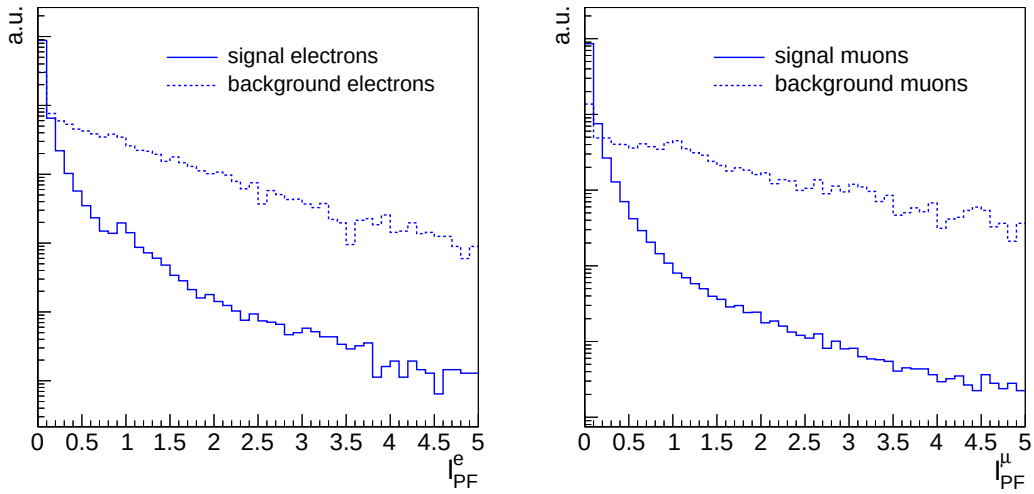

Figure 4.6: Distribution of the combined PF isolation variable $\mathcal{I}^{\ell}_{\mathrm{PF}}$, for signal and background electrons (left) and muons (right). The electrons and muons are required to be selected as loose leptons and to satisfy the identification requirements.

As discussed in Section 3.2.2, isolation is particularly sensitive to undesirable energy deposits from pileup interactions. While the charged hadron is naturally protected by a vertex constraint, the component of neutral particles (photons and neutral hadrons) is corrected for its pileup contribution by subtracting the $p_{\mathrm{T}}^{\mathrm{PU}}$ quantity (constraining the difference to stay non-negative). Electrons and muons here use two different approaches:

- Electrons use the FastJet technique, like at HLT. The subtracted quantity is:

$$p_{\mathrm{T}}^{\mathrm{PU}}(\mathrm{e}) \equiv \rho \times A_{\mathrm{eff}} \ , \tag{4.5}$$

  where the effective area A_{eff} again requires a prior computation, which is done

here in five bins of $|\eta^{e}|$.

- Muons instead use the so-called $\Delta\beta$ correction, which is defined as:

$$p_{T}^{PU}(\mu) \equiv 0.5 \times \sum_{i} p_{T}^{PU,i} \ , \tag{4.6}$$

where i runs over the momenta of the charged hadron PF candidates that do not originate from the primary vertex. This formula relies on the assumption that the pileup contribution of neutral particles to the isolation cone is about half of that of charged particles associated to pileup.

Besides this pileup correction, the photon isolation component is also slightly modified with respect to its default definition. The sum of photon momenta indeed disregards any photon that is selected by the algorithm that recovers final-state radiation, which will be described in Section 4.4.

Although the aforementioned isolation strategy is similar to that used in Run I, two important points have been reoptimized, namely the cut threshold and the size of the cone in which PF candidates are considered. The latter is defined by the maximal angular distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ between the PF candidate and the lepton direction at the primary vertex, which used to be $\Delta R = 0.4$. The optimization was performed taking signal electrons from a mixture of Higgs boson production mechanisms weighted by their cross sections instead of just ggH, since other, more crowded topologies can justify a narrower cone. Among a few cone size hypotheses, an angular radius of $\Delta R = 0.3$ was chosen because it was found to provide some improvement for electrons in the efficiency region that is relevant to the analysis, as illustrated in the ROC curve of Fig. 4.7 (left), done using only ggH. The operating point is defined as 0.35, and the A_{eff} is computed specifically for $\Delta R = 0.3$. As regards muons (Fig. 4.7, right), both cone sizes give similar performance in this region, and the same ΔR and operating point are chosen. The behaviour of the new $\mathcal{I}^{\ell}_{\text{PF}}$ variable in data relative to simulation is illustrated in Fig. 4.8 using the Z+1ℓ control region.

4.2 Lepton momentum calibration

As explained in Section 2.2.2, the determination of the momentum of electrons relies on a combination from ECAL and tracker, while that of muons involves the tracker and the muon chambers. Although the ECAL calibrations and the tracker and muon system alignment rely on the best knowledge of the detector conditions, some small discrepancies remain between data and simulation, e.g. because of residual misalignment and imperfect corrections of the ECAL crystal transparency losses. The scale and resolution of lepton momenta thus have to be calibrated, which is done in bins of p_{T}^{ℓ} and η^{ℓ}, exploiting some dilepton resonances. Besides benefiting the overall accuracy of the analysis, this calibration is of utmost importance for the Higgs boson mass measurement.

First, the scale of electrons in data is calibrated using a $Z \to e^{+}e^{-}$ control sample, correcting the momenta so as to place the peak of the reconstructed dielectron mass spectrum at the nominal Z boson mass. To account for the time-dependent loss of the transparency of crystals, the correction is derived as a function of time. Second, a smearing of the electron energies is applied in simulation so as to make the $Z \to e^{+}e^{-}$ mass resolution in simulation match that observed in data. To this end, a pseudorandom

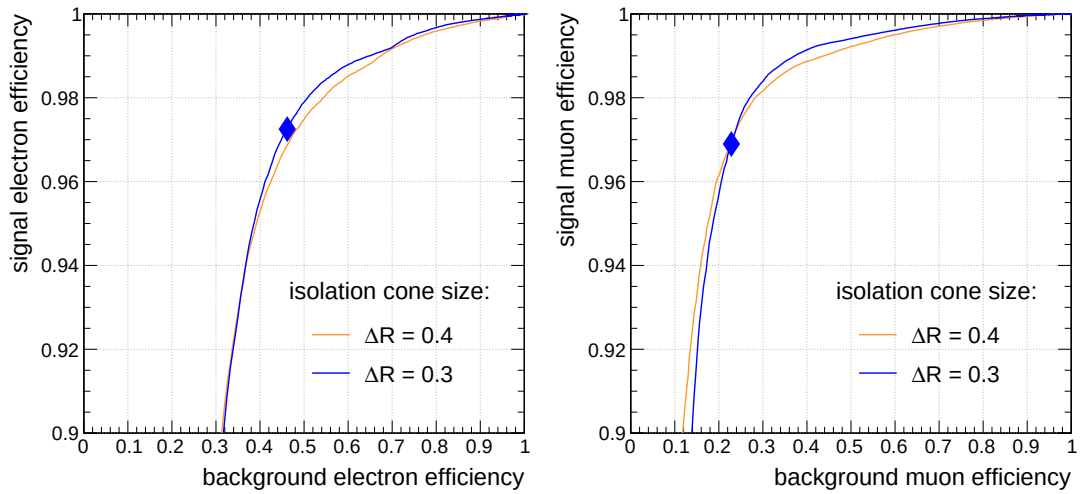

Figure 4.7: ROC curves of combined PF isolation for electrons (left) and muons (right). The electrons and muons are required to be selected as loose leptons and to satisfy the identification requirements. Two versions of the variable are compared, for two sizes of the isolation cone: $\Delta R = 0.4$ as used in Run I, and $\Delta R = 0.3$ as used in this analysis as of 2015. The new operating point is shown by the diamond markers.

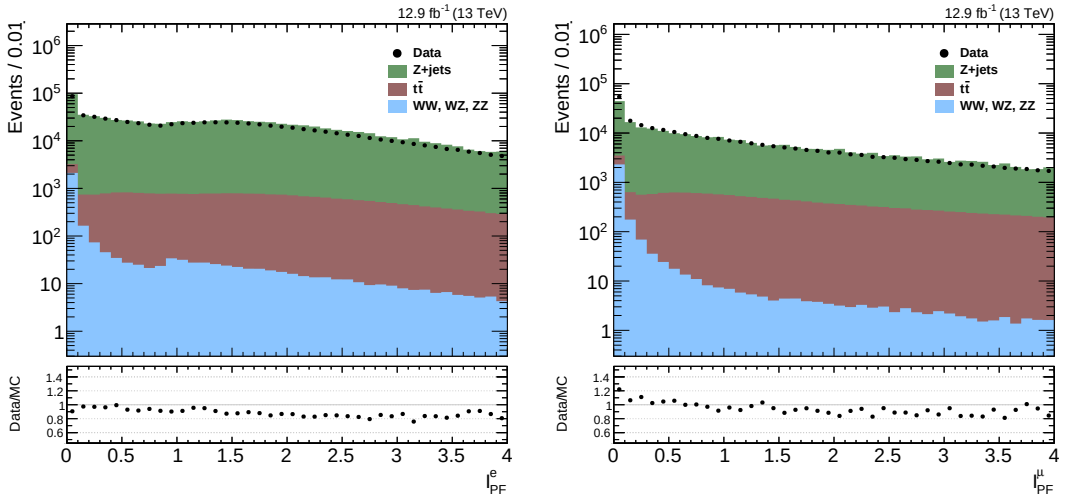

Figure 4.8: Distribution of the combined PF isolation variable $\mathcal{I}^{\ell}_{\mathrm{PF}}$ in 2016 data and simulation, for electrons (left) or muons (right) selected as the third lepton of a $\mathrm{Z}(\to \ell\ell) + 1\ell$ control region.

Gaussian multiplicative factor is applied, with Gaussian parameters varying in bins of $p_{\mathrm{T}}^{\mathrm{e}}$ and $|\eta^{\mathrm{e}}|$. The impact of these corrections on the Z boson lineshape for different categories of electrons in data and simulation can be seen comparing the top row and bottom row of Fig. 4.9.

Muons, on the other hand, are calibrated using J/ψ and $\mathrm{Z} \to \mu^{+}\mu^{-}$ decays with a method recently developed in CMS in the context of a W-like measurement of the Z-boson mass [66]. Using a Kalman Filter, the curvature of muon track in the silicon

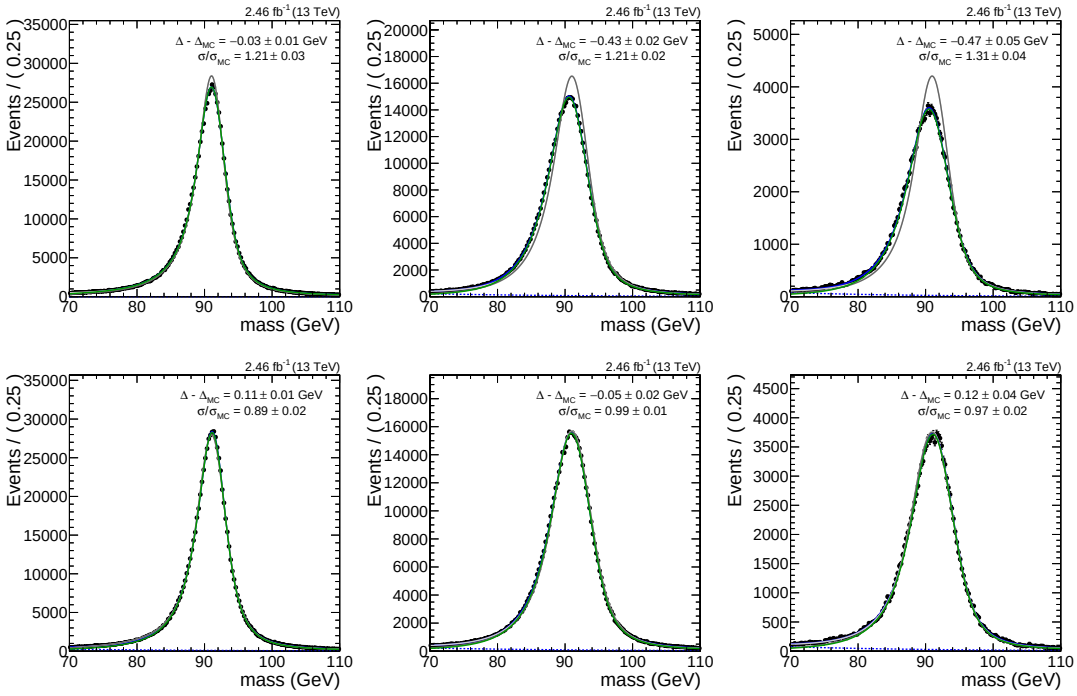

Figure 4.9: Fits to the dielectron invariant mass distribution from $Z \to e^+e^-$ events in data (green curves) and simulation (grey curves), for events with both electrons in the barrel (left), one electron in the barrel and one in the endcaps (middle), and both electrons in the endcaps (right). The dashed blue curve represents the background model. Plots in the top row are uncorrected, and plots in the bottom row have the momentum scale corrections and smearing applied.

tracker is corrected for three effects: small variations of the magnetic field, residual misalignment, and the imperfect modelling of the material. Similarly to electrons, the muon momentum resolution in simulation is smeared to match that of data. Figure 4.10 shows the effect of these corrections on the Z boson lineshape in data and simulation.

Per-lepton momentum uncertainties, i.e. the uncertainties on the measurement of individual electron and muon momenta, constitute another observable that is also updated during the lepton calibration procedure. As a closure test of both the corrections of uncertainties and the smearing of momenta in simulation, Fig. 4.11 displays the measured dilepton mass resolutions in data and simulation, in bins of the resolution predicted from per-lepton momentum uncertainties. This is explained in more detail in Section 5.3.3, where the per-lepton momentum uncertainties are again calibrated for the computation of the analysis-specific per-event mass uncertainties.

4.3 Lepton efficiency measurements

The efficiency of reconstructing and selecting signal leptons is a cornerstone of the analysis in several respects. Not only does it need to be optimized with respect to background leptons, but it is also very important to accurately measure it in data and simulation, and to correct for the possible discrepancies. Indeed, the analysis ultimately consists in comparing predictions from simulation to data, and any sizeable mismodelling can propagate to the predicted event yields and bias the physics results.

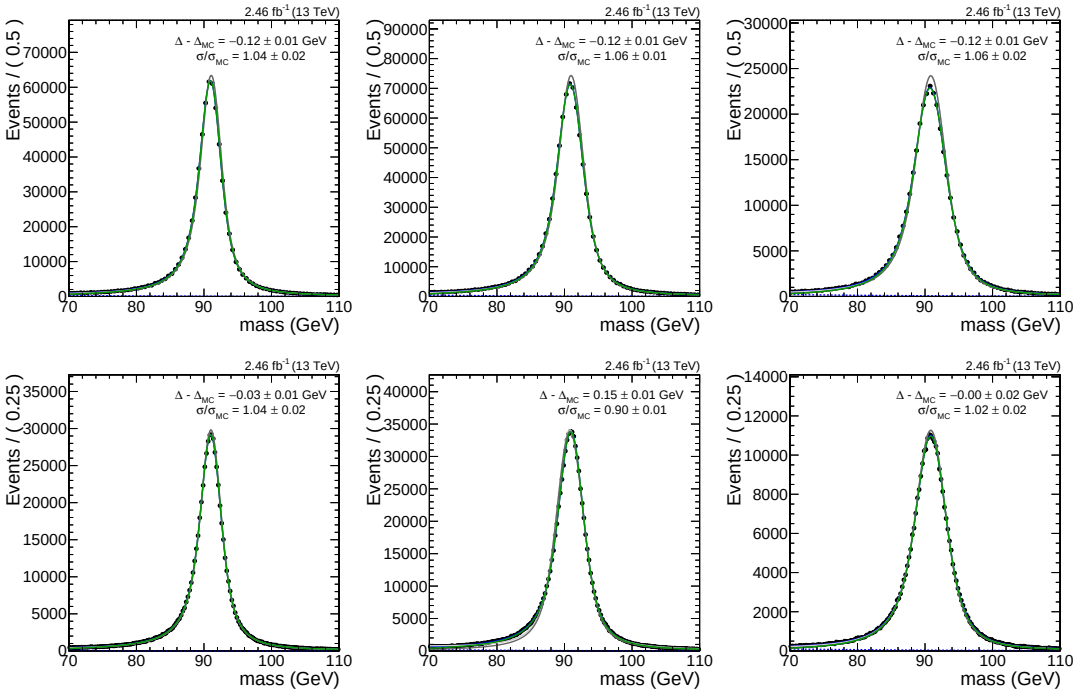

Figure 4.10: Fits to the dimuon invariant mass distribution from $Z \to \mu^+\mu^-$ events in data (green curves) and simulation (grey curves), for events with both muons in the barrel (left), one muon in the barrel and one in the endcaps (middle), and both muons in the endcaps (right). The dashed blue curve represents the background model. Plots in the top row are uncorrected, and plots in the bottom row have the momentum scale corrections and smearing applied.

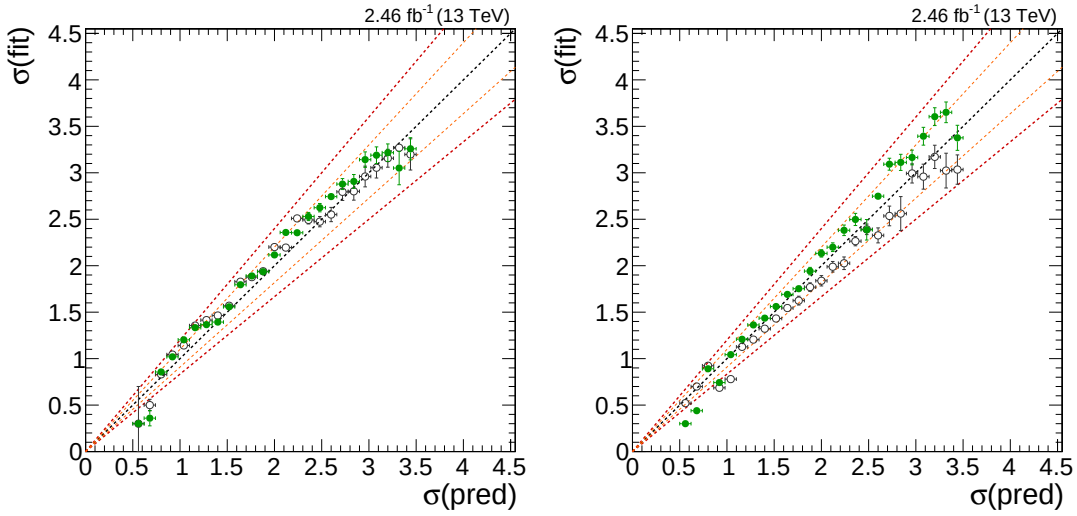

Figure 4.11: Measured vs. predicted mass resolution for $Z \to e^+e^-$ (left) and $Z \to \mu^+\mu^-$ (right) events in data (filled markers) and simulation (empty markers), in bins of the predicted resolution, after applying the lepton momentum calibrations. The orange and red dashed lines represent the ±10% and ±20% envelopes, respectively.

The selection of events for the signal region relies on no fewer than four leptons, which means that the impact of any uncontrolled discrepancy in lepton-level efficiencies is taken to the power four in event-level efficiencies.

To bring this under control, the efficiency ϵ of the general reconstruction (Section 2.2.2) and analysis-specific selection (Section 4.1) of signal electrons and muons is measured in several bins of p_{T}^{ℓ} and η^{ℓ}, with the same definitions and data taking conditions as in the analysis. This measurement is done in data (ϵ_{data}) and in simulation (ϵ_{MC}) with the same method, leading to the definition of a per-lepton efficiency scale factor as

$$\mathrm{SF}_{\ell}(p_{\mathrm{T}}^{\ell}, \eta^{\ell}) \equiv \frac{\epsilon_{\mathrm{data}}(p_{\mathrm{T}}^{\ell}, \eta^{\ell})}{\epsilon_{\mathrm{MC}}(p_{\mathrm{T}}^{\ell}, \eta^{\ell})} \,. \tag{4.7}$$

A per-event data-to-simulation scale factor is then defined as the product of the scale factors of the four selected leptons:

$$\mathrm{SF}_{4\ell} \equiv \prod_{\ell=1}^{4} \mathrm{SF}_{\ell}(p_{\mathrm{T}}^{\ell}, \eta^{\ell}) \,, \tag{4.8}$$

which is then used to reweight the simulated samples on a per-event basis. Furthermore, systematic uncertainties on this number also need to be thoroughly assessed.

4.3.1 Tag-and-Probe technique

Efficiency measurements in data rely on the Tag-and-Probe (T&P) technique [67], which exploits the $\mathrm{Z} \to \ell^{+}\ell^{-}$ candle to build an unbiased source of leptons. First, a lepton is selected with tight identification requirements and is considered as the *tag*. Second, another lepton of same flavour, opposite charge, and as loose selection as possible is sought for in a certain mass window around m_{Z^0} (e.g. $60 < m_{\mathrm{e^+e^-}} < 120\,\mathrm{GeV}$ for electrons) and considered as the *probe*. The efficiency of a selection cut is then estimated as $\epsilon = \mathrm{N_P}/(\mathrm{N_P} + \mathrm{N_F})$, where $\mathrm{N_P}$ and $\mathrm{N_F}$ are the numbers of *passing* and *failing* probes, respectively. Studies have shown that efficiencies estimated this way are almost insensitive to the selection criteria applied to the tag.

Despite being linked to tags by the Z resonance, probes are generally not a pure source of leptons and are contaminated by the multi-jet background, especially at low p_{T}. As a result, a cut-and-count approach is not enough to reliably estimate the efficiencies. Instead, this study relies on a simultaneous fit of the dilepton invariant mass distribution of passing and failing probes to the sum of the Z peak and underlying background (Drell-Yan continuum and fake leptons), which provides the number of passing and failing actual signal leptons. An example of such a fit is shown in Fig. 4.12 for a certain category of electrons, which will be discussed in Section 4.3.2.

Although efficiencies in simulation could be simply estimated by a direct method, they here rely on the same Tag-and-Probe approach as data, since the priority is to avoid any bias.

4.3.2 Electron efficiencies

The efficiency of selecting electrons is factorized as

$$\epsilon = \epsilon_{\mathrm{reco.}} \times \epsilon_{\mathrm{sel.\,|\,reco.}} \,. \tag{4.9}$$

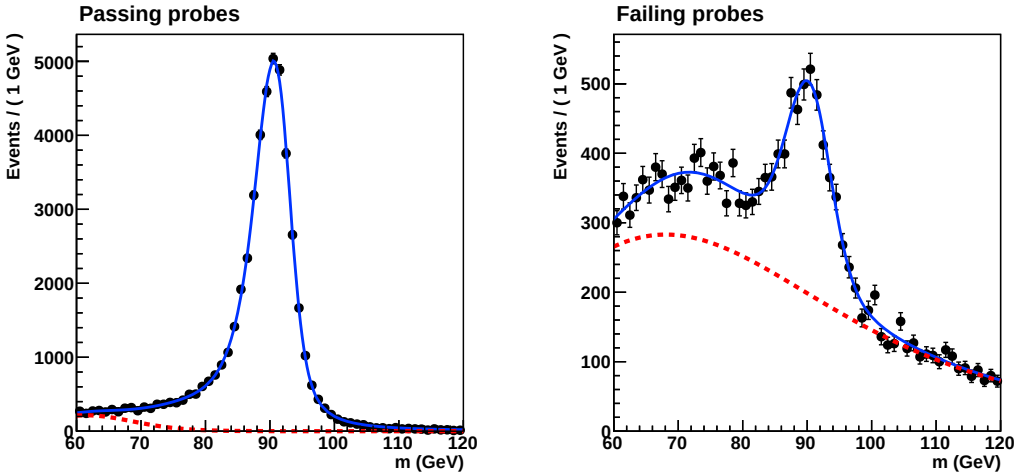

Figure 4.12: Illustration of the Tag-and-Probe method for the measurement of the efficiency of the analysis-specific electron selection (identification, impact parameter and isolation requirements) relative to electron reconstruction. Fits to the 2015 experimental data are shown for probe non-crack electrons that pass (left) or fail (right) the selection, in a kinematic window defined as $20 < p_{\mathrm{T}}^{\mathrm{e}} < 30\,\mathrm{GeV}$ and $0.8 < \left|\eta^{\mathrm{SC}}\right| < 1.479$. Black points are 2015 data, the red dashed line represents the background model, and the blue line represents the signal plus background model.

The reconstruction efficiency $\epsilon_{\mathrm{reco.}}$ is computed centrally in CMS and will not be described in detail here. Its associated scale factors (Eq. 4.7) range from 0.96 in the central barrel to 1.01 at large η^{SC}. The second factor $\epsilon_{\mathrm{sel.\,|\,reco.}}$ is specific to the $\mathrm{H}\to 4\ell$ analysis and represents the efficiency of the identification, impact parameter and isolation cuts described in Section 4.1, with respect to reconstructed electrons.

Data samples for the T&P method are selected with an isolated single-electron HLT path, while simulated events are taken from a LO Drell-Yan+jets sample. Tag electrons are required to have $p_{\mathrm{T}}^{\mathrm{e}} > 30\,\mathrm{GeV}$, to be in the geometrical detector acceptance, and to pass a tight identification operating point. They are also geometrically matched to the triggering electron of the HLT path, and to the generator-level electron in simulation. Despite this tight selection, background electrons from QCD multijet production are still present, particularly in low-$p_{\mathrm{T}}^{\mathrm{e}}$ bins, and in the set of failing probes because the measured efficiencies are large, which makes it hard to fit for the Z peak there. Consequently, a loose cut on tracker isolation is introduced to purify the set of probes.

The Z peak is modelled with a Breit-Wigner function convolved with a Crystal Ball function, while the background component in data is described by an error function times an exponential. Figure 4.12 shows the example of the fits in the $20 < p_{\mathrm{T}}^{\mathrm{e}} < 30\,\mathrm{GeV}$ region, where an accumulation of events can be seen in the left-hand side tail of $m_{\mathrm{e^+e^-}}$ distribution for failing probes. This effect was understood to arise from three contributions of similar sizes:

- electrons located near cracks between the EB and EE and between modules of the ECAL barrel (see Section 2.2.1), which were then handled separately (defining two collections: *crack electrons* and *non-crack electrons*),
- electrons experiencing final-state radiation, which were thus treated in the same

way as explained in Section 4.4,

- electrons with a large amount of bremsstrahlung radiation that was not accounted for during reconstruction.

A Gaussian function is added to the signal model in medium-$p_{\mathrm{T}}^{\mathrm{e}}$ bins in order to model this residual tail. Its ratio to the whole signal component is first determined from the fit in simulation, and then fixed for the fit in data, while other parameters are left floating.

After extracting $\epsilon_{\mathrm{sel.\,|\,reco.}}$ from the simultaneous fit of passing and failing probes, the data-to-simulation scale factors are computed from Eq. 4.7. They are presented in Fig. 4.13 as a function of $p_{\mathrm{T}}^{\mathrm{e}}$ and η^{SC}, showing that selection efficiencies are slightly overestimated in simulation, with scale factors mostly ranging from 0.97 to 0.99, except for some low-$p_{\mathrm{T}}^{\mathrm{e}}$ bins in the endcaps where they reach 1.02 for non-crack electrons.

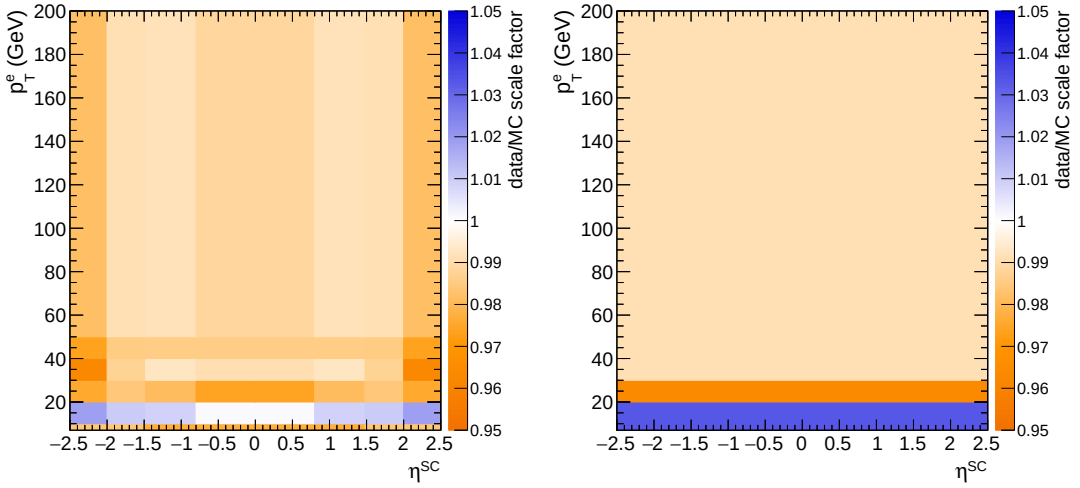

Figure 4.13: Data-to-simulation scale factors for electron selection efficiencies (identification, impact parameter and isolation, relative to reconstruction) in the 2015 data sample, computed using the Tag-and-Probe method as a function of the electron transverse momentum $p_{\mathrm{T}}^{\mathrm{e}}$ and supercluster pseudorapidity η^{SC}, for non-crack (left) and crack (right) probes.

The uncertainties on the scale factors include following contributions:

- statistical uncertainties obtained from the fits,
- variations obtained from alternative modellings of signal and background shapes,
- a $\pm 5\%$ variation of the minimum bias cross section used for pileup reweighting in simulation,
- the comparison between two alternative Drell-Yan + jets samples,
- the effect of dropping the aforementioned tracker isolation cut applied to probes.

The total systematic uncertainty is of the order of 10% in low-$p_{\mathrm{T}}^{\mathrm{e}}$ bins and around 1–2% at high $p_{\mathrm{T}}^{\mathrm{e}}$. It is propagated to the event-level uncertainties described in Section 5.7.3.

4.3.3 Muon efficiencies

Measurements of muon efficiencies also rely on the T&P method, and generally use a $Z \to \mu^+\mu^-$ sample collected with an isolated single-muon HLT path. The overall efficiency is broken down into four factors, which are measured separately:

- The efficiency to reconstruct a muon track in the inner tracker, whereby the probes are tracks reconstructed in the muon system only. For this factor and the next one, the efficiency of low-p_T muons is measured with a $J/\psi \to \mu^+\mu^-$ sample instead of the Z sample, because it has a better purity in this kinematic regime. Here, the scale factor is 0.99 below 10 GeV, and 1 above this.

- The efficiency for a muon to be reconstructed as a global or tracker muon, and to pass the identification requirements. Here, the probes are tracks reconstructed in the inner tracker. The corresponding scale factors range from 0.985 to 1 and are independent of p_T^μ and of pileup; they have a complex η^μ-dependency due to the geometry of the muon system.

- The efficiency for reconstructed, identified muons to also pass all impact parameter requirements of the analysis (i.e. cuts on d_{xy}, d_z, and SIP_{3D}). The corresponding scale factors are p_T^μ-independent as well, and the largest deviations from unity (≈ 0.99) are found in the endcaps.

- The efficiency for such muons to also pass isolation. Here, final-state radiation results in a component of off-peak signal events in the distribution of failing probes, albeit to a lesser extent than for electrons. To account for this effect, the signal model uses a template from simulated events, convolved with a Gaussian function to cover possible resolution and scale differences between data and simulation. Data and simulation are found to be in very good agreement.

Multiplying these four efficiency factors provides the overall data-to-simulation scale factors for muons, which are illustrated in Fig. 4.14 as a function of p_T^μ and η^μ. Here again, the final correction is negative: about -1% or less in most bins, and about -2% for muons below 10 GeV or with $|\eta^\mu| > 2$. The systematic uncertainties on these scale factors are estimated by varying the event selection and the models for the signal and background shapes used in the fits, and from recomputing the reconstruction and isolation efficiencies with a cut-and-count approach instead of the T&P fits. This all leads to uncertainties of 1–2% for muons below 10 GeV or with $|\eta^\mu| > 2$, and about 0.5% or less in the rest of the phase space.

4.4 FSR photon recovery

Leptons from Z bosons decays can radiate a high-energy photon, a phenomenon called *final-state radiation* (FSR), which degrades the accuracy of the information extracted from four-lepton candidates. It is therefore important to collect and identify these photons, and to associate them to their parent lepton, to fully reconstruct the Higgs boson decay system. The FSR recovery algorithm is designed to discriminate such photons from background photons that can originate from initial-state radiation or pileup interactions. This is achieved exploiting the particular kinematics of FSR photons, which are mostly emitted with a direction nearly collinear with their parent lepton, and tend to be isolated from other particles.

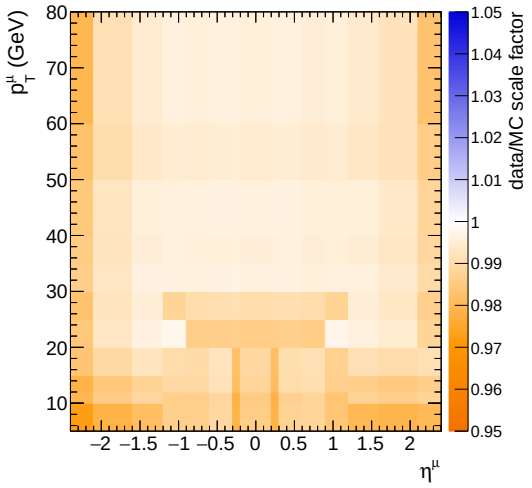

Figure 4.14: Data-to-simulation scale factors for overall muon efficiencies (tracking, reconstruction, identification, impact parameter and isolation) in the 2015 data sample, computed using the Tag-and-Probe method as a function of the muon transverse momentum p_{T}^{μ} and pseudorapidity η^{μ}.

FSR photon candidates are sought for among photons reconstructed by the PF algorithm within $|\eta^{\gamma}| < 2.4$, passing $p_{\mathrm{T}}^{\gamma} > 2\,\mathrm{GeV}$. They are required to pass an isolation criterion $\mathcal{I}^{\gamma} < 1.8$, whereby the relative PF isolation $\mathcal{I}^{\gamma}$ for photon candidates is defined as

$$\mathcal{I}^{\gamma} \equiv \frac{1}{p_{\mathrm{T}}^{\gamma}} \left(\sum_{\substack{\text{photons}}} p_{\mathrm{T}} + \sum_{\substack{\text{neutral} \\ \text{hadrons}}} p_{\mathrm{T}} + \sum_{\substack{\text{charged} \\ \text{hadrons}}} p_{\mathrm{T}} \right) , \tag{4.10}$$

using a cone of radius $\Delta\mathrm{R} = 0.3$, and including the contribution from pileup vertices in the sums. FSR candidates that are involved in the supercluster of a loose electron are then discarded. The remaining ones are associated to the closest loose lepton in the event. To reject background photons, a twofold kinematic cut is applied:

$$\Delta\mathrm{R}(\gamma, \ell)/(E_{\mathrm{T}}^{\gamma})^2 < 0.012 \quad \text{and} \quad \Delta\mathrm{R}(\gamma, \ell) < 0.5 \, . \tag{4.11}$$

In the rare cases where a lepton has more than one associated FSR candidate, only the lowest-$\Delta\mathrm{R}(\gamma, \ell)/(E_{\mathrm{T}}^{\gamma})^2$ one is retained.

Since the FSR photons are often located in the isolation cone of their lepton, they tend to make it fail the isolation requirement. Hence, all selected FSR photons are explicitly subtracted from the isolation sums, not only of their associated leptons (see Eq. 4.4), but also of all loose leptons in the event.

This FSR recovery algorithm was optimized using a ggH sample, and its efficiency of selecting genuine reconstructed FSR photons in 2e2μ events was 57% in the 2015 analysis, for a purity of 81%. The procedure was found to affect 1.5%, 3.8%, and 6.4% of 4e, 2e2μ, and 4μ events, respectively. The low number for electrons is partly caused by an inefficiency in the way that candidate photons already collected in an electron supercluster are discarded from the procedure. This point was reoptimized in 2016, allowing the algorithm to now affect 4.3% of 4e events.

4.5 Jet selection and flavour tagging

The present thesis work aims at probing subdominant Higgs boson production modes, with some of them including jets, e.g. VBF, VH where the associated vector boson decays hadronically, and $\mathrm{t\bar{t}H}$. To this end, the number and kinematics of jets is exploited to build discriminants (see Section 5.3) and to categorize the final selected events (see Section 5.4). This section describes the selection of these jets, and two methods that aim at identifying the flavour of the initiating parton, which I have introduced in the analysis to aid production mode extraction.

4.5.1 Jet selection

Jets reconstruction is described in Section 2.2.2. The jet energy corrections have been computed centrally for each data taking period, and are applied to jets in both simulation and data. Besides the correction of the jet energy scale, a smearing of the jet energy resolution is applied in simulation to match that observed in data.

Kinematic requirements are then applied to jet candidates: they must have a transverse momentum $p_{\mathrm{T}}^{\mathrm{jet}}$ larger than 30 GeV and be found in the pseudorapidity range $|\eta^{\mathrm{jet}}| < 4.7$. In order to reject the background from fake, badly reconstructed, or noise jets, some loose identification criteria centrally recommended for CMS analyses are also applied. These have been reoptimized for 13 TeV collisions, and include requirements on the number of constituents in the jet, on the multiplicity of charged and neutral particles, and on the fraction of the jet energy carried by the different types of constituents. Dedicated identification criteria for the rejection of pileup jets were also tested, but they have not been applied yet, for lack of optimization of the new Run II algorithm.

Electrons and photons can also be reconstructed as jets, which requires to explicitly remove the corresponding jets from the jet collection. This *jet cleaning* procedure is implemented as follows: to be considered for the analysis, the selected jets are required to be geometrically separated from any of the leptons that pass the complete selection described in Section 4.1 and from their associated FSR photons by an angular distance of $\Delta\mathrm{R}(\ell/\gamma, \mathrm{jet}) > 0.4$.

4.5.2 b tagging

Many measurements and searches at CMS rely on the accurate identification of jets originating from the hadronization of bottom quarks (b jets). This has motivated the development and optimization of b jet identification techniques, referred to as *b tagging*. These mostly rely on the fact that the B hadrons present in b jets have relatively large masses (5–6 GeV) and a long lifetime, and thus decay a few millimetres away from the primary interaction vertex. Thanks to the very good position resolution of the pixel tracker, the impact parameters of the resulting displaced tracks are measured, and a secondary vertex may be reconstructed. Tagging algorithms usually condense the extracted information into one classifier variable, the b tagger, which is computed for every jet, and takes larger values for jets that are most likely to be b jets.

In the present analysis, processes such as $\mathrm{t\bar{t}H}$ and ZH, $\mathrm{Z} \to \mathrm{b\bar{b}}$ involve two b jets. As will be shown in Section 5.4, this fact makes b tagging an important ingredient of Higgs boson production mode extraction, because the number of b-tagged jets found in each event is used as a criterion for event categorization.

The b tagger which is used here is called *Combined Secondary Vertex* (CSVv2) and is described in detail in Ref. [68]. This classifier is a Run II-reoptimized version of the CSV tagger that was used in CMS Run I results. It combines the information on displaced tracks with information related to secondary vertices associated to the jet, using a multivariate regression. The output variables takes values between 0 (light-flavour-like jet) and 1 (b-like jet) and is illustrated in Fig. 4.15 in data and simulation in the Z+1ℓ control region, exhibiting a moderately good agreement, and illustrating the fact that b-like jets dominate in the $\mathrm{t\bar{t}}$ sample. The $\mathrm{H} \to 4\ell$ analysis uses a medium-efficiency operating point (0.800), which was tuned by the CMS b-tagging group so as to have a misidentification probability of around 1% for light-flavour jets with $p_{\mathrm{T}}^{\mathrm{jet}} > 30\,\mathrm{GeV}$.

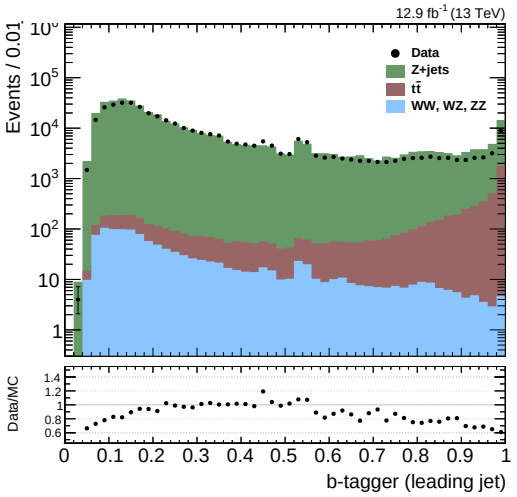

Figure 4.15: Distribution of the CSVv2 tagger in 2016 data and simulation, for the highest-p_{T} jet passing the selection of Section 4.5.1 in a Z+1ℓ control region.

The imperfect modelling of the distributions of the b tagger causes some discrepancies of the b-jet identification efficiency and misidentification probability between data and simulation. To correct for these differences, scale factors have to be applied to the simulation; they are defined as the ratio of the efficiencies of tagging genuine b jets in data and simulation:

$$\mathrm{SF} \equiv \frac{\epsilon_{\mathrm{data}}}{\epsilon_{\mathrm{MC}}}. \tag{4.12}$$

They are provided centrally for every tagger, as a function of the jet p_{T}, η and flavour (i.e. b jet, c jet or light-flavour jet). On that basis, a method to correct the number of b-tagged jets in a given event consists in updating the b-tagging status on a jet-by-jet basis, depending on the scale factor value and using a random number generator. When $\mathrm{SF} < 1$, one just has to consider a fraction

$$f = 1 - \mathrm{SF} \tag{4.13}$$

of the originally b-tagged jets as untagged. If, on the other hand, $\mathrm{SF} > 1$, one needs to

tag a fraction of the previously untagged jets, namely:

$$f = \frac{1 - \mathrm{SF}}{1 - 1/\epsilon_{\mathrm{MC}}}. \tag{4.14}$$

This formula involves the tagging efficiency in simulation, ϵ_{MC}, which has to be computed beforehand in the phase space region of the analysis.

4.5.3 Quark-gluon tagging

Pairs of jets produced in VBF and VH processes originate from quark jets, while jets produced in association with the ggH process are gluon-induced. Differences in jet substructure between quark and gluon induced jets can be exploited to build discriminating variables, thus helping to extract these processes of interest. Such methods have been developed in CMS since Run I and are referred to as *quark-gluon tagging*.

The quark-gluon tagger used for 13 TeV analyses combines three input variables:

- *mult*: the jet *particle multiplicity*, which is larger in gluon jets than in light-quark jets,
- $axis2$: the minor axis of the jet ellipse on the $\eta - \phi$ plane, which accounts for the fact that gluon jets are less collimated than quark jets,
- $p_{\mathrm{T}}D$: the jet *fragmentation function*, defined as

$$p_{\mathrm{T}}D \equiv \frac{\sqrt{\sum_i p_{\mathrm{T},i}^2}}{\sum_i p_{\mathrm{T},i}} \ . \tag{4.15}$$

  It describes the energy spectrum of the jets constituents, and is softer for gluon jets than for quark jets.

Probability density functions (pdfs) are constructed from a training sample for these three variables, separately for light quark jets and gluon jets. Global pdfs for light quarks ($\mathcal{Q}$) and gluons ($\mathcal{G}$) are then defined by multiplying all three quark or gluon pdfs ($f_{\mathcal{Q}}^i$ and $f_{\mathcal{G}}^i$):

$$\mathcal{Q}(\mathrm{J}) \equiv \prod_i f_{\mathcal{Q}}^i(x_i) \ , \tag{4.16}$$

$$\mathcal{G}(\mathrm{J}) \equiv \prod_i f_{\mathcal{G}}^i(x_i) \ , \tag{4.17}$$

where $\mathrm{J} = (x_1, x_2, x_3) = (mult, axis2, p_{\mathrm{T}}D)$ denotes the set of input variables of a given jet. These quantities are used to construct a likelihood discriminant as

$$\mathcal{D}^{\mathrm{q/g}}(\mathrm{J}) \equiv \left[1 + \frac{\mathcal{G}(\mathrm{J})}{\mathcal{Q}(\mathrm{J})}\right]^{-1} \ . \tag{4.18}$$

This variable is illustrated in Fig. 4.16 in data and simulation in the Z+1ℓ control region. It takes values between 0 (gluon-like jets) and 1 (quark-like jets) and may be interpreted as the probability for a jet to originate from a quark. Its usage for production mode discrimination will be described in Section 5.3.2.

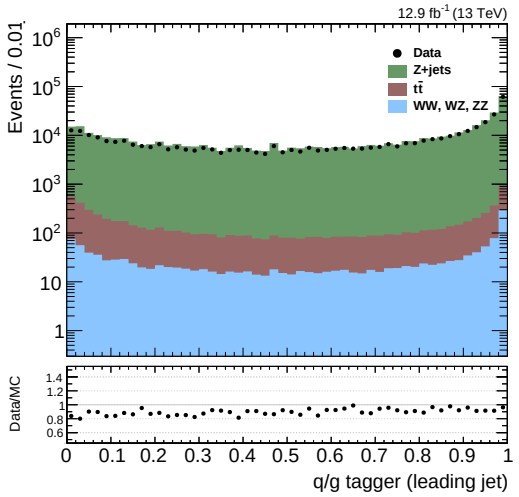

Figure 4.16: Distribution of the quark-gluon tagger $\mathcal{D}^{q/g}$ in 2016 data and simulation, for the highest-p_T jet passing the selection of Section 4.5.1 in a Z+1ℓ control region.

4.6 Conclusion

At this point of the analysis workflow, the following ingredients have been defined:

- a collection of identified, isolated, primary-vertex-compatible, calibrated electrons and muons,
- a collection of isolated photons consistent with final-state radiation,
- a collection of well-identified jets that do not overlap with these leptons and photons, and are accompanied by flavour-tagging variables.

Moreover, a good control over the efficiencies of the aforementioned selection cuts in data relative to simulation is achieved, and event-level scale factors are applied when necessary.

The way how all these pieces come together to build event-level quantities and extract the H$\to$4ℓ signal is discussed in Chapter 5.

Chapter 5

Probing the four-lepton final state

Having defined particle-level ingredients of the analysis, the extraction of the $H \to 4\ell$ signal can now be discussed, along with all successive steps that lead to the measurements of physics parameters. This includes defining and optimizing the signal region selection, building observables that discriminate the Higgs boson signal against backgrounds and different types of signals against one another, categorizing events, estimating the contribution of background processes to the signal region, modelling signals and performing the measurements. This chapter presents this core workflow of the $H \to ZZ \to 4\ell$ analysis in a way that does not exploit collision data yet. Data-to-simulation comparisons of event-level variables and final results are instead discussed later in Chapters 6 and 7, for the 2015 and 2016 analysis, respectively.

Building on the solid selection strategy inherited from Run I, my contribution consisted in extending the analysis so that it handles rare production mechanisms of the H(125) boson, such as vector boson fusion and associated VH and $t\bar{t}H$ production. Besides some changes to the event selection (Section 5.2), I focused on exploiting novel discriminants based on kinematics and jet tagging (5.3) to define a new event categorization (5.4). I then integrated the new categories into the rest of the analysis flow, including reducible background estimation (5.5.2), signal modelling (5.6), and the assessment of systematic uncertainties (5.7.3).

5.1 Simulated samples

From the optimization of the event selection to the production of results, various aspects of the analysis rely on Monte Carlo (MC) simulated samples for the involved signal and background processes. The production of a simulated sample includes several steps:

1. Generation of the hard physics process, which provides a list of events featuring a collection of 4-vectors of the outgoing particles. Different generator programs are used depending on the process. For all processes, the NNPDF 3.0 parton distribution functions are used [69].

2. Simulation of hadronization, jet fragmentation, and showering, and addition of the underlying event: this is done with PYTHIA 8 [70].

3. Addition of the overlapping pp interactions (pileup), with a profile close to what is expected in data taking. The simulated events are later reweighted to match the distribution of the number of interactions per bunch crossing to that observed in data.

4. Processing through a detailed simulation of the CMS detector based on GEANT4 [71, 72]. This provides a raw detector output for every event.

5. Reconstruction of the events with the same algorithms as used for data, with a similar version of the CMS software to that used to reconstruct the targeted data sets. The results of the HLT are also emulated during this step.

5.1.1 Signal samples

MC simulated samples for the signal processes are listed in Table 5.1. The five main production mechanisms of the SM Higgs boson are generated at NLO in perturbative QCD (pQCD) for $\sqrt{s} = 13\,\text{TeV}$ with the POWHEG 2.0 [73, 74, 75] generator: ggH [76], VBF [77], WH and ZH, for which the MINLO HVJ extension of POWHEG 2.0 is used [78], and $\mathrm{t\bar{t}H}$ [79]. The samples are reweighted according to state-of-the-art cross sections calculations that include some higher orders, as mentioned in Section 1.2.1. The decay of the Higgs boson to four leptons is modelled with JHUGEN V6 [80, 81], which properly accounts for interference effects associated with permutations of identical leptons in the 4e, 4μ and 4τ final states. Adding $2\mathrm{e}2\mu$, $2\mathrm{e}2\tau$, and $2\mu2\tau$, six final states are included at generator level in total, even though the analysis does not use reconstructed tau leptons. Nevertheless, owing to the existence of leptonic decays of the taus, a small amount of events involving tau pairs are reconstructed as 4ℓ ($\ell = \mathrm{e}, \mu$) events.

Table 5.1: List of signal MC samples used in the analysis. In this table, ℓ means e, μ, or τ, and X denotes the set of any other possible particles in the considered final states. Information on generator software is quoted, along with the production cross section at 13 TeV for $m_\mathrm{H} = 125\,\text{GeV}$ (see Section 1.2.1) times branching fraction of $\mathrm{H} \to 4\ell$ (2.76×10^{-4}). In the case of ZH and $\mathrm{t\bar{t}H}$, the latter number is replaced by the efficiency of the generator-level filter.

process		generator	$\sigma \cdot BR$
ggH	$\mathrm{gg} \to \mathrm{H} \to \mathrm{ZZ} \to 4\ell$	POWHEG 2.0 + JHUGEN V6	13.41 fb
VBF	$\mathrm{qq'} \to \mathrm{Hqq'} \to \mathrm{ZZqq'} \to 4\ell\mathrm{qq'}$	POWHEG 2.0 + JHUGEN V6	1.044 fb
WH	$\mathrm{q\bar{q}} \to \mathrm{W^+H} \to \mathrm{W^+ZZ} \to 4\ell + \mathrm{X}$	POWHEG 2.0 + MINLO HWJ + JHUGEN V6	0.232 fb
WH	$\mathrm{q\bar{q}} \to \mathrm{W^-H} \to \mathrm{W^-ZZ} \to 4\ell + \mathrm{X}$	POWHEG 2.0 + MINLO HWJ + JHUGEN V6	0.147 fb
ZH	$\mathrm{q\bar{q}} \to \mathrm{ZH} \to \mathrm{ZZZ} \to 4\ell + \mathrm{X}$	POWHEG 2.0 + MINLO HZJ + JHUGEN V6	0.668 fb
$\mathrm{t\bar{t}H}$	$\mathrm{gg} \to \mathrm{t\bar{t}H} \to \mathrm{t\bar{t}ZZ} \to 4\ell + \mathrm{X}$	POWHEG 2.0 + JHUGEN V6	0.393 fb

In the case of ZH and $\mathrm{t\bar{t}H}$, pairs of prompt leptons can also originate from the decay of the associated Z bosons and $\mathrm{t\bar{t}}$ pairs. Therefore, instead of only considering the decay to four leptons, all other decays of the $\mathrm{H} \to \mathrm{ZZ}$ system are allowed at first, and a generator-level filter then requires the presence of two pairs of opposite-sign same leptons in the final state, regardless of whether they originate from the H boson or from the associated Z or $\mathrm{t\bar{t}}$. In other words, part of the studied ZH and $\mathrm{t\bar{t}H}$ events feature a $\mathrm{H} \to 2\ell2\mathrm{X}$ decay where X either stands for a quark or for a neutrino. These events are often selected by the analysis, forming a non-resonant contribution to the signal region.

In order to parameterize the analysis expectations as a function of the Higgs boson mass m_H, signal samples are generated for a variety of values of m_H. Low-mass studies of the H(125) boson use five points: 120, 124, 125, 126, and 130 GeV. Moreover, to build signal models in studies of the H(125) that use the off-shell region or in new resonance searches, some higher-mass samples of ggH and VBF mechanisms are generated and processed through full simulation, for masses up to 3 TeV, and for the wide width $\Gamma_\mathrm{H}^\mathrm{SM}(m_\mathrm{H})$ of a SM Higgs boson at such masses. Signal modelling is discussed in detail in Section 5.6.

5.1.2 Background samples

The simulated background samples are listed in Table 5.2.

The dominant SM background to the Higgs boson signal arises from the production of ZZ or Zγ^* pairs through quark-antiquark annihilation ($q\bar{q} \to ZZ$) or gluon fusion ($gg \to ZZ$). This irreducible background is directly estimated from simulation. Additional scale factors that account for missing higher-order corrections will be discussed in Section 5.5.1.

Table 5.2: List of MC samples used for the estimation of irreducible backgrounds and for the study of control regions. Information on generator software is quoted, along with the production cross sections at 13 TeV times relevant branching fractions. These cross sections correspond to the phase space of event generation, and do not include K-factors for ZZ backgrounds (see Section 5.5.1). In this table, ℓ means e, μ, or τ.

process	additional information	generator	$\sigma \cdot BR$
$q\bar{q} \to ZZ \to 4\ell$		POWHEG 2.0	1.256 pb
$q\bar{q} \to ZZ \to 4\ell$		MadGraph5_aMC@NLO	1.212 pb
$gg \to ZZ \to 4e$	background only, no interference	MCFM	0.00159 pb
$gg \to ZZ \to 4\mu$			0.00159 pb
$gg \to ZZ \to 4\tau$			0.00159 pb
$gg \to ZZ \to 2e2\mu$			0.00319 pb
$gg \to ZZ \to 2e2\tau$			0.00319 pb
$gg \to ZZ \to 2\mu2\tau$			0.00319 pb
$Z/\gamma^* \to \ell\ell$ + jets	$m_{\ell\ell} > 50$ GeV	MadGraph5_aMC@NLO	6104 pb
$Z/\gamma^* \to \ell\ell$ + jets	$10 < m_{\ell\ell} < 50$ GeV		18610 pb
$WZ \to 3\ell\nu$		MadGraph5_aMC@NLO	5.29 pb
$t\bar{t}$		MadGraph5_aMC@NLO	815.96 pb
$t\bar{t} \to 2\ell2\nu2b$		POWHEG 2.0	87.31 pb

The $q\bar{q} \to ZZ$ process is generated at NLO in pQCD using POWHEG 2.0 [82], with the same settings as signal samples. Since this simulation spans a large range of ZZ invariant masses, the QCD factorization and renormalization scales are chosen to be dynamical, equal to m_{ZZ}.
The $gg \to ZZ$ process is generated at LO with MCFM [83, 84]. Here, to match the $gg \to H \to ZZ$ transverse momentum spectra predicted by POWHEG at NLO, the showering is performed with different PYTHIA 8 settings, only allowing emissions up to the parton-level scale.

Additional background samples of WZ, Drell-Yan + jets, $t\bar{t}$ are generated using either MadGraph5_aMC@NLO [85] or POWHEG 2.0. As will be explained in Section 5.5.2, the contribution of these reducible backgrounds is actually estimated from control regions in data, but the simulated samples are used for the optimization and validation of the methods. Drell-Yan + jets samples are also used in lepton-level optimization studies as a source of background leptons, as explained in Section 4.1.

5.1.3 Sample and event reweighting

To compute expected yields from simulated samples, and compare expected distributions to collision data, simulated events have to be properly weighted. This includes three contributions:

- The baseline procedure to obtain an expectation of event yields from separate MC samples consists in weighting each of them by the product of the process cross section and the targeted integrated luminosity, divided by the total number of simulated events.
- Some generator programs associate weights to the events they deliver, which need to be applied to properly modelling the physics process. This point is key when using MadGraph5_aMC@NLO and recent powheg processes, which can involve events with both positive and negative weights.
- Although the modelling of pileup in simulated samples is performed according to what is expected for data taking, the samples have to be reweighted a posteriori to match the actual pileup profile in data. The average number of additional pileup interactions is either computed from the number of reconstructed primary vertices or from the measured instantaneous luminosity per bunch crossing, and its distributions in simulation and data are used to compute pileup weights for a particular data taking period, and in a region of the phase space similar to that of the analysis. An illustration of the impact of these weights is presented in Fig. 5.1 for 2015 data. This procedure is found to have a sub-percent impact on expected yields in the signal region, which shows that the analysis is pileup-robust.

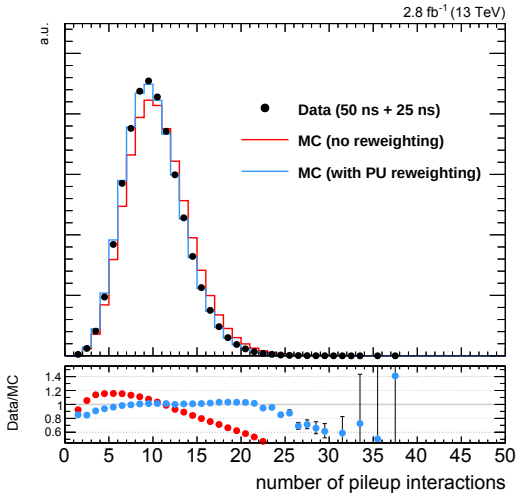

Figure 5.1: Distribution of the number of pileup interactions in 2015 simulation before and after pileup reweighting, compared to that in 2015 data. Simulated distributions use the Drell-Yan + jets, $t\bar{t}$ and ZZ processes. Events are required to pass the trigger requirement of the $H \to 4\ell$ analysis, and to contain a Z candidate (as defined in Section 5.2.2).

All weights are treated as independent, and the sum of products of generator weights and pileup weights for all simulated events is used as their denominator. The final event

weight is then multiplied by the data-to-simulation scale factor $SF_{4\ell}$ defined in Eq. 4.8, in order to correct for discrepancies in lepton reconstruction and selection efficiencies. This scales the expected yields down by a few percents, depending on the final state and data taking period.

5.2 Event selection

The $H \to ZZ \to 4\ell$ process is characterized by its low cross section, and by its very clean signature of four isolated leptons originating from the same vertex and compatible with the decay of a pair of Z bosons. As a result, only a handful of signal events are expected per inverse femtobarn of collision data, but an appropriate event selection can achieve excellent signal-to-background ratio.

This section defines the two main ingredients of event selection, which are technically independent: the *trigger* requirement, which is designed to be as close as possible to 100% efficient for the signal, and the *signal region* selection, where all care has to be taken to optimize the tradeoff between signal and backgrounds. The latter optimization was already made in Run I, and it has been moderately amended in Run II to comply with the new goal of unravelling production mechanisms. This step is fully independent from event categorization, which will be described in Section 5.4.

5.2.1 Trigger requirement

By definition, the events of the recorded collision data sample are known to have fired at least one path of the HLT menu. This menu is also emulated in MC samples, so that the same trigger requirement can be applied as in data.

As described in detail in Chapter 3, a dedicated collection of HLT paths has been designed and optimized to cover the phase space of the 4ℓ signal (the exact paths that are used depend on data taking conditions and will be listed before discussing the results). The trigger requirement is defined as follows: any event where at least one of these 4ℓ-dedicated paths is fired is always accepted, without tying the three possible final states to particular paths. The latter point is motivated by the hunt for production modes: for instance, the triple-electron path is not tied to the 4e final state, because a 2e2μ VH event can have additional true electrons coming from the decay of the associated W or Z boson.

Trigger efficiency measurement

The efficiency of the trigger requirement needs to be determined in collision data and simulation in order to correct for possible mismatches, either caused by mismodelling effects or by the fact that the menu used in data taking is steadily being optimized and diverges from the version that was included in MC samples. In general, efficiencies of leptonic HLT paths in collision data are meant to be measured for each leg with the Tag-and-Probe technique. However, this method is not applicable here, because many different paths need to be combined, and there are no Level-1 seeds with thresholds low enough to measure the efficiencies of the lowest-p_T legs of triple lepton triggers anyway. A different Tag-and-Probe-like method is used instead, based on 4ℓ events that pass the signal region selection, and using the HLT paths in a more global approach.

The tag is first defined as one of the four reconstructed leptons that geometrically matches a trigger-level lepton that fires a tight single-muon or single-electron path, and

the other three leptons are used as probes. Since more than one lepton may satisfy the tag requirement, up to four combinations of one tag and three probes are available in each 4ℓ event, and all of them are counted in the denominator of the efficiency. A set of probes is then counted in the numerator whenever the three corresponding trigger-level leptons fire at least one of the 4ℓ-dedicated paths. Such a method slightly underestimates the efficiency, because only three leptons are probed instead out of four; in particular, triplet of leptons of the 2e2μ channel cannot simultaneously probe the double-electron and double-muon paths. To monitor this bias, the efficiency in simulated samples is measured both with the same procedure and in a direct way (i.e. as the fraction of 4ℓ events that satisfy the trigger requirement among those that are reconstructed and selected in the signal region).

Results of the T&P-like and direct measurements are presented in Table 5.3 for 2015 data and MC simulation. Both data and $q\bar{q} \to ZZ$ simulation exhibit a nearly 100% efficiency in the 4μ final state and 97–99% for 4e and 2e2μ, and discrepancies are covered by the statistical uncertainties in data. These are of the order of 1–2% in the 4e and 2e2μ channel, which is almost as large as the overall amount of inefficiency, and also covers the bias with respect to the direct method. Additional studies of the dependency on probe lepton momenta led to similar levels of agreement, even in the low-p_T range where inefficiencies are larger. It was therefore decided to not apply any scale factor, but to use a flat systematic uncertainty that covers the discrepancies. It is defined as $^{+1\%}_{-4\%}$, $^{+0\%}_{-1\%}$ and $\pm 1\%$ in the 4e, 4μ and 2e2μ final states, respectively. As will be explained in Section 5.7.3, these numbers are then absorbed in the uncertainty on lepton trigger, reconstruction and selection.

Table 5.3: Efficiency of the set of HLT paths of the $H \to 4\ell$ analysis in 2015 data, measured either directly from MC simulation, or with the Tag-and-Probe-like method.

Final state	4e	4μ	2e2μ
ggH simulation (direct method)	$0.986^{+.003}_{-0.003}$	$0.998^{+.001}_{-0.001}$	$0.990^{+.002}_{-0.002}$
$q\bar{q} \to ZZ$ simulation (direct method)	$0.992^{+.004}_{-0.007}$	$0.998^{+.002}_{-0.003}$	$0.998^{+.001}_{-0.003}$
$q\bar{q} \to ZZ$ simulation (T&P-like method)	$0.994^{+.003}_{-0.004}$	$0.999^{+.001}_{-0.001}$	$0.977^{+.004}_{-0.004}$
data (T&P-like method)	$0.974^{+.019}_{-0.037}$	$1.000^{+.000}_{-0.010}$	$0.990^{+.007}_{-0.014}$

5.2.2 ZZ candidate selection

The signal region definition is mostly determined by the selection of one four-lepton candidate (or ZZ candidate) per event, on which any subsequent physics result will rely. This involves the following steps.

1. The event is required to contain at least one reconstructed primary vertex passing some quality requirements, such as a small radius and collisions restricted along the z axis. This step is nearly 100% efficient.

2. *Selected leptons* are defined as the reconstructed leptons that pass all selection cuts defined in Section 4.1, namely the impact parameter (4.1.1) identification (4.1.2 and 4.1.2) and FSR-corrected isolation (4.1.3) requirements.

3. *Z candidates* are defined as pairs of selected leptons of opposite charge and matching flavour (e^+e^-, $\mu^+\mu^-$) that satisfy $12 < m_{\ell\ell(\gamma)} < 120\,\text{GeV}$, where the $m_{\ell\ell(\gamma)}$ invariant mass includes the selected FSR photons if any.

4. *ZZ candidates* are defined as pairs of non-overlapping Z candidates. The Z candidate with invariant mass $m_{\ell\ell(\gamma)}$ closest to the nominal Z boson mass m_{Z^0} is denoted as Z_1, and the other one is denoted as Z_2.

5. ZZ candidates are required to satisfy the following list of requirements (unless otherwise stated, the selected FSR photons are included in invariant mass computations).

   - *Ghost removal*: all six pairs that can be built with the four leptons must be separated by $\Delta R > 0.02$.
   - *Lepton* p_T: Two of the four selected leptons must pass $p_{T,1} > 20\,\text{GeV}$ and $p_{T,2} > 10\,\text{GeV}$.
   - *Z_1 invariant mass*: $m_{Z_1} > 40\,\text{GeV}$.
   - *QCD suppression*: all four opposite-sign pairs that can be built with the four leptons (regardless of lepton flavour) must satisfy $m_{\ell^+\ell'^-} > 4\,\text{GeV}$. Here, selected FSR photons are not used in computing $m_{\ell^+\ell'^-}$, since a QCD-induced low mass dilepton resonance (e.g. J/ψ) may have photons nearby (e.g. from a π^0).
   - *Alternative pairing check*: defining Z_a and Z_b as the mass-sorted alternative pairing Z candidates (Z_a being the one closest to the nominal Z boson mass), the ZZ candidate is excluded if $m_{Z_b} < 12\,\text{GeV}$ while m_{Z_a} is closer to the nominal Z boson mass than m_{Z_1} is. This cut discards 4μ and 4e candidates where the alternative pairing looks like an on-shell Z plus a low-mass $\ell^+\ell^-$ resonance.
   - *Four-lepton invariant mass*: $m_{4\ell} > 70\,\text{GeV}$.

6. The events that contain at least one selected ZZ candidate form the signal region.

5.2.3 ZZ candidate choice

Some of the events selected in the signal region have more than four selected leptons, and can thus have more than one selected ZZ candidate. In particular, such a situation is expected to happen in VH and $t\bar{t}H$ events where the associated W, Z or $t\bar{t}$ particles decay leptonically. One candidate then has to be chosen, and is referred to as the *best ZZ candidate*. If this choice is incorrect, or if none of the selected ZZ candidates includes all four leptons from the $H \to 4\ell$ decay (because one lepton is not in the acceptance, not reconstructed or not selected), the event contributes to the non-resonant signal component, along with the $H \to 2\ell 2X$ events that were mentioned in Section 5.1.1.

For the 2015 analysis, the ZZ candidate choice strategy is similar to that used in the Run I analysis [23]. When more than one ZZ candidate is selected, the one with the Z_1 mass closest to the nominal Z boson mass is chosen. If two or more candidates include the same Z_1 and just differ by their Z_2 candidate, the Z_2 that has the largest scalar sum of lepton transverse momenta is chosen. This strategy is known not to be optimal: for instance, in $ZH \to 6\ell$ events, the associated Z boson is on mass shell, which often results

in the best ZZ candidate erroneously being made of this associated Z and the on-shell Z of the H boson decay.

To correct for this effect, a new strategy is introduced in the 2016 analysis, whereby the best ZZ candidate is now defined as the one with the largest value of $\mathcal{D}_{\text{bkg}}^{\text{kin}}$, a kinematic discriminant which will be defined in Section 5.3.1. When different ZZ candidates involve the same four leptons, they have identical values of $\mathcal{D}_{\text{bkg}}^{\text{kin}}$, and the candidate with the Z_1 mass closest to the nominal Z boson mass is again chosen. This overall change of logic is facilitated by the fact that the ZZ candidate selection flow described in Section 5.2.2 has been reordered for Run II, by applying all kinematic cuts before making a choice, instead of choosing a candidate during the cut flow. The $\mathcal{D}_{\text{bkg}}^{\text{kin}}$-based choice naturally favours ZZ candidates that have a kinematic configuration most consistent with a Higgs boson decays, a change that greatly benefits VH and $\text{t}\bar{\text{t}}\text{H}$ events, as numerically shown in Table 5.4.

Table 5.4: Fraction of events where the best ZZ candidate has four correctly assigned leptons, among selected events that have at least 5 selected leptons including the 4 leptons from the $\text{H}\to 4\ell$ decay. The correspondence of selected leptons with the $\text{H}\to 4\ell$ decay products is tested via a geometrical matching ($\Delta\text{R} < 0.1$) to the generator-level leptons.

	Run I strategy	Run II strategy (largest $\mathcal{D}_{\text{bkg}}^{\text{kin}}$)
WH	54%	77%
ZH	19%	50%
$\text{t}\bar{\text{t}}\text{H}$	52%	74%

Other strategies were investigated, for example choosing ZZ candidates for which $m_{4\ell}$ is found in a certain window around 125 GeV. The largest-$\mathcal{D}_{\text{bkg}}^{\text{kin}}$ logic was favoured because it is less likely to artificially bring background events in the vicinity of the Higgs boson peak.

5.2.4 Event selection efficiency

The efficiency for the analysis to trigger and select the $\text{H}\to 4\ell$ signal in the signal region is studied in Fig. 5.2 for the ggH process ($m_{\text{H}} = 125\,\text{GeV}$), for the three final states. About half of the 4ℓ events (4e: 46.9%, 4μ: 52.8%, $2\text{e}2\mu$: 49.3%) are found to have their four leptons in the detector acceptance, which is defined as $p_{\text{T}}^{\text{e}} > 7\,\text{GeV}$, $|\eta^{\text{e}}| < 2.5$ for electrons, and $p_{\text{T}}^{\mu} > 5\,\text{GeV}$, $|\eta^{\mu}| < 2.4$ for muons. Since electrons are more difficult to reconstruct than muons and compete against much more background objects, the efficiency difference between the three final states widens as the selection progresses: the final efficiency of the full analysis selection (including signal region and trigger) is 22.4%, 40.2% and 29.2% for the 4e, 4μ and $2\text{e}2\mu$ final states, respectively. To give a per-lepton overview of the impacts of acceptance and analysis selection, Fig. 5.3 compares the p_{T}^{ℓ} and $|\eta^{\ell}|$ distributions of the four leptons at generator level and after the full analysis selection is applied.

5.3 Observables

Any physics analysis where a signal is extracted has to rely on some well-chosen variables that discriminate it from its backgrounds. And more generally, measuring a physics

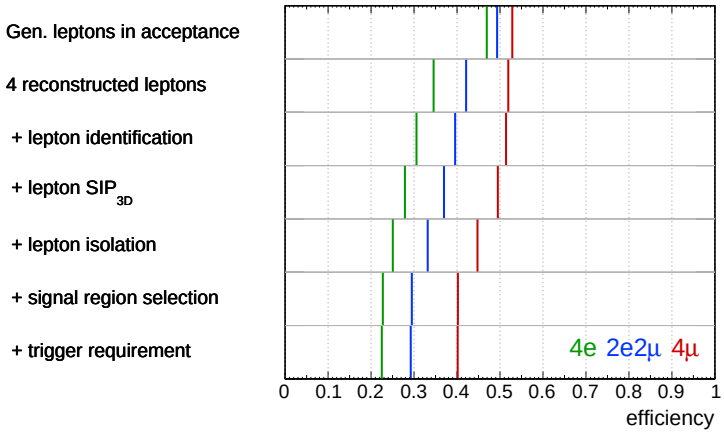

Figure 5.2: Efficiencies of different steps of the analysis selection, using a $gg \to H(125) \to ZZ \to 4\ell$ simulated sample, comparing three subsets of events corresponding to the generated 4e, 2e2μ and 4μ final states. The successive steps include the efficiency for the four generated leptons to be in the η^ℓ and p_T^ℓ acceptance, the efficiency for the event to contain two pairs of opposite-sign same-flavour reconstructed leptons, with increasingly stringent selection requirements up to the signal region selection, and the efficiency for the event to be both in the signal region and triggered.

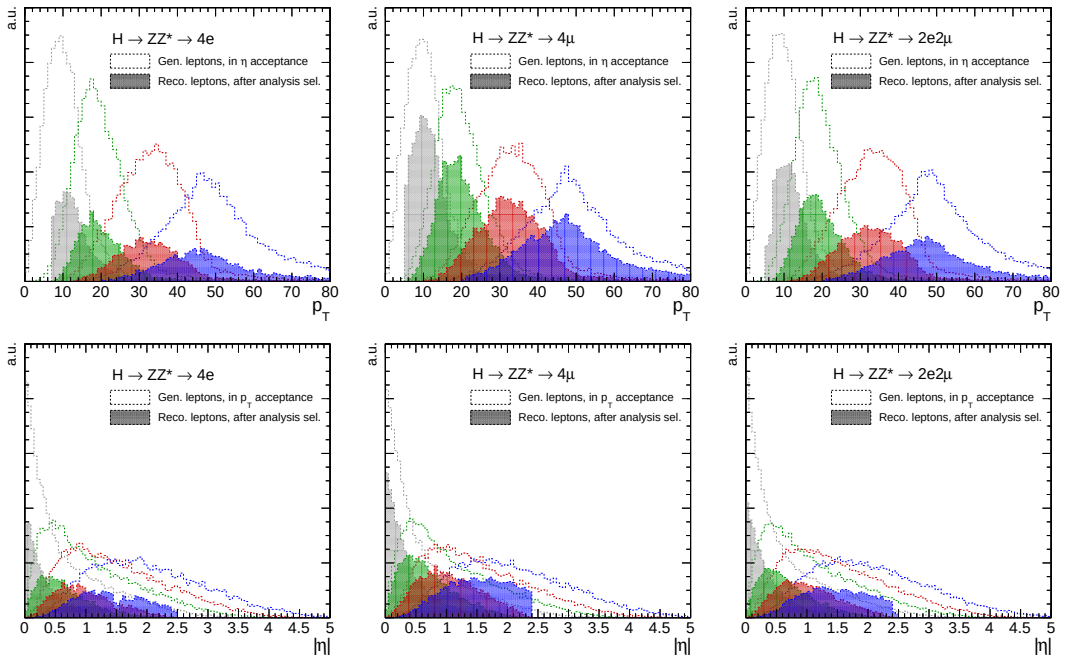

Figure 5.3: Distributions of the transverse momenta p_T^ℓ of the four p_T^ℓ-ordered leptons (top), and of the absolute pseudorapidities $|\eta^\ell|$ of the four $|\eta^\ell|$-ordered leptons (bottom) in the $gg \to H(125) \to ZZ \to 4\ell$ process, for the 4e (left), 4μ (middle) and 2e2μ (right) final states. Empty histograms show distributions for the four generated leptons, for events where all four leptons are in the $|\eta^\ell|$ acceptance defined as $|\eta^e| < 2.5$ and $|\eta^\mu| < 2.4$ (top), or in the p_T^ℓ acceptance defined as $p_T^e > 7\,\text{GeV}$ and $p_T^\mu > 5\,\text{GeV}$ (bottom). Shaded histograms show distributions for the reconstructed leptons of the best ZZ candidate, for events passing the full analysis selection.

parameter such as a mass or a coupling strength requires to find the observables that are most sensitive to it.

In both respects, the $H \to ZZ \to 4\ell$ analysis is in a privileged position, because it involves a complete reconstruction of a final state featuring many particles, ranging from just four leptons in the ggH case to more than ten objects (particles or jets) in the $t\bar{t}H$ events that fall entirely in the detector acceptance. The numerous associated observables are of two kinds:

- *kinematic properties*, i.e. any information that can be extracted from the four-momenta of these objects, such as angles and invariant masses,
- *internal properties*, such as jet flavour taggers and lepton momentum resolution, and their per-event counterparts.

These variables can be used either in selection cuts or in fits. For practical reasons, and in the case of fits, for computational reasons, this manifold information generally has to be condensed into a smaller number of combined observables.

In this analysis, kinematic discriminants related to either production or decay of the Higgs boson are mostly formed using matrix-element calculations. Compared to the commonly used multivariate techniques, this strategy has the advantage of extracting the information in a transparent way that is directly related to the physics of the targeted process, and does not depend on a prior training, thus avoiding the associated caveats. The matrix-element computations are performed with the Matrix Element Likelihood Approach (MELA) package [19, 80, 81, 86], whereby JHUGen and MCFM matrix elements are used for the signal and the background, respectively.

5.3.1 Decay discriminants

The decay part of the kinematic information is used for different purposes, such as the separation of the Higgs boson signal from the SM ZZ backgrounds, and the measurements of its mass, width and its spin/parity quantum numbers. In this analysis, two main decay-related observables are defined from the best ZZ candidate:

Four-lepton invariant mass

The most emblematic observable is the four-lepton invariant mass $m_{4\ell}$, which makes the Higgs boson signal appear as a narrow resonance peak near $m_{4\ell} = 125\,\text{GeV}$, on top of a background that is almost flat in this region. This variable directly benefits from the efforts in optimizing the lepton momentum resolution of CMS. It will be illustrated in detail in Chapters 6 and 7.

Kinematic discriminant

The second observable arises from the observation that the kinematic configuration of the four-lepton decay of a scalar boson is fundamentally different from that of the irreducible ZZ background. The four-momenta of $H \to 4\ell$ decay products carry eight independent degrees of freedom that fully describe this configuration in the centre-of-mass frame of the 4ℓ system, up to an arbitrary rotation around the beam axis [80, 81, 86]. A possible set of eight variables is given by:

- $m_{4\ell}$, the already mentioned four-lepton invariant mass,

- m_{Z_1} and m_{Z_2}, the invariant masses of the two constituent Z candidates,
- θ^*, the angle between the beam axis and the Z_1 direction in the 4ℓ rest frame,
- Φ and Φ_1, the two azimuthal angles between the three planes defined by the H boson, Z_1, and Z_2 decay products in the 4ℓ rest frame,
- θ_1 and θ_2, the angles between the Z_1 and Z_2 boost directions and the momenta of their daughter particles in the Z_1 and Z_2 rest frames, respectively.

The definition of the five angles is illustrated in Fig. 5.4 (left). Leaving $m_{4\ell}$ aside here, the seven remaining variables are collectively referred to as $\vec{\Omega}^{H\to 4\ell}$. A discriminant sensitive to the 4ℓ kinematics is then defined as

$$\mathcal{D}^{\text{kin}}_{\text{bkg}} \equiv \left[1 + \frac{\mathcal{P}^{q\bar{q}}_{\text{bkg}}(\vec{\Omega}^{H\to 4\ell}|m_{4\ell})}{\mathcal{P}^{gg}_{\text{sig}}(\vec{\Omega}^{H\to 4\ell}|m_{4\ell})}\right]^{-1} , \tag{5.1}$$

where the denominator is the probability for the Higgs boson signal, and the numerator is the probability for the dominant $q\bar{q}\to 4\ell$ background process, calculated from the LO matrix element squared with JHUGen and MCFM, respectively [19, 33].

As evidenced by the distributions of Fig. 5.4 (right), the $\mathcal{D}^{\text{kin}}_{\text{bkg}}$ variable provides significant discrimination between signal and ZZ background, and it only characterizes the decay of the Higgs boson, not its production mechanism. This partly motivates the choice to not describe the system in the laboratory frame here; another reason why the transverse momentum $p_{\text{T}}^{4\ell}$ of the 4ℓ system is not taken into account is to reduce the associated QCD uncertainties.

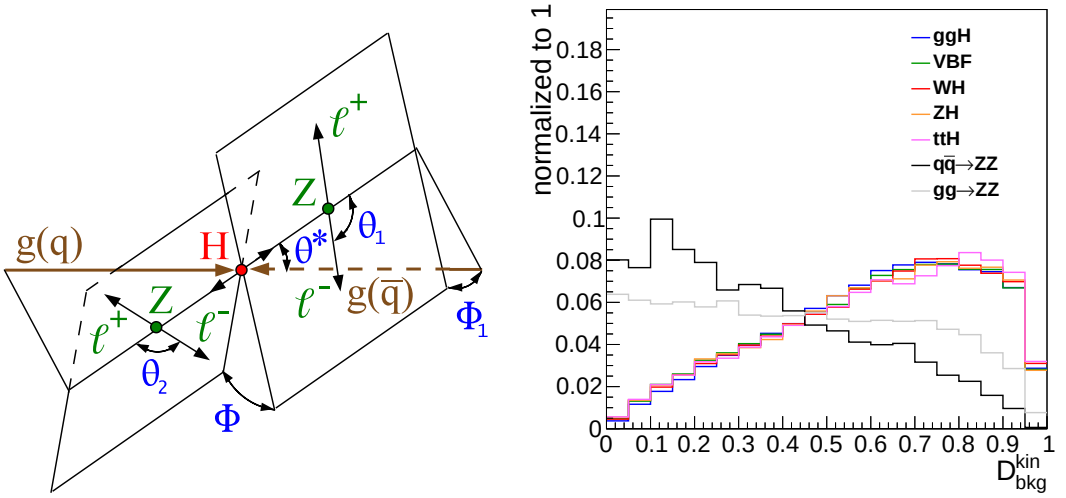

Figure 5.4: (left) Illustrations of the five angles that characterize the orientation of Higgs boson production and decay as $gg/q\bar{q} \to H \to ZZ \to 4\ell$. The θ^*, Φ_1, and Φ angles are defined in the H rest frame, whereas θ_1 and θ_2 are defined in the Z_1 and Z_2 rest frames, respectively. (right) Distribution of the decay kinematic discriminant $\mathcal{D}^{\text{kin}}_{\text{bkg}}$ for the main production mechanisms of the H(125) signal and for the $gg\to ZZ$ and $q\bar{q}\to ZZ$ background processes. Events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130$ GeV window.

5.3.2 Production discriminants

Identifying and extracting the five Higgs boson production mechanisms examined in this work (ggH, VBF, WH, ZH and $\mathrm{t\bar{t}H}$) requires to examine other objects than the four-lepton system, namely the selected jets (defined in Section 4.5.1), and the *additional leptons*, i.e. the selected leptons that are not involved in the best ZZ candidate. Various observables are defined from these objects and will be used as inputs to the event categorization described in Section 5.4.

Numbers of additional objects

The simplest production-related observables arise from counting the additional objects. This analysis uses four such observables, which are illustrated in Fig. 5.5:

- the number of selected jets, whereby VBF and hadronic VH events most often have one or two jets, and $\mathrm{t\bar{t}H}$ more than three;
- the number of selected b-tagged jets, mostly allowing extraction of $\mathrm{t\bar{t}H}$ events, but also of ZH events where the associated Z decays to $\mathrm{b\bar{b}}$;
- the number of additional leptons, the presence of which characterizes leptonic VH and leptonic $\mathrm{t\bar{t}H}$ events;
- the number of pairs of opposite-sign same-flavour additional leptons, which arise in leptonic ZH and fully leptonic $\mathrm{t\bar{t}H}$.

Matrix-element production discriminants

When selected jets are present, their kinematic properties can be encoded in matrix-element based observables that discriminate vector boson fusion and hadronic VH processes from ggH events produced in association with jets. The production-related kinematic information associated to VBF and hadronic-VH candidate events is fully described at leading order by five angles [86], which are collectively referred to as $\vec{\Omega}^{\mathrm{H+JJ}}$ and illustrated in Fig. 5.6. Compared to Section 5.3.1 and Fig. 5.4 (left), the angles are now defined from the point of view of the production HVV vertex, with the $\mathrm{Z_1}$–$\mathrm{Z_2}$ direction replacing the beam axis, and the $\mathrm{Vq\bar{q}}$ vertices replacing the $\mathrm{Z_1}$ and $\mathrm{Z_2}$ decays. The transverse momentum of the combined $4\ell + 2\,\mathrm{jets}$ system is not included in the input observables in order to reduce the associated uncertainties.

In events containing at least two selected jets, a discriminant sensitive to the VBF topology can be built as

$$\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-2j}} \equiv \left[1 + \frac{\mathcal{P}_{\mathrm{H+JJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{VBF}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}\right]^{-1} \tag{5.2}$$

where the $\mathcal{P}_{\mathrm{VBF}}$ and $\mathcal{P}_{\mathrm{H+JJ}}$ probabilities for VBF and $\mathrm{ggH} + 2\,\mathrm{jets}$ are obtained from the JHUGen matrix elements [41, 37]. The two highest-p_{T} jets are used when there are more than two. Distributions of $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-2j}}$ are shown for various processes in Fig. 5.7 (left). This discriminant equally well separates VBF from the $\mathrm{ggH} + 2\,\mathrm{jets}$ signal and from the $\mathrm{gg/q\bar{q}} \to 4\ell + 2\,\mathrm{jets}$ background, because the latter processes have similar jet kinematics.

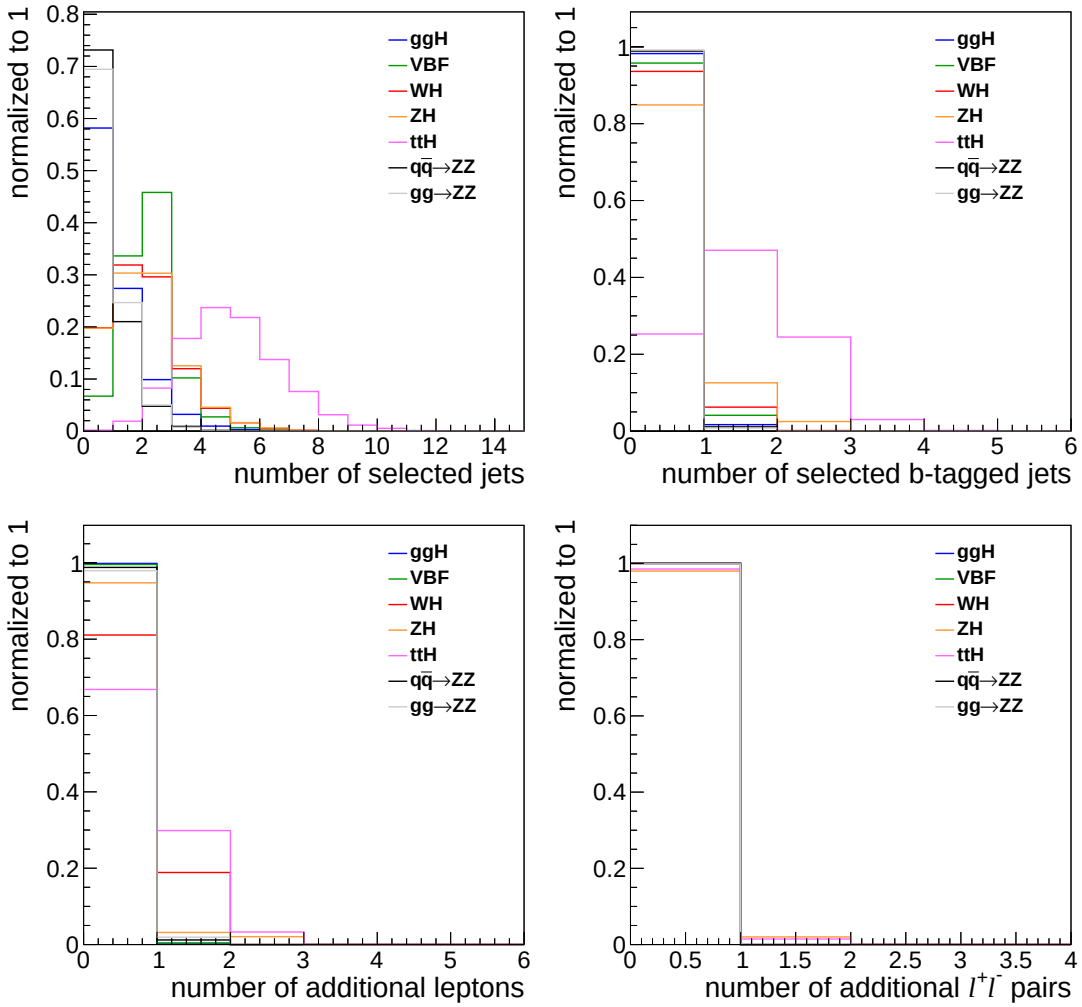

Figure 5.5: Normalized distributions of the number of selected jets (top left), of selected b-tagged jets (top right) of additional leptons (bottom left), and of pairs of opposite-sign additional leptons (bottom right), for the 5 main production processes of the H(125) signal and the ZZ irreducible backgrounds. Events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window.

Simulation shows that $\approx 40\%$ of the VBF events have fewer than two selected jets, which either means that a VBF jet is out of the detector acceptance, or that it is not reconstructed or fails the selection requirements defined in Section 4.5.1. The remaining jet alone still holds valuable information, given its forward position. A VBF signal probability can be constructed in events containing exactly one selected jet by simply integrating $\mathcal{P}_{\text{VBF}}$ over the pseudorapidity of the unobserved jet while constraining the transverse momentum of the $4\ell+2$ jets system to be zero. Together with the probability for ggH + 1 jet computed from JHUGEN, this leads to the discriminant:

$$\mathcal{D}^{\text{ME}}_{\text{VBF}-1\text{j}} \equiv \left[1 + \frac{\mathcal{P}_{\text{H+J}}(\vec{\Omega}^{\text{H+J}}|m_{4\ell})}{\int d\eta_{\text{J}} \mathcal{P}_{\text{VBF}}(\vec{\Omega}^{\text{H+JJ}}|m_{4\ell})}\right]^{-1} \tag{5.3}$$

which is illustrated in Fig. 5.7 (right) for signal and background processes.

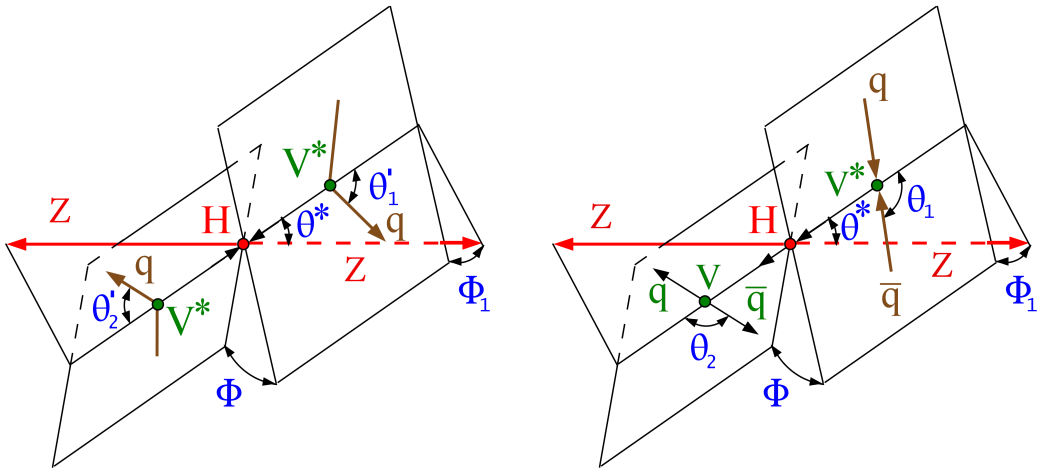

Figure 5.6: Illustrations of the five angles that fully describe the kinematic configuration of the Higgs boson production chain at leading order in VBF (left) and hadronic-VH (right) processes. The θ^*, Φ_1, and Φ angles are defined in the H rest frame; θ_1' and θ_2' are defined in the laboratory frame; and θ_1 and θ_2 are defined in the V* and V rest frames, respectively.

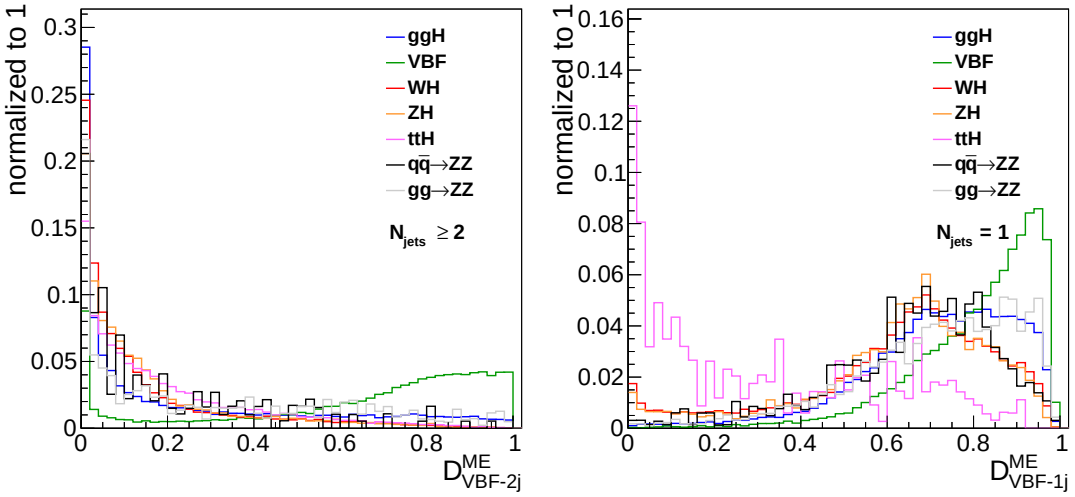

Figure 5.7: Distributions of the matrix-element based production discriminants sensitive to VBF production $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-2j}}$ and $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-1j}}$, in events with at least two or exactly one selected jet respectively, for the main production mechanisms of the H(125) signal and the $\mathrm{gg}\to\mathrm{ZZ}$ and $\mathrm{q\bar{q}}\to\mathrm{ZZ}$ background processes. Events are required to pass the full analysis selection and to be in a $118<m_{4\ell}<130\,\mathrm{GeV}$ window.

Finally, two discriminants separating hadronic VH processes from ggH + 2 jets are defined for events with at least two selected jets, as

$$\mathcal{D}^{\mathrm{ME}}_{\mathrm{WH-hadr.}} \equiv \left[1+\frac{\mathcal{P}_{\mathrm{HJJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{WH-hadr.}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}\right]^{-1}, \tag{5.4}$$

$$\mathcal{D}^{\mathrm{ME}}_{\mathrm{ZH-hadr.}} \equiv \left[1+\frac{\mathcal{P}_{\mathrm{HJJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{ZH-hadr.}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}\right]^{-1}, \tag{5.5}$$

where $\mathcal{P}_{\mathrm{WH-hadr.}}$ and $\mathcal{P}_{\mathrm{ZH-hadr.}}$ are again computed from JHUGEN matrix elements. These discriminants are very similar, the main differing point being the mass of the W and Z bosons. Distributions are compared in Fig. 5.8 for various production modes.
By contrast to VBF, hadronic VH events with one lost jet exhibit little outstanding kinematic information and are therefore not exploited in matrix elements. Matrix element probabilities for leptonic-VH processes have also been studied, but they are not exploited in the analysis because the sole presence of additional selected leptons was found to provide enough information to extract such processes.

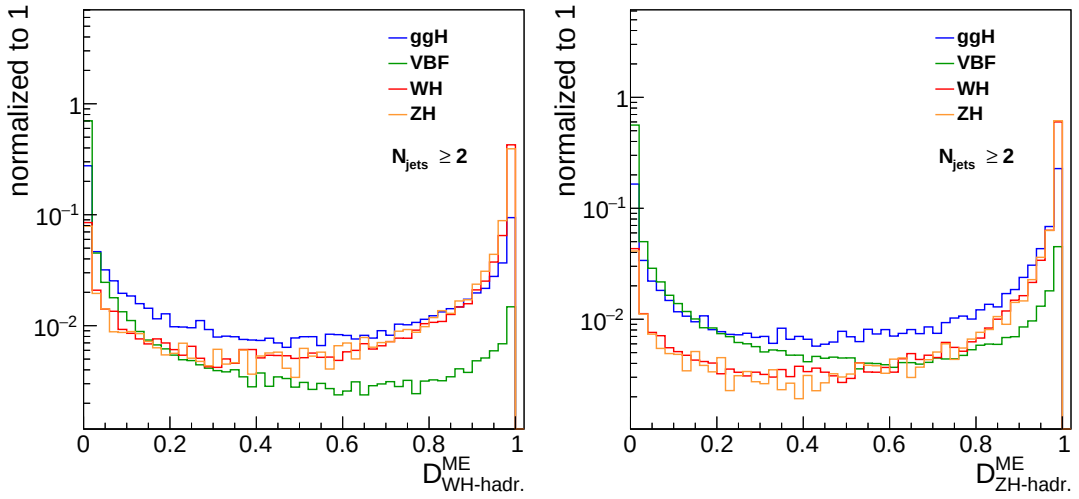

Figure 5.8: Distributions of the matrix-element based production discriminants sensitive to VH production $\mathcal{D}^{\mathrm{ME}}_{\mathrm{WH-hadr.}}$ and $\mathcal{D}^{\mathrm{ME}}_{\mathrm{ZH-hadr.}}$, in events with at least two selected jets, for the main production mechanisms of the H(125) signal. Events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\mathrm{GeV}$ window.

Quark-gluon taggers

Since jets produced next to the ggH process are typically gluon-induced, as opposed to pairs of jets from VBF and VH processes, quark-gluon tagging can be exploited at event level and provides information that is independent from that of jet 4-vectors.

The likelihood discriminant defined in Section 4.5.3 can be adapted to these processes that feature two tagging jets. Using again the two highest-p_{T} selected jets J_1 and J_2, the tagger reads:

$$\mathcal{D}^{\mathrm{q/g}}_{\mathrm{2jets}} \equiv \left[1 + \frac{\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)}{\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)}\right]^{-1} . \tag{5.6}$$

Events containing exactly one jet can exploit the one-jet version of the tagger:

$$\mathcal{D}^{\mathrm{q/g}}_{\mathrm{1jet}} \equiv \left[1 + \frac{\mathcal{G}(\mathrm{J})}{\mathcal{Q}(\mathrm{J})}\right]^{-1} . \tag{5.7}$$

Combined production discriminants

For computational and practical reasons, it makes sense to combine information from the matrix element method and from quark-gluon tagging into one single variable per

process. For instance, to separate VBF from gluon fusion in two-jet events, one could write

$$\left[1 + \frac{\mathcal{P}_{\mathrm{H+JJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{VBF}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})} \times \frac{\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)}{\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)}\right]^{-1} , \tag{5.8}$$

relying on the assumption that the $\mathcal{P}_{\mathrm{H+JJ}}\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)$ and $\mathcal{P}_{\mathrm{VBF}}\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)$ quantities behave as products of uncorrelated probabilities. However, there are several reasons why this assumption does not hold. On the one hand, the involved factors were found to be correlated, albeit moderately; on the other hand, the $\mathcal{Q}$ and $\mathcal{G}$ are products of correlated pdfs, as explained in Section 4.5.3. As a result, within the quark-gluon discriminant, the information shared by the three input variables is given more weight than the more independent information, while the quark-gluon discriminant itself is then given too much weight compared to the matrix-element discriminant.

This non-optimal handling of correlations can be partially bypassed by applying some transformation functions to $\mathcal{G}/\mathcal{Q}$ ratios, in order to give them less weight in the combined tagger. Since they are products of correlated variables, I tested a variety of $x \longmapsto x^{1/n}$ fractional exponent functions. Figure 5.9 presents the outcome of such tests. For the four situations already envisaged with matrix-element discriminants, ROC curves are presented for individual matrix-element and quark-gluon taggers, and for four versions of the combined taggers (untransformed, square, cube and quartic root).

The 2-jet-based extraction of VBF exhibits the most complicated behaviour, since its untransformed combined tagger does not even perform better than matrix elements alone. Its transformed versions however do perform better, and very similarly to one another. For the 1-jet based extraction of VBF and the extraction of hadronic WH and ZH, the untransformed combined taggers outperform both their components, and are further improved by the fractional exponent transformation functions, which are again similar to one another. The best transformation depends on the desired efficiency. For overall consistency, a cube root transformation is chosen for all four taggers.

Based on the above observations, the following four combined taggers are built for the subsequent analysis:

- to separate VBF from gluon fusion in events with at least 2 selected jets:

$$\mathcal{D}^{\mathrm{comb.}}_{\mathrm{VBF-2j}} \equiv \left[1 + \frac{\mathcal{P}_{\mathrm{H+JJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{VBF}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})} \times \left(\frac{\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)}{\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)}\right)^{1/3}\right]^{-1} \tag{5.9}$$

- to separate VBF from gluon fusion in events with exactly 1 selected jet:

$$\mathcal{D}^{\mathrm{comb.}}_{\mathrm{VBF-1j}} \equiv \left[1 + \frac{\mathcal{P}_{\mathrm{H+J}}(\vec{\Omega}^{\mathrm{H+J}}|m_{4\ell})}{\int d\eta_{\mathrm{J}} \mathcal{P}_{\mathrm{VBF}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})} \times \left(\frac{\mathcal{G}(\mathrm{J})}{\mathcal{Q}(\mathrm{J})}\right)^{1/3}\right]^{-1} \tag{5.10}$$

- to separate hadronic WH from gluon fusion in events with at least 2 selected jets:

$$\mathcal{D}^{\mathrm{comb.}}_{\mathrm{WH-hadr.}} \equiv \left[1 + \frac{\mathcal{P}_{\mathrm{HJJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{WH-hadr.}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})} \times \left(\frac{\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)}{\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)}\right)^{1/3}\right]^{-1} \tag{5.11}$$

- to separate hadronic ZH from gluon fusion in events with at least 2 selected jets:

$$\mathcal{D}^{\mathrm{comb.}}_{\mathrm{ZH-hadr.}} \equiv \left[1 + \frac{\mathcal{P}_{\mathrm{HJJ}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})}{\mathcal{P}_{\mathrm{ZH-hadr.}}(\vec{\Omega}^{\mathrm{H+JJ}}|m_{4\ell})} \times \left(\frac{\mathcal{G}(\mathrm{J}_1)\mathcal{G}(\mathrm{J}_2)}{\mathcal{Q}(\mathrm{J}_1)\mathcal{Q}(\mathrm{J}_2)}\right)^{1/3}\right]^{-1} \tag{5.12}$$

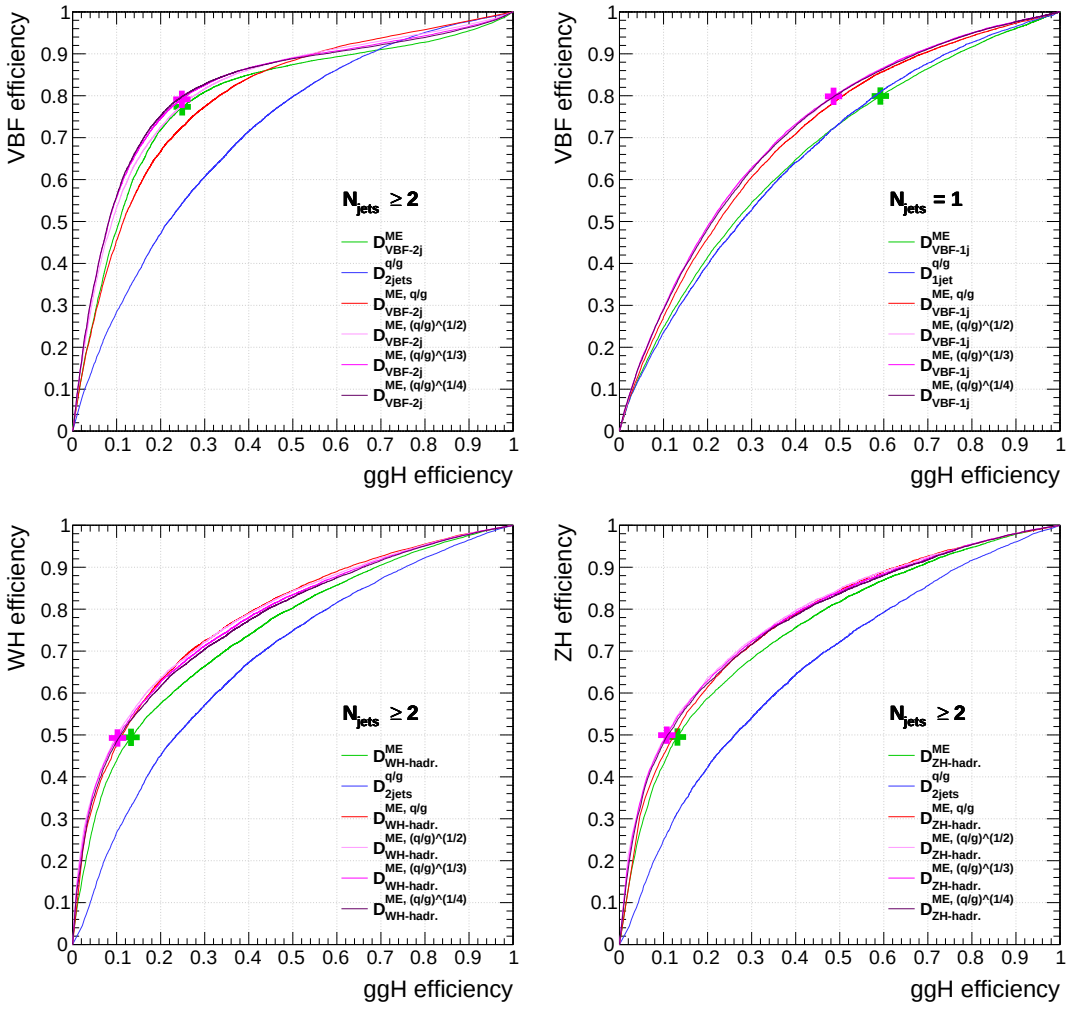

Figure 5.9: Performance comparison of matrix-element based, quark-gluon based and combined discriminants. Four cases are studied: 2-jet VBF vs. gluon fusion H + 2 jets (top left), 1-jet VBF vs. gluon fusion H + 1 jet (top right), hadronic WH vs. gluon fusion H + 2 jets (bottom left), and hadronic ZH vs. gluon fusion H + 2 jets (bottom right). The combined discriminants are shown in an untransformed version, and with three transformations of the $\mathcal{G}/\mathcal{Q}$ ratio (square, cube and quartic root). Samples are generated with $m_{\mathrm{H}} = 125\,\mathrm{GeV}$; events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\mathrm{GeV}$ window. The cross markers represent operating points for event categorization, which will be defined in Section 5.4.

5.3.3 Per-event mass uncertainties

The uncertainty $\sigma_{m_{4\ell}}$ on the measurement of the four-lepton invariant mass $m_{4\ell}$ can be estimated from the individual lepton momentum uncertainties on a per-event basis. The relative uncertainty $\mathcal{D}_{\mathrm{m}} \equiv \sigma_{m_{4\ell}}/m_{4\ell}$ is then used as an observable: since it can vary a lot from one event to another, it helps the Higgs boson mass measurement, by giving more importance to the observed events that have a low uncertainty.

The uncertainty on the muon momentum measurement relies on the full uncertainty matrix as obtained from the muon track fit, while the momentum uncertainty of electrons

is estimated from the combination of the ECAL and tracker measurement, neglecting the uncertainty on the track direction. When some FSR photons are selected, the uncertainty on the p_{T}^{γ} is taken from a parameterization obtained from simulation. To compute the per-event mass uncertainty, the contribution of the transverse momentum of each individual lepton or FSR photon to the overall $m_{4\ell}$ uncertainty is calculated separately, and $\sigma_{m_{4\ell}}$ is defined as the sum in quadrature of these ≥ 4 contributions (all directional uncertainties are neglected). This simple method was shown to agree at a $<1\%$ level with an analytical propagation of per-lepton uncertainties to $\sigma_{m_{4\ell}}$ that includes all correlations.

These per-event mass uncertainties then need to be calibrated to match the observed mass resolution. This is done by correcting the lepton-level momentum uncertainties which then propagate to $\sigma_{m_{4\ell}}$ (this comes on top of the calibrations that were described in Section 4.2). The correction is performed both in data and simulation, using $\mathrm{Z}\to\mathrm{e}^{+}\mathrm{e}^{-}$ and $\mathrm{Z}\to\mu^{+}\mu^{-}$ event samples. The lineshape of these resonances is modelled with a Breit-Wigner function convolved with a double Crystal Ball plus exponential function. Events are divided into categories based on the p_{T}^{ℓ} and η^{ℓ} of the leptons, and the dilepton mass resolution predicted by the above technique is compared to the resolution extracted from a fit to the dilepton mass in each category, so that a correction to lepton momentum uncertainties can be extracted by an iterative procedure. After the corrections are derived, a closure test of the agreement between the predicted and fitted 4ℓ mass resolution is performed in data and in simulation, in bins of the predicted $m_{4\ell}$ resolution, confirming that the calibration brings it close to the fitted one. A systematic uncertainty of $\pm 20\%$ on the 4ℓ mass resolution is assigned to cover the residual differences.

5.4 Event categorization

The strategy for the extraction of Higgs boson production modes is based on *event categorization*, a step that consists in distributing all selected events among mutually exclusive event categories, each of which is enriched in a particular process or set of processes. The available information is exploited in two ways:

- Production-related observables, i.e. the numbers of associated objects and matrix-element or combined discriminants defined in Section 5.3.2, are only used as criteria to define the event categories, with the goal of maximizing both the purities and expected yields of said categories.
- Decay-related observables, i.e. $m_{4\ell}$ and $\mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}}$ defined in Section 5.3.1, are used to build two-dimensional pdfs in each of these event categories, in order to extract the processes via an unbinned likelihood fit. This point will be explained in Sections 5.6 and 5.7.

Focusing on the first point, three different event categorization schemes were successively designed in this thesis work, targeting different rounds of public results.

5.4.1 Categorization for the 2015 analysis

To produce the public results based on 2015 data which will be reported in Chapter 6, a limited integrated luminosity was available. Thus, a simple two-category scheme based was used, focusing on vector boson fusion only. The categories were defined as follows:

1. at least 2 selected jets and $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}} > 0.5 \longrightarrow$ **VBF-tagged category**;
2. remaining events $\longrightarrow$ **Untagged category**.

Compared to the categorization used in Run I [23], the addition of the cut on $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$ enhances the purity of the VBF-tagged category. Figure 5.10 illustrates the contents of each category in terms of the five main Higgs boson production modes, and the expected event yields in 2015 data.

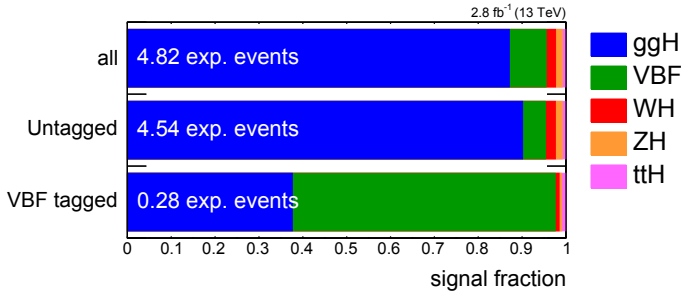

Figure 5.10: Signal process composition of the full selected event sample (line 1) and of the two event categories (lines 2 and 3), for the categorization of the 2015 analysis. Samples are generated with $m_{\text{H}} = 125\,\text{GeV}$; events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window.

5.4.2 Categorization for the 2016 analysis

The larger data sets that are being collected in 2016 call for the introduction of a more complete categorization strategy, which does not only target VBF, but also associated VH and $\text{t}\bar{\text{t}}\text{H}$ production. This section describes the version of the event categorization that was used in the public results of the 2016 analysis reported in Chapter 7. All the observables defined in Section 5.3.2 are now exploited, except those related to quark/gluon tagging.

In order to design pure event categories, a bottom-up approach is used. First, numerous baskets of selected events are defined according to all possible combinations of the following criteria:

- 0, 1, 2, 3, or ≥ 4 selected jets,
- 0, 1, or ≥ 2 selected b-tagged jets,
- 0, 1, or ≥ 2 additional leptons,
- 0, or ≥ 1 pairs of opposite-sign same-flavour additional leptons,
- cuts on $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$, $\mathcal{D}^{\text{ME}}_{\text{WH}-\text{hadr.}}$ and $\mathcal{D}^{\text{ME}}_{\text{ZH}-\text{hadr.}}$ (for events with at least two selected jets),
- a cut on $\mathcal{D}^{\text{ME}}_{\text{VBF}-1\text{j}}$ (for events with exactly one selected jet).

Second, the composition of every basket is evaluated in terms of relative fractions of production processes. For instance, baskets using a cut on $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$ are dominated by VBF, while $\text{t}\bar{\text{t}}\text{H}$ tends to populate baskets featuring b-tagged jets and/or leptons, etc.

Third, the baskets are merged with one another, so as to form only one event category per process of interest. In general, each basket is assigned to the category whose process most contributes to it. This procedure results in six categories, the definitions of which are given by the following criteria, applied in this very order:

1. exactly 4 leptons + (2 or 3 jets, of which at most 1 b-tag, OR at least 4 jets + no b-tag) + $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}} > 1.043 - \frac{460}{m_{4\ell}+634}$
   $\longrightarrow$ **VBF-2jet-tagged category**

2. exactly 4 leptons + (2 or 3 jets, of which at most 1 b-tag, OR at least jets + no b-tag) + ($\mathcal{D}^{\text{ME}}_{\text{WH}-\text{hadr.}} > 0.959$ OR $\mathcal{D}^{\text{ME}}_{\text{ZH}-\text{hadr.}} > 0.9946$)
   OR: exactly 4 leptons + 2 or 3 jets, of which at least 2 b-tags
   $\longrightarrow$ **VH-hadronic-tagged category**

3. no more than 3 jets + no b-tagged jet in the event + (exactly 1 additional lepton OR 1 pair of additional opposite-sign same-flavour leptons)
   OR no jet + at least 1 additional lepton
   $\longrightarrow$ **VH-leptonic-tagged category**

4. at least 4 jets, of which at least 1 b-tagged,
   OR: at least 1 additional lepton
   $\longrightarrow$ **t$\bar{\text{t}}$H-tagged category**

5. exactly 4 leptons + exactly 1 jet + $\mathcal{D}^{\text{ME}}_{\text{VBF}-1\text{j}} > 0.699$
   $\longrightarrow$ **VBF-1jet-tagged category**

6. remaining events
   $\longrightarrow$ **Untagged category**

Since VBF extraction relies on two very different discriminants $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$ and $\mathcal{D}^{\text{ME}}_{\text{VBF}-1\text{j}}$, it was decided to have both a VBF-2jet-tagged and a VBF-1jet-tagged category instead of a single VBF-tagged one. Moreover, instead of having a WH-tagged and a ZH-tagged category, a VH-hadronic-tagged one and a VH-leptonic-tagged one are defined. Indeed, on the one hand, $\text{ZH} \to 6\ell$ events rarely have 6 selected leptons and tend to cluster with $\text{WH} \to 5\ell\nu$ events, and on the other end, $\mathcal{D}^{\text{ME}}_{\text{WH}-\text{hadr.}}$ and $\mathcal{D}^{\text{ME}}_{\text{ZH}-\text{hadr.}}$ are so similar that they do not really help separating WH from ZH either. The VH-hadronic-tagged category thus relies on the logical 'or' of both discriminants, and also has a b-tagging-based criterion that selects $\text{Z}(\to \text{b}\bar{\text{b}})\text{H}$.

The cuts on the production discriminants use a set of four operating points, which are illustrated as green markers on Fig. 5.9. These thresholds were tuned so as to have a good compromise between purity and yield in the corresponding categories. This was found to require tighter cuts for WH and ZH than for VBF, because the signature of hadronic VH is less distinctive. For $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$, the threshold was made $m_{4\ell}$-dependent so as to have the same VBF tagging efficiency in the separate high-mass studies.

Figure 5.11 reports the relative contributions of every signal process in the six event categories, together with the expected event yields for the 2016 data sample of Chapter 7. At this point, less than one event is expected in the VH-hadronic-tagged, VH-leptonic-tagged, and t$\bar{\text{t}}$H-tagged categories, but such a six-category scheme is meant to be ultimately used with the full Run II data sample, i.e. with one order of magnitude more data. Figures 5.12 and 5.13 describe how the various signal and irreducible background processes are individually distributed among the six categories, while Fig. 5.14 illustrates the S/(S+B) ratios.

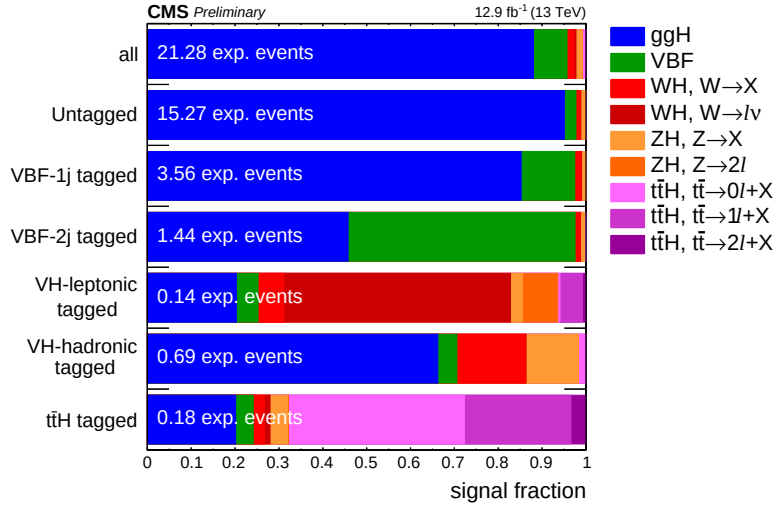

Figure 5.11: Signal process composition of the full selected event sample (line 1) and of the six event categories (other lines), for the categorization of the 2016 analysis. The WH, ZH and $t\bar{t}H$ processes are split according to the generated decay of the associated W, Z, or $t\bar{t}$ system, whereby the X symbol stands for a set of any other particles than electrons or muons. Samples are generated with $m_H = 125\,\text{GeV}$; events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window.

5.4.3 Improved categorization from combined discriminants

To produce future $H \to ZZ \to 4\ell$ results, a new version of the event categorization is under study. Compared to the six categories of Section 5.4.2, the only change consists in replacing the four cuts on matrix-element production discriminants by cuts on the discriminants that combine information from the matrix element method and from quark-gluon tagging:

$$\begin{aligned}
\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}} &> 1.043 - \frac{460}{m_{4\ell} + 634} &\longrightarrow\quad & \mathcal{D}^{\text{comb.}}_{\text{VBF}-2\text{j}} > 0.391 \\
\mathcal{D}^{\text{ME}}_{\text{VBF}-1\text{j}} &> 0.699 &\longrightarrow\quad & \mathcal{D}^{\text{comb.}}_{\text{VBF}-1\text{j}} > 0.72 \\
\mathcal{D}^{\text{ME}}_{\text{WH}-\text{hadr.}} &> 0.959 &\longrightarrow\quad & \mathcal{D}^{\text{comb.}}_{\text{WH}-\text{hadr.}} > 0.973 \\
\mathcal{D}^{\text{ME}}_{\text{ZH}-\text{hadr.}} &> 0.9946 &\longrightarrow\quad & \mathcal{D}^{\text{comb.}}_{\text{ZH}-\text{hadr.}} > 0.996
\end{aligned}$$

The operating points of the four new taggers are illustrated as pink markers on Fig. 5.9. The additional discrimination power is used to enhance VBF efficiency in the case of $\mathcal{D}^{\text{comb.}}_{\text{VBF}-2\text{j}}$, and to improve ggH rejection in the case of $\mathcal{D}^{\text{comb.}}_{\text{VBF}-1\text{j}}$, $\mathcal{D}^{\text{comb.}}_{\text{WH}-\text{hadr.}}$, and $\mathcal{D}^{\text{comb.}}_{\text{ZH}-\text{hadr.}}$.

Figure 5.15 shows the updated version of the category compositions and expected yields, which can be compared to Fig. 5.11. As expected, the VBF-2jet-tagged, VBF-1jet-tagged, and VH-hadronic-tagged categories have become more pure, either because they accept more signal (VBF-2jet) or because they reject more background (VBF-1jet and VH-hadronic).

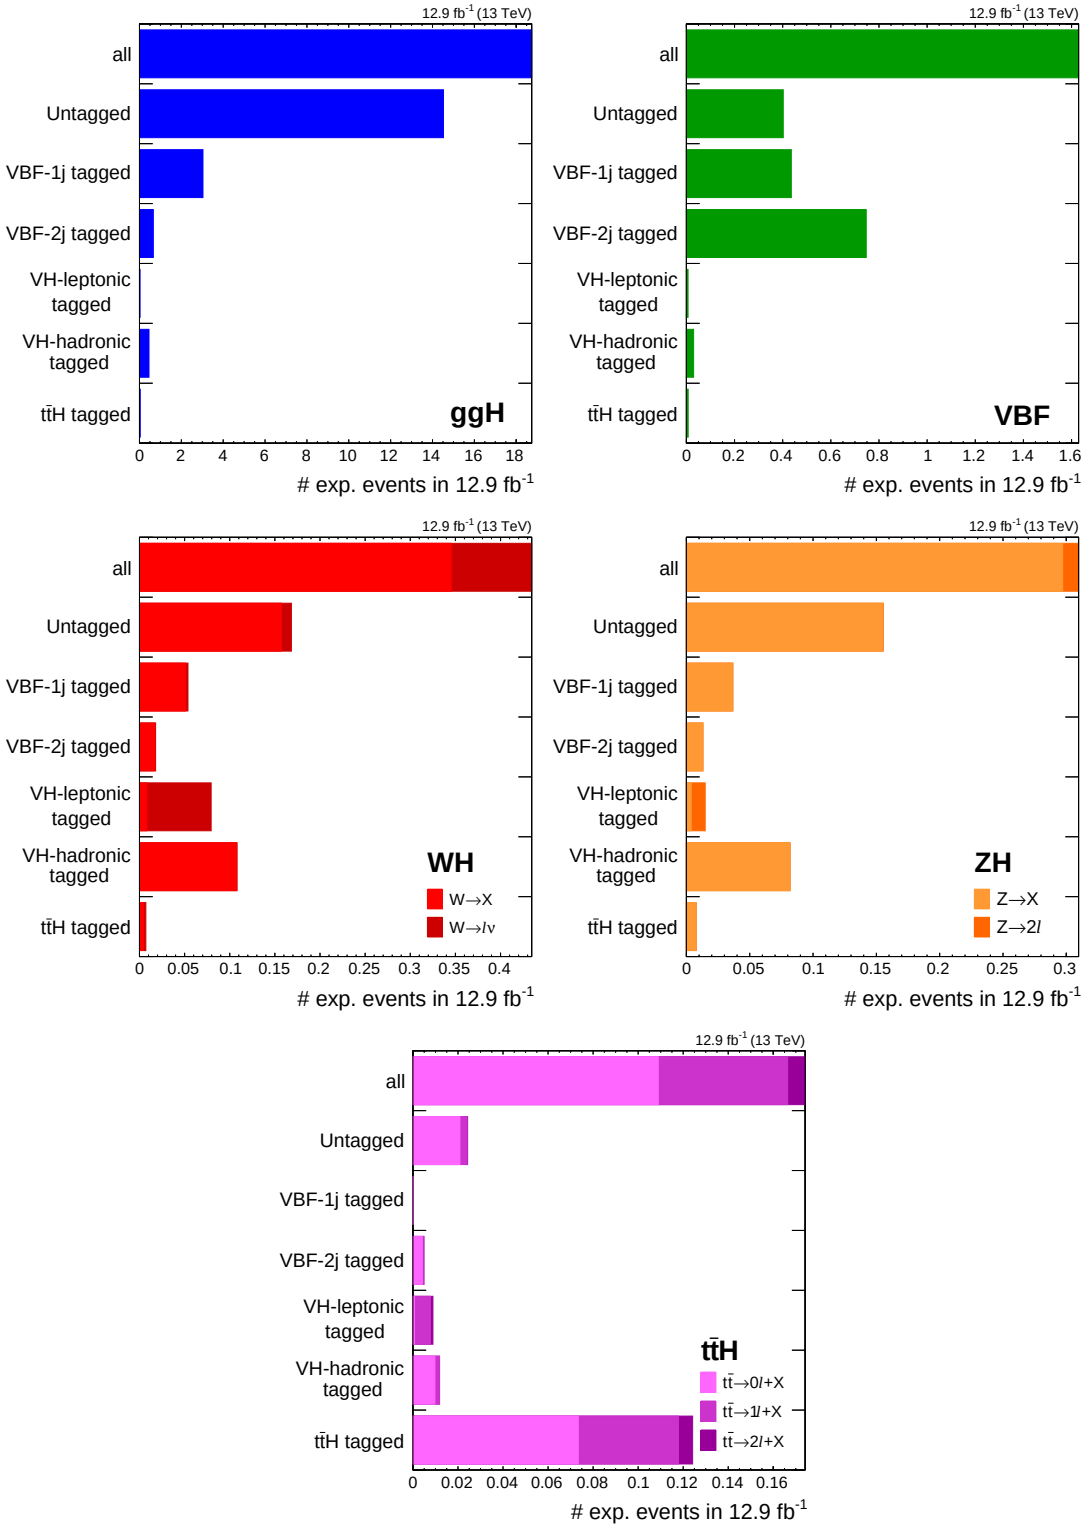

Figure 5.12: Distributions of signal processes in the six categories, for the categorization of the 2016 analysis. The ggH, VBF, WH, ZH, and $t\bar{t}H$ processes are presented in that order. The three latter processes are split according to the generated decay of the associated W, Z, or $t\bar{t}$ system, whereby the X symbol stands for a set of any other particles than electrons or muons. Samples are generated with $m_H = 125\,\text{GeV}$; events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window. The first line of each plot shows the expected yield of the full selected event sample.

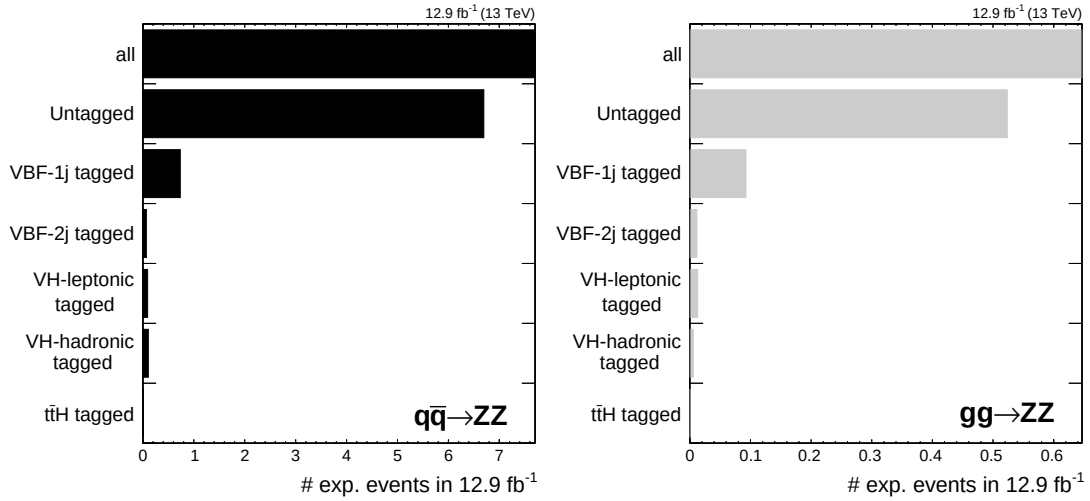

Figure 5.13: Distributions of the $q\bar{q} \to ZZ$ (left) and $gg \to ZZ$ (right) irreducible background processes in the six categories, for the categorization of the 2016 analysis. Events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window. The first line of each plot shows the expected yield of the full selected event sample.

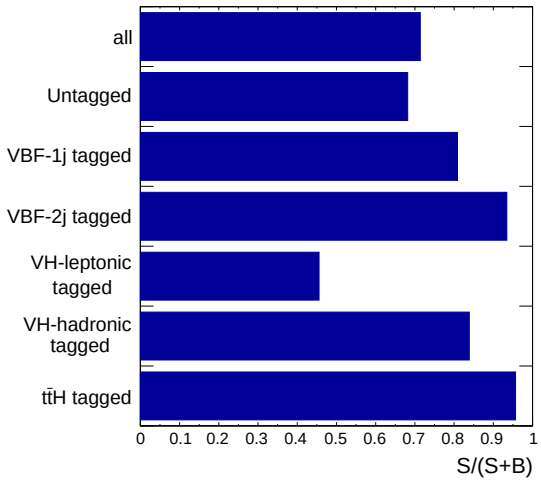

Figure 5.14: Ratio of signal to signal plus irreducible background yields in the full selected event sample (line 1) and in the six event categories (other lines), for the categorization of the 2016 analysis. Signal samples are generated with $m_{\text{H}} = 125\,\text{GeV}$; events are required to pass the full analysis selection and to be in a $118 < m_{4\ell} < 130\,\text{GeV}$ window.

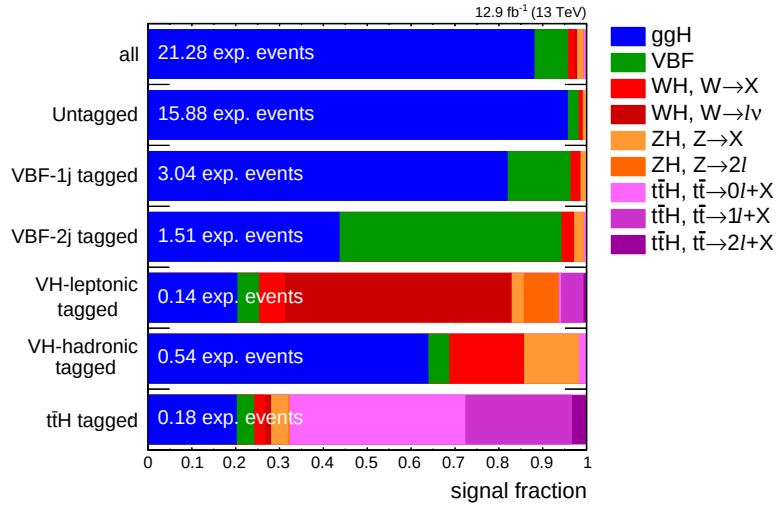

Figure 5.15: Signal process composition of the full selected event sample (line 1) and of the six event categories (other lines), for the improved categorization that uses combined discriminants. Colours, symbols, and event selection are defined as in Fig. 5.11.

5.5 Background estimation

On a conceptual level, the extraction of a signal requires to estimate two types of backgrounds.

- Physics processes which have the exact same final state as the targeted signal are called *irreducible backgrounds*. They cannot be distinguished from the signal looking at one single event: they only differ from it in the statistical properties of their kinematics.
- Events that pass the analysis selection as a result of object misreconstruction or misidentification are referred to as the *reducible backgrounds*. They materialize the non-perfect performance of the detector and of the analysis techniques.

5.5.1 Irreducible backgrounds

Processes that have two pairs of opposite-sign same-flavour prompt leptons in their final state are irreducible backgrounds to the $H \to ZZ \to 4\ell$ signal. Two such processes are studied here: the production of ZZ pairs via $q\bar{q}$ annihilation ($q\bar{q} \to ZZ$) and via gluon fusion ($gg \to ZZ$). They are collectively referred to as ZZ backgrounds, and their contribution to the signal region is estimated using MC simulation, both for expected yields and for distributions of observables.

However, the ZZ background contributions to the signal region have to be corrected for some higher-order contributions that are not taken into accounts by MC generators. This is done by multiplying the expected event yields by the so-called K-factors. These are applied differentially, usually as a function of the generated-level m_{ZZ}, so that the final irreducible background expectation reflects the state-of-the-art knowledge of the production cross sections for the given process.

The $q\bar{q} \to ZZ$ background is generated with POWHEG at NLO, but a full computation of its differential cross section exists at NNLO [87], although it is not yet available in

a partonic-level event generator. Therefore, a NNLO/NLO K-factor is applied, based on cross sections obtained using the same parton distribution functions (PDFs) and renormalization and factorization scales as the POWHEG sample. This K-factor is applied as a function of m_{ZZ} and varies between 1.0 and 1.2 (Fig. 5.16, left).
Besides this QCD K-factor, an additional NLO/LO electroweak correction is applied to the $q\bar{q} \to ZZ$ background as a function of the initial state quark flavour and kinematics. This correction is applied in the region $m_{ZZ} > 2m_{Z^0}$ for which it has been computed, and the average effect is a decrease of the predicted yield by about 2% at 200 GeV, and 14% at 1 TeV, as illustrated in Fig. 5.16 (right).

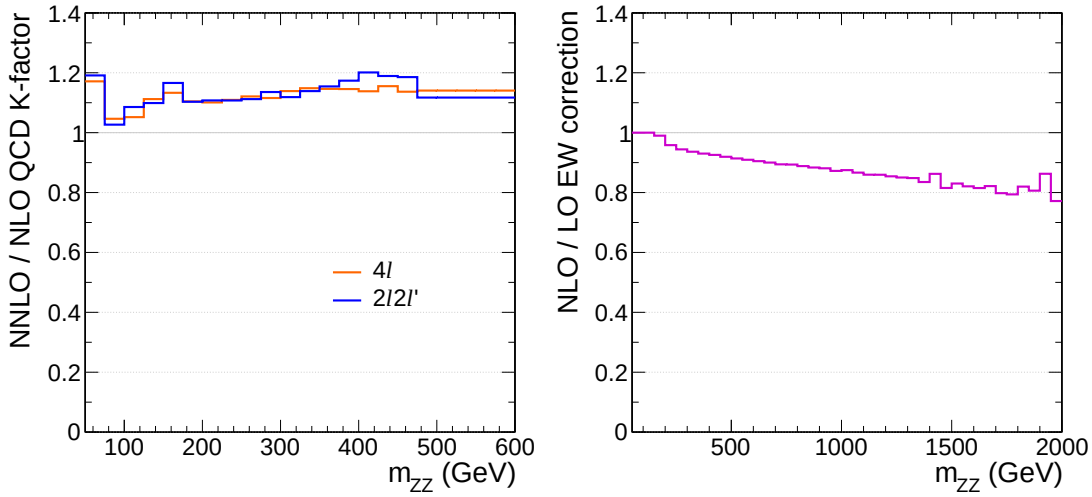

Figure 5.16: (left) NNLO/NLO QCD K-factor for the $q\bar{q} \to ZZ$ background, as a function of generator-level m_{ZZ}, for the 4ℓ and $2\ell 2\ell'$ final states. (right) Average impact of the NLO/LO electroweak correction for the $q\bar{q} \to ZZ$ background on its yields, as a function of generator-level m_{ZZ}.

The $gg \to ZZ$ background is generated at LO with MCFM 7.0 [83]. Although no exact calculation exists beyond the LO, it has been recently shown [88] that the soft collinear approximation is able to describe the background cross section and the interference term at NNLO. Further calculations also show that the K-factors are very similar at NLO for the signal and background [89] terms, and at NNLO for the signal and interference terms [90], suggesting that the same K-factor can be used for the $gg \to H \to ZZ$ signal and for the $gg \to ZZ$ background [91]. This analysis relies on a NNLO/LO K-factor computed for the signal process as a function of m_{ZZ}, using the HNNLO v2 program [92, 93, 94]. The NNLO and LO $gg \to H \to 2\ell 2\ell'$ cross sections are computed for a fixed small H boson decay width of 4.07 MeV across the whole m_{ZZ} range (Fig. 5.17, left), and the K-factor is defined as their ratio. It varies from about 2.0 to 2.6 and is 2.27 at $m_{ZZ} = 125$ GeV, as shown in Fig. 5.17 (right). As intended, it is then applied to the background process ($gg \to ZZ$), in which case an additional systematic uncertainty of 10% on its determination is used in the analysis.

After these corrections are applied, analytical shapes of the $m_{4\ell}$ distribution of the irreducible background are defined, in order to be used as inputs to H(125)-related measurements. Histogram templates of $m_{4\ell}$ in the $105 < m_{4\ell} < 140$ GeV range are fitted to empirical functions, namely Bernstein polynomials of degree 3, separately for $q\bar{q} \to ZZ$

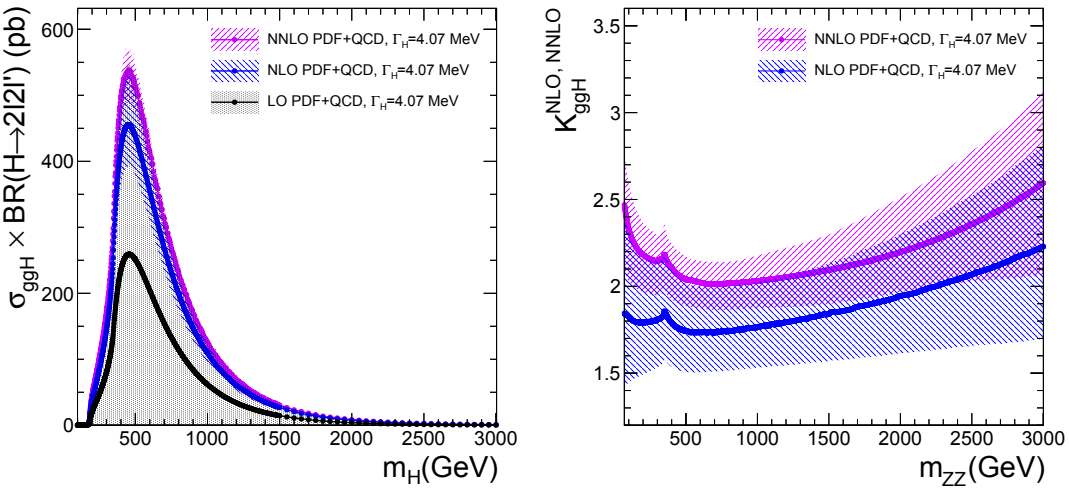

Figure 5.17: (left) $gg \to H \to 2\ell 2\ell'$ cross sections at NNLO, NLO and LO as a function of m_H, at the fixed small decay width of 4.07 MeV. (right) Corresponding NNLO/LO and NLO/LO K-factors. The former is applied to the $gg \to ZZ$ background, as a function of generator-level m_{ZZ}. Uncertainties on the PDF set and QCD scales are displayed as hatched bands.

and for $gg \to ZZ$, and for every final state and event category. The same histogram templates also provide the expected yields to be used as measurement inputs.

5.5.2 Reducible backgrounds

In the $H \to 4\ell$ analysis, reducible backgrounds arise from events where at least one of the four selected leptons is not a genuine prompt lepton. This can be due to secondary leptons from heavy-flavour jets found in processes such as $Z + b\bar{b}$ and $t\bar{t}$ production, or, more generally, to decays of heavy-flavour hadrons, in-flight decays of light mesons within jets, or decays of charged hadrons overlapping with π^0 decays that are misidentified as leptons, which can happen in various processes such as Z + jets, $Z\gamma$ + jets, WW + jets and WZ + jets. In what follows, all reducible backgrounds are collectively denoted as Z+X.

The contribution of the Z+X background to the signal region cannot be reliably estimated using MC simulation. It indeed involves several physics processes which are not simulated in samples large enough to correctly populate the four-lepton signal region, let alone in low-yield event categories. Moreover, it also critically relies on such higher-order effects as the modelling of additional jets, which is not simulated accurately enough yet. As a consequence, the Z+X background is estimated from data-driven methods, using data events collected in *control regions*, i.e. regions of the phase space that are orthogonal to the signal region and more populated than it.

Two independent strategies are used: the *same-sign method* (SS) and the *opposite-sign method* (OS), with control regions also orthogonal to each other. In both cases, these regions rely on so-called $Z + \ell\ell$ candidates, where two of the four leptons satisfy relaxed identification requirements, and the expected yield in the signal region is obtained by reweighting that of the control regions according to lepton *fake rates*, i.e. the fraction of such relaxed leptons that also pass the full selection criteria. The two methods then

differ in the exact definitions of the control regions, and rely on different strategies to handle electrons from photon conversions.

Control region definition

In this section, the selected leptons defined in Section 5.2.2 are referred to as *passing leptons*, and denoted as P. Two types of relaxed leptons are then defined.

- *Loose leptons* (denoted as L) were defined in Section 4.1.1. Compared to selected leptons, they are allowed to fail identification (Section 4.1.2) or isolation (Section 4.1.3) requirements, but can also pass them.
- *Failing leptons* (denoted as F) are required to fail identification or isolation, or both.

In other words, P and F are disjoint sets, the reunion of which is L.

Three control regions are defined by altering the signal region selection described in Section 5.2.2 in the following way.

- The SS method uses the *2P2L$_{SS}$ region*, where $\mathrm{Z}+\ell\ell$ candidates are formed by combining a Z candidate with a pair of *same-sign* loose leptons.
- The OS method relies on the *2P2F* and *3P1F regions*, where $\mathrm{Z}+\ell\ell$ candidates are formed by combining a Z candidate with a pair of *opposite-sign* leptons, where either two (2P2F) or exactly one (3P1F) of said leptons is a failing lepton, respectively.

In these regions, the Z_1 is defined as the normal Z candidate, while the Z_2 is the relaxed one. While FSR recovery is treated the same way as in the signal region, the jet cleaning procedure is extended to also involve the loose and failing leptons of selected $\mathrm{Z}+\ell\ell$ candidates and their associated photons. The 2P2L$_{\mathrm{SS}}$ region uses slight adaptations to kinematic cuts: the QCD suppression cut only applies to the three possible opposite-sign pairs, while the alternative pairing check only involves the pairing where the Z_a is an opposite-sign pair.

To avoid overlaps, events selected in the signal region are not used to define control region candidates, and no more than one such candidate is chosen in each event for each method (SS and OS). Nevertheless, one event can still provide candidates for both the 2P2F and 3P1F regions.

Lepton fake rates

The lepton fake rates f_{e} and f_{μ} are determined from data, separately for the SS and OS method, using a special $\mathrm{Z}+1\ell$ control region.

A Z_1 candidate is first chosen among pairs of same-flavour, opposite-charge selected leptons passing $p_{\mathrm{T}} > 20/10\,\mathrm{GeV}$. For the SS method, a $40 < m_{\ell^+\ell^-} < 120\,\mathrm{GeV}$ requirement is applied, which is tightened to $|m_{\ell^+\ell^-} - m_{\mathrm{Z}^0}| < 7\,\mathrm{GeV}$ for the OS method in order to reduce the contribution of asymmetric photon conversions. The Z_1 candidate is chosen as the pair with invariant mass closest to m_{Z_1}. Moreover, the presence of exactly one additional loose lepton is required. The QCD suppression cut $m_{\ell\ell} > 4\,\mathrm{GeV}$ is applied using the loose lepton and the opposite-sign lepton from the Z_1 candidate. To suppress contamination from WZ and $\mathrm{t\bar{t}}$ processes, the missing transverse energy is required to be less than 25 GeV.

Finally, the additional loose electron or muon is used as a probe to measure the fake rates f_e and f_μ, which are simply defined as the fractions of such leptons that are also passing leptons. The fake rates are measured in bins of p_T^ℓ, separately for the barrel and endcap regions.

Method with same-sign leptons

The SS method makes use of the 2P2L$_{SS}$ region. The expected number of Z+X events in the signal region is computed as:

$$N_{SR}^{Z+X} = r_{OS/SS} \cdot \sum_{2P2L_{SS}} \tilde{f}(\ell_3) \cdot \tilde{f}(\ell_4)\,, \tag{5.13}$$

which uses following definitions:

- ℓ_3 and ℓ_4 are the two loose leptons.
- $\tilde{f}$ is a corrected version of the fake rate that takes the contribution of photon conversions into account. The average fraction r_{miss} of loose electrons with tracks having one missing hit in the pixel detector is used as an estimator of this contribution. The dependency of f on r_{miss} is studied by varying the requirement on $|m_{\ell^+\ell^-} - m_{Z^0}|$ and $|m_{\ell^+\ell^- e_{loose}} - m_{Z^0}|$ in order to obtain samples with different FSR contributions. $f(r_{miss})$ is found to be linear, and the value of r_{miss} measured in the 2P2L$_{SS}$ region is chosen to define $\tilde{f}$.
- $r_{OS/SS}$ is the ratio of the number of events in 2P2L$_{OS}$ and 2P2L$_{SS}$ regions, where 2P2L$_{OS}$ region is defined analogously to 2P2L$_{SS}$ but with opposite-sign loose leptons. It is computed from simulation, separately for every final state.

Method with opposite-sign leptons

The OS method exploits both the 2P2F and 3P1F regions, and proceeds as follows.

- The 2P2F control region is used to estimate backgrounds which intrinsically have only two prompt leptons, such as Z + jets and $t\bar{t}$. Their contribution to the signal region is obtained by weighting each event of the 2P2F region by a factor $\frac{f(\ell_3)}{1-f(\ell_3)}\frac{f(\ell_4)}{1-f(\ell_4)}$, where ℓ_3 and ℓ_4 are the failing leptons.
- The 3P1F control region is both populated by the previous processes and by backgrounds that have three genuine prompt leptons, such as WZ + jets and Zγ + jets where the photon converts to e^+e^-. Their contribution is estimated by weighting each event of the 3P1F region by a factor $\frac{f(\ell_4)}{1-f(\ell_4)}$, where ℓ_4 is the failing lepton.
- The contamination of 2P2F-type processes in the 3P1F region is estimated as $\sum_i(\frac{f(\ell_3)}{1-f(\ell_3)} + \frac{f(\ell_4)}{1-f(\ell_4)})$ and contributes to the final weighted sum of the 3P1F events with the component $\sum_i(2\frac{f(\ell_3)}{1-f(\ell_3)}\frac{f(\ell_4)}{1-f(\ell_4)})$, which is subtracted from the Z+X estimate.
- The contamination from true ZZ events in the 3P1F region N_{3P1F}^{ZZ} is estimated from simulation, and the expected yield in the signal region is reduced by a factor $1 - N_{3P1F}^{ZZ}/N_{3P1F}$, where N_{3P1F} is the number of events in the 3P1F region.

Putting these ingredients together, the expected number of Z+X events in the signal region can be written as

$$N_{\mathrm{SR}}^{\mathrm{Z+X}} = \left(1 - \frac{N_{\mathrm{3P1F}}^{\mathrm{ZZ}}}{N_{\mathrm{3P1F}}}\right) \sum_{\mathrm{3P1F}} \frac{f(\ell_4)}{1-f(\ell_4)} \; - \sum_{\mathrm{2P2F}} \frac{f(\ell_3)}{1-f(\ell_3)} \frac{f(\ell_4)}{1-f(\ell_4)} \,. \tag{5.14}$$

Uncertainties

The predicted yield $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ in the signal region is affected by the following sources of uncertainty:

- A statistical component arises from the limited size of the data sample in the control regions, and in the regions where fake rates are computed. It is typically in the range 1–10%.
- An important systematic uncertainty comes from the fact that the composition in background processes (Drell-Yan + jets, $\mathrm{t\bar{t}}$, WZ, $\mathrm{Z}\gamma^{(*)}$) of the $\mathrm{Z}+1\ell$ region used to compute the fake rates typically differs from that of control regions where they are applied. To estimate this uncertainty, per-process fake rates are first measured in the $\mathrm{Z}+1\ell$ region in simulation for individual background processes. Two overall fake rates can then be computed as weighted averages, according to the background composition of either the $\mathrm{Z}+1\ell$ or the 2P2F region as obtained from simulation. The difference between the two results defines the uncertainty, which affects $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ at the level of 30–40%. It is computed for the OS method, but it is also applied to predictions of the SS method.
- For the SS method, an additional component arises from the statistical uncertainty of the measurement of r_{miss} in the $\mathrm{2P2L_{SS}}$ region, which propagates to the electron fake rates and to $N_{\mathrm{SR}}^{\mathrm{Z+X}}$.

These three contributions are eventually added in quadrature to define the total uncertainty.

Combination, categories, and shape

Each of the two methods provides an independent prediction $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ for the Z+X yield in the signal region. Since the fake rates and the composition of control regions depends on the final state, the prediction is performed separately for each final state: 4e, 4μ, $2\mathrm{e}2\mu$ and $2\mu2\mathrm{e}$.[1] The final inclusive estimate is obtained by weighting the individual mean values of both methods according to their corresponding variances. The uncertainty range on this combined estimate is the envelope that covers the uncertainty ranges of both methods.

To predict the Z+X yield in every event category, only the SS method is used, and the per-category $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ numbers are then slightly rescaled according to ratio of combined and SS-based inclusive numbers, in each of the four final states. Since event categorization relies on associated objects (jets and additional selected leptons), the definitions of categories of the signal region are naturally extended to categories of the $\mathrm{2P2L_{SS}}$ region, with the $\mathrm{Z}+\ell\ell$ candidate now playing the role of 4ℓ candidate. However, these control region categories have two slight differences. First, the jet cleaning procedure introduced

[1] For the sake of Z+X background estimation, the usual $2\mathrm{e}2\mu$ final state is split into two separate final states $2\mathrm{e}2\mu$ and $2\mu2\mathrm{e}$, depending whether the Z_1 candidate is made of electrons or muons, respectively.

in Section 4.5.1 is now also performed with respect to loose leptons, so that there is no overlap with any lepton of the $Z + \ell\ell$ candidate. Second, since no event of the signal region can be used for control regions, events containing at least two additional selected leptons that can form a ZZ candidate with two leptons of the best ZZ candidate do not contribute to the $2P2L_{SS}$ region. This point conceptually deteriorates the Z+X estimation in the VH-leptonic-tagged or $t\bar{t}H$-tagged categories, but it is irrelevant in practice because events with ≥ 6 leptons are a small part of these categories, and the uncertainties are large with the available data. It will nevertheless need to be improved as the collected integrated luminosity increases.

Besides the yield predictions, distributions of observables for the Z+X background are determined from the corresponding distributions of $Z + \ell\ell$ related observables in the control regions. In particular, an analytical parameterization of the $m_{4\ell}$ shape of Z+X needs is done: it is obtained by combining the inclusive OS and SS predictions in bins of $m_{4\ell}$, and fitting the resulting distributions with Landau functions [95], separately for every final state. At this point, a shape uncertainty can be defined from the difference between $m_{4\ell}$ distributions from both methods, but since this difference varies slowly with $m_{4\ell}$, the uncertainty is absorbed in the already conservative combined uncertainty on predicted yields. As regards categories, it was shown that the shapes obtained from the inclusive sample describe the exclusive distributions of the Untagged and VBF-1jet-tagged category well. Other categories do not contain enough data to check this, but the same shapes are used in them.

The numerical results of this entire procedure and the resulting shapes will be presented in Chapters 6 and 7 for the 2015 and 2016 analyses, respectively.

5.6 Signal modelling

5.6.1 Low-mass signal model

The $m_{4\ell}$ histogram templates that are obtained from simulation are not used directly in the measurements: analytical shapes are used instead, which has the virtue of smoothing out the irregularities due to the finite number of simulated events. Moreover, in order to measure m_H, or to measure another physics parameter for a certain value of m_H, a continuous parameterization of signal predictions as a function of m_H first needs to be done, based on the few m_H values for which a MC sample has been simulated (see Section 5.1). For measurements related to the on-shell H(125), this analysis exploits the $105 < m_{4\ell} < 140\,\text{GeV}$ window, and the parameterization is built for the $118 < m_H < 130\,\text{GeV}$ range using five points of m_H: 120, 124, 125, 126 and 130 GeV.

First, the *shape* of the signal pdf for $m_{4\ell}$ has to be described by an analytical function. For such low values of m_H, the narrow-width resonance hypothesis holds, which means that the theoretical signal lineshape can be described as a relativistic Breit-Wigner function. It then has to be convolved with an empirical function that accounts for the effects of experimental resolution, such as tails due to bremsstrahlung, final-state radiation, and energy leakage in the ECAL. A double-sided Crystal Ball (dCB) function is usually chosen, which both accounts for the Gaussian resolution of the core of the $m_{4\ell}$ distribution, and for the two asymmetric non-Gaussian tails described by power laws. In practice, given the narrowness of the theoretical lineshape, the signal pdf was

found to be well described by a dCB function alone:

$$\mathcal{P}(\,m_{4\ell}\,|\,m_{\mathrm{H}}\,) \;=\; f_{\mathrm{dCB}}(\,m_{4\ell}\,|\,m_{\mathrm{H}}\,) \;=\; N\cdot\begin{cases} A\cdot(B+|\xi|)^{-n_L}\,, & \text{for } \xi<\alpha_L \\ \exp\left(-\xi^2/2\right), & \text{for } \alpha_L\leq\xi\leq\alpha_R \\ A\cdot(B+|\xi|)^{-n_R}\,, & \text{for } \xi>\alpha_R \end{cases} \tag{5.15}$$

where $\xi = (m_{4\ell} - m_{\mathrm{H}} - \Delta m_{\mathrm{H}}) \,/\, \sigma_m$. This function has six independent parameters: the variance σ_m of the Gaussian core, the systematic mass shift Δm_{H} of the peak, and two parameters for each tail: the powers n_L and n_R control their prominence, while α_L and α_R define the position of the boundary between the core and the tails. The values of A and B ensue from requiring the continuity of the function and its first derivative, while the normalizing constant N is determined later.

The dependency of each of the six parameters p_i on m_{H} has to be determined for every final state and event category. To this end, a linear approximation is used:

$$p_i(m_{\mathrm{H}}) = C_i + D_i \times (m_{\mathrm{H}} - 125\,\mathrm{GeV})\;. \tag{5.16}$$

For every final state, the C_i parameters are first obtained from the shapes in the ggH production mode in the Untagged category for $m_{\mathrm{H}} = 125\,\mathrm{GeV}$. This is illustrated in Fig. 5.18 for the three final states, showing that resolution is better for muons than for electrons. Then, the D_i parameters are determined from a simultaneous fit to all other m_{H} points. Studies show that these pdf from ggH distributions are also valid for other categories and for VBF.

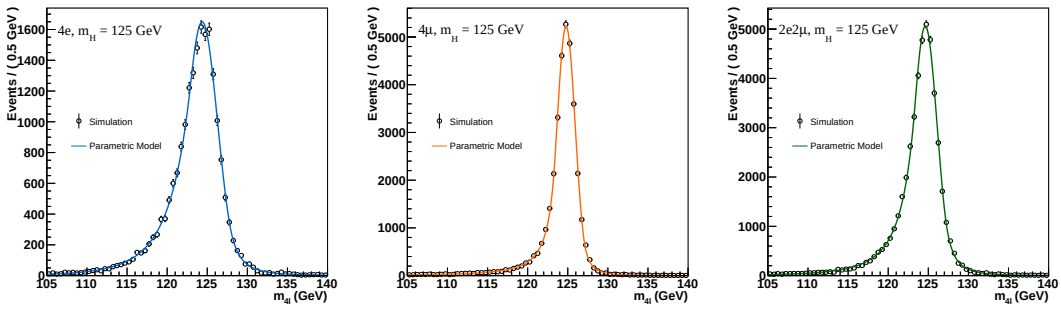

Figure 5.18: MC simulated $m_{4\ell}$ distributions after full event selection, and fitted pdf $f_{\mathrm{dCB}}(m_{4\ell}|m_{\mathrm{H}})$, for the ggH process generated with $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, in the Untagged Category. The three final states are compared: 4e (left), 4μ (centre), and $2e2\mu$ (right).

For the WH, ZH, and $\mathrm{t\bar{t}H}$ processes, the Higgs boson peak is accompanied by a non-resonant component from events where the four leptons are misassigned (as mentioned in Section 5.2.3). A Landau function is thus added to the dCB to perform the fit in the $m_{\mathrm{H}} = 125\,\mathrm{GeV}$ case, adding two more parameters, which are then fixed for the simultaneous fit. The relative normalization of both components is finally adjusted for every category. An example is given for the ZH process in Fig. 5.19.

Finally, the dependency of the expected signal *yields* on m_{H} is also estimated from simulation, independently of the shape, in the same $105 < m_{4\ell} < 140\,\mathrm{GeV}$ window. For every combination of production modes, final states and event categories, a fit to a second order polynomial is performed using the five m_{H} points from 120 to 130 GeV.

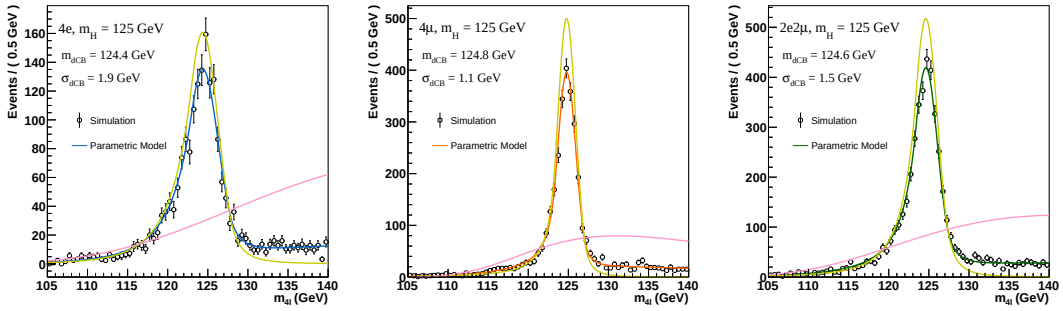

Figure 5.19: MC simulated $m_{4\ell}$ distributions after full event selection, and fitted pdf $f_{\mathrm{dCB}}(m_{4\ell}|m_{\mathrm{H}})$, for the ZH process generated with $m_{\mathrm{H}} = 125\,\mathrm{GeV}$, in the Untagged Category. The pre-fit pdfs for the resonant and non-resonant component are shown as yellow and pink curves, respectively. The three final states are compared: 4e (left), 4μ (centre), and $2\mathrm{e}2\mu$ (right).

However, when the triplet (process, final state, category) has a low number of events in the MC simulation, a first-order polynomial is used instead, to get a more robust result. Examples of such fits are presented in Fig. 5.20 for ggH in the Untagged category and ZH in the VBF-2jet-tagged category.

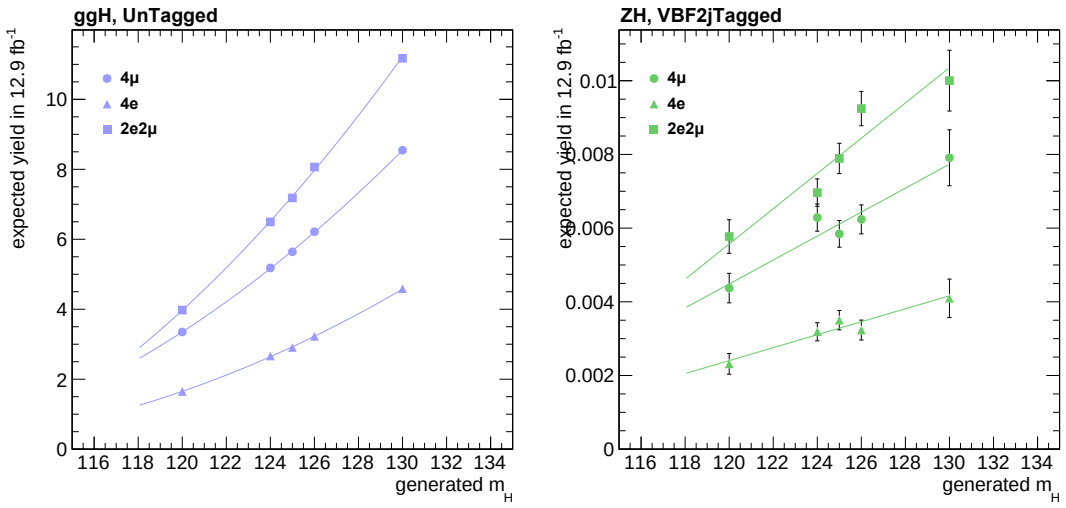

Figure 5.20: Fits of the m_{H} dependency of the expected signal yields for $12.9\,\mathrm{fb}^{-1}$ in the $105 < m_{4\ell} < 140\,\mathrm{GeV}$ window after full event selection, shown for the $118 < m_{\mathrm{H}} < 130\,\mathrm{GeV}$ interval. Two examples are shown: ggH in the Untagged category (left) to illustrate a case with a large number of simulated events, and ZH in the VBF-2jet-tagged category (right) for a less populated case.

5.6.2 High-mass signal model

Signal parameterization at masses beyond the 125 GeV peak is relevant to two different studies. First, studying the off-shell tail of the H(125) boson allows the measurement of its decay width Γ_{H}, which is sensitive to the ratio of on-shell and off-shell cross sections,

as will be discussed in Section 7.5.4. Second, a search is performed for an additional resonance X of spin-parity $J^P = 0^+$ decaying to four leptons, with mass m_X between 130 GeV and 2.5 TeV, arbitrary width Γ_X and cross section σ_X, a study which is reported in Section 7.5.5. Both these studies require to not only parameterize the $H(m_H, \Gamma_H)$ and $X(m_X, \Gamma_X, \sigma_X)$ signals, but also to carefully account for the interference between the signals and backgrounds in the $gg \to 4\ell$ and $qq' \to 4\ell$ processes, and between both types of signal.

Given the many parameters at play and the interference, this situation cannot be handled like the low-mass H(125) study. Here, the description of the signal model is split into two parts: the theoretical modelling of the generator-level four-lepton invariant mass $m_{4\ell}^{\text{gen.}}$, and the discussion of detector effects leading to the reconstructed-level $m_{4\ell}$. The signal pdf is then built as the convolution of the theoretical lineshape $\mathcal{P}_{\text{theo.}}$ and the experimental resolution function $\mathcal{R}$, multiplied by the experimental efficiency function $\mathcal{E}$:

$$\mathcal{P}(m_{4\ell}) = \mathcal{E}(m_{4\ell}) \times \left[\mathcal{R}(m_{4\ell} \,|\, m_{4\ell}^{\text{gen.}}) \otimes \mathcal{P}_{\text{theo.}}(m_{4\ell}^{\text{gen.}} \,|\, m_H, \Gamma_H, m_X, \Gamma_X, \sigma_X) \right] . \quad (5.17)$$

To model the theoretical pdf $\mathcal{P}_{\text{theo.}}(m_{4\ell}^{\text{gen.}} \,|\, m_H, \Gamma_H, m_X, \Gamma_X, \sigma_X)$, a general parameterization of the $gg/qq' \to \text{bkgd} + H(125)^* + X(m_X) \to 4\ell$ process is used. Gluon fusion, VBF and associated VH/VX productions are considered, while $t\bar{t}H$ is neglected because it is highly suppressed above the $2m_Z$ threshold. All contributing amplitudes for 4ℓ production from ZZ backgrounds and H(125) and $X(m_X)$ resonances, as well as their interference, are included. The mass and width of both the H and X resonances are included as general parameters of the model, and the relative contribution from ggH and VBF + VH is floated. This novel approach is supported by the MC simulation from the MCFM + JHUGen framework, including NNLO corrections to the cross section based on MCFM using the HNNLO program [92, 93, 94] (consistent with the K-factor discussed in Section 5.5.1), and the implementation is done within the MELA package [80, 81, 86]. Such developments were part of the LHC Higgs Cross Section Group Yellow Report 4 preparation [10]; an example is shown in Fig. 5.21, showing the various contributions to $\mathcal{P}_{\text{theo.}}$ for $m_X = 450$ GeV and $\Gamma_X = 46.8$ GeV. The interference of X with the backgrounds and off-shell tail of the H(125) boson is an important effect when Γ_X is large, and the parameterization naturally scales this interference with σ_X.

The parameterization of the resolution function $\mathcal{R}(m_{4\ell} \,|\, m_{4\ell}^{\text{gen.}})$ is extracted separately for each final state (4e, 4μ, 2e2μ), from the full simulation of the $gg \to X \to 4\ell$ process based on POWHEG + JHUGen, for a wide width $\Gamma_X = \Gamma_H^{\text{SM}}(m_H)$. Samples for 54 mass points from 120 GeV to 3 TeV are used, and the distribution of $m_{4\ell} - m_{4\ell}^{\text{gen.}}$ is studied in a narrow window around m_X. $\mathcal{R}$ is modelled with a dCB function, and the mass dependency of its six parameters is parameterized assembling polynomial functions that cover several mass windows. Other mechanisms than gluon fusion are found to have a similar resolution and are assigned the same parameterization.

The multiplicative efficiency function $\mathcal{E}(m_{4\ell})$ accounts both for detector acceptance and the efficiency of event reconstruction and selection. Like $\mathcal{R}$, it is computed separately for each final state, for a similar set of mass points of the POWHEG + JHUGen full simulation, but gluon fusion and VBF have to be treated separately. The efficiency parameterization is shown in Fig 5.22 for these two mechanisms in the three final states.

Several closure tests have showed that the $m_{4\ell}$ distributions obtained from full simulation are in good agreement with the signal model of Eq. 5.17, both for the modelling

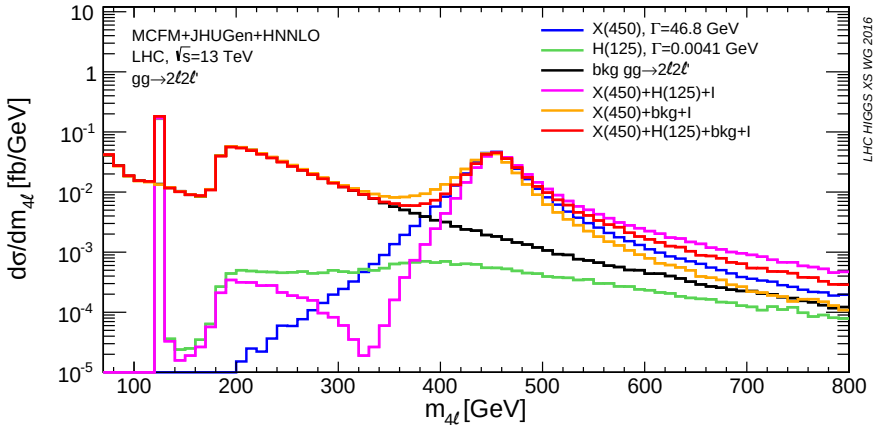

Figure 5.21: Differential cross section of the process $gg \to ZZ/Z\gamma^*/\gamma^*\gamma^* \to 2\ell 2\ell'$ (where $\ell, \ell' = e, \mu$, or τ) as a function of $m_{4\ell}$, generated with the MCFM+JHUGEN framework including the NNLO corrections calculated with MCFM + HNNLO. The plot shows distributions in the presence of a hypothetical X(450) resonance, for several combinations of signal and background components, including the interference I of contributing amplitudes [10].

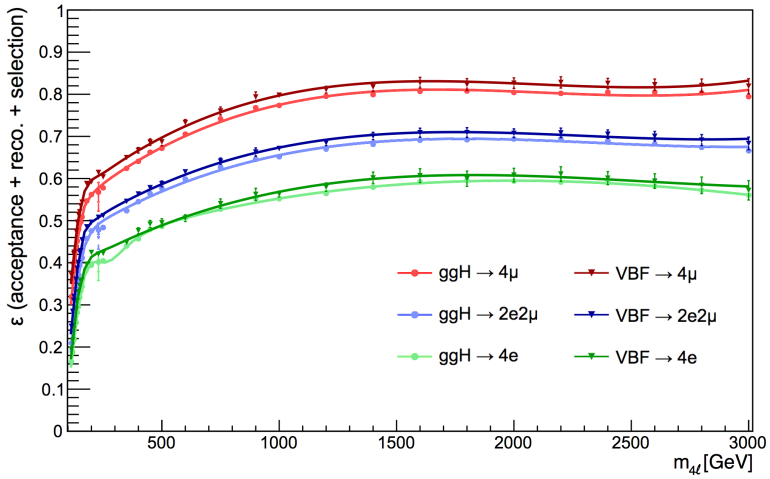

Figure 5.22: Parameterization of the experimental efficiency function $\mathcal{E}(m_{4\ell})$ for a SM-like signal with $J^P = 0^+$ produced via gluon fusion (solid circles) or vector boson fusion (solid triangles), for the three final states. The markers show the results from full simulation, while the curves show the analytical fits.

the off-shell tail of the H(125) boson for the width measurement, or for the arbitrary $X(m_X, \Gamma_X)$ resonance.

5.7 Measurement strategies

The $H \to 4\ell$ channel is suited to a variety of measurements. When studying the previously discovered H(125) boson, one can first estimate the significance of the observed signal, and then measure the values (central values plus uncertainty intervals) of some

parameters of interest, such as the signal strength, couplings, and mass (using the on-shell region only), and the decay width (exploiting both the on-shell and high-mass off-shell region). By contrast, in the high-mass search for a new resonance decaying to four leptons, unless a sizeable excess is observed, the results take the form of exclusion limits on the cross section of the hypothesized process.

This section describes the statistical methods, fit models and systematic uncertainties that are used to derive all these results from the inputs defined in the previous sections.

5.7.1 Framework for statistical interpretation

To quantify or reject the presence of a signal in data and to estimate parameters, a well-defined common statistical procedure is used at the LHC. This methodology was developed in early Run I by the CMS and ATLAS collaborations in the context of the LHC Higgs Combination Group [96], in order to coordinate searches for the SM Higgs boson [16].

General likelihood function

The most widespread model relies on one single *signal strength* parameter μ that multiplies the expected SM Higgs boson cross section such that $\sigma = \mu \cdot \sigma_{\mathrm{SM}}$. All production modes are scaled by this same factor, while branching fractions are assumed to be preserved. Each independent source of systematic uncertainty is assigned a *nuisance parameter* θ_i, the full set of which is denoted as θ. The expected background and SM Higgs boson signal yields are functions of these parameters, and can be written as $b(\theta)$ and $\mu \cdot s(\theta)$, respectively. Systematic uncertainties usually reflect the possible deviations of a quantity from the input value $\tilde{\theta}_i$ provided by a separate measurement. On that basis, the *likelihood* can be defined, given the data (either actual data or a pseudo-experiment) and the measurements of $\tilde{\theta}$, as:

$$\mathcal{L}(\mathrm{data}, \tilde{\theta}\,|\,\mu, \theta) \equiv \prod_c \mathcal{L}_c(\,\mathrm{data}\,|\,\mu \cdot s(\theta) + b(\theta)\,) \cdot \prod_i p_i(\tilde{\theta}_i|\theta_i)\ . \tag{5.18}$$

The first product in Eq. 5.18 runs over all channels considered in the analysis. The $\mathcal{L}_c$ functions generally stand for products of Poisson probabilities for the number of events in every bin of their channel. However, the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis uses data in an unbinned way, by exploiting the individual per-event values of a set of observables $\mathcal{O}$. In this case, the likelihood is expressed as the product over all N observed events of the pdfs of $\mathcal{O}$ for signal and background events ($f_s(\mathcal{O}|\theta)$ and $f_b(\mathcal{O}|\theta)$), weighted by the total expected signal and background rates $S(\theta)$ and $B(\theta)$:

$$\mathrm{Poisson}(\,\mathrm{data}\,|\,\mu \cdot s + b\,) = N^{-1} \prod_{e=1}^{N} (\mu S f_s(\mathcal{O}_e) + B f_b(\mathcal{O}_e)) \cdot e^{-(\mu S + B)}\ , \tag{5.19}$$

where dependencies on θ are implied. The set of channels c, the choice of relevant observables $\mathcal{O}$ and the construction of the pdfs f_s and f_b will be described in Section 5.7.2.

The factors $p_i(\tilde{\theta}_i|\theta_i)$ in Eq. 5.18 represent pdfs for the measurements associated to each nuisance parameter, i.e. the probabilities of measuring a value $\tilde{\theta}_i$ given the true value θ_i. The different sources of systematic uncertainties and their treatment will be described in Section 5.7.3.

Quantifying an excess

Whether one tries to re-observe the Higgs boson discovered in Run I or to find another resonance, one needs an objective way of testing the presence or absence of a signal. This requires to build a quantity called the *test statistic*, which encodes the information about the observed data, expected signal and background, and uncertainties into a single variable. Using the pdfs Poisson(data $|\, \mu \cdot s + b$) and $p(\tilde{\theta}|\theta)$, with values of θ that maximize $\mathcal{L}$, one can generate pseudo-data sets in order to build expected distributions of the test statistic for the background-only or signal+background hypotheses. Comparing them to its value in data then helps ruling on the presence or absence of a signal.

When trying to quantify the statistical significance of an excess over the background-only expectation, the following test statistic is used:

$$q_0 \equiv -2 \ln \frac{\mathcal{L}(\text{data}, \tilde{\theta} \,|\, 0, \hat{\theta}_0)}{\mathcal{L}(\text{data}, \tilde{\theta} \,|\, \hat{\mu}, \hat{\theta})} \,, \quad \text{with the constraint } \; \hat{\mu} \geq 0 \,, \tag{5.20}$$

where $\hat{\mu}$ and $\hat{\theta}$ maximize the likelihood in the denominator, and $\hat{\theta}_0$ maximizes the numerator under the background-only hypothesis ($\mu = 0$). By definition, q_0 is positive in the presence of a signal-like excess and zero otherwise.

The significance of the excess is quantified by the local p-value p_0, defined as the probability for q_0 to be at least as large as the value q_0^{obs} observed in data, under the background-only hypothesis:

$$p_0 \equiv \mathbb{P}(\, q_0 \geq q_0^{\text{obs}} \,|\, b \,) \,, \tag{5.21}$$

In other words, p_0 characterizes the probability for a local background fluctuation to resemble the hypothesized signal at least as much as the data does. It is independent of the normalization of the expected signal, but information on the magnitude of the excess relative to the expectation is instead provided by the best-fit signal strength $\hat{\mu}$.

The p-value is usually converted into a significance Z via the Gaussian one-sided tail integral:

$$p_0 = \int_Z^{+\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-x^2/2\right) \, \mathrm{d}x \,, \tag{5.22}$$

whereby the conventional 5σ threshold for a discovery corresponds to $p_0 = 2.8 \times 10^{-7}$.

In practice, the results on significance will be presented as a scan of the observed local p-value throughout the probed signal parameter space, here an interval of the Higgs boson mass m_{H}. This is usually compared to the scan of mean expected local p-values at m_{H} for a SM Higgs boson with the corresponding mass m_{H}. It should be noted that the probability for a statistical fluctuation to arise anywhere in this parameter space is larger than the probability of observing one at a certain point. This so-called *look-elsewhere effect* is usually corrected for by computing a *global* p-value, which will not be discussed here.

Setting exclusion limits

When no significant excess is observed, one usually sets upper limits on the considered signal hypothesis. To this end, another test statistic is used, this time using the

signal+background model in the numerator:

$$q_\mu \equiv -2\ln\frac{\mathcal{L}(\text{data},\tilde{\theta}\,|\,\mu,\hat{\theta}_\mu)}{\mathcal{L}(\text{data},\tilde{\theta}\,|\,\hat{\mu},\hat{\theta})}\,, \quad \text{with the constraint } \; 0\leq\hat{\mu}\leq\mu\;, \tag{5.23}$$

where $\hat{\theta}_\mu$ maximizes the numerator under the hypothesis of a signal strength μ. The constraint $\hat{\mu}\geq 0$ prevents an unphysical negative signal strength, while the constraint $\hat{\mu}\leq\mu$ guarantees a one-sided confidence interval: this way, an excess larger than expected for a signal with strength μ is not counted as evidence against this hypothesis. Unlike earlier test statistics that were used at LEP or the Tevatron, the definition of Eq. 5.23 has the advantage that the expected distributions of q_μ under the signal+background and background-only hypotheses are known analytically in the asymptotic limit of a large number of background events [97].

Exclusion limits are then computed with the modified frequentist construction $\mathrm{CL_s}$ [98, 99, 100, 96]. The probabilities for q_μ to be at least as large as the value q_μ^{obs} observed in data, under the signal+background and background-only hypotheses, are respectively defined as:

$$\begin{aligned}\mathrm{CL_{s+b}}\,(\mu) &\equiv \mathbb{P}(\,q_\mu\geq q_\mu^{\text{obs}}\,|\,\mu\cdot s+b\,)\;,\\ \mathrm{CL_b}\,(\mu) &\equiv \mathbb{P}(\,q_\mu\geq q_\mu^{\text{obs}}\,|\,b\,)\;,\end{aligned} \tag{5.24}$$

and the $\mathrm{CL_s}$ quantity is defined as their ratio:

$$\mathrm{CL_s}\,(\mu)\equiv\frac{\mathrm{CL_{s+b}}\,(\mu)}{\mathrm{CL_b}\,(\mu)}\,. \tag{5.25}$$

The signal is excluded at a 95% confidence level (CL) if $\mathrm{CL_s}\,(1)<0.05$, which is a more conservative condition than if $\mathrm{CL_{s+b}}$ were used, because it protects against downward fluctuations of the background. To quote an upper limit on μ at the 95% CL, the value of μ is adjusted until $\mathrm{CL_s}\,(\mu)=0.05$.

While the μ model is mainly used for the SM-like Higgs boson, the exclusion limits of Section 7.5.5 for a new resonance decaying to four leptons will be reported as a function of its production cross section σ_{X}, which is a similar model.

Measuring parameters

As soon as a Higgs boson was discovered near 125 GeV, part of the focus has shifted from searches to measuring its properties. Some of these measurements, e.g. spin-parity tests, consisted in discriminating between two specific models, and are not considered again in this early Run II analysis. Only measurements of continuous physics parameters are performed here, such as the signal strength, couplings, mass, and width. Two parameters or more can be measured simultaneously.

Denoting the vector of such parameters as π, a likelihood function $\mathcal{L}(\text{data},\tilde{\theta}\,|\,\pi,\theta)$ is defined. This generalizes the likelihood of Eq. 5.18 to another signal model, the details of which depends on the measurement and will be expanded on in later sections. The strategy then consists in maximizing $\mathcal{L}$ with respect to π and θ, to obtain the best fit values $\hat{\pi}$ for the parameters of interest, and $\hat{\theta}$ for the nuisance parameters.

The *negative log-likelihood function* is then defined as:

$$-\,2\Delta\mathrm{lnL}\equiv -2\ln\frac{\mathcal{L}(\text{data},\tilde{\theta}\,|\,\pi,\hat{\theta}_\pi)}{\mathcal{L}(\text{data},\tilde{\theta}\,|\,\hat{\pi},\hat{\theta})}\,, \tag{5.26}$$

where $\hat{\theta}_\pi$ maximizes the numerator for a fixed set of values π of the parameters of interest. In this likelihood-based method (which is one of several possible statistical methodologies), the determination of confidence intervals relies on the Wilks theorem [101], which states that for a model with n parameters of interest, the distribution of $-2\Delta\text{lnL}$ approaches a χ^2 with n degrees of freedom in the limit of a large data sample. For example, when measuring one single parameter, the 68% and 95% CL intervals are deduced from the conditions $-2\Delta\text{lnL} \leq 1$ and ≤ 3.84, respectively. When measuring one parameter in a model that has more than one, the method can be applied either fixing the other ones to some nominal value, or profiling them, i.e. leaving them floating in the fit.

In general, the results are both quoted as central values with 68% CL intervals, and displayed graphically as scans of $-2\Delta\text{lnL}$. Expected results can also be provided for some nominal values of the parameters. This would in principle require to generate a large number of pseudo-experiments and determine their median outcome, but a very good approximation is provided by the Asimov data set [97], i.e. one single representative data set in which the observed rates and distributions coincide with predictions under the nominal set of nuisance parameters.

5.7.2 Multi-dimensional pdfs

As mentioned when discussing Eq. 5.18, the $\text{H}\to\text{ZZ}\to 4\ell$ analysis is an unbinned shape analysis, whereby the likelihoods are constructed from the pdfs of some well-chosen observables, in a number of distinct channels. The number and nature of the observables depends on the measurement, but there cannot be more than three at a time, for computational and statistical reasons. The one-, two-, and three-dimensional (1D, 2D, 3D) pdfs need to be constructed for the five considered signal processes (ggH, VBF, WH, ZH and $\text{t}\bar{\text{t}}\text{H}$) and the three backgrounds ($\text{q}\bar{\text{q}}\to\text{ZZ}$, $\text{gg}\to\text{ZZ}$ and Z+X).

For every considered measurement, Table 5.5 reports the set of observables, the name of the associated pdf, and the set of channels used in the fit. These choices are discussed in what follows.

Table 5.5: Observables, pdfs and channels used in the physics measurements of the $\text{H}\to\text{ZZ}\to 4\ell$ analysis.

measurement	pdf and observables	channels
signal significance	$f^{\mu}_{\text{1D}}(m_{4\ell})$ $f^{\mu}_{\text{2D}}(m_{4\ell}, \mathcal{D}^{\text{kin}}_{\text{bkg}})$	2 categories × 3 final states (2015) or 6 categories × 3 final states (2016)
signal strengths	$f^{\mu}_{\text{2D}}(m_{4\ell}, \mathcal{D}^{\text{kin}}_{\text{bkg}})$	
mass	$f^{m}_{\text{1D}}(m_{4\ell})$ $f^{m}_{\text{2D}}(m_{4\ell}, \mathcal{D}_{\text{m}})$ $f^{m}_{\text{3D}}(m_{4\ell}, \mathcal{D}_{\text{m}}, \mathcal{D}^{\text{kin}}_{\text{bkg}})$	3 final states
width resonance searches	$f^{\text{h.m.}}_{\text{1D}}(m_{4\ell})$	2 categories × 3 final states

Significance and signal strengths

The simplest way of assessing the significance of the excess of events caused by the Higgs boson signal is with a one-dimensional fit to the four-lepton invariant mass distribution. The likelihood function then just relies on the 1D mass pdf that has been discussed in Sections 5.5 and 5.6.1 for the backgrounds and the signal:

$$f^{\mu}_{\mathrm{1D}}(m_{4\ell}) = \mathcal{P}(m_{4\ell}) \ . \tag{5.27}$$

This model can be greatly improved by exploiting the second decay-related discriminating variable, namely the matrix-element-based discriminant $\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$, as a second dimension:

$$f^{\mu}_{\mathrm{2D}}(m_{4\ell}, \mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}) = \mathcal{P}(m_{4\ell}) \times \mathcal{P}(\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}|m_{4\ell}) \ , \tag{5.28}$$

a pdf which is also used for signal strength measurements. The second factor controls the shape of $\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$ and its correlation with $m_{4\ell}$, and is based on 2D histogram templates, to which a smoothing procedure is applied. Their conditional part is implemented by normalizing each bin of the $m_{4\ell}$ dimension to one. As a result, $\mathcal{P}(\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}|m_{4\ell})$ only provides information on the kinematic discriminant given a certain mass, but no information on the mass itself. The 2D templates are built separately for the 3 final states (4μ, 4e, $2\mathrm{e}2\mu$), from simulation for the signal, $\mathrm{q\overline{q}}\to\mathrm{ZZ}$ and $\mathrm{gg}\to\mathrm{ZZ}$ components and from the control region for Z+X, using the SS method.

Eventually, based on the event categories defined in Section 5.4 and on the three final states, the f^{μ} pdfs are built for $2 \times 3 = 6$ channels in the 2015 analysis, and $6 \times 3 = 18$ channels in the 2016 analysis.

Mass measurement

Besides $m_{4\ell}$ and $\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$, the measurement of the Higgs boson mass also exploits the additional information from per-event mass uncertainties described in Section 5.3.3. A 3D pdf is thus built, together with 1D and 2D versions:

$$\begin{array}{lcl} f^{m}_{\mathrm{1D}}(m_{4\ell}) & = & \mathcal{P}(m_{4\ell}) \ , \\ f^{m}_{\mathrm{2D}}(m_{4\ell}, \mathcal{D}_{\mathrm{m}}) & = & \mathcal{P}(m_{4\ell}) \times \mathcal{P}(\mathcal{D}_{\mathrm{m}}) \ , \\ f^{m}_{\mathrm{3D}}(m_{4\ell}, \mathcal{D}_{\mathrm{m}}, \mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}) & = & \mathcal{P}(m_{4\ell}) \times \mathcal{P}(\mathcal{D}_{\mathrm{m}}) \times \mathcal{P}(\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}|m_{4\ell}) \ . \end{array} \tag{5.29}$$

The templates for the $\mathcal{D}_{\mathrm{m}}$ variable are built for each final state, from simulation for the signal and irreducible backgrounds, and from the control region for Z+X. While they should in principle be conditional upon the other variables, $\mathcal{D}_{\mathrm{m}}$ is actually not correlated to $\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$, and it is found to have little dependency on $m_{4\ell}$. Therefore, one-dimensional $\mathcal{P}(\mathcal{D}_{\mathrm{m}})$ templates are used, which also helps dealing with the limited size of samples. Their shapes are fitted with the sum of a Landau [95] and a Gaussian function.

No event categorization is used for the mass measurement, since the expected contribution from associated production modes (which have a different $m_{4\ell}$ lineshape) is still very small. There are thus 3 channels, one for each final state.

Off-shell-related measurements

For measurements that use the high-mass region, i.e. the measurement of the decay width of the H(125) boson and the search for a new resonance, only information on the four-lepton invariant mass is used for the moment. The 1D pdf $f^{\mathrm{h.m.}}_{\mathrm{1D}}(m_{4\ell})$ is based on the signal model described in Section 5.6.2.

The event categorization is also simplified here: since the contribution of production mechanisms other than VBF and ggH is expected to be negligible at high mass, all categories other than the VBF-2jet-tagged category are combined into one broad Untagged category. As always, the 3 final states are considered separately, for a total of $2 \times 3 = 6$ channels.

5.7.3 Systematic uncertainties

Systematic uncertainties here refer to all sources of uncertainties that can make the prediction of signal and background yields and shapes imperfect, and thus bias their comparison to the data. This includes theoretical uncertainties, related to cross sections and generators, and experimental uncertainties, either related to the detector response (efficiencies, calibration, etc.) or to the data control regions. In the statistical model (Section 5.7.1), each independent source of uncertainty is assigned a nuisance parameter θ_i, which is profiled during the minimization. It can affect either one or multiple processes, and in different ways depending on the final state or the event category. The exact way and magnitude have to be determined and quantified from an auxiliary study.

Two classes of uncertainties are considered here. First, *normalization uncertainties* only affect the expected yields of some processes. Their associated nuisance parameters simply scale the yield in a multiplicative way, and their pdf $p_i(\tilde{\theta}_i|\theta_i)$ follow a log-normal distribution. Second, *shape uncertainties* alter the shapes of the pdfs of some of the considered observables. This is modelled by considering a family of alternative templates, which is usually governed by one single nuisance parameter that itself follows a Gaussian pdf.

The following sources of uncertainty are taken into account in this $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ analysis. All items of this list are treated as uncorrelated to one another, and unless otherwise stated, they are normalization uncertainties.

- The measurement of the *integrated luminosity* of the data sample is affected by a normalization uncertainty, which is defined for all CMS analyses for a given data taking period.
- Experimental uncertainties on the *efficiency of trigger and lepton reconstruction and selection* are covered by one flat number for every final state, defined as the quadrature sum of the per-final-state uncertainty on trigger efficiency described in Section 5.2.1, and that on scale factors for lepton reconstruction and selection efficiencies described in Section 4.3. Per-lepton uncertainties are taken to the four-lepton-candidate level via a toy-based estimation, using a ggH sample and assuming that the leptons are uncorrelated. The final uncertainties are symmetrized, and the procedure is made more conservative by fixing a floor value of 4% (1% per lepton), which only changes the value for the 4μ final state.
- A variety of uncertainties are related to theoretical computations:
  - A *QCD uncertainty* is applied to every signal and irreducible background sample in an uncorrelated way. The renormalization and factorization scales are varied between 0.5 and 2 times their nominal value, while keeping their ratio between 0.5 and 2. This has an impact of 3–10% on global process normalizations and also causes some anti-correlated migrations of events between different categories.

  - The impact of the *modelling of hadronization and the underlying event* has been studied for VBF-2jet tagging in the ggH process, based on dedicated generator-level samples with varied PYTHIA parameters for the underlying event tune and hadronization scale. The resulting uncertainty is found to be 6% and is added linearly to the ggH QCD uncertainty.
  - The uncertainty from the choice of a set of *parton distribution functions* is defined by taking the root mean square of the variation when using different replicas of the default NNPDF set. It is determined independently for different sets of processes grouped by initial state, namely $gg(\to H) \to ZZ$, $qq' \to VBF/VH$, $q\bar{q} \to ZZ$, and $t\bar{t}H$. For the moment, this uncertainty is handled inclusively and not yet per category.
  - The *electroweak corrections* for the $q\bar{q} \to ZZ$ background prediction are assigned an $m_{4\ell}$-dependent uncertainty of 1–15 %. For measurements exploiting the full mass range, it is applied as a shape uncertainty, while on-shell-only measurements just use a normalization uncertainty (1–8%, depending on the category) to account for event migrations from high mass.
  - The $gg \to ZZ$ background yield is applied an additional normalization uncertainty of 10% (besides the QCD uncertainty) to account for the fact that its NNLO/LO K-factor was actually computed for a $gg \to H \to ZZ$ signal, as explained in Section 5.5.1.
  - A systematic uncertainty of 2% on the *branching ratio* of $H \to ZZ \to 4\ell$ is applied to the yields of all signal processes.

- Uncertainties on the *reducible background estimation* have been described in Section 5.5.2. The final numbers include a component for statistical limitations of control regions, and some purely systematic components. The uncertainties are treated as uncorrelated between final states, since each of them can be viewed as a separate measurement.

- The statistical limitations of the $\mathcal{P}(\mathcal{D}^{\text{kin}}_{\text{bkg}}|m_{4\ell})$ template for Z+X are covered by a shape uncertainty. Since the corresponding template of the $q\bar{q} \to ZZ$ background is found to have a similar shape, the uncertainty on the Z+X one is defined from the difference between these two templates, and the mirror of the difference.

- The uncertainty on the *lepton energy scale* follows from the calibrations described in Section 4.2, and is crucial to the measurement of the Higgs boson mass. This is a shape uncertainty on $m_{4\ell}$, whereby the mean of the dCB function (i.e. $m_H + \Delta m_H$ with the notation of Eq. 5.15) varies by 0.3% for electrons and 0.04% for muons.

- Another shape uncertainty affects the 4ℓ *mass resolution* (σ_m in Eq. 5.15): it is fixed to 20% for all three final states, in order to cover the residual differences between the predicted and fitted resolution, as explained in Section 5.3.3.

- Two sources of experimental uncertainty only affect the event categorization, i.e. they do not impact the overall normalization or shapes, but they cause part of the selected events to migrate from one category to another. Both are fully correlated for different processes, and anti-correlated between categories.

  - The imprecise knowledge of the *jet energy scale* affects the number of selected jets in an event, i.e. the primary categorization variable. Varying the jet energy correction factor up and down leads to uncertainties ranging from 2

to 18% depending on the process and category, with largest impacts for the VBF-2jet-tagged and $t\bar{t}H$-tagged categories.

– For the categorization of the 2016 analysis, the data-to-simulation scale factors for *b-tagging efficiency* (described in Section 4.5.2) are assigned some uncertainties. Their upward and downward variations can modify the number of b-tagged jets per event, altering the expected yields of the $t\bar{t}H$-tagged category (which most involves b-tagging) by 2 to 8% depending on the process. The impact on other categories is 0–1% for most processes and up to 5% for $t\bar{t}H$.

The numerical values of part of these systematic uncertainties depend on the data taking period they are used for. Therefore, they will be summarized in Sections 6.3 and 7.3, just before the corresponding results are discussed.

5.8 Conclusion

With the whole analysis strategy being now defined, unblinding collision data at event level and producing physics results involves the following steps:

1. Definition of the data sample and choice of a set of trigger paths among those defined in Chapter 3.

2. Application of data-set-dependent lepton calibrations (Section 4.2) to both data and simulation; computation of data-to-simulation trigger and lepton efficiency scale factors (Section 4.3).

3. Unblinding of the control regions, and data-driven estimation of the reducible background, following the procedure presented in Section 5.5.2.

4. Unblinding of the signal region, and checks of the data-to-simulation agreement of variables related to the event selection and observables defined in Section 5.3.

5. Unblinding of the physics measurements prepared according to the methodology of Section 5.7.

Chapters 6 and 7 will now present the production of two successive sets of results, based on 2015 and 2016 data, respectively.

Chapter 6

Evidence for the Higgs boson in 2015 data

As a first look into 13 TeV data, this chapter presents a set of preliminary results based on the pp collisions recorded by the CMS detector in 2015. The data that were accumulated during this commissioning-oriented year correspond to a moderate integrated luminosity of $2.8\,\text{fb}^{-1}$. It is nevertheless more than enough to extract the first ever set of 13 TeV-based high energy physics results. At this point, ambitious goals were already set for the $\text{H} \to 4\ell$ analysis, namely to obtain the first evidence of the H(125) boson at this unexplored centre-of-mass energy, and to perform the first property measurements wherever possible.

The results of this chapter were produced within the CMS $\text{H} \to \text{ZZ} \to 4\ell$ working group in the fall of 2015 and the winter of 2016, and were first presented on 16^{th} March 2016, at the *51^{st} Rencontres de Moriond on Electroweak interactions and unified theories* in La Thuile. As one of the two co-editors of the CMS Physics Analysis Summary (PAS) [102] in which all 4ℓ-related results were reported, I was involved in all aspects of the analysis, from running the full chain of analysis selection defined in Chapters 4 and 5 to producing the observed and expected distributions for all main variables and contributing to the measurements.

6.1 Recorded data sets

The year 2015 marked the beginning of Run II of the LHC and the first exploration of the $\sqrt{s} = 13\,\text{TeV}$ centre-of-mass energy in pp collisions. The need to commission the full apparatus resulted in the first runs being carried out at low luminosities. The first collisions were delivered in June 2015 with a proton bunch spacing of 50 ns, and CMS used a modified version of its Run I Level-1 trigger, together with the 50 ns-dedicated version of the Run II HLT menu. In August, the LHC moved to 25 ns bunches and CMS then switched to the new Stage 1 Level-1 trigger, which was used until the end of 2015, with increasingly stringent prescales and HLT menus as the LHC ramped up. The highest instantaneous luminosity that was reached in the 2015 runs was $5.2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$ and was accompanied by an average number of 16 pileup interactions per event.

The data sample that is exploited in this chapter reflects the progression of the LHC: it is divided into four different run periods: two at a bunch spacing of 50 ns, and two at 25 ns, the last of which contains most of the total integrated luminosity. For

computational reasons, the recorded events of each period are split into *primary data sets*, each of which contains all events collected by a particular subset of all HLT paths that involve similar physics objects. The 2015 analysis relies on multilepton HLT paths, plus a single-electron path: it therefore exploits the DoubleMuon, DoubleEG, MuonEG, and SingleElectron primary data sets. This results in a total of 16 data sets, which are listed in Table 6.1, while the complete list of HLT paths is given in Table 6.2. The multilepton paths that are used here are the ones that I designed for the 5e33 and 7e33 scenarios (see Section 3.2.3). The 2015 pileup distribution was shown in Fig. 5.1: the average number of pileup interactions for this 2015 sample is ≈ 10.

Table 6.1: List of data sets used in the $H \to ZZ \to 4\ell$ analysis for 2015 data, with information on the corresponding run periods.

data set	run period	integrated luminosity	bunch spacing
/DoubleMuon/Run2015B-16Dec2015-v1 /DoubleEG/Run2015B-16Dec2015-v1 /MuonEG/Run2015B-16Dec2015-v1 /SingleElectron/Run2015B-16Dec2015-v1	Run2015B	$43.9\,\text{pb}^{-1}$	50 ns
/DoubleMuon/Run2015C_50ns-16Dec2015-v1 /DoubleEG/Run2015C_50ns-16Dec2015-v1 /MuonEG/Run2015C_50ns-16Dec2015-v1 /SingleElectron/Run2015C_50ns-16Dec2015-v1	Run2015C	$27.3\,\text{pb}^{-1}$	
/DoubleMuon/Run2015C_25ns-16Dec2015-v1 /DoubleEG/Run2015C_25ns-16Dec2015-v1 /MuonEG/Run2015C_25ns-16Dec2015-v1 /SingleElectron/Run2015C_25ns-16Dec2015-v1		$17.7\,\text{pb}^{-1}$	25 ns
/DoubleMuon/Run2015D-16Dec2015-v1 /DoubleEG/Run2015D-16Dec2015-v2 /MuonEG/Run2015D-16Dec2015-v1 /SingleElectron/Run2015D-16Dec2015-v1	Run2015D	$2.67\,\text{fb}^{-1}$	

Table 6.2: List of HLT paths used in the $H \to ZZ \to 4\ell$ analysis for 2015 data, with their respective L1 seeds and the primary data sets to which they belong.

HLT path	L1 seed	primary data set
`HLT_Ele17_Ele12_CaloIdL_TrackIdL_IsoVL_DZ` `HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL`	`L1_DoubleEG_15_10` `L1_TripleEG_14_10_8`	DoubleEG

`HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ`
`HLT_Mu17_TrkIsoVVL_TkMu8_TrkIsoVVL_DZ`
`HLT_TripleMu_12_10_5`
	`L1_DoubleMu_10_3p5`
	`L1_DoubleMu_10_3p5`
	`L1_TripleMu_5_5_3`
		DoubleMuon

`HLT_Mu8_TrkIsoVVL_Ele17_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu17_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu8_DiEle12_CaloIdL_TrackIdL`
`HLT_DiMu9_Ele9_CaloIdL_TrackIdL`
	`L1_Mu5_EG15`
	`L1_Mu12_EG10`
	`L1_Mu6_DoubleEG10`
	`L1_DoubleMu7_EG7`
		MuonEG

`HLT_Ele23_WPLoose_Gsf`	`(multiple)`	SingleElectron

The data sets are *re-reconstructed*, which means that the standard CMS event reconstruction was run again in the Winter of 2016 with improved calibrations with respect to those used during the data taking in 2015. Moreover, a certification workflow ensures that only good-quality data, recorded with a fully operational CMS detector, are used for physics. For instance, some problems with the cooling system of the superconducting solenoid of CMS led to turn it off in some runs. This resulted in 23% of the data being recorded without any magnetic field, thus being unusable for many physics analyses, including $\mathrm{H} \to 4\ell$. In contrast, although 15% of the remaining data were affected by a low-occupancy problem in the HF calorimeter, this was shown not to hamper the $\mathrm{H} \to 4\ell$ analysis, and the corresponding runs are exploited here. The data sample that is finally used in this analysis corresponds to an integrated luminosity of $2.8\,\mathrm{fb}^{-1}$.

6.2 Results of background estimation

The control regions of the 2015 data sample are studied in order to estimate the reducible background yield $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ in the signal region, according to the SS and OS methods described in Section 5.5.2. Both results are shown in Table 6.3: the predicted Z+X yields from the two methods are found to be overall compatible, with a 1σ tension in the 4e final state. Their central values and uncertainty ranges are combined as described in Section 5.5.2, and the result is also given in Table 6.3.

Figure 6.1 illustrates the combination of the $m_{4\ell}$ distributions of Z+X from the SS and OS methods, showing the sum of all final states, which are separately fitted to Landau functions. The resulting analytical shapes are then used in both event categories.

Table 6.3: Estimated reducible background yield $N_{\mathrm{SR}}^{\mathrm{Z+X}}$ in the signal region, as obtained from the SS and OS methods and from their combination, for $2.8\,\mathrm{fb}^{-1}$ of 2015 data. The result is given for every final state, for the sum of the two event categories. The quoted uncertainties include statistical and systematic sources.

Final state	4e	4μ	$2\mathrm{e}2\mu$
SS method	1.7 ± 0.9	2.3 ± 0.7	3.4 ± 1.2
OS method	3.1 ± 1.1	2.0 ± 0.7	3.1 ± 1.1
Combination	$2.2^{+2.0}_{-1.2}$	2.1 ± 0.9	3.2 ± 1.3

6.3 Systematic uncertainties

Table 6.4 summarizes the values of the systematic uncertainties used in the measurements of this 2015 analysis. Details about the definition and computation of these uncertainties are given in Section 5.7.3.

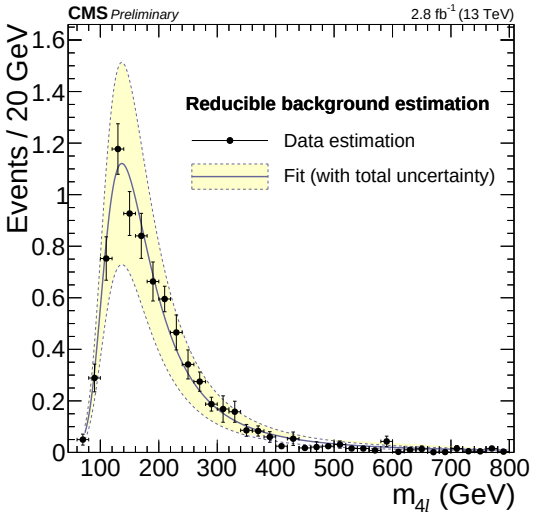

Figure 6.1: Predicted distribution of the four-lepton invariant mass $m_{4\ell}$ for the Z+X background in the signal region, for the sum of all final states and event categories. Points with vertical uncertainty bars represent the per-bin prediction obtained from combining results of the SS and OS methods, while the curve is the sum of the analytical shapes of each final state. The yellow band shows the total uncertainty on the prediction.

Table 6.4: Summary of the systematic uncertainties affecting measurements of the 2015 $H \to ZZ \to 4\ell$ analysis. The uncertainties either affect the shape (S) or only the normalization (N) of predicted distributions. (p) means that the value of the uncertainty is process-dependent.

source of uncertainty	type	magnitude (4e/4μ/2e2μ)
luminosity	N	2.7 %
trigger+reconstruction+selection efficiencies	N	10% / 4% / 5%
QCD scale, inclusive (p)	N	3 – 10 %
QCD scale + underlying event, in categories (p)	N	5 – 25 %
PDF set (p)	N	3 – 5 %
NLO/LO EW corrections for $q\bar{q} \to ZZ$	N	1 – 15 %
NNLO/LO K-factor for $gg \to ZZ$	N	10 %
$H \to ZZ \to 4\ell$ branching ratio	N	2 %
Z+X reducible background prediction	N	$^{+90\%}_{-55\%}$ / 40% / 40%
lepton energy scale	S	0.3% / 0.04% / 0.17%
4ℓ mass resolution	S	20 %
jet energy scale for VBF-tagged category (p)	N	2 – 9 %

6.4 Results of event selection

This section presents the outcome of event selection in the signal region, for two classes of measurement inputs: the predicted and observed event yields, and the distributions of the main kinematic variables in data and simulation. These simple comparisons contribute to checking that every element of the analysis behaves as expected before performing the measurements.

6.4.1 Event yields

The number of expected background and signal events and the number of events observed in the 2015 data sample are reported in Table 6.5, for each final state and for the inclusive sample. The $q\bar{q} \to ZZ$ and $gg \to ZZ$ estimations are estimated from MC simulation, while the Z+X component comes from the data-driven estimation reported in Section 6.2. The quoted uncertainties include systematic, plus, in the case of Z+X, statistical sources. 109 events are observed in the full signal region ($m_{4\ell} > 70\,\text{GeV}$), where $118.3^{+8.3}_{-7.8}$ (background only) or $123.7^{+8.6}_{-8.2}$ (background and H(125) signal) are expected. The downward fluctuation in data is driven by the 4e and 2e2μ final states, while 4μ exhibits an upward fluctuation.

Table 6.5: Number of observed events compared to the expected background and signal yields for each final state, in the full signal region ($m_{4\ell} > 70\,\text{GeV}$). Uncertainties include statistical and systematic sources.

Final state	4e	4μ	2e2μ	4ℓ
$q\bar{q} \to ZZ$	$18.3^{+1.9}_{-1.8}$	31.1 ± 2.0	$42.6^{+3.5}_{-3.3}$	$92.0^{+6.7}_{-6.4}$
$gg \to ZZ$	3.9 ± 0.6	5.9 ± 0.8	9.0 ± 1.3	$18.8^{+2.6}_{-2.5}$
Z+X	$2.2^{+2.0}_{-1.2}$	2.1 ± 0.9	3.2 ± 1.3	$7.5^{+2.5}_{-2.0}$
Sum of backgrounds	$24.4^{+3.0}_{-2.4}$	39.1 ± 2.5	$54.8^{+4.4}_{-4.2}$	$118.3^{+8.3}_{-7.8}$
Signal ($m_{\text{H}} = 125\,\text{GeV}$)	1.1 ± 0.1	1.9 ± 0.2	2.5 ± 0.2	5.5 ± 0.5
Total expected	$25.5^{+3.1}_{-2.5}$	$40.9^{+2.6}_{-2.5}$	$57.3^{+4.5}_{-4.4}$	$123.7^{+8.6}_{-8.2}$
Observed	17	49	43	109

Table 6.6 shows a similar breakdown of expected and observed yields, focusing this time on a $118 < m_{4\ell} < 130\,\text{GeV}$ window around the mass of the Higgs boson discovered in Run I. To obtain a Z+X prediction in this mass range, the prediction of the full signal region is scaled according to the relevant fraction of the integral of the full $m_{4\ell}$ analytical shape. 8 events are observed between 118 and 130 GeV, while $2.80^{+0.25}_{-0.22}$ (background only) or $7.62^{+0.58}_{-0.56}$ (background and H(125) signal) are expected. The characteristics of this excess are discussed in the next section. Table 6.7 shows the breakdown of expected yields into event categories instead of final states, in the same $118 < m_{4\ell} < 130\,\text{GeV}$ range. All 8 data events are Untagged.

6.4.2 Event distributions

This section shows an overview of the distributions of the three main observables used in the event categorization and in the measurements, confronting the data to expectations from signal and background processes.

Figure 6.2 shows the distribution of the reconstructed four-lepton invariant mass in the full signal region, for the sum of the three final states and the two categories. The vertical bars on data points are asymmetric Poisson uncertainties that cover the 68% probability interval around the central value [103]. The observed distribution agrees with the expectation within the statistical uncertainties over the whole spectrum: at low mass, in the region of the $Z \to 4\ell$ resonance peak; at high mass, for $m_{4\ell} > 2m_{Z^0}$

Table 6.6: Number of observed events compared to the expected background and signal yields for each final state, for the mass range $118 < m_{4\ell} < 130\,\text{GeV}$. Uncertainties include statistical and systematic sources.

Final state	4e	4μ	$2e2\mu$	4ℓ
$q\bar{q} \to ZZ$	0.33 ± 0.03	0.75 ± 0.05	0.92 ± 0.07	2.00 ± 0.14
$gg \to ZZ$	0.04 ± 0.01	0.08 ± 0.01	0.07 ± 0.01	$0.18^{+0.03}_{-0.02}$
Z+X	$0.17^{+0.15}_{-0.09}$	0.19 ± 0.08	0.26 ± 0.10	$0.62^{+0.20}_{-0.16}$
Sum of backgrounds	$0.54^{+0.16}_{-0.10}$	1.02 ± 0.09	1.25 ± 0.13	$2.80^{+0.25}_{-0.22}$
Signal ($m_H = 125\,\text{GeV}$)	$0.91^{+0.11}_{-0.10}$	1.70 ± 0.15	2.21 ± 0.22	$4.82^{+0.44}_{-0.45}$
Total expected	$1.45^{+0.21}_{-0.16}$	2.72 ± 0.20	3.45 ± 0.29	$7.62^{+0.58}_{-0.56}$
Observed	1	3	4	8

Table 6.7: Number of observed events compared to the expected background and signal yields for each event category, for the mass range $118 < m_{4\ell} < 130\,\text{GeV}$.

Category	Untagged	VBF-tagged	Total
$q\bar{q} \to ZZ$	2.00 ± 0.14	0.01 ± 0.00	2.00 ± 0.14
$gg \to ZZ$	$0.18^{+0.03}_{-0.02}$	0.00 ± 0.00	$0.18^{+0.03}_{-0.02}$
Z+X	$0.60^{+0.19}_{-0.15}$	$0.02^{+0.01}_{-0.00}$	$0.62^{+0.20}_{-0.16}$
Sum of backgrounds	$2.77^{+0.25}_{-0.21}$	0.03 ± 0.01	$2.80^{+0.25}_{-0.22}$
Signal ($m_H = 125\,\text{GeV}$)	$4.54^{+0.43}_{-0.44}$	0.28 ± 0.04	$4.82^{+0.44}_{-0.45}$
Total expected	$7.32^{+0.56}_{-0.55}$	0.31 ± 0.04	$7.62^{+0.58}_{-0.56}$
Observed	8	0	8

where the two Z bosons are produced on shell; and in the 125 GeV region, where the SM Higgs boson is visible as a narrow peak.

Focusing on the low-mass region, Fig. 6.3 shows the correlations of both kinematic discriminants versus $m_{4\ell}$: the decay-related $\mathcal{D}^{\text{kin}}_{\text{bkg}}$ discriminant which is used in the fits, and the production-related $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$ discriminant which is used to define the VBF-tagged category. The few events of the H(125) peak turn out to cluster at medium values of $\mathcal{D}^{\text{kin}}_{\text{bkg}}$, leaving the region of larger values empty. The only event of the peak which has two selected jets (a 4μ event) has $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}} \simeq 0.28$, and is therefore not VBF-tagged.

Many more distributions and control variables were studied and are not shown here, but Chapter 7 will give more details with the 2016 data sample, which has 4.6 times more data.

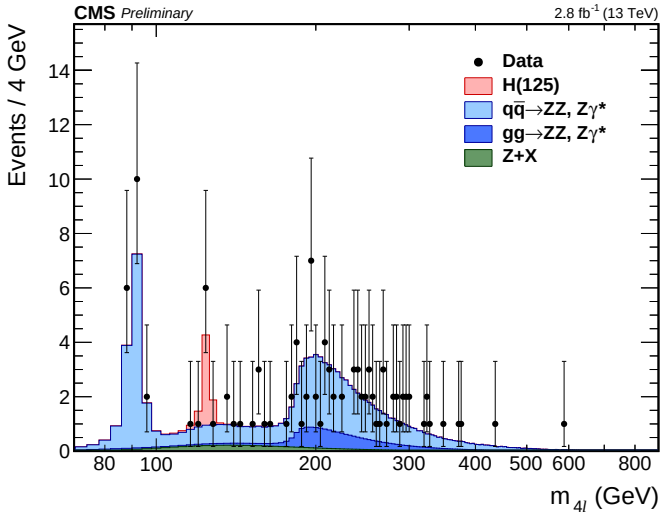

Figure 6.2: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the signal region. Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data. No events are observed with $m_{4\ell} > 600\,\text{GeV}$.

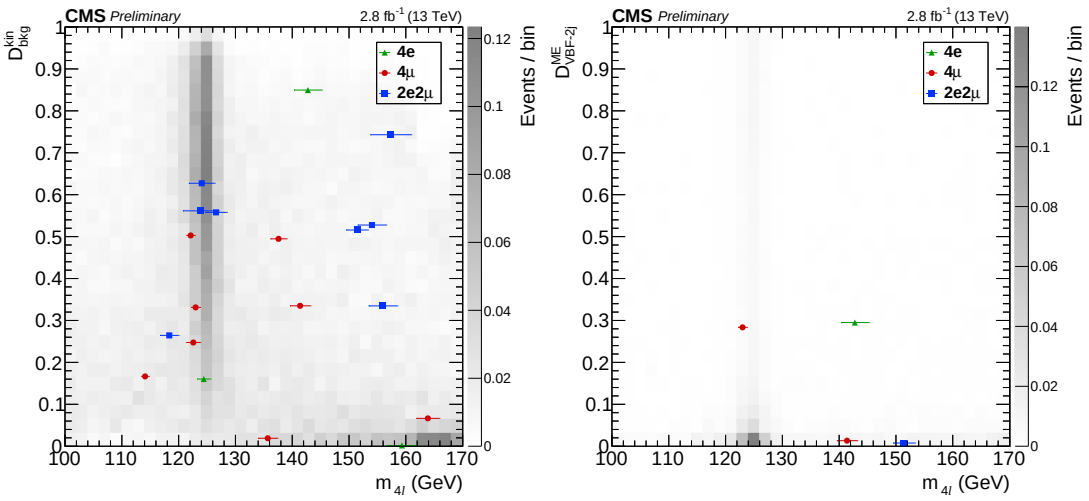

Figure 6.3: Distributions of kinematic discriminants versus the four-lepton reconstructed mass $m_{4\ell}$ in the mass region $100 < m_{4\ell} < 170\,\text{GeV}$: (left) decay discriminant $\mathcal{D}^{\text{kin}}_{\text{bkg}}$ versus $m_{4\ell}$, and (right) VBF-sensitive production discriminant $\mathcal{D}^{\text{ME}}_{\text{VBF}-2\text{j}}$ versus $m_{4\ell}$ for events with two selected jets. The grey scale represents the expected relative density of ZZ background plus H(125) signal, which are normalized to the SM expectation. The points show the data and the horizontal bars represent the measured per-event mass uncertainties (see Section 5.3.3).

6.5 Measurements

6.5.1 Significance

The significance of the excess of events observed near 125 GeV is estimated following the statistical procedure outlined in Section 5.7.1. Figure 6.4 (left) shows the local p-value and significance relative to the SM background-only expectation as a function of m_{H}, for the range $118 < m_{\mathrm{H}} < 130$ GeV, with a breakdown into final states. The minimum of the p-value is reached at 123.4 GeV and corresponds to a local significance of 3.0 standard deviations (σ), while 3.1σ are expected at this mass for a SM Higgs boson. At $m_{\mathrm{H}} = 125.09$ GeV, which corresponds to the Run I LHC combined measurement of the Higgs boson mass [27], the observed significance is 2.5σ, for an expectation of 3.4σ.

It is worth noticing that when the one-dimensional model $f^{\mu}_{\mathrm{1D}}(m_{4\ell})$ is used instead of the nominal two-dimensional version $f^{\mu}_{\mathrm{2D}}(m_{4\ell}, \mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}})$, the observed significances are larger, with 3.4σ observed (2.6σ expected) at 123.4 GeV and 2.7σ observed (2.9σ expected) at 125.09 GeV, as illustrated in Fig. 6.4 (right). This peculiar difference is due to the fact that the on-peak events do not populate the large-$\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$ region, as was noted from Fig. 6.3 (left). As a result, the events appear as less signal-like in the 2D template than in the 1D one, despite the gain in expected sensitivity.

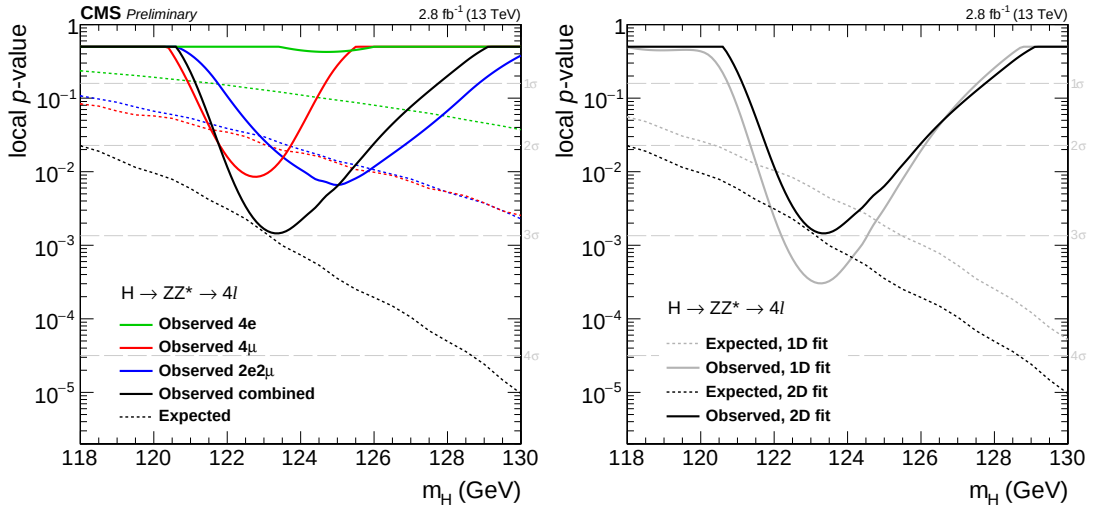

Figure 6.4: p-value and significance of the local fluctuation with respect to the SM expectation as a function of the Higgs boson mass. Dashed lines show the median expected local p-values for the local mass hypothesis. (left) Main result using $m_{4\ell}$ and $\mathcal{D}^{\mathrm{kin}}_{\mathrm{bkg}}$, (right) comparison to the result using $m_{4\ell}$ only.

The dedicated measurement of the Higgs boson mass was not performed yet in this 2015 analysis, but m_{H} was nevertheless measured as a cross-check, from the same 2D model as used to estimate the significance. The result is $m_{\mathrm{H}} = 123.4^{+0.8}_{-0.7}$ GeV, which is compatible with the Run I LHC measurement of 125.09 ± 0.24 GeV at the 1.6σ level.

6.5.2 Signal strength and production modes

As explained in Section 5.7.1, the signal model used to estimate the significance is also used to measure the signal strength, from a simultaneous unbinned maximum likelihood

fit to all six channels (2 categories × 3 final states). The combined measured signal strength is $\mu = \sigma/\sigma_{\mathrm{SM}} = 0.82^{+0.57}_{-0.43}$ for a mass hypothesis of $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$.

A trivial possible alteration of the baseline signal model consists in having separate parameters for each channel. For instance, the signal strength can be measured in each event category, as illustrated in Fig. 6.5 (left). The result is $0.89^{+0.62}_{-0.46}$ in the Untagged category, and $0.0^{+1.7}_{-0.0}$ in the VBF-tagged category, in which no event is observed. The latter value is bounded by zero from below because it was chosen to constrain all signal strength parameters to be non-negative. The measurements are consistent with SM expectations within their uncertainties, which are large and dominated by their statistical component with the limited 2015 data set.

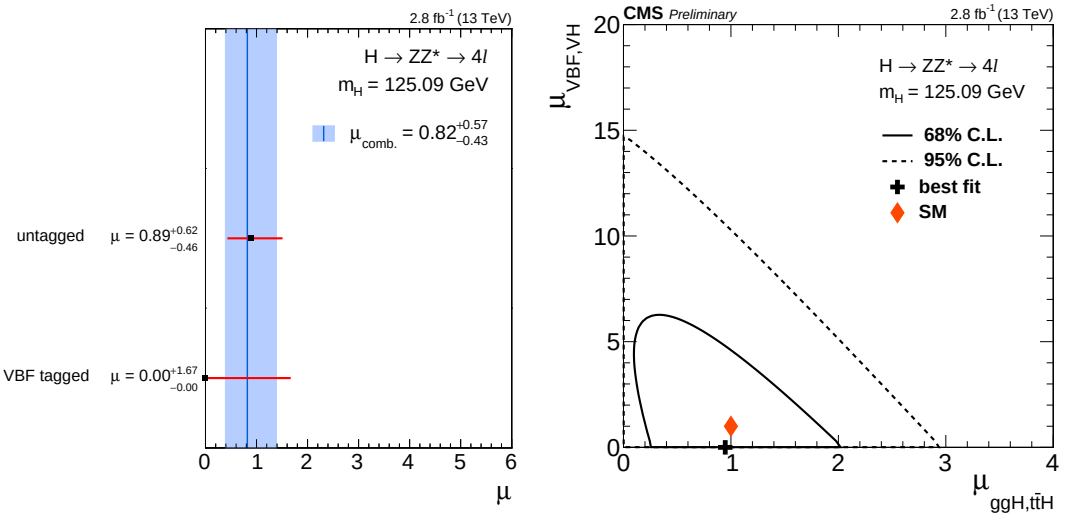

Figure 6.5: (left) Observed values of the signal strength $\mu = \sigma/\sigma_{SM}$ for the two event categories, compared to the combined one shown as a vertical line. The horizontal bars and the filled band indicate the $\pm 1\sigma$ uncertainties. The uncertainties include both statistical and systematic sources. (right) Result of the 2D likelihood scan for the $\mu_{\mathrm{ggH,t\bar{t}H}}$ and $\mu_{\mathrm{VBF,VH}}$ signal-strength modifiers. The solid and dashed contours show the 68% and 95% CL regions, respectively. The cross indicates the best-fit values, and the diamond represents the expected values for the SM Higgs boson.

However, the information from event categorization can already be used to constrain Higgs boson production mechanisms. A simple widespread model relies on two parameters of interest, namely two signal strength modifiers defined as scale factors for the fermion and vector-boson induced contributions to the expected SM cross section:

- $\mu_{\mathrm{ggH,t\bar{t}H}}$ scales the Higgs boson production cross sections in modes related to couplings to fermions (i.e. gluon fusion, where the loop mostly involves top and bottom quarks, and $\mathrm{t\bar{t}H}$, where a Higgs boson is radiated off a top quark).
- $\mu_{\mathrm{VBF,VH}}$ scales the cross sections in modes related to couplings to electroweak gauge bosons (i.e. VBF, WH and ZH, all of which rely on the vertex involving one Higgs boson and two gauge bosons).

In other words, the $\mu \cdot s + b$ model is replaced by $\mu_{\mathrm{ggH,t\bar{t}H}} \cdot s_{\mathrm{ggH,t\bar{t}H}} + \mu_{\mathrm{VBF,VH}} \cdot s_{\mathrm{VBF,VH}} + b$. The two-dimensional fit is performed assuming a mass of $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$, leading to the measurements of $\mu_{\mathrm{ggH,t\bar{t}H}} = 0.95^{+0.64}_{-0.49}$ and $\mu_{\mathrm{VBF,VH}} = 0.0^{+2.5}_{-0.0}$. The 68% and 95%

CL contours in the $(\mu_{\mathrm{ggH,t\bar{t}H}}, \mu_{\mathrm{VBF,VH}})$ plane are shown in Fig. 6.5 (right), illustrating the good compatibility of the result with the SM expectation.

6.6 Conclusion

Based on $2.8\,\mathrm{fb}^{-1}$ of pp data recorded in 2015, the first CMS results in the $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ channel at $\sqrt{s} = 13\,\mathrm{TeV}$ were produced, proving that every aspect of the analysis was already under control with the upgraded apparatus, software, and analysis methods. The first evidence for the previously discovered Higgs boson at this new centre-of-mass energy is obtained, with an observed (expected) significance of 3.0σ (3.1σ) at the minimum of the p-value and 2.5σ (3.4σ) at $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$, meaning that the sensitivity is already comparable to that of the contribution of the $\mathrm{H}\to 4\ell$ channel to the discovery of the Higgs boson in July 2012 [19]. The measured signal strength is $\mu = 0.82^{+0.57}_{-0.43}$ at $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$, and a simple Run I-like event categorization is used to measure the signal strength modifiers associated with fermions and vector bosons, which are $\mu_{\mathrm{ggH,t\bar{t}H}} = 0.95^{+0.64}_{-0.49}$ and $\mu_{\mathrm{VBF,VH}} = 0.0^{+2.5}_{-0.0}$, respectively.

While all these preliminary results are consistent with the expectations for the SM Higgs boson within the uncertainties, the 2015 integrated luminosity is still too limited to claim a rediscovery in the 4ℓ channel, let alone to establish subdominant production modes: such endeavours require the high-luminosity data sample which is being collected as of 2016.

Chapter 7

Measurement of Higgs boson properties and search for a new resonance in 2016 data

With the LHC performance ramping up, the complete 2016 data sample was expected to bring one order of magnitude more integrated luminosity than the 2015 one did, opening prospects for fruitful analyses in the scalar sector at $\sqrt{s} = 13\,\text{TeV}$. This chapter thus reports a second round of results, based on a partial 2016 data sample that was unblinded for the summer conferences. This sample was recorded from April to mid-July 2016 and corresponds to an integrated luminosity of $12.9\,\text{fb}^{-1}$.

Owing to the increase of Higgs boson production cross sections reported in Chapter 1, this luminosity is already enough to reach a similar expected sensitivity as CMS did with the $\approx 25\,\text{fb}^{-1}$ of the full Run I data set (at 7 TeV and 8 TeV). This allows for a full standalone rediscovery of the H(125) boson in the 4ℓ channel and for an extensive study of its properties. In particular, it now makes sense to deploy the full event categorization to probe all main production mechanisms, which is done here via the measurement of a set of signal strength modifiers. The mass and decay width of the boson are also measured, and a search for an additional scalar resonance decaying to four leptons is performed in a large mass range. Besides running the analysis selection chain to provide the inputs to all results, I specifically performed the measurements related to signal strength and production mode extraction.

The results presented in this chapter are part of an even larger set of Higgs boson measurements in the four-lepton channel, which are exposed in a new PAS referenced as [104]. They were made public on 4^{th} and 5^{th} August 2016 during the *38^{th} International Conference on High Energy Physics* in Chicago, at which I presented the overall analysis strategy and the on-shell-related measurements.

7.1 Recorded data sets

As of 2016, all LHC pp collisions are being delivered with a bunch spacing of 25 ns. CMS has now switched to the new Stage 2 Level-1 trigger setup, which had been commissioned in parallel to the Stage 1 one during the latest 2015 runs, and will be used in the remainder of Run II. The data sample used in this chapter spans three run periods and requires the highest level of data certification, for a total integrated luminosity of

$12.9\,\text{fb}^{-1}$. This time, five different primary data sets are exploited in the $\text{H}\to\text{ZZ}\to 4\ell$ analysis, for a total of 15 data sets, the full list of which is given in Table 7.1.

Since the preliminary results of this chapter were produced during the data taking, they exploit the *prompt reconstruction* workflow, which means that the data underwent the standard CMS event reconstruction as soon as they were recorded, with preliminary cosmics-based detector calibrations. The data will then be studied to better master the detector alignment and fine-tune the calibrations, so that a re-reconstruction can be performed on raw data and used in the final publications, for the benefit of precision measurements such as that of the Higgs boson mass.

The trigger strategy for this 2016 analysis relies on a mixture of two sets of HLT paths: depending on the run, either the paths designed for the 7e33 scenario or those of the 1.4e34 scenario were used, as dictated by peak instantaneous luminosities. These paths have been slightly modified to be interfaced with the Stage 2 trigger. Although the impact on $\text{H}\to\text{ZZ}\to 4\ell$ trigger efficiencies is small, it was decided to complement the standard set of multilepton paths with some additional double-electron, single-electron and single-muon[1] paths that were originally designed for other unrelated analyses, in order to bring the efficiency as close as possible to 100%. The complete list of HLT paths is given in Table 7.2.

7.2 Results of background estimation

Before examining the signal region for the considered data sample, control regions are studied in order to estimate the reducible background yield $N_{\text{SR}}^{\text{Z+X}}$ in the signal region with the SS and OS methods.

Figure 7.1 illustrates the contents of the three control regions defined in Section 5.5.2, comparing their observed $m_{4\ell}$ distributions to the simulation of the main processes that contribute to them. Such data involving background leptons are not perfectly modelled, but the simulated distributions are not directly used in background estimation.

The observed yields of the 2P2L_{SS} region are presented in more detail in Table 7.3, for each final state and each of the six event categories defined in Section 5.4.2. Every channel is populated, with lower yields when the two loose leptons are reconstructed muons (indeed, fewer physics objects mimic prompt muons than prompt electrons). The lowest yields are observed in the VH-leptonic-tagged category, but are still enough for the SS method to be applied. The $\text{Z}+\text{X}$ predicted yields $N_{\text{SR}}^{\text{Z+X}}$ for the signal region are then also presented in Table 7.3, as computed from Eq. 5.13, and with their total uncertainty. The dominant contribution to this uncertainty is the systematic component from the process composition of the control region, which is estimated inclusively.

The combination of inclusive predictions from the SS and OS methods is presented in Table 7.4. Predicted $\text{Z}+\text{X}$ yields from the two methods are found to be in good agreement within their uncertainties, and are combined as described in Section 5.5.2. The detailed SS-based expected yields of Table 7.3 are then rescaled by the ratio of this combined inclusive prediction to the SS one.

Finally, Fig. 7.2 illustrates the combination of the $m_{4\ell}$ distributions from the SS and OS methods, showing the sum of the four final states (4e, 4μ, and distinct $2\text{e}2\mu$ and

[1] Hence the use of a fifth primary data set called SingleMuon, which was not exploited in the 2015 analysis.

Table 7.1: List of data sets used in the $H \to ZZ \to 4\ell$ analysis for early 2016 data, with information on the corresponding run periods.

data set	run period	integrated luminosity	bunch spacing
/DoubleEG/Run2016B-PromptReco-v2 /DoubleMuon/Run2016B-PromptReco-v2 /MuonEG/Run2016B-PromptReco-v2 /SingleElectron/Run2016B-PromptReco-v2 /SingleMuon/Run2016B-PromptReco-v2	Run2016B	$5.8\,\text{fb}^{-1}$	25 ns
/DoubleEG/Run2016C-PromptReco-v2 /DoubleMuon/Run2016C-PromptReco-v2 /MuonEG/Run2016C-PromptReco-v2 /SingleElectron/Run2016C-PromptReco-v2 /SingleMuon/Run2016C-PromptReco-v2	Run2016C	$2.6\,\text{fb}^{-1}$	
/DoubleEG/Run2016D-PromptReco-v2 /DoubleMuon/Run2016D-PromptReco-v2 /MuonEG/Run2016D-PromptReco-v2 /SingleElectron/Run2016D-PromptReco-v2 /SingleMuon/Run2016D-PromptReco-v2	Run2016D	$4.3\,\text{fb}^{-1}$	

Table 7.2: List of HLT paths used in the $H \to ZZ \to 4\ell$ analysis for early 2016 data, with their respective L1 seeds and the primary data sets to which they belong.

HLT path	L1 seed	primary data set

`HLT_Ele17_Ele12_CaloIdL_TrackIdL_IsoVL_DZ`
`HLT_Ele23_Ele12_CaloIdL_TrackIdL_IsoVL_DZ`
`HLT_DoubleEle33_CaloIdL_GsfTrkIdVL`
`HLT_Ele16_Ele12_Ele8_CaloIdL_TrackIdL`
	`L1_DoubleEG_15_10`
	`L1_DoubleEG_22_10`
	`(multiple)`
	`L1_TripleEG_14_10_8`
		DoubleEG

`HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL`
`HLT_Mu17_TrkIsoVVL_TkMu8_TrkIsoVVL`
`HLT_TripleMu_12_10_5`
	`L1_DoubleMu_11_4`
	`L1_DoubleMu_11_4`
	`L1_TripleMu_5_5_3`
		DoubleMuon

`HLT_Mu8_TrkIsoVVL_Ele17_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu8_TrkIsoVVL_Ele23_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu17_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu23_TrkIsoVVL_Ele12_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu23_TrkIsoVVL_Ele8_CaloIdL_TrackIdL_IsoVL`
`HLT_Mu8_DiEle12_CaloIdL_TrackIdL`
`HLT_DiMu9_Ele9_CaloIdL_TrackIdL`
	`L1_Mu5_EG15`
	`L1_Mu5_EG20`
	`L1_Mu12_EG10`
	`L1_Mu20_EG10`
	`L1_SingleMu*`
	`L1_Mu6_DoubleEG10`
	`L1_DoubleMu7_EG7`
		MuonEG

`HLT_Ele25_eta2p1_WPTight`
`HLT_Ele27_WPTight`
`HLT_Ele27_eta2p1_WPLoose_Gsf`
	`L1_SingleEG*`
	`L1_SingleEG*`
	`L1_SingleEG*`
		SingleElectron

`HLT_IsoMu20 OR HLT_IsoTkMu20` `HLT_IsoMu22 OR HLT_IsoTkMu22`	`L1_SingleMu*` `L1_SingleMu*`	SingleMuon

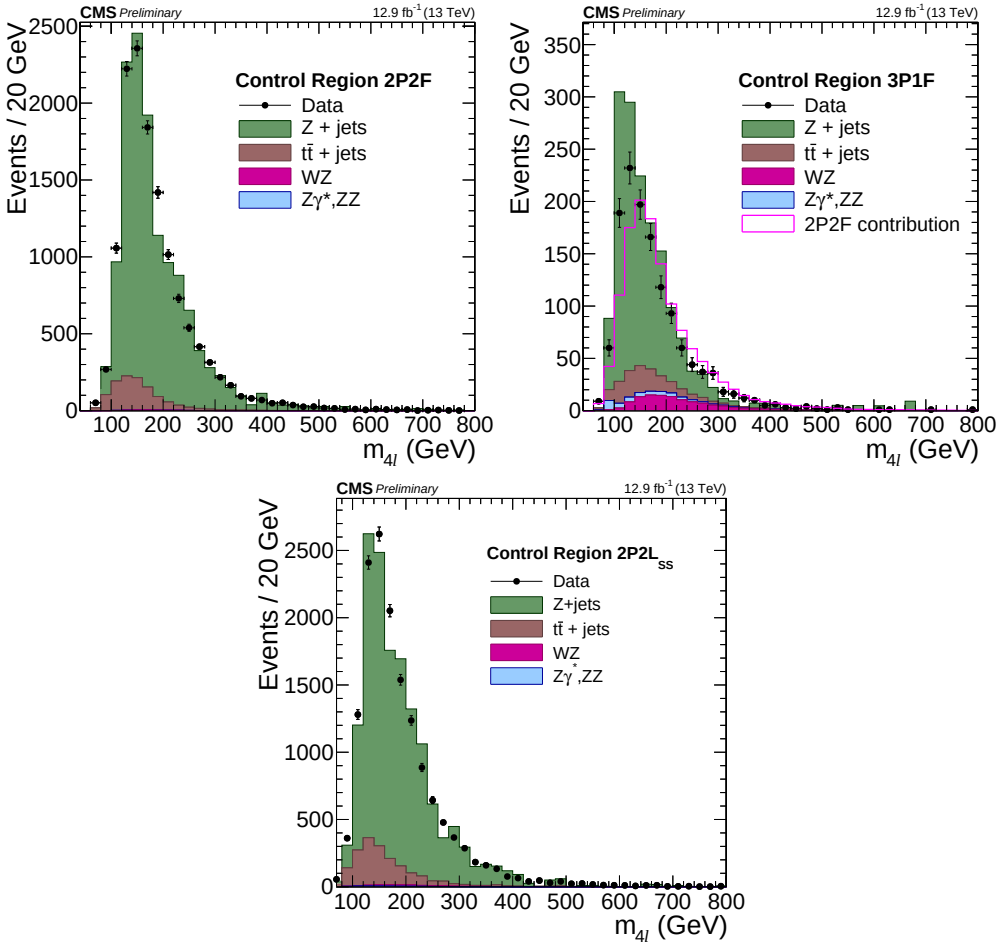

Figure 7.1: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the 2P2F (top left), 3P1F (top right) and 2P2L$_{SS}$ (bottom) control regions, for the sum of all final states and event categories. Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions from simulation, for the main contributing processes. The empty pink histogram in the 3P1F plot represents the predicted contribution of 2P2F-type processes to this region, which has to be subtracted when estimating the Z+X background (see Eq. 5.14).

2μ2e) that are separately fitted to Landau functions. The resulting analytical shapes are later used in each of the six event categories.

7.3 Systematic uncertainties

The values of the systematic uncertainties used in the measurements of this 2016 analysis are summarized in Table 7.5. Details about the definition and computation of these uncertainties are given in Section 5.7.3.

Compared to the values of the 2015 analysis shown in Table 6.4, there is a sizeable increase of the uncertainties on lepton efficiencies. This is driven by the lepton reconstruction component, where the uncertainties are increased mostly because of a problem of inefficiency in tracking reconstruction in 2016 data, caused by events with highly ionizing particles. The larger electron uncertainties are also caused by GSF tracking

Table 7.3: Results of the SS method for Z+X background estimation in 12.9 fb^{-1} of 2016 data, for every final state and event category. Two numbers are quoted in every case: the number of events in the 2P2L$_{\text{SS}}$ region, and the expected yield $N_{\text{SR}}^{\text{Z+X}}$ in the signal region for the full range of $m_{4\ell}$. The quoted uncertainties include statistical and systematic sources.

	4e		2μ2e		4μ		2e2μ	
inclusive	5973	$9.75^{+4.16}_{-3.95}$	7245	$11.73^{+5.22}_{-4.96}$	789	$10.95^{+3.96}_{-3.94}$	714	$10.48^{+3.69}_{-3.67}$
Untagged	5057	$8.25^{+3.53}_{-3.34}$	6044	$9.84^{+4.38}_{-4.16}$	641	$8.89^{+3.20}_{-3.19}$	589	$8.60^{+3.04}_{-3.02}$
VBF-1jet-tagged	430	$0.69^{+0.31}_{-0.29}$	558	$0.91^{+0.41}_{-0.39}$	52	$0.76^{+0.29}_{-0.29}$	40	$0.62^{+0.24}_{-0.24}$
VBF-2jet-tagged	257	$0.43^{+0.19}_{-0.18}$	347	$0.49^{+0.22}_{-0.21}$	54	$0.77^{+0.30}_{-0.30}$	33	$0.51^{+0.20}_{-0.20}$
VH-leptonic-tagged	38	$0.07^{+0.03}_{-0.03}$	64	$0.12^{+0.06}_{-0.05}$	8	$0.09^{+0.05}_{-0.05}$	11	$0.18^{+0.09}_{-0.09}$
VH-hadronic-tagged	90	$0.12^{+0.05}_{-0.05}$	121	$0.17^{+0.07}_{-0.07}$	14	$0.21^{+0.09}_{-0.09}$	17	$0.26^{+0.11}_{-0.11}$
$\text{t}\bar{\text{t}}$H-tagged	101	$0.20^{+0.09}_{-0.08}$	111	$0.20^{+0.10}_{-0.09}$	20	$0.27^{+0.12}_{-0.12}$	24	$0.32^{+0.13}_{-0.13}$

Table 7.4: Estimated reducible background yield $N_{\text{SR}}^{\text{Z+X}}$ in the signal region, as obtained from the SS and OS methods and from their combination, for 12.9 fb^{-1} of 2016 data. The result is given for every final state, for the sum the six event categories. Unlike in Table 7.3, 2μ2e and 2e2μ are here merged into one 2e2μ final state, like for the rest of the analysis. The quoted uncertainties include statistical and systematic sources.

Final state	4e	4μ	2e2μ
SS method	$9.8^{+4.2}_{-4.0}$	$10.9^{+4.0}_{-3.9}$	$22.2^{+8.3}_{-8.2}$
OS method	$9.9^{+3.1}_{-3.4}$	$9.7^{+3.3}_{-3.4}$	$19.1^{+6.9}_{-6.8}$
Combination	$9.8^{+4.2}_{-4.0}$	$10.2^{+4.6}_{-3.9}$	$20.4^{+11.1}_{-8.1}$

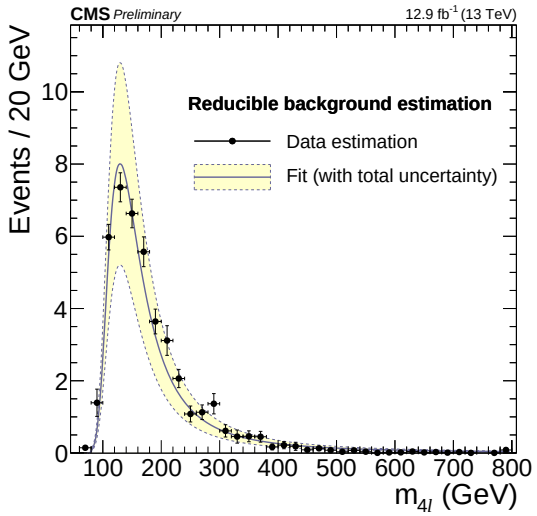

Figure 7.2: Predicted distribution of the four-lepton invariant mass $m_{4\ell}$ for the Z+X background in the signal region, for the sum of all final states and event categories. Points with vertical uncertainty bars represent the per-bin prediction obtained from combining results of the SS and OS methods, while the curve is the sum of the analytical shapes of each final state. The yellow band shows the total uncertainty on the prediction.

efficiency not yet being measured for $p_{\mathrm{T}}^{\mathrm{e}} < 20\,\mathrm{GeV}$. On the other hand, the inclusive Z+X estimation has become more precise in the 4e channel, thanks to the larger data sample. The uncertainty on the integrated luminosity is larger than in 2015, and other differences arise from the use of the six new event categories, for which event migrations have been studied in detail.

Table 7.5: Summary of the systematic uncertainties affecting measurements of the 2016 $\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ analysis. The uncertainties either affect the shape (S) or only the normalization (N) of predicted distributions. (p) means that the value of the uncertainty is process-dependent, (c) means that it is category-dependent.

source of uncertainty	type	magnitude (4e/4μ/2e2μ)
luminosity	N	6.2 %
trigger+reconstruction+selection efficiencies	N	11% / 6% / 9%
QCD scale, inclusive (p)	N	3 – 10 %
QCD scale + underlying event, in categories (p, c)	N	1 – 20 %
PDF set (p)	N	3 – 5 %
NLO/LO EW corrections for $\mathrm{q\overline{q}}\to\mathrm{ZZ}$ (c)	N / S	1 – 15 %
NNLO/LO K-factor for $\mathrm{gg}\to\mathrm{ZZ}$	N	10 %
$\mathrm{H}\to\mathrm{ZZ}\to 4\ell$ branching ratio	N	2 %
Z+X reducible background prediction (c)	N	45% / $^{+40\%}_{-35\%}$ / $^{+45\%}_{-40\%}$
lepton energy scale	S	0.3% / 0.04% / 0.17%
4ℓ mass resolution	S	20 %
jet energy scale (p, c)	N	1 – 18 %
b-tagging efficiency (p, c)	N	0 – 8 %

7.4 Results of event selection

7.4.1 Event yields

The number of expected background and signal events and the number of events observed in the 2016 data sample are reported in Table 7.6, for each final state and for the inclusive sample. As usual, the $\mathrm{q\overline{q}}\to\mathrm{ZZ}$ and $\mathrm{gg}\to\mathrm{ZZ}$ estimations are obtained from MC simulation, while the Z+X component comes from the data-driven estimation reported in Section 7.2. The observed event rates are again found to be compatible with the expectation: 533 events are observed, where $494.4^{+44.6}_{-43.3}$ (background only) or $518.9^{+46.6}_{-45.3}$ (background and H(125) signal) are expected. The slight upward fluctuation of this observed yield originates from the 4e final state.

The yields for the $118 < m_{4\ell} < 130\,\mathrm{GeV}$ range are shown in Table 7.7. 33 events are observed in total, where $13.77^{+1.41}_{-1.38}$ (background only) or $35.38^{+3.43}_{-3.45}$ (background and H(125) signal) are expected. The observed yield agrees well with the signal+background expectation in every final state.

The breakdown into event categories is shown in Table 7.8 for the same $m_{4\ell}$ range. Beyond the dominant Untagged category, two VBF-2jet-tagged events and one VH-hadronic-tagged event are observed, in accordance with their respective expectations. A downward fluctuation is obtained in the VBF-1jet-tagged category, where about five

Table 7.6: Number of observed events compared to the expected background and signal yields for each final state, in the full signal region ($m_{4\ell} > 70\,\text{GeV}$). Uncertainties include statistical and systematic sources.

Final state	4e	4μ	$2e2\mu$	4ℓ
$q\bar{q}\to ZZ$	$71.3^{+8.4}_{-8.0}$	$132.6^{+11.6}_{-11.6}$	$173.3^{+20.6}_{-19.3}$	$377.2^{+36.6}_{-35.7}$
$gg\to ZZ$	$14.8^{+2.6}_{-2.5}$	$25.3^{+3.9}_{-3.9}$	$36.8^{+6.5}_{-6.5}$	$76.9^{+12.6}_{-12.6}$
Z+X	$9.8^{+4.2}_{-4.0}$	$10.2^{+3.8}_{-3.8}$	$20.4^{+7.9}_{-7.7}$	$40.4^{+9.7}_{-9.5}$
Sum of backgrounds	$95.9^{+10.9}_{-10.4}$	$168.0^{+14.1}_{-14.0}$	$230.5^{+25.5}_{-24.2}$	$494.4^{+44.6}_{-43.3}$
Signal ($m_H = 125\,\text{GeV}$)	$4.6^{+0.6}_{-0.6}$	$8.7^{+1.0}_{-1.0}$	$11.2^{+1.5}_{-1.5}$	$24.5^{+2.9}_{-3.0}$
Total expected	$100.5^{+11.4}_{-10.8}$	$176.7^{+14.8}_{-14.7}$	$241.7^{+26.6}_{-25.2}$	$518.9^{+46.6}_{-45.3}$
Observed	111	178	244	533

Table 7.7: Number of observed events compared to the expected background and signal yields for each final state, for the mass range $118 < m_{4\ell} < 130\,\text{GeV}$. Uncertainties include statistical and systematic sources.

Final state	4e	4μ	$2e2\mu$	4ℓ
$q\bar{q}\to ZZ$	$1.37^{+0.16}_{-0.15}$	$3.09^{+0.27}_{-0.27}$	$3.90^{+0.46}_{-0.43}$	$8.36^{+0.81}_{-0.79}$
$gg\to ZZ$	$0.16^{+0.03}_{-0.03}$	$0.32^{+0.05}_{-0.05}$	$0.30^{+0.05}_{-0.05}$	$0.77^{+0.12}_{-0.12}$
Z+X	$0.90^{+0.38}_{-0.37}$	$1.40^{+0.52}_{-0.51}$	$2.34^{+0.91}_{-0.89}$	$4.64^{+1.11}_{-1.09}$
Sum of backgrounds	$2.42^{+0.42}_{-0.40}$	$4.81^{+0.59}_{-0.59}$	$6.54^{+1.03}_{-1.00}$	$13.77^{+1.41}_{-1.38}$
Signal ($m_H = 125\,\text{GeV}$)	$3.90^{+0.53}_{-0.54}$	$7.92^{+0.88}_{-0.93}$	$9.80^{+1.34}_{-1.36}$	$21.61^{+2.63}_{-2.71}$
Total expected	$6.32^{+0.78}_{-0.76}$	$12.73^{+1.21}_{-1.24}$	$16.34^{+1.92}_{-1.90}$	$35.38^{+3.43}_{-3.45}$
Observed	5	12	16	33

events are expected and only one is observed. The VH-leptonic- and $t\bar{t}H$-tagged categories are empty in this mass range, but this is consistent with the expectation of 0.3 event.

7.4.2 Event distributions

Similarly to Section 6.4.2 with 2015 data, this section presents the distributions of the main kinematic variables for 2016 data, comparing data to expectations from signal and background processes. Many more distributions are shown this time than for the 2015 analysis. Again, simulation provides the normalization and shapes of the H(125) signal and the ZZ backgrounds, while the Z+X background is data-driven. For the $m_{4\ell}$ variable, the Z+X normalization and analytical shapes come from the OS+SS combination described in Section 7.2, while for variables other than $m_{4\ell}$, the distributions of Z+X are first extrapolated from the $2P2L_{SS}$ region using the SS method, and then renormalized

Table 7.8: Number of observed events compared to the expected background and signal yields for each event category, for the mass range $118 < m_{4\ell} < 130\,\text{GeV}$.

Category	Untag.	VBF-1j.	VBF-2j.	VH-lep.	VH-had.	$t\bar{t}H$	Total
$q\bar{q} \to ZZ$	7.27	0.82	0.06	0.10	0.11	0.01	8.36
$gg \to ZZ$	0.62	0.11	0.01	0.01	0.01	0.00	0.77
Z+X	3.83	0.32	0.24	0.05	0.08	0.10	4.64
Sum of backgrounds	11.73	1.25	0.32	0.16	0.20	0.11	13.77
Signal ($m_H = 125\,\text{GeV}$)	15.51	3.62	1.45	0.14	0.70	0.19	21.61
Total expected	27.24	4.87	1.77	0.30	0.90	0.30	35.38
Observed	29	1	2	0	1	0	33

to the combined yield. In all 1D plots, the vertical bars on data points are asymmetric Poisson uncertainties that cover the 68% probability interval around the central value.

The distribution of the reconstructed four-lepton invariant mass $m_{4\ell}$ in the full signal region is shown in Fig. 7.3, for the sum of the three final states and the six categories, and in different intervals for better readability. The observed distributions still agree well with the expectation over the whole spectrum, and are now smoothed by the much larger amount of data than in the 2015-based results of Chapter 6. The low-mass part of the spectrum is shown again in Fig. 7.4 for each separate final state, and in Fig. 7.5 for each event category.

The reconstructed invariant masses of lepton pairs selected as Z_1 and Z_2 are shown in Fig. 7.6 with their correlation, both for the full signal region and for $118 < m_{4\ell} < 130\,\text{GeV}$. They also exhibit good agreement between data and expectation. The decay discriminant $\mathcal{D}_{\text{bkg}}^{\text{kin}}$ and its correlation with $m_{4\ell}$ are shown in Fig. 7.7, together with information about the final state and event category of every observed event. Signal-like events appear as a cluster of events at medium values of $\mathcal{D}_{\text{bkg}}^{\text{kin}}$ at the 125 GeV peak.

The four matrix-element based production discriminants used in event categorization, namely $\mathcal{D}_{\text{VBF}-2\text{j}}^{\text{ME}}$, $\mathcal{D}_{\text{VBF}-1\text{j}}^{\text{ME}}$, $\mathcal{D}_{\text{WH}-\text{hadr.}}^{\text{ME}}$, and $\mathcal{D}_{\text{ZH}-\text{hadr.}}^{\text{ME}}$, are shown in Fig. 7.8 in the $118 < m_{4\ell} < 130\,\text{GeV}$ mass window, for events with at least two selected jets (except $\mathcal{D}_{\text{VBF}-1\text{j}}^{\text{ME}}$ which is for events with exactly one selected jet). Their correlations with $m_{4\ell}$ are shown in Fig. 7.9, together with the operating points used to define categories.

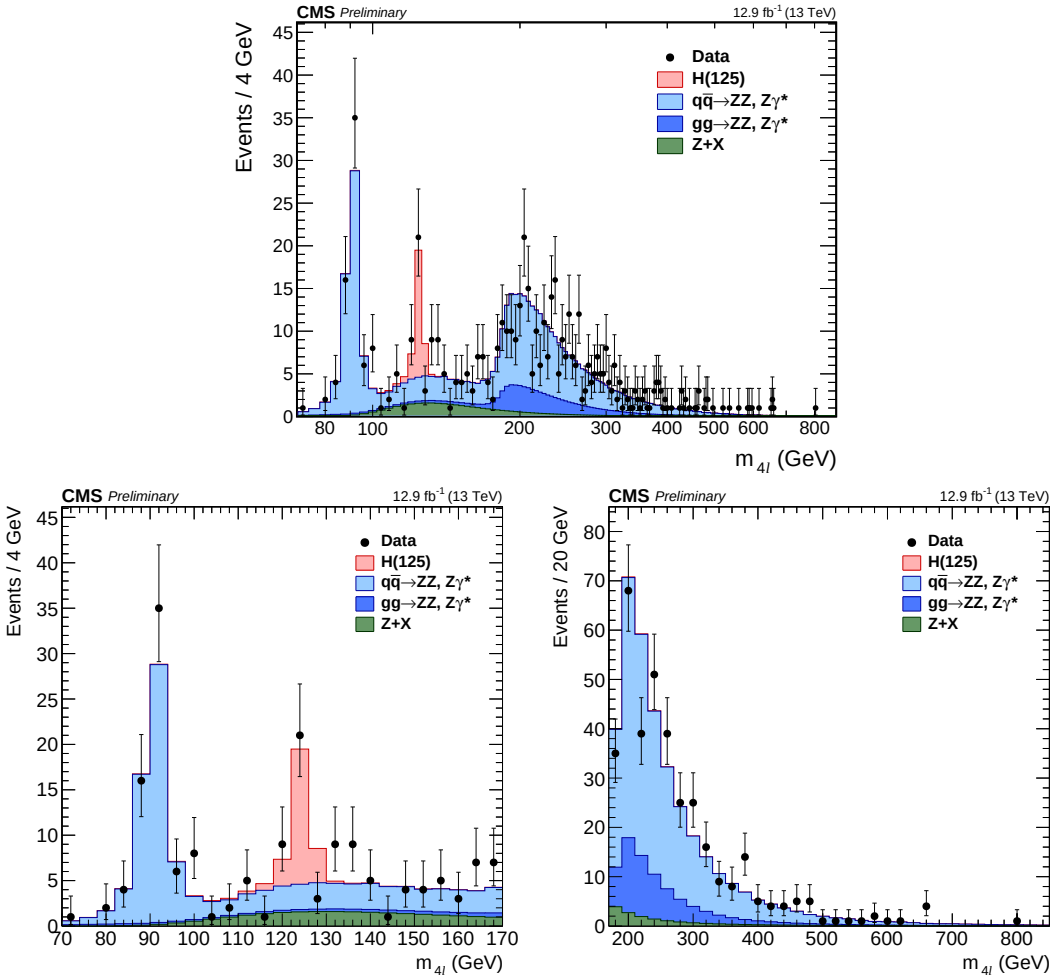

Figure 7.3: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the full mass range (top), the low-mass range (bottom left), and the high-mass range (bottom right). Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data. No events are observed with $m_{4\ell} > 850\,\text{GeV}$.

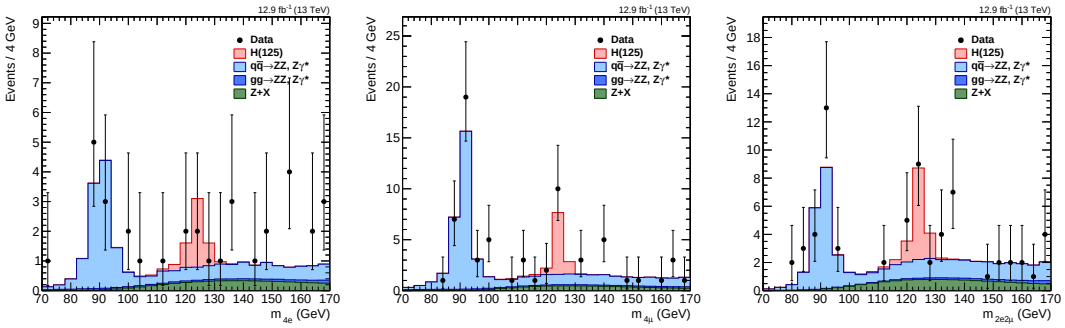

Figure 7.4: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the low-mass range for the 4e (left), 4μ (centre) and $2e2\mu$ (right) final states. Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data.

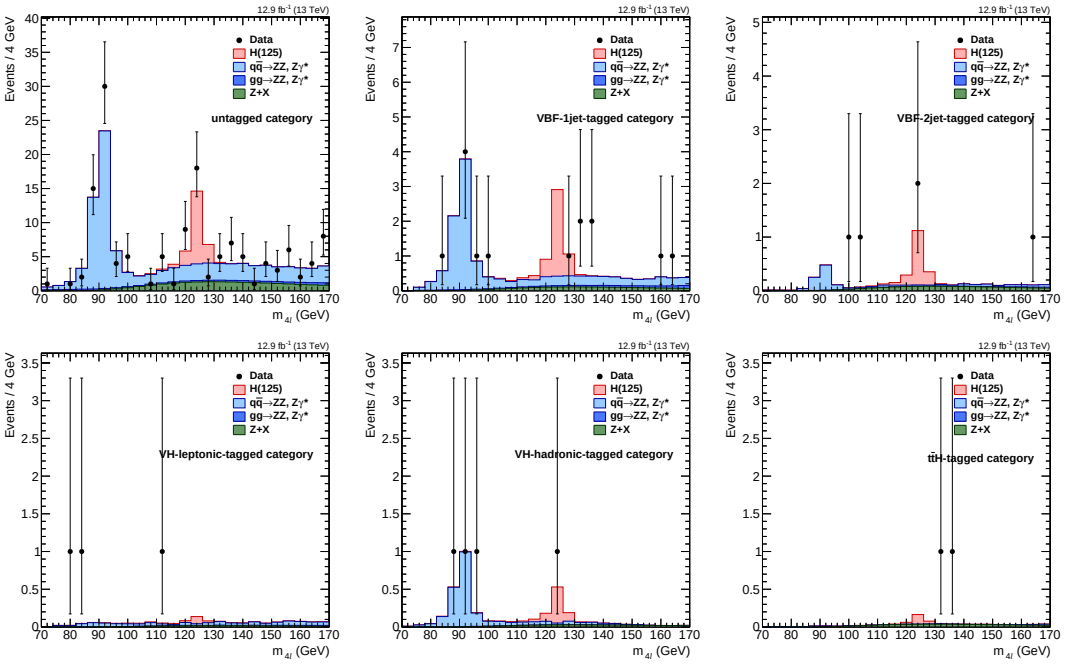

Figure 7.5: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the low-mass range for the six event categories. Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data.

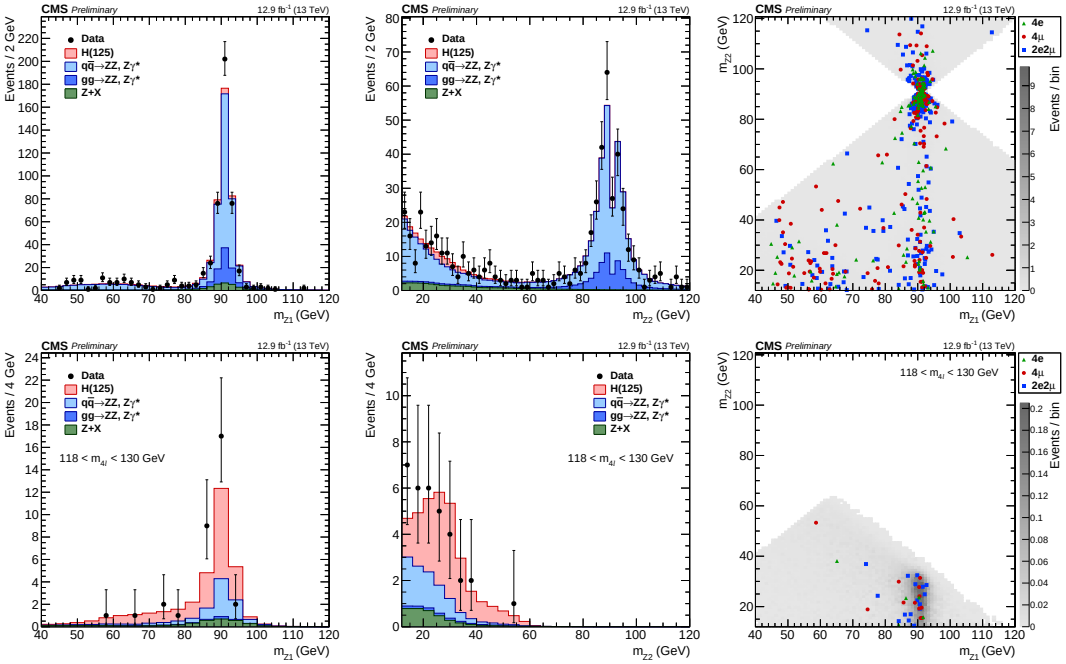

Figure 7.6: Distribution of the Z_1 (left) and Z_2 (centre) reconstructed invariant masses and correlation between the two (right), for the full mass range (top row) and the mass region $118 < m_{4\ell} < 130\,\text{GeV}$ (bottom row). The stacked histograms and the grey scale represent expected distributions, and points represent the data. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data.

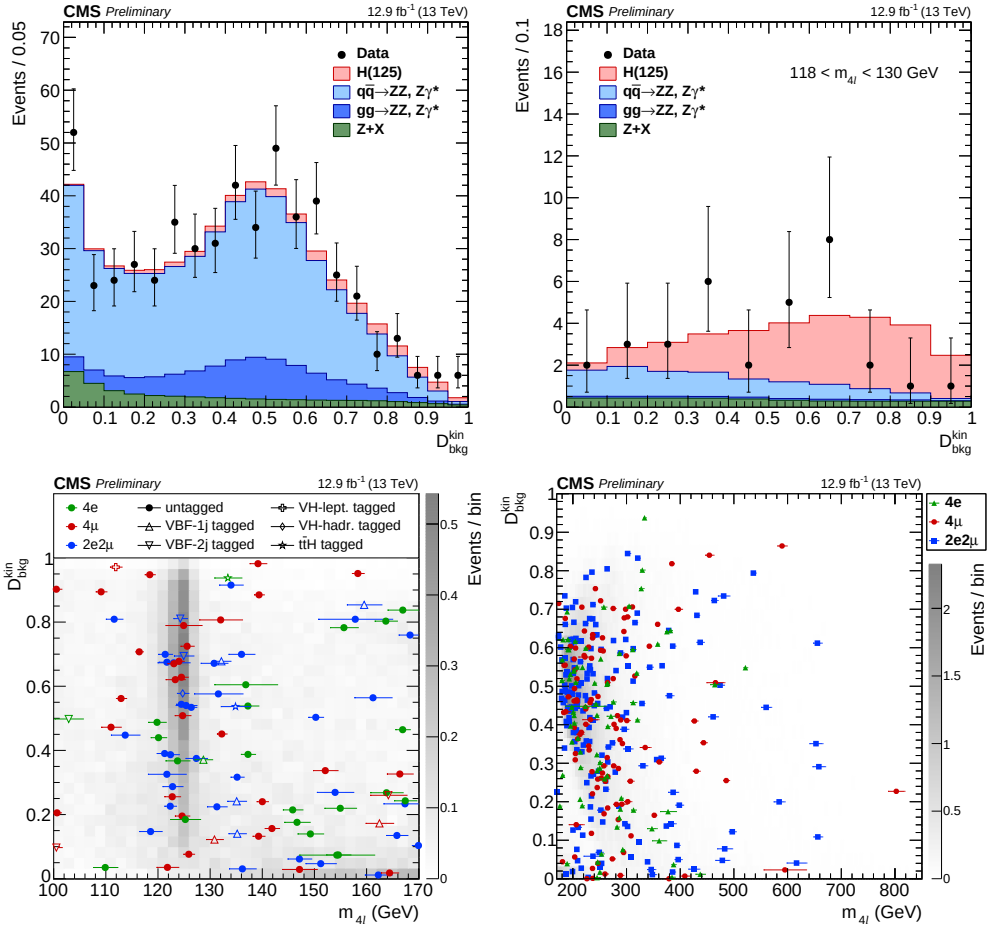

Figure 7.7: Top row: Distribution of the kinematic discriminant $\mathcal{D}_{\text{bkg}}^{\text{kin}}$ in the full mass range (left) and in the mass region $118 < m_{4\ell} < 130\,\text{GeV}$ (right). Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data. Bottom row: Distribution of $\mathcal{D}_{\text{bkg}}^{\text{kin}}$ versus the four-lepton reconstructed mass $m_{4\ell}$ in the low-mass region (left) and in the high-mass region (right). The grey scale represents the expected relative density of ZZ background plus H(125) signal. The points show the data and the horizontal bars represent the measured per-event mass uncertainties.

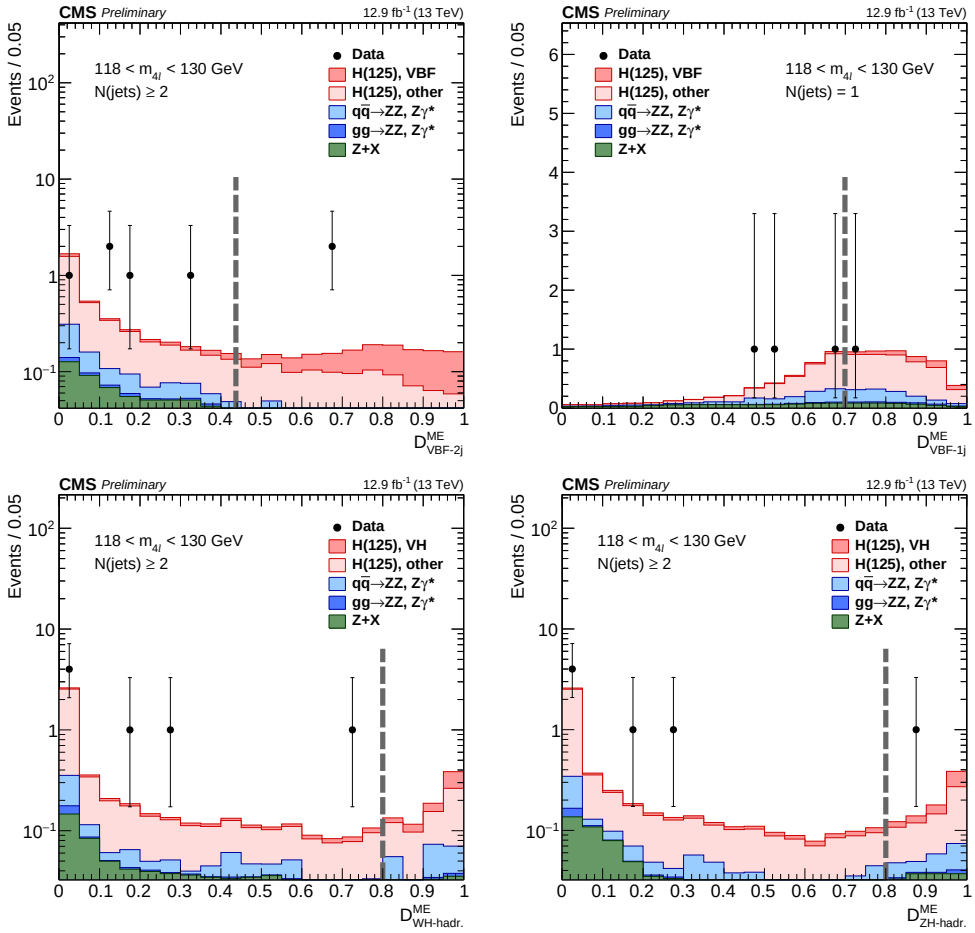

Figure 7.8: Distribution of the four production discriminants used for event categorization, in the mass region $118 < m_{4\ell} < 130$ GeV. $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-2j}}$, $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-1j}}$, $\mathcal{D}^{\mathrm{ME}}_{\mathrm{WH-hadr.}}$, and $\mathcal{D}^{\mathrm{ME}}_{\mathrm{ZH-hadr.}}$ are shown in that order, and the latter two are rescaled for visualization purposes. Points represent the data, with statistical uncertainties shown as vertical bars, and stacked histograms represent expected distributions. The H(125) signal and the ZZ backgrounds are normalized to the SM expectation, the Z+X background to the estimation from data. The vertical grey dashed lines denote the operating points used in the event categorization.

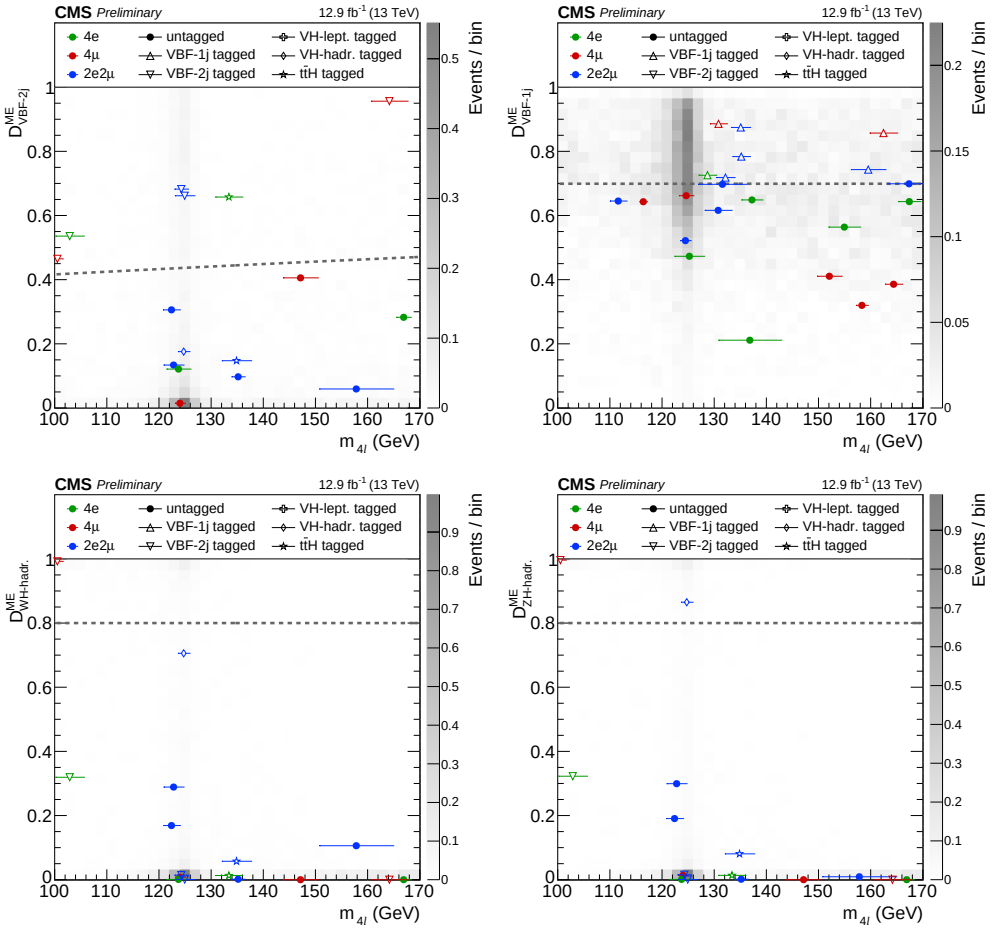

Figure 7.9: Distribution of the four production discriminants used for event categorization versus the four-lepton reconstructed mass $m_{4\ell}$ in the low-mass region. $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-2j}}$, $\mathcal{D}^{\mathrm{ME}}_{\mathrm{VBF-1j}}$, $\mathcal{D}^{\mathrm{ME}}_{\mathrm{WH-hadr.}}$, and $\mathcal{D}^{\mathrm{ME}}_{\mathrm{ZH-hadr.}}$ are shown in that order, and the latter two are rescaled for visualization purposes. The grey scale represents the expected relative density of ZZ background plus H(125) signal. The points show the data and the horizontal bars represent the measured per-event mass uncertainties. The grey dashed lines denote the operating points used in the event categorization.

7.5 Measurements

7.5.1 Significance

As in Section 6.5.1, the local p-value and significance relative to the SM background-only expectation are presented in Fig. 7.10 (left) as a function of m_{H}, for the range $118 < m_{\mathrm{H}} < 130\,\mathrm{GeV}$, with their breakdown into final states. With the full six-category scheme, 18 channels are now used in total. The minimum of the p-value is reached at 124.3 GeV and corresponds to a local significance of 6.4σ, for an expectation of 6.3σ for the SM Higgs boson. This means that a standalone rediscovery is achieved in 13 TeV data. At the Run I measured mass $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$, the observed significance is 6.2σ, where 6.5σ are expected. As expected, the 2e2μ and 4μ channels contribute most to the excess.

Figure 7.10 (right) shows that this time, the two-dimensional model consistently leads to larger observed significances than the one-dimensional one. Indeed, with the larger data sample, the $\mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}}$ distribution near the peak now better matches the expectation than in 2015 data, albeit still with a moderate deficit at large values.

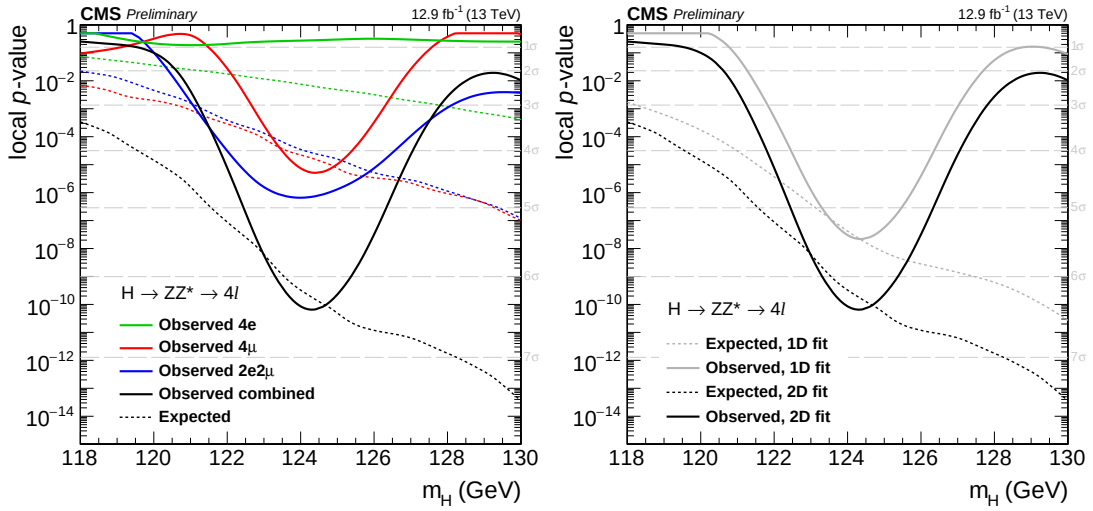

Figure 7.10: p-value and significance of the local fluctuation with respect to the SM expectation as a function of the Higgs boson mass. Dashed lines show the median expected local p-values for the local mass hypothesis. (left) Main result using $m_{4\ell}$ and $\mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}}$, (right) comparison to the result using $m_{4\ell}$ only.

7.5.2 Signal strength and production modes

The combined signal strength for a mass hypothesis of $m_{\mathrm{H}} = 125.09\,\mathrm{GeV}$ is measured to be $\mu = \sigma/\sigma_{\mathrm{SM}} = 0.99^{+0.33}_{-0.26}$. As in Section 6.5.2, the signal model is modified to measure separate signal strengths in every event category, but the fit is still simultaneous. The six per-category results are shown in Fig. 7.11 and compared to the combined result. All the parameters are again constrained to non-negative values. Although the signal strength in the VBF-1jet-tagged category is measured to be $0.00^{+0.16}_{-0.00}$, it was checked to be compatible with 1 at the 2.2σ level. This behaviour is related to the fact the event deficit in this category tends to pull its best-fit parameter value toward unphysical negative values, and the non-negative constraint makes the 68% CL interval small. In all

other categories, the observed values are fully consistent with 1 within the uncertainties. The constraint is very loose in the VH-leptonic-tagged and $t\bar{t}H$-tagged categories, which do not have much sensitivity yet: in both cases, 0.3 event is expected within $118 < m_H < 130\,\text{GeV}$ and none is observed.

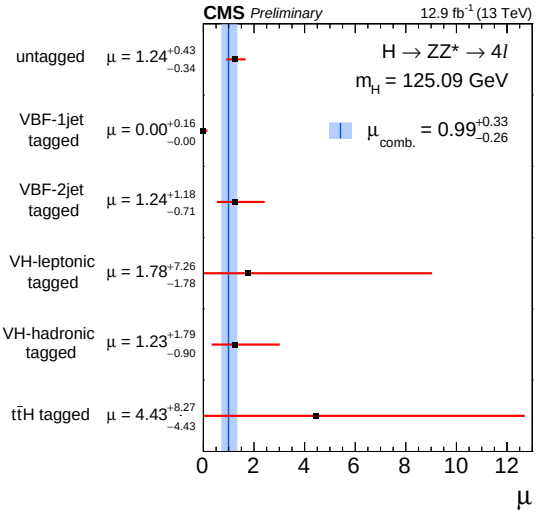

Figure 7.11: Observed values of the signal strength $\mu = \sigma/\sigma_{SM}$ for the six event categories, compared to the combined one shown as a vertical line. The horizontal bars and the filled band indicate the $\pm 1\sigma$ uncertainties. The uncertainties include both statistical and systematic sources.

The Higgs boson production mechanisms are first probed with the same two-parameter model as introduced in Section 6.5.2. The two-dimensional fit for $\mu_{\text{VBF,VH}}$ and $\mu_{\text{ggH,t}\bar{\text{t}}\text{H}}$ at a mass of $m_H = 125.09\,\text{GeV}$ leads to the measurements of $\mu_{\text{ggH,t}\bar{\text{t}}\text{H}} = 1.00^{+0.39}_{-0.32}$ and $\mu_{\text{VBF,VH}} = 0.91^{+1.56}_{-0.91}$. The new 68% and 95% CL contours in the $(\mu_{\text{ggH,t}\bar{\text{t}}\text{H}}, \mu_{\text{VBF,VH}})$ plane are shown in Fig. 7.12 (left), bounding a much narrower area, and highlighting the very good compatibility of the measurement with the SM expectation (1,1).

A second, more precise signal model is also introduced in order to control the cross sections of all considered Higgs boson production modes individually. This time, four signal strength modifiers are defined: μ_{ggH}, μ_{VBF}, μ_{VH}, and $\mu_{t\bar{t}H}$, each of which scales the production cross section of the process named in its subscript. One single μ_{VH} parameter is used to control both WH and ZH, because the analysis still has little sensitivity to these two modes separately. The model therefore assumes that the ratio of HWW to HZZ coupling strengths is as predicted by the SM, an assumption that was already made in considering W boson fusion and Z boson fusion together in the VBF process. The results of the likelihood scans for μ_{ggH}, μ_{VBF}, μ_{VH} and $\mu_{t\bar{t}H}$ are reported in Fig. 7.12 (right) and compared to the combined signal strength. Although the constraints on VH and $t\bar{t}H$ are still moderate with the data collected so far, they constitute the first step to achieving standalone observations of subdominant production modes in the $H \to ZZ \to 4\ell$ channel, and to other precision measurements that will exploit them by the end of Run II. Beyond this particular analysis, these results will also contribute to the upcoming CMS-wide combinations with other Higgs boson decay channels, whereby the couplings will be constrained with a much greater precision.

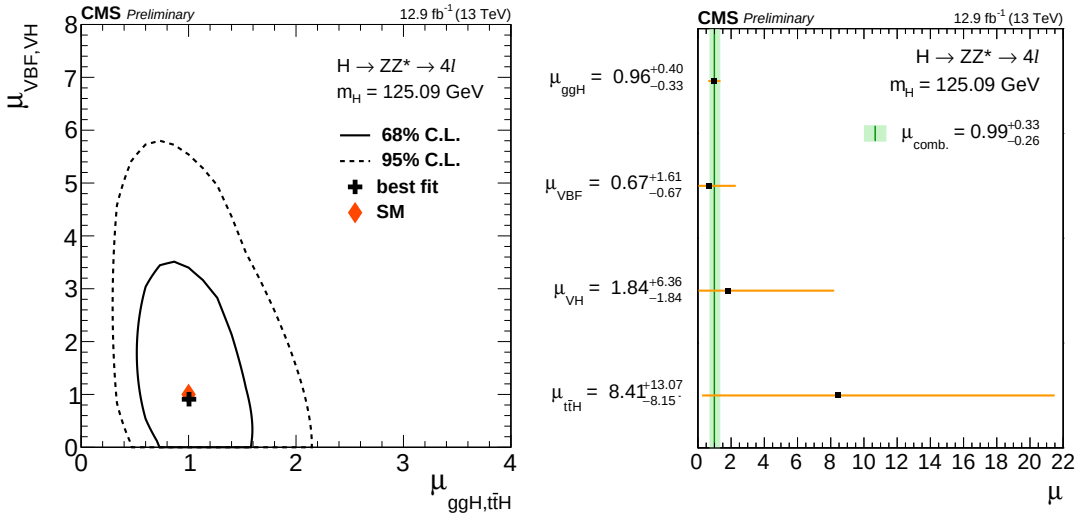

Figure 7.12: (left) Result of the 2D likelihood scan for the $\mu_{\mathrm{ggH,t\bar{t}H}}$ and $\mu_{\mathrm{VBF,VH}}$ signal-strength modifiers. The solid and dashed contours show the 68% and 95% CL regions, respectively. The cross indicates the best-fit values, and the diamond represents the expected values for the SM Higgs boson. (right) Results of likelihood scans versus the four signal strength modifiers corresponding to the main Higgs boson production modes, compared to the combined signal strength shown as a vertical line. The horizontal bars and the filled band indicate the $\pm 1\sigma$ uncertainties. The uncertainties include both statistical and systematic sources.

7.5.3 Mass measurement

Besides the $m_{4\ell}$ and $\mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}}$ variables, the measurement of the mass of the rediscovered Higgs boson exploits the additional information from per-event relative mass uncertainties $\mathcal{D}_{\mathrm{m}}$ (see Section 5.3.3). The $\mathrm{H}\to 4\ell$ analysis is particularly suited to such a strategy, owing to the low number of events, and to the use of two flavours of leptons resulting in a wide spread of per-event uncertainties.

Scans of the negative log-likelihood versus m_{H} are performed profiling the signal strength μ along with the nuisance parameters. These scans are shown in Fig. 7.13 (left), for the combination of the three final states, comparing the outcome of the 1D, 2D, and 3D models described in Section 5.7.2. The full expected and observed values are reported in Table 7.9 with their total uncertainties.

Table 7.9: Expected and observed results of the measurement of the Higgs boson mass for the combination of the 4e, 4μ, and $2\mathrm{e}2\mu$ final states, for the 1D, 2D, and 3D fit models. The expected results use a generated mass of 125 GeV. The quoted uncertainties include statistical and systematic sources.

Fit model	Expected result for $m_{\mathrm{H}}^{\mathrm{gen.}} = 125\,\mathrm{GeV}$	Observed result
$f_{\mathrm{1D}}^{m}(m_{4\ell})$	$125.01 \pm 0.62\,\mathrm{GeV}$	$124.31^{+0.46}_{-0.45}\,\mathrm{GeV}$
$f_{\mathrm{2D}}^{m}(m_{4\ell}, \mathcal{D}_{\mathrm{m}})$	$125.00 \pm 0.56\,\mathrm{GeV}$	$124.52^{+0.47}_{-0.47}\,\mathrm{GeV}$
$f_{\mathrm{3D}}^{m}(m_{4\ell}, \mathcal{D}_{\mathrm{m}}, \mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}})$	$124.98 \pm 0.54\,\mathrm{GeV}$	$124.50^{+0.48}_{-0.46}\,\mathrm{GeV}$

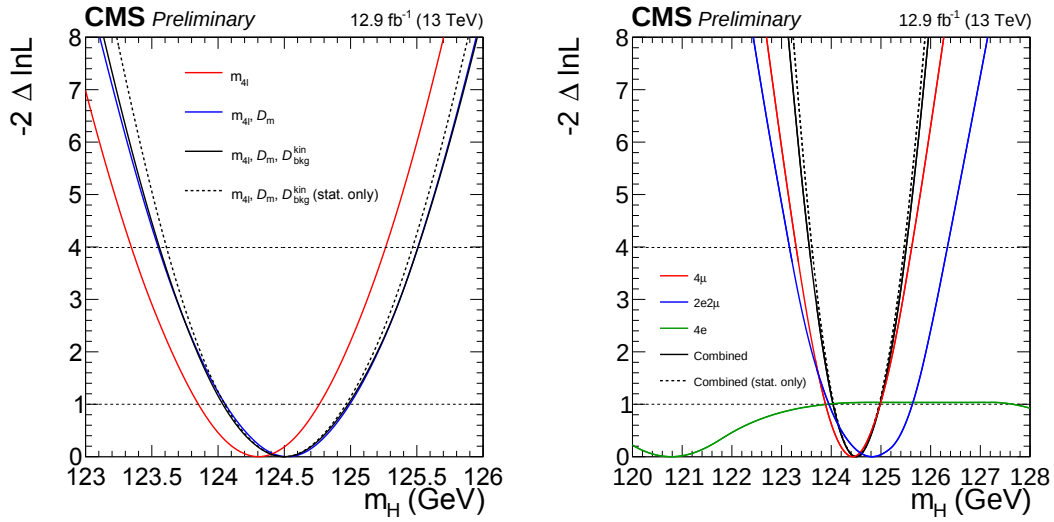

Figure 7.13: Scans of the negative log-likelihood versus m_{H}, comparing the 1D, 2D, and 3D measurement (left), and showing the 3D measurement for the 3 final states and their combination (right). The solid lines represent scans with the full uncertainties from statistical and systematic sources, while the dashed lines include only statistical uncertainties.

Including the per-event mass uncertainties (i.e. moving from 1D to 2D) brings an improvement of about 8% to the expected relative uncertainty of the measurement, and including $\mathcal{D}_{\mathrm{bkg}}^{\mathrm{kin}}$ (i.e. moving to 3D) improves it further by about 3%. The observed value is lower with the 1D fit than with the 2D and 3D ones, which can be understood from Fig. 7.7 (bottom left): the events observed on the left-hand side of the Higgs boson peak tend to have slightly larger per-event mass uncertainties, and get therefore downweighted when $\mathcal{D}_{\mathrm{m}}$ is used.

Figure 7.13 (right) presents the likelihood scan separately for the three final states, showing that they are statistically compatible. The systematic uncertainty is dominated by the limited knowledge of the lepton momentum scale, which has an uncertainty of 0.04% in the 4μ channel and 0.3% in the 4e channel. As a result, the sensitivity is driven by the 4μ channel, which also has the best selection efficiency and $m_{4\ell}$ resolution. Furthermore, the 4e channel has a slight downward fluctuation of on-peak yields in 2016 data (see Table 7.7), and its observed events do not form a very sharp peak (see Figs. 7.4, left, and 7.7, bottom left), which explains the peculiar shape of the 4e likelihood.

To conclude, the final measurement of the Higgs boson mass is the one provided by the 3D fit:

$$m_{\mathrm{H}} = 124.50^{+0.48}_{-0.46}(\mathrm{tot.})\,\mathrm{GeV} = 124.50^{+0.47}_{-0.45}(\mathrm{stat.})^{+0.13}_{-0.11}(\mathrm{syst.})\,\mathrm{GeV}\ , \tag{7.1}$$

where the total uncertainty is split into statistical and systematic contributions[2]. This result is well compatible with the Run I LHC combined measurement of the Higgs boson mass [27] ($m_{\mathrm{H}} = 125.09 \pm 0.21(\mathrm{stat.}) \pm 0.11(\mathrm{syst.})\,\mathrm{GeV}$) and with the standalone Run I

[2]The statistical contribution to the total uncertainty is estimated by performing the likelihood scan neglecting all systematic sources. The signal strength is still profiled, and its uncertainty is included in the statistical component. The systematic uncertainty is finally computed as the quadrature difference between the total and statistical one.

CMS measurement in the $H \to 4\ell$ channel [23] ($m_H = 125.6 \pm 0.4(\text{stat.}) \pm 0.2(\text{syst.})$ GeV), albeit with a downward fluctuation with respect to them.

7.5.4 Width measurement

Measuring the decay width Γ_H of the discovered Higgs boson is key to determining its consistency with the SM, in particular to probing possible couplings to new undiscovered particles. Γ_H can be directly measured at the resonance peak in the $H \to 4\ell$ and $H \to \gamma\gamma$ channels, as was reported by CMS in Ref. [23, 26] for the Run I data set. However, such a measurement is limited by the experimental resolution (of the order of the GeV), and is thus only sensitive to values at least three orders of magnitude larger than the expected width of around 4 MeV for a SM Higgs boson of mass near 125 GeV.

It was then proposed that much better precision can be achieved exploiting the off-shell Higgs boson production away from its resonance peak [105, 106, 84]. In the ggH production mode, the $H \to ZZ$ decay amplitude is enhanced as of the vicinity of the Z-boson pair production threshold, with the region $m_{4\ell} > 2m_Z$ representing about a tenth of the total cross section. It was shown that a measurement of the relative off-shell and on-shell production in the $H \to ZZ$ channel gives access to Γ_H, under the assumption that the coupling ratios are unchanged, or, in other words, that no BSM particles or interactions modify the Higgs boson couplings either in production or in decay. This measurement strategy was first exploited by CMS with its Run I data set: the Higgs boson width was measured in the $H \to ZZ$ channel using the 4ℓ and $2\ell 2\nu$ final states [35], a result which was then combined with a similar measurement from the $H \to WW$ channel, providing a combined observed (expected) upper limit of $\Gamma_H < 13$ MeV (26 MeV) at a 95% CL [107]. A generalized analysis of the 4ℓ channel was also done including anomalous couplings of the Higgs boson to two electroweak bosons via the parameter $f_{\Lambda Q}$ [37].

The first such measurement of the width Γ_H in 13 TeV data is presented here, exploiting the invariant mass distribution throughout the range $105 < m_{4\ell} < 1600$ GeV. Besides Γ_H, the chosen signal model has three other parameters: the Higgs boson mass m_H, which is either simultaneously measured, or left floating in the fit, and the strengths $\mu_{ggH,t\bar{t}H}$ and $\mu_{VBF,VH}$ of fermion- and vector-boson-induced production mechanisms, which are profiled independently. These can be constrained thanks to the splitting of selected events into a VBF-2jet-tagged and an Untagged category, as mentioned in Section 5.7.2. The modelling of the pdf $\mathcal{P}(m_{4\ell} \,|\, m_H, \Gamma_H, \mu_{VBF,VH}, \mu_{ggH,t\bar{t}H})$ for the $gg/qq' \to 4\ell$ processes relies on the general parameterization described in Section 5.6.2. Interference effects between signal and background are particularly relevant in the off-shell region, but they are also in the on-shell region in the case where the width is as large as the experimental resolution.

It is also instructive to still perform a direct measurement of Γ_H, leaving out the off-shell region and the aforementioned assumptions on BSM contributions. This is done here by limiting the exploited mass range to $105 < m_{4\ell} < 140$ GeV, thus only relying on the shape of the $m_{4\ell}$ distribution at the Higgs boson peak.

Figure 7.14 illustrates the simultaneous measurement of the Higgs boson width Γ_H and mass m_H for the two types of fits: either using the full mass range or the on-shell region only. Since the measurement of the Higgs boson mass is only driven by the events of the 125 GeV peak, one can see that the precision on m_H is similar in both cases. The

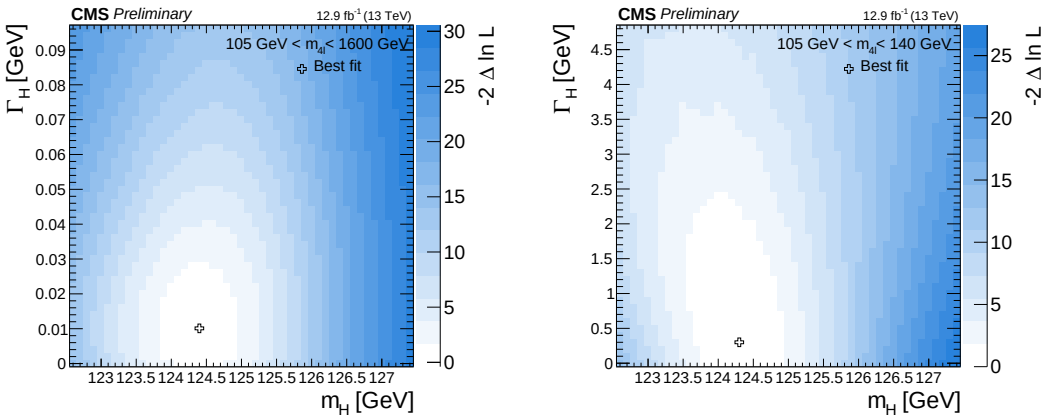

Figure 7.14: Scans of the negative log-likelihood versus m_{H} and Γ_{H}, using the full mass range $105 < m_{4\ell} < 1600\,\mathrm{GeV}$ (left) or the on-shell region $105 < m_{4\ell} < 140\,\mathrm{GeV}$ (right).

precision on Γ_{H}, on the other hand, improves by two orders of magnitude when the off-shell region is included.

Figure 7.15 shows 1D likelihood scans versus Γ_{H} where m_{H} is floated, also using either the full or on-shell-only range. Numerical values for the corresponding observed and expected results are summarized in Table 7.10. The observed (expected) limit on the width at a 95% CL is $\Gamma_{\mathrm{H}} < 41\,\mathrm{MeV}$ (32 MeV). The expected sensitivity is thus slightly better than in the CMS Run I 4ℓ-only result, which was $\Gamma_{\mathrm{H}} < 33\,\mathrm{MeV}$ (42 MeV) [35], although the latter measurement was also exploiting a decay kinematic discriminant, which the present Run II result does not use yet.

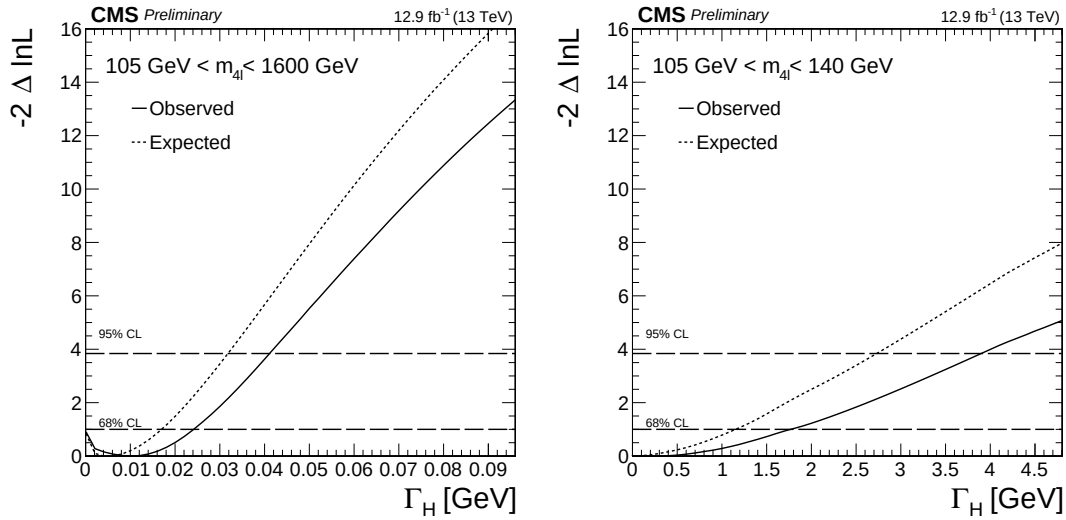

Figure 7.15: Scans of the observed and expected negative log-likelihood versus Γ_{H}, with m_{H} floated, using the full mass range $105 < m_{4\ell} < 1600\,\mathrm{GeV}$ (left) or the on-shell region $105 < m_{4\ell} < 140\,\mathrm{GeV}$ (right).

Table 7.10: Measurement results of the width Γ_H of the H(125) boson, with the best fit values and 68% CL intervals (quoted as uncertainties), and the 95% CL intervals (quoted as ranges in square brackets). The expected results are obtained for the SM signal production cross sections ($\mu_{VBF,VH} = \mu_{ggH,t\bar{t}H} = 1$), and for $m_H = 125\,\text{GeV}$ and $\Gamma_H = \Gamma_H^{SM}(125\,\text{GeV}) = 0.0041\,\text{GeV}$. All results are quoted in GeV.

mass range	Expected results	Observed results
$105 < m_{4\ell} < 1600\,\text{GeV}$	$0.004^{+0.013}_{-0.004}$ $[0.000, 0.032]$	$0.010^{+0.014}_{-0.010}$ $[0.000, 0.041]$
$105 < m_{4\ell} < 140\,\text{GeV}$	$0.0^{+1.1}_{-0.0}$ $[0.0, 2.7]$	$0.3^{+1.4}_{-0.0}$ $[0.0, 3.9]$

7.5.5 High-mass search

The $ZZ \to 4\ell$ channel is well suited to searches for new bosonic resonances, thanks to its excellent experimental resolution and to such resonances often having large decay branching fractions to ZZ at high mass. It was already exploited in Run I in the CMS search for a high-mass Higgs boson in the $145 < m_{4\ell} < 1000\,\text{GeV}$ range, together with other $H \to ZZ$ and $H \to WW$ channels. The results were then interpreted both for a heavy SM-like Higgs boson and for an electroweak-singlet extension of the SM [41].

Now that the increase of the centre-of-mass energy to 13 TeV gives better access to hypothetical high-mass particles, a more general search is performed here for a scalar resonance X decaying to $ZZ \to 4\ell$, with a resonance mass $130 < m_X < 2500\,\text{GeV}$, exploiting the invariant mass distribution throughout the range $100 < m_{4\ell} < 3000\,\text{GeV}$. Any value of the decay width Γ_X of the resonance is allowed, from the narrow-width approximation (denoted as $\Gamma_X = 0$) to an arbitrarily large width (with $\Gamma_X < m_X$). Gluon fusion, vector boson fusion and associated VX production are considered. VX is tied to VBF according to expectations of their relative cross sections, and the fraction of VBF+VX production is denoted as f_{VBF}, so that the fraction of gluon fusion is $1 - f_{VBF}$. The modelling of the pdf $\mathcal{P}(\, m_{4\ell} \,|\, m_X,\, \Gamma_X,\, \sigma_X\,)$ for the $gg/qq' \to 4\ell$ processes relies on the general parameterization described in Section 5.6.2, including the SM backgrounds, $X(m_X)$ signal, off-shell tail of H(125) and interference between various components, all treated as one process for either gluon fusion or VBF + VX.

The results take the form of upper limits on the resonance production cross section σ_X as a function of m_X and Γ_X, whereby f_{VBF} can be either fixed to a certain value or left unconstrained in the fit. Constraints on f_{VBF} are possible because the events are split into a VBF-2jet-tagged and Untagged category, as explained in Section 5.7.2. The cross section σ_X is reported for the pure signal process $pp \to X \to ZZ \to 4\ell$ excluding the interference. For a wide resonance with $m_X > 2m_Z$, σ_X includes the off-shell production of $X^* \to ZZ$ above this $2m_Z$ threshold. The systematic uncertainties are similar to those used in the other measurements.

Figure 7.16 shows the observed and expected 95% CL exclusion limits on σ_X as a function of m_X, for a few values of Γ_X, either profiling the fraction f_{VBF}, or fixing it to 1. No excess above 2σ is observed throughout the studied mass range in either case. The shape of the observed limits can be linked to the fluctuations of the $m_{4\ell}$ distribution shown in Fig. 7.3. As expected, the smaller the hypothesized width, the

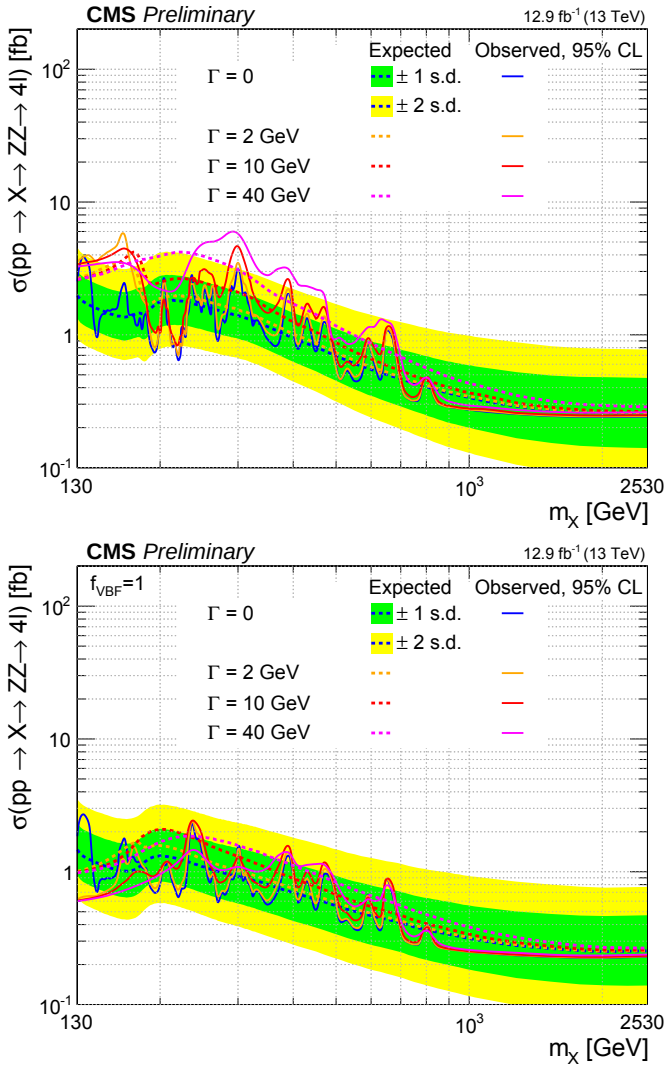

Figure 7.16: Expected and observed upper limits at a 95% CL on the pp $\to$ X $\to$ ZZ $\to 4\ell$ cross section σ_{X} as a function of the resonance mass m_{X}, shown for several values of its decay width Γ_{X}, where the fraction f_{VBF} of VBF and VH production is either floated (top) or fixed to 1 (bottom).

more stringent the constraints on σ_{X}. Constraints are also tighter under the assumption of pure electroweak production ($f_{\mathrm{VBF}} = 1$) because of the larger signal-to-background ratio in the VBF-2jet-tagged event category.

7.6 Conclusion

Increasingly precise Higgs boson results in the four-lepton final state have been produced with a first subset of 2016 data, using a pp collision data sample corresponding to an integrated luminosity of 12.9 fb^{-1}, i.e. 4.6 times the luminosity used in the 2015 analysis.

The Higgs boson discovered in Run I is now rediscovered at $\sqrt{s} = 13$ TeV, with an observed (expected) significance of 6.2σ (6.5σ) at $m_{\mathrm{H}} = 125.09$ GeV. The measured signal strength is $\mu = 0.99^{+0.33}_{-0.26}$, and the measured signal strength modifiers associated with

fermions and vector bosons are $\mu_{\mathrm{ggH,t\bar{t}H}} = 1.00^{+0.39}_{-0.32}$ and $\mu_{\mathrm{VBF,VH}} = 0.91^{+1.56}_{-0.91}$, respectively. Four signal strength modifiers associated to individual production modes are also measured; in particular, the VH and $\mathrm{t\bar{t}H}$ mechanisms are probed independently for the first time in the four-lepton final state. The Higgs boson mass is measured to be $m_{\mathrm{H}} = 124.50^{+0.48}_{-0.46}\,\mathrm{GeV}$, while its width is constrained to $\Gamma_{\mathrm{H}} < 41\,\mathrm{MeV}$ at a 95% CL. All results are consistent with the expectations for the SM Higgs boson within their uncertainties, which are dominated by the statistical component with the current data set.

A search for an additional scalar resonance is carried out for masses up to 2.5 TeV and for various widths; no significant excess is observed and upper limits at a 95% CL are set on the resonance production cross section under various assumptions.

Conclusions

After a fruitful Run I period rewarded by the discovery of a standard model Higgs boson like particle and first property measurements, Run II of the LHC is now unfolding, with vast prospects for direct and indirect searches for physics beyond the SM. From the trigger level to the statistical interpretation, every item of the CMS $\mathrm{H} \to \mathrm{ZZ} \to 4\ell$ analysis has now been redesigned and optimized to extract this rare signal. A new event categorization has been introduced to distinguish production-specific topologies, with an improved selection for vector boson fusion and the first study of associated VH and associated $\mathrm{t\bar{t}H}$ production with this decay channel.

Two sets of CMS preliminary four-lepton-based results at $\sqrt{s} = 13\,\mathrm{TeV}$ have already been released, respectively based on 2015 and early 2016 data. The Higgs boson is rediscovered at the new collision energy, and the first hints for vector boson fusion and associated VH production begin to show up, while no other significant excess is observed in the remainder of the $m_{4\ell}$ spectrum. All measurements performed so far show that the properties of the boson are still consistent, within the current precision, with the expectations for a standard model with a minimal scalar sector, i.e. with only one physical H boson. The analysis performed for exclusive production modes will allow for a better sensitivity to couplings in the future combination of all Higgs boson production and decay channels.

While this thesis work is coming to a close, the exploration of 13 TeV data is only just beginning. At the time when the last presented results were frozen, the LHC had delivered one third of the total luminosity foreseen for 2016, and about one tenth of that foreseen for the remainder of Run II, which will last until 2018. In a context when searches for supersymmetry are rendered more difficult by the increasing constraints on its phase space, the thirst of the particle physics community for signs of new physics is huge. While the episodic hints for new resonances attract much scrutiny, precision tests in the just discovered scalar sector constitute a privileged portal to new physics. Four-lepton-based analyses will keep exploring both of these directions in the upcoming years, with the quest for deviations of Higgs boson properties from the standard model, and the search for new resonances at high mass.

Bibliography

[1] S. L. Glashow, "Partial Symmetries of Weak Interactions", *Nucl. Phys.* **22** (1961) 579–588, doi:10.1016/0029-5582(61)90469-2.

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

[3] A. Salam, "Weak and electromagnetic interactions", in *Elementary particle physics: relativistic groups and analyticity*, N. Svartholm, ed., p. 367. Almqvist & Wiksell, 1968. Proceedings of the eighth Nobel symposium.

[4] F. Englert and R. Brout, "Broken Symmetry and the Mass of Gauge Vector Mesons", *Phys. Rev. Lett.* **13** (1964) 321, doi:10.1103/PhysRevLett.13.321.

[5] P. W. Higgs, "Broken Symmetries and the Masses of Gauge Bosons", *Phys. Rev. Lett.* **13** (1964) 508, doi:10.1103/PhysRevLett.13.508.

[6] G. Guralnik, C. Hagen, and T. Kibble, "Global Conservation Laws and Massless Particles", *Phys. Rev. Lett.* **13** (1964) 585, doi:10.1103/PhysRevLett.13.585.

[7] J. Goldstone, "Field Theories with Superconductor Solutions", *Nuovo Cim.* **19** (1961) 154–164, doi:10.1007/BF02812722.

[8] A. Djouadi, "The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the standard model", *Phys. Rept.* **457** (2008) 1–216, doi:10.1016/j.physrep.2007.10.004, arXiv:hep-ph/0503172.

[9] C. Anastasiou et al., "High precision determination of the gluon fusion Higgs boson cross-section at the LHC", *JHEP* **05** (2016) 058, doi:10.1007/JHEP05(2016)058, arXiv:1602.00695.

[10] LHC Higgs Cross Section Working Group, "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector", 2016. arXiv:1610.07922.

[11] LHC Higgs Cross Section Working Group, "Handbook of LHC Higgs Cross Sections: 3. Higgs Properties", 2013. arXiv:1307.1347.

[12] G. Ridolfi, "Search for the Higgs boson: Theoretical perspectives", in *Results and perspectives in particle physics. Proceedings, 15th Rencontres de physique de la Vallée d'Aoste, La Thuile, Italy, March 4-10, 2001*, pp. 291–304. arXiv:hep-ph/0106300.

[13] ALEPH, DELPHI, L3, OPAL Collaborations, and LEP Working Group for Higgs Boson Searches, "Search for the standard model Higgs boson at LEP", *Phys. Lett. B* **565** (2003) 61–75, doi:10.1016/S0370-2693(03)00614-2, arXiv:hep-ex/0306033.

[14] CDF and D0 Collaborations, "Combination of Tevatron searches for the standard model Higgs boson in the W+W- decay mode", *Phys. Rev. Lett.* **104** (2010) 061802, doi:10.1103/PhysRevLett.104.061802, arXiv:1001.4162.

[15] ALEPH, CDF, D0, DELPHI, L3, OPAL, SLD Collaborations, the LEP Electroweak Working Group, the Tevatron Electroweak Working Group, and the SLD Electroweak and Heavy Flavour Groups, "Precision Electroweak Measurements and Constraints on the Standard Model", CERN PH-EP-2010-095, 2010. arXiv:1012.2367.

[16] CMS Collaboration, "Combined results of searches for the standard model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV", *Phys. Lett. B* **710** (2012) 26–48, doi:10.1016/j.physletb.2012.02.064, arXiv:1202.1488.

[17] ATLAS Collaboration, "Combined search for the Standard Model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector", *Phys. Rev. D* **86** (2012) 032003, doi:10.1103/PhysRevD.86.032003, arXiv:1207.0319.

[18] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC", *Phys. Lett. B* **716** (2012) 1–29, doi:10.1016/j.physletb.2012.08.020, arXiv:1207.7214.

[19] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC", *Phys. Lett. B* **716** (2012) 30–61, doi:10.1016/j.physletb.2012.08.021, arXiv:1207.7235.

[20] CMS Collaboration, "Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **06** (2013) 081, doi:10.1007/JHEP06(2013)081, arXiv:1303.4571.

[21] CDF and D0 Collaborations, "Evidence for a particle produced in association with weak bosons and decaying to a bottom-antibottom quark pair in Higgs boson searches at the Tevatron", *Phys. Rev. Lett.* **109** (2012) 071804, doi:10.1103/PhysRevLett.109.071804, arXiv:1207.6436.

[22] CDF and D0 Collaborations, "Higgs Boson Studies at the Tevatron", *Phys. Rev. D* **88** (2013), no. 5, 052014, doi:10.1103/PhysRevD.88.052014, arXiv:1303.6346.

[23] CMS Collaboration, "Measurement of the properties of a Higgs boson in the four-lepton final state", *Phys. Rev. D* **89** (2014), no. 9, 092007, doi:10.1103/PhysRevD.89.092007, arXiv:1312.5353.

[24] ATLAS Collaboration, "Measurement of the Higgs boson mass from the $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ channels with the ATLAS detector using 25 fb^{-1} of pp collision data", *Phys. Rev. D* **90** (2014), no. 5, 052004, doi:10.1103/PhysRevD.90.052004, arXiv:1406.3827.

[25] CMS Collaboration, "Observation of the diphoton decay of the Higgs boson and measurement of its properties", *Eur. Phys. J. C* **74** (2014), no. 10, 3076, doi:10.1140/epjc/s10052-014-3076-z, arXiv:1407.0558.

[26] CMS Collaboration, "Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the standard model predictions using proton collisions at 7 and 8 TeV", *Eur. Phys. J. C* **75** (2015), no. 5, 212, doi:10.1140/epjc/s10052-015-3351-7, arXiv:1412.8662.

[27] ATLAS and CMS Collaborations, "Combined Measurement of the Higgs Boson Mass in *pp* Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments", *Phys. Rev. Lett.* **114** (2015) 191803, doi:10.1103/PhysRevLett.114.191803, arXiv:1503.07589.

[28] CMS Collaboration, "Study of the Mass and Spin-Parity of the Higgs Boson Candidate Via Its Decays to Z Boson Pairs", *Phys. Rev. Lett.* **110** (2013), no. 8, 081803, doi:10.1103/PhysRevLett.110.081803, arXiv:1212.6639.

[29] ATLAS Collaboration, "Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s} = 7$ and 8 TeV in the ATLAS experiment", *Eur. Phys. J. C* **76** (2016), no. 1, 6, doi:10.1140/epjc/s10052-015-3769-y, arXiv:1507.04548.

[30] ATLAS and CMS Collaborations, "Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **08** (2016) 045, doi:10.1007/JHEP08(2016)045, arXiv:1606.02266.

[31] ATLAS Collaboration, "Measurements of Higgs boson production and couplings in the four-lepton channel in pp collisions at center-of-mass energies of 7 and 8 TeV with the ATLAS detector", *Phys. Rev. D* **91** (2015), no. 1, 012006, doi:10.1103/PhysRevD.91.012006, arXiv:1408.5191.

[32] ATLAS Collaboration, "Evidence for the spin-0 nature of the Higgs boson using ATLAS data", *Phys. Lett. B* **726** (2013) 120–144, doi:10.1016/j.physletb.2013.08.026, arXiv:1307.1432.

[33] CMS Collaboration, "Constraints on the spin-parity and anomalous HVV couplings of the Higgs boson in proton collisions at 7 and 8 TeV", *Phys. Rev. D* **92** (2015), no. 1, 012004, doi:10.1103/PhysRevD.92.012004, arXiv:1411.3441.

[34] ATLAS Collaboration, "Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector", *Eur. Phys. J. C* **75** (2015), no. 10, 476, doi:10.1140/epjc/s10052-015-3685-1, arXiv:1506.05669. [Erratum: Eur. Phys. J. C 76, no.3, 152 (2016)].

[35] CMS Collaboration, "Constraints on the Higgs boson width from off-shell production and decay to Z-boson pairs", *Phys. Lett. B* **736** (2014) 64–85, doi:10.1016/j.physletb.2014.06.077, arXiv:1405.3455.

[36] ATLAS Collaboration, "Constraints on the off-shell Higgs boson signal strength in the high-mass ZZ and WW final states with the ATLAS detector", *Eur. Phys. J. C* **75** (2015), no. 7, 335, doi:10.1140/epjc/s10052-015-3542-2, arXiv:1503.01060.

[37] CMS Collaboration, "Limits on the Higgs boson lifetime and width from its decay to four charged leptons", *Phys. Rev. D* **92** (2015), no. 7, 072010, doi:10.1103/PhysRevD.92.072010, arXiv:1507.06656.

[38] ATLAS Collaboration, "Fiducial and differential cross sections of Higgs boson production measured in the four-lepton decay channel in *pp* collisions at $\sqrt{s}$=8 TeV with the ATLAS detector", *Phys. Lett. B* **738** (2014) 234–253, doi:10.1016/j.physletb.2014.09.054, arXiv:1408.3226.

[39] CMS Collaboration, "Measurement of differential and integrated fiducial cross sections for Higgs boson production in the four-lepton decay channel in pp collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **04** (2016) 005, doi:10.1007/JHEP04(2016)005, arXiv:1512.08377.

[40] ATLAS Collaboration, "Search for an additional, heavy Higgs boson in the $H \to ZZ$ decay channel at $\sqrt{s} = 8$ TeV in *pp* collision data with the ATLAS detector", *Eur. Phys. J. C* **76** (2016), no. 1, 45, doi:10.1140/epjc/s10052-015-3820-z, arXiv:1507.05930.

[41] CMS Collaboration, "Search for a Higgs boson in the mass range from 145 to 1000 GeV decaying to a pair of W or Z Bosons", *JHEP* **10** (2015) 144, doi:10.1007/JHEP10(2015)144, arXiv:1504.00936.

[42] The LHC Study Group, "The Large Hadron Collider: Accelerator project", Technical Report CERN-AC-93-03-LHC, 1993.

[43] The LHC Study Group, "The Large Hadron Collider: Conceptual design", Technical Report CERN-AC-95-05-LHC, 1995.

[44] L. Evans and P. Bryant, "LHC Machine", *JINST* **3** (2008) S08001, doi:10.1088/1748-0221/3/08/S08001.

[45] W. J. Stirling. Private communication, 2016.

[46] C. De Melis, "The CERN accelerator complex. Complexe des accélérateurs du CERN". OPEN-PHO-ACCEL-2016-009, 2016.

[47] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.

[48] P. Adzic et al., "Energy resolution of the barrel of the CMS electromagnetic calorimeter", *JINST* **2** (2007) P04004, doi:10.1088/1748-0221/2/04/P04004.

[49] CMS Collaboration, "Performance of CMS muon reconstruction in *pp* collision events at $\sqrt{s} = 7$ TeV", *JINST* **7** (2012) P10002, doi:10.1088/1748-0221/7/10/P10002, arXiv:1206.4071.

[50] CMS Collaboration, "Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV", *JINST* **10** (2015) P06005, doi:10.1088/1748-0221/10/06/P06005, arXiv:1502.02701.

[51] W. Adam, R. Frühwirth, A. Strandlie, and T. Todorov, "Reconstruction of electrons with the Gaussian sum filter in the CMS tracker at LHC", *J. Phys. G* **31** (2005) N9, doi:10.1088/0954-3899/31/9/N01, arXiv:physics/0306087.

[52] S. Baffioni et al., "Electron reconstruction in CMS", *Eur. Phys. J. C* **49** (2007) 1099–1116, doi:10.1140/epjc/s10052-006-0175-5.

[53] P. Billoir, "Progressive track recognition with a Kalman like fitting procedure", *Comput. Phys. Commun.* **57** (1989) 390–394, doi:10.1016/0010-4655(89)90249-X.

[54] R. Fruhwirth, "Application of Kalman filtering to track and vertex fitting", *Nucl. Instrum. Meth.* **A262** (1987) 444–450, doi:10.1016/0168-9002(87)90887-4.

[55] CMS Collaboration, "Particle-flow Event Reconstruction in CMS and performance for jets, taus, and $E_{\mathrm{T}}^{\mathrm{miss}}$", CMS Physics Analysis Summary CMS-PAS-PFT-09-001, 2009.

[56] CMS Collaboration, "Commissioning of the Particle-flow Event Reconstruction with the first LHC collisions recorded in the CMS detector", CMS Physics Analysis Summary CMS-PAS-PFT-10-001, 2010.

[57] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014) P10009, doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569.

[58] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_t jet clustering algorithm", *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189.

[59] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual", *Eur. Phys. J. C* **72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097.

[60] CMS Collaboration, "Determination of jet energy calibration and transverse momentum resolution in CMS", *JINST* **6** (2011) P11002, doi:10.1088/1748-0221/6/11/P11002, arXiv:1107.4277.

[61] M. Cacciari and G. P. Salam, "Pileup subtraction using jet areas", *Phys. Lett. B* **659** (2008) 119, doi:10.1016/j.physletb.2007.09.077, arXiv:0707.1378.

[62] M. Cacciari, G. P. Salam, and G. Soyez, "The Catchment Area of Jets", *JHEP* **04** (2008) 005, doi:10.1088/1126-6708/2008/04/005, arXiv:0802.1188.

[63] CMS Collaboration, "The CMS trigger system", 2016. arXiv:1609.02366. Submitted to *JINST*.

[64] B. Kreis et al., "Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger", *JINST* **11** (2016), no. 01, C01051, doi:10.1088/1748-0221/11/01/C01051, arXiv:1511.05855.

[65] CMS Collaboration, "Particle-flow commissioning with muons and electrons from J/Psi and W events at 7 TeV", CMS Physics Analysis Summary CMS-PAS-PFT-10-003, 2010.

[66] CMS Collaboration, "W-like measurement of the Z boson mass using dimuon events collected in pp collisions at $\sqrt{s}$ =7 TeV", CMS Physics Analysis Summary CMS-PAS-SMP-14-007, 2016.

[67] CMS Collaboration, "Measurement of the Inclusive W and Z Production Cross Sections in pp Collisions at $\sqrt{s} = 7$ TeV", *JHEP* **10** (2011) 132, doi:10.1007/JHEP10(2011)132, arXiv:1107.4789.

[68] CMS Collaboration, "Identification of b quark jets at the CMS Experiment in the LHC Run 2", CMS Physics Analysis Summary CMS-PAS-BTV-15-001, 2016.

[69] R. D. Ball et al., "Parton distributions for the LHC run II", *Journal of High Energy Physics* **2015** (2015), no. 4, 1–148, doi:10.1007/JHEP04(2015)040.

[70] T. Sjöstrand et al., "An Introduction to PYTHIA 8.2", *Comput. Phys. Commun.* **191** (2015) 159–177, doi:10.1016/j.cpc.2015.01.024, arXiv:1410.3012.

[71] GEANT4 Collaboration, "GEANT4: a simulation toolkit", *Nucl. Instrum. Meth. A* **506** (2003) 250, doi:10.1016/S0168-9002(03)01368-8.

[72] J. Allison et al., "Geant4 developments and applications", *IEEE Trans. Nucl. Sci.* **53** (2006) 270, doi:10.1109/TNS.2006.869826.

[73] S. Alioli, P. Nason, C. Oleari, and E. Re, "NLO vector-boson production matched with shower in POWHEG", *JHEP* **07** (2008) 060, doi:10.1088/1126-6708/2008/07/060, arXiv:0805.4802.

[74] P. Nason, "A new method for combining NLO QCD with shower Monte Carlo algorithms", *JHEP* **11** (2004) 040, doi:10.1088/1126-6708/2004/11/040, arXiv:hep-ph/0409146.

[75] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with parton shower simulations: the POWHEG method", *JHEP* **11** (2007) 070, doi:10.1088/1126-6708/2007/11/070, arXiv:0709.2092.

[76] E. Bagnaschi, G. Degrassi, P. Slavich, and A. Vicini, "Higgs production via gluon fusion in the POWHEG approach in the SM and in the MSSM", *JHEP* **02** (2012) 088, doi:10.1007/JHEP02(2012)088, arXiv:1111.2854.

[77] P. Nason and C. Oleari, "NLO Higgs boson production via vector-boson fusion matched with shower in POWHEG", *JHEP* **02** (2010) 037, doi:10.1007/JHEP02(2010)037, arXiv:0911.5299.

[78] G. Luisoni, P. Nason, C. Oleari, and F. Tramontano, "$HW^{\pm}/HZ$ + 0 and 1 jet at NLO with the POWHEG BOX interfaced to GoSam and their merging within MiNLO", *JHEP* **10** (2013) 083, doi:10.1007/JHEP10(2013)083, arXiv:1306.2542.

[79] H. B. Hartanto, B. Jager, L. Reina, and D. Wackeroth, "Higgs boson production in association with top quarks in the POWHEG BOX", *Phys. Rev. D* **91** (2015), no. 9, 094003, doi:10.1103/PhysRevD.91.094003, arXiv:1501.04498.

[80] Y. Gao et al., "Spin determination of single-produced resonances at hadron colliders", *Phys. Rev. D* **81** (2010) 075022, doi:10.1103/PhysRevD.81.075022, arXiv:1001.3396.

[81] S. Bolognesi et al., "On the spin and parity of a single-produced resonance at the LHC", *Phys. Rev. D* **86** (2012) 095031, doi:10.1103/PhysRevD.86.095031, arXiv:1208.4018.

[82] P. Nason and G. Zanderighi, "W^+W^-, WZ and ZZ production in the POWHEG-BOX-V2", *Eur. Phys. J. C* **74** (2014), no. 1, 2702, doi:10.1140/epjc/s10052-013-2702-5, arXiv:1311.1365.

[83] J. M. Campbell and R. K. Ellis, "MCFM for the Tevatron and the LHC", *Nucl. Phys. Proc. Suppl.* **205** (2010) 10, doi:10.1016/j.nuclphysbps.2010.08.011, arXiv:1007.3492.

[84] J. M. Campbell, R. K. Ellis, and C. Williams, "Bounding the Higgs width at the LHC using full analytic results for $gg \to e^-e^+\mu^-\mu^+$", *JHEP* **04** (2014) 060, doi:10.1007/JHEP04(2014)060, arXiv:1311.3589.

[85] J. Alwall et al., "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations", *JHEP* **07** (2014) 079, doi:10.1007/JHEP07(2014)079, arXiv:1405.0301.

[86] I. Anderson et al., "Constraining anomalous HVV interactions at proton and lepton colliders", *Phys. Rev. D* **89** (2014) 035007, doi:10.1103/PhysRevD.89.035007, arXiv:1309.4819.

[87] M. Grazzini, S. Kallweit, and D. Rathlev, "ZZ production at the LHC: fiducial cross sections and distributions in NNLO QCD", *Phys. Lett. B* **750** (2015) 407–410, doi:10.1016/j.physletb.2015.09.055, arXiv:1507.06257.

[88] M. Bonvini et al., "Signal-background interference effects in $gg \to H \to WW$ beyond leading order", *Phys. Rev. D* **88** (2013) 034032, doi:10.1103/PhysRevD.88.034032, arXiv:1304.3053.

[89] K. Melnikov and M. Dowling, "Production of two Z-bosons in gluon fusion in the heavy top quark approximation", *Phys. Lett. B* **744** (2015) 43, doi:10.1016/j.physletb.2015.03.030, arXiv:1503.01274.

[90] C. S. Li, H. T. Li, D. Y. Shao, and J. Wang, "Soft gluon resummation in the signal-background interference process of $gg(\to h^*) \to ZZ$ ", *JHEP* **08** (2015) 065, doi:10.1007/JHEP08(2015)065, arXiv:1504.02388.

[91] G. Passarino, "Higgs CAT", *Eur. Phys. J. C* **74** (2014) 2866, doi:10.1140/epjc/s10052-014-2866-7, arXiv:1312.2397.

[92] S. Catani and M. Grazzini, "An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC", *Phys. Rev. Lett.* **98** (2007) 222002, doi:10.1103/PhysRevLett.98.222002, arXiv:hep-ph/0703012.

[93] M. Grazzini, "NNLO predictions for the Higgs boson signal in the H $\to$ WW $\to \ell\nu\ell\nu$ and H $\to$ ZZ $\to 4\ell$ decay channels", *JHEP* **02** (2008) 043, doi:10.1088/1126-6708/2008/02/043, arXiv:0801.3232.

[94] M. Grazzini and H. Sargsyan, "Heavy-quark mass effects in Higgs boson production at the LHC", *JHEP* **09** (2013) 129, doi:10.1007/JHEP09(2013)129, arXiv:1306.4581.

[95] L. Landau, "On the energy loss of fast particles by ionization", *J. Phys. (USSR)* **8** (1944) 201–205.

[96] ATLAS and CMS Collaborations, LHC Higgs Combination Group, "Procedure for the LHC Higgs boson search combination in Summer 2011", Technical Report CMS-NOTE-2011-005/ATL-PHYS-PUB-2011-11, 2011.

[97] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics", *Eur. Phys. J. C* **71** (2011) 1554, doi:10.1140/epjc/s10052-011-1554-0, arXiv:1007.1727. [Erratum: Eur. Phys. J. C 73 (2013) 2501].

[98] A. L. Read, "Modified frequentist analysis of search results (the CL_s method)", CERN Report CERN-OPEN-2000-005, 2000.

[99] T. Junk, "Confidence level computation for combining searches with small statistics", *Nucl. Instrum. Meth. A* **434** (1999) 435–443, doi:10.1016/S0168-9002(99)00498-2, arXiv:hep-ex/9902006.

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

[101] S. S. Wilks, "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses", *Annals Math. Statist.* **9** (1938), no. 1, 60–62, doi:10.1214/aoms/1177732360.

[102] CMS Collaboration, "Studies of Higgs boson production in the four-lepton final state at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-HIG-15-004, 2016.

[103] R. D. Cousins, "Why isn't every physicist a Bayesian?", *Am. J. Phys.* **63** (1995) 398, doi:10.1119/1.17901.

[104] CMS Collaboration, "Measurements of properties of the Higgs boson and search for an additional resonance in the four-lepton final state at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-HIG-16-033, 2016.

[105] F. Caola and K. Melnikov, "Constraining the Higgs boson width with ZZ production at the LHC", *Phys. Rev. D* **88** (2013) 054024, doi:10.1103/PhysRevD.88.054024, arXiv:1307.4935.

[106] N. Kauer and G. Passarino, "Inadequacy of zero-width approximation for a light Higgs boson signal", *JHEP* **08** (2012) 116, doi:10.1007/JHEP08(2012)116, arXiv:1206.4803.

[107] CMS Collaboration, "Search for Higgs boson off-shell production in proton-proton collisions at 7 and 8 TeV and derivation of constraints on its total decay width", *JHEP* **09** (2016) 051, doi:10.1007/JHEP09(2016)051, arXiv:1605.02329.

université PARIS-SACLAY | ÉCOLE DOCTORALE PHENIICS

Title: Measurements of Higgs boson properties in the four-lepton final state at $\sqrt{s} = 13\,\text{TeV}$ with the CMS experiment at the LHC

Keywords: *Higgs boson, CMS experiment, LHC collider, standard model, weak interaction*

Abstract:

This thesis reports a study of Higgs boson production in proton-proton (pp) collisions at $\sqrt{s} = 13\,\text{TeV}$ recorded with the CMS detector at the CERN Large Hadron Collider (LHC), exploiting the decay channel into a pair of Z bosons that in turn decay into pairs of electrons or muons ($\text{H} \to \text{ZZ} \to 4\ell$, $\ell = \text{e}, \mu$).

This work is carried out in the context of the beginning of Run II of the LHC, a new data-taking period that started in 2015, following a two-year shutdown. This restart is marked by an increase of the centre-of-mass energy from 8 TeV to 13 TeV and an upgrade of collider parameters that increases the luminosity, setting new constraints on the triggering, reconstruction and analysis of pp collision events. Therefore, considerable effort is devoted to the improvement and reoptimization of the CMS trigger system for Run II, focusing on the reconstruction and selection of electrons and on the preparation of multilepton trigger paths that preserve a maximal efficiency for the $\text{H} \to \text{ZZ} \to 4\ell$ channel.

Secondly, the offline algorithms for electron and muon selection are optimized and their efficiencies are measured in data, while the selection logic of four-lepton candidates is improved. In order to extract rare production modes of the Higgs boson such as vector boson fusion, VH associated production, and $\text{t}\bar{\text{t}}\text{H}$ associated production, a new classification of selected events into exclusive categories is introduced, using discriminants based on matrix-element calculations and jet flavour tagging.

Results of the analysis of first 13 TeV data are presented for two data sets recorded in 2015 and early 2016, corresponding to integrated luminosities of $2.8\,\text{fb}^{-1}$ and $12.9\,\text{fb}^{-1}$, respectively. A standalone rediscovery of the Higgs boson in the four-lepton channel is achieved at the new energy. The signal strength relative to the standard model prediction, the mass and decay width of the boson, and a set of parameters describing the contributions of its main predicted production modes are measured. All results are in good agreement with standard model expectations for a 125 GeV Higgs boson within the uncertainties, which are dominated by their statistical component with the current data set. Finally, a search for an additional high-mass resonance decaying to four leptons is performed, and no significant excess is observed.

Titre : Mesure des propriétés du boson de Higgs dans l'état final à quatre leptons à $\sqrt{s} = 13\,\text{TeV}$ avec l'expérience CMS au LHC

Mots-clés : *boson de Higgs, expérience CMS, collisionneur LHC, modèle standard, interaction faible*

Résumé :

Cette thèse présente une étude de la production de boson de Higgs dans les collisions proton-proton (pp) à $\sqrt{s} = 13\,\text{TeV}$ enregistrées avec le détecteur CMS au Grand collisionneur de hadrons (LHC) du CERN, exploitant la voie de désintégration en une paire de bosons Z qui se désintègrent à leur tour en paires d'électrons ou de muons ($\text{H} \to \text{ZZ} \to 4\ell$, $\ell = \text{e}, \mu$).

Ce travail s'inscrit dans le contexte du début du Run II du LHC, une nouvelle période de prise de données qui a commencé en 2015 après une interruption de deux ans. Ce redémarrage est marqué par une augmentation de l'énergie dans le centre de masse de 8 TeV à 13 TeV et par une mise à niveau des paramètres du collisionneur augmentant la luminosité, ce qui place des contraintes inédites sur le déclenchement, la reconstruction et l'analyse des données de collisions pp. Un effort important est donc consacré à l'amélioration et la réoptimisation du système de déclenchement de CMS pour le Run II, en mettant l'accent sur la reconstruction et la sélection des électrons et sur la préparation de voies de déclenchement multi-leptons préservant une efficacité maximale pour le canal $\text{H} \to \text{ZZ} \to 4\ell$.

Dans un second temps, les algorithmes de sélection hors-ligne des électrons et des muons sont optimisés et leurs efficacités sont mesurées dans les données, tandis que la logique de sélection des candidats à quatre leptons est améliorée. Afin d'extraire des modes de production rares du boson de Higgs tels que la fusion de bosons vecteurs, la production associée VH et la production associée $\text{t}\bar{\text{t}}\text{H}$, une nouvelle répartition des événements sélectionnés en catégories exclusives est introduite, fondée sur des discriminants utilisant le calcul d'éléments de matrice et l'étiquetage de saveur des jets.

Les résultats de l'analyse des premières données à 13 TeV sont présentés pour des lots de données enregistrés en 2015 et au début de 2016, qui correspondent à des luminosités intégrées respectives de $2.8\,\text{fb}^{-1}$ et $12.9\,\text{fb}^{-1}$. Le boson de Higgs est redécouvert de façon indépendante à la nouvelle énergie. L'intensité du signal relative à la prédiction du modèle standard, la masse et la largeur de désintégration du boson sont mesurées, ainsi qu'un jeu de paramètres contrôlant les contributions des principaux modes de production attendus. Tous les résultats sont en bon accord avec les prévisions du modèle standard pour un boson de Higgs à 125 GeV, aux incertitudes de mesure près, ces dernières étant dominées par la composante statistique avec l'échantillon de données actuel. Enfin, une recherche de nouvelle résonance se désintégrant en quatre leptons à haute masse est effectuée, et aucun excès significatif n'est observé.

Université Paris-Saclay
Espace Technologique – Immeuble Discovery
Route de l'Orme aux Merisiers RD 128 – 91190 Saint-Aubin – France