

Dark matter searches in invisible Higgs boson decays with the CMS experiment

Analysing Higgs boson decays to invisible particles in $t\bar{t}H$ and VH final states and combining all final states using Run 1 and Run 2 CMS data



SCHOOL OF PHYSICS
UNIVERSITY OF BRISTOL

ROBERT STEPHEN WHITE

June 22, 2023

A DISSERTATION SUBMITTED TO THE UNIVERSITY OF BRISTOL IN ACCORDANCE
WITH THE REQUIREMENTS FOR AWARD OF THE DEGREE OF DOCTOR OF
PHILOSOPHY IN THE FACULTY OF SCIENCE.



Abstract

Dark matter (DM) is a well-known but as-of-yet unidentified phenomenon that constitutes more than a quarter of the universe. It is dark in that it neither absorbs nor emits light; it is electromagnetically inert. Its effects, however, are observable: spiral galaxies behave in such a way that their rotation is more like what is expected in the presence of substantially more mass than that which is visible. In the Bullet Cluster, where two galaxies are colliding, gravitational lensing techniques show how the centre of mass of these galaxies is consistent with that of two large bodies of DM-dominated mass.

A high energy environment such as that of the Large Hadron Collider (LHC) at CERN would be ideal to produce DM particles with a mass $\mathcal{O}(\text{TeV})$, where proton-proton collisions take place and purpose-built detectors such as the Compact Muon Solenoid (CMS) detector can record the collisions and reconstruct the particles produced. One of the primary aims in the current frontier of high energy physics is to detect DM directly, although so far this has not been achieved. In the LHC, efforts to discover DM rely on production from standard model (SM) particles, including via models where the Higgs boson acts as a mediator. With the CMS detector, 165 fb^{-1} data has been recorded across the periods 2011–2013 (Run 1) and 2015–2018 (Run 2), with which in light of the absence of DM discovery limits are set on the branching fraction $\mathcal{B}(\text{H} \rightarrow \text{inv})$, the probability of decays of Higgs bosons to *invisible* particles not detected in the CMS detector.

In the standard model, given the only invisible decays of the Higgs boson are to neutrinos, $\mathcal{B}(\text{H} \rightarrow \text{inv})$ is approximately 0.12%. Using 2016–2018 data, $\mathcal{B}(\text{H} \rightarrow \text{inv})$ is set at 54% observed (39% expected) at the 95% confidence level (CL) in the combined $t\bar{t}\text{H}$ and VH channel. Using both Run 1 and Run 2 data from the CMS detector, this is set at 15% observed (8% expected) at the 95% CL, from a combination of channels in which a Higgs boson is produced in association with another or other particles.

Dedication and acknowledgements

I would like to thank first and foremost, my supervisors! First, Henning Flächer, whose constant guidance and drive has pushed me to be a better, more rigorous scientist, and has helped pull to shore the Higgs-to-invisible analysis during the year-and-a-half-long review. This thesis would also be very pessimistic and misleading without your feedback as well. I wish you luck finding Dark Matter with LZ, whether at SURF or maybe Boulby eventually, because we're not finding any in high-energy colliders...can't believe DM might be discovered for the first time just outside Middlesbrough! The invaluable knowledge of Sudan Paramesvaran on the CMS detector has also aided my understanding of the Collaboration I've been working for since 2019. Your proofreading and sanity checks for the analysis (and thesis) have been a huge help, and pushing forward my contributions to AutoDQM will hopefully set me up for a second paper in 2023. Your aid in postdoc applications was certainly a huge factor in helping me get a new job!

I would like to thank Nick Wardle at Imperial, whose indispensable knowledge of the HiggsCombine tool has been huge for producing a statistically sound analysis, and has been a fountain of information regarding Higgs Physics! Likewise, I thank Olivier Davignon of LLR, who was the original architect of the fit procedure for the analysis amongst other major contributions, and whose rigour and understanding has been integral to producing our Higgs-to-invisible result long after you moved away from the project. Additional thanks go to Andreas Albert, who really synthesised the combination and produced a repository into which our channels fit smoothly during the handover, and Ben Krikler, who left some time ago but whose efforts with F.A.S.T made for a very sleek and interactive workflow (barring fast-plotter), and who introduced me to the AutoDQM group in the early days of my PhD. Also Aaron Bundock's work on the *new fit* was key to reshaping our workflow and producing correct results (which still receives eleventh-hour tweaks!) Exploring the tunnels (and nicking useless crap down there) made my possibly final visit to CERN as a student so much more enriching!

I would have nothing to show for in the analysis without the work of those before me: Eshwen Bhal and David Anthony, whose major contributions and constant help have been crucial throughout a very taxing and relentless analysis, but for whom the publication of our paper has hopefully been made worthwhile. Also thanks to Maciek Glowacki, a West Yorkshire prodigy and fellow United supporter whose knowledge of the computing and ML world and setting up the CHIP pipeline has been a massive help (and whose banter made the LTA office worthwhile working in!) The (ultra-)legacy of

Hinv is now in the capable hands!

For the help on the AutoDQM project (which will be my life for the next couple of months) I'd like to thank Chad Freer, Indara Suarez, and Andrew Brinkerhoff, who have overseen the project and provided me with expertise on the academic world and how to write a proper professional CV and statement! I look forward to working with you all over the summer!

A special thanks to Katie Walkingshaw Pass, who again was a rock throughout the PhD - diolch am y crancod! And to the rest of the Bristol gang - Jake, Charlotte, Mei, Ben, Lakshan, Nikos, Chris, and more recently Florian, Liam and Shahzad - who throughout my time in Bristol has made my life so much more fun! The biggest draw to the Bristol group is always the social side, so thanks to Joel, Kostas, Jonas, David C, Paras, and especially thanks to Martisse, all of whom have made my years at Bristol (and its many pubs) a joy and made the group so lively and inclusive!

For making my life liveable during lockdown, I'd like to thank Tyler Woodburn, an ||ledge|| who has cooked for me on countless occasions, and now a place to stay in my final months in Bristol. Also, to Señor Harry Manocha, with whom I had so many good times, not least clearing a shed! And to George Wickett and Ahmed Taha, the rest of the Ladz Inc. gang who during my days at Coronation Road helped make isolating over Christmas bearable and forget about the actual slum we lived in.

Further, I'd like to pay homage to my time in Meyrin, and the LTA family who couldn't wait to explore the world after lockdown on LTA: first and foremost Ethan Simpson, who I lived with the whole 15 months out in Switzerland and who was there through the highs and lows of the drama and mountains and tunnels - can't thank you enough for saving me from wolves in Morillon! Shout outs to Zak, Luke, Emil, Jake, Holly, Riz, Jenny, Abbie (TBC), Estifa'a, Jack, Ryun, Neža and Tim and Dom and Josh (the 4000er crew), Ben, Red Ben, Rousso, Levi, Giovanna (a presto!), Anja [1], Rebecca, Eva, Gareth, Will and Johnny (survivors of the Erlebnisklettersteig), Sammy, Kathryn, Meirin, Fionn, Ryan, Martin, Clara, Ynyr, Iza, and Harry and Eimear and Adam and Sav (the social secs), whose place will forever remain in my heart (and on the LTA wall)!!

Also to the Alps, my true love in life and what made the whole PhD worthwhile! LTA will probably be the best time of my life - but I'll be conquering more passes soon!

And finally to my family and truly my parents, the OG doctors in the house, for being there throughout - to whom I can only apologise I'm not a chemist...

Author's declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: DATE:

Contents

1	Introduction	9
2	The Standard Model and the Dark Sector	13
2.1	The Standard Model	13
2.2	Gauge Theories in the SM	16
2.2.1	Electromagnetism	17
2.2.2	Quantum Chromodynamics	19
2.2.3	Electroweak Unification	22
2.3	Electroweak Symmetry Breaking	24
2.3.1	SM masses and couplings	29
2.4	The parton model	30
3	The CMS Detector	33
3.1	The Large Hadron Collider	33
3.1.1	Luminosity	36
3.1.2	Pileup	37
3.2	Compact Muon Solenoid	39
3.2.1	Tracker	41
3.2.2	Electromagnetic Calorimeter	43
3.2.3	Hadronic Calorimeter	45
3.2.4	Muon Chamber and Yoke	47
3.3	Triggering and Data Acquisition	49
3.3.1	Level 1 Trigger	49
3.3.2	Layer 2 Calorimeter and DQM Overview	51
3.3.3	High Level Trigger	54
3.3.4	Data Reduction and Storage	55
3.4	AutoDQM	56
3.4.1	AutoDQM ML	60

4	Event reconstruction and simulation	65
4.1	Software	66
4.2	Simulated samples	67
4.3	Event reconstruction	71
5	Analysis strategy	78
5.1	Trigger selections and efficiencies	80
5.1.1	Trigger efficiency	81
5.2	Offline event selection	83
5.3	Event categorisation	86
5.3.1	ℓ_{lost} and $\mathbf{Z} \rightarrow \mathbf{inv}$ background estimation	88
5.3.2	QCD multijet background estimation	92
5.3.3	Signal region composition	97
6	Statistical interpretation	100
6.1	Theoretical systematic uncertainties	102
6.1.1	Systematic uncertainties in t quark samples	106
6.2	Experimental systematic uncertainties	110
6.2.1	Lepton and photon identification and isolation	111
6.2.2	Trigger efficiency and pre-firing	111
6.2.3	Object tagging	112
6.2.4	Luminosity and pileup	112
6.2.5	QCD prediction	113
6.2.6	Jet energy uncertainties	113
6.3	Statistical uncertainty	115
6.4	Statistical interpretation	116
7	Fit Results	118
8	Combination with other Higgs boson production channels	130
8.1	Selection overlap	133
8.2	Systematic uncertainties	136
8.2.1	Theoretical uncertainties	137
8.2.2	Experimental uncertainties	138
8.3	Limits and evaluation of results	141
8.4	Dark matter interpretation	144
9	Summary	149

A	Look-up values	151
B	Statistical analysis	153
B.1	Goodness-of-fit tests	156
B.2	Parameter fitting and maximum likelihood	158
B.3	Interval estimates and limit setting	159
B.4	Nuisance parameters and systematic uncertainty	163
C	Supporting material for the $t\bar{t}H$ and VH analysis	166
D	Supporting material for the combination	175

Chapter 1

Introduction

The universe is comprised principally of three components: *baryonic*, or normal matter from which the world around us is formed; dark energy, the descriptor for the phenomenon by which the accelerating rate of expansion of the universe is explained; and dark matter (DM), the *invisible* mass that maintains the shape of spiral galaxies and constitutes 85% of all matter in the universe, according to the best parametrisation of the cosmological Big Bang model. DM is invisible in the sense that it neither emits nor absorbs electromagnetic radiation, hence why it does not appear visibly on galactic space surveys like ordinary matter. In a similar vein, were DM to exist on Earth, it would be difficult to detect directly: it weakly interacts with normal matter, therefore would not be picked up in particle detectors, while instrumentation to detect gravitational effects is not sufficiently sensitive on the subatomic level.

Several candidates for DM sources have been proposed over time, ranging from massive compact halo objects (MACHOs) on a macroscopic level to fast-moving neutrinos on a microscopic scale. Both would explain the very low interaction cross-section and the gravitation-only interaction with SM matter, based on observation. Although not excluded entirely, MACHOs are unlikely to constitute all contributions to DM in the universe, while fast-moving neutrinos are too light and would wash out the large scale structure of the universe, meaning galaxies would not form. The lead candidate for DM effectively describes the known properties of DM: weakly-interacting massive particles (WIMPs), which are elementary non-baryonic particles outside of the SM that perform well in modelling the large scale structure of the universe.

Means of detecting WIMPs include direct and indirect methods. Direct detection typically uses underground vessels (which are naturally shielded from background radiation such as cosmic rays) containing cryogenically-cooled inert liquids, where the recoil of the nuclei in the liquid due to their interaction with latent DM is measured

through the emission of scintillating light induced by the nuclear recoil. The energy scale for direct detection is $\mathcal{O}(\text{keV})$, and so far has set strong upper limits on the DM-nucleon interaction cross-section for small WIMP masses. Indirect methods involve searches for particles produced in the annihilation or decay of DM candidates in space, assuming a DM model that couples readily to the Standard Model (SM). The signature for this annihilation or decay would be the emission of high energy γ rays that can be detected in γ ray telescopes both on Earth and in orbit.

The final method of detecting WIMPs is to produce them in a laboratory setting, for which discoveries or exclusion limits can be set on WIMP masses $\mathcal{O}(100 \text{ GeV} - 10 \text{ TeV})$. Higgs bosons decaying to invisible particles ($H \rightarrow \text{inv}$) is an ideal signature for DM production: the Higgs boson [2–7] is sufficiently massive to couple to heavy particles like WIMP candidates, and from the detection of other particles produced in association with the Higgs decay in an event, information about the missing energy from the invisible Higgs decay can be inferred.

Searches like these are one of the frontiers of contemporary high-energy physics (HEP), however no observation has yet been made in particle colliders. The CERN *Large Hadron Collider* (LHC) provides ideal conditions for these searches, where as of 2022 proton-proton (pp) collision data at centre-of-mass energy $\sqrt{s} = 13.6 \text{ TeV}$ is being collected in detectors such as the *Compact Muon Solenoid* (CMS). The LHC is effectively a factory for producing 125 GeV Higgs bosons. In the SM, the Higgs can only decay invisibly via $H \rightarrow ZZ^* \rightarrow 4\nu$, with an expected branching ratio of 0.12%. Typical high energy particle detectors such as CMS will ideally detect all SM particles that are not neutrinos. Some DM models place the $H \rightarrow \text{inv}$ branching ratio, $\mathcal{B}(H \rightarrow \text{inv})$, at $\mathcal{O}(10\%)$, therefore measuring an upper limit on it will give an indication of the contribution due to Higgs decays into DM candidates.

The tightest exclusions on $\mathcal{B}(H \rightarrow \text{inv})$ are obtained in vector boson fusion (VBF) searches using data collected at $\sqrt{s} = 13 \text{ TeV}$, yielding exclusion limits of $\mathcal{B}(H \rightarrow \text{inv}) = 0.18$ (0.10 expected) at the 95% confidence level (CL) with 165 fb^{-1} of CMS data [8] from all runs of the LHC, and $\mathcal{B}(H \rightarrow \text{inv}) = 0.15$ (0.10 expected) with 139 fb^{-1} of *Toroidal LHC ApparatuS* (ATLAS) data [9] from Run 2 (2015–2018).

The main focus of this thesis is to analyse the Higgs boson production channels for which a $t\bar{t}$ pair ($t\bar{t}H$), or a vector boson (VH) with a resolved topology, where the presence of the V boson is inferred from well separated decay products, is also produced, and the final state of these associated objects is entirely hadronic [10]. The exclusion limits calculated in these channels are then combined with the CMS VBF production channel, as well as Higgs production in gluon-gluon fusion (ggH) [11], $t\bar{t}H$

where the $t\bar{t}$ decays into at least one charged lepton [12–14], and VH channels with varying final states [11, 15] including the topology with merged decay products arising from boosted V bosons (to complement the resolved VH topology studied here). The combination is performed using all 165 fb^{-1} CMS Run 1 (2011–2014) [16] and Run 2 data for which $H \rightarrow \text{inv}$ analyses have been published. The VBF, VH, $t\bar{t}H$, and ggH topologies are illustrated as representative tree level diagrams in Fig. 1.1.




Figure 1.1: Representative tree level diagrams of the four Higgs production channels, (clockwise) VBF, VH, ggH, and $t\bar{t}H$, in which $H \rightarrow \text{inv}$ searches are performed by the CMS Collaboration.

In order to introduce the analysis, an overview of particle physics in the SM is provided in Ch. 2. The technical aspects of the CMS detector and data collection are discussed in Ch. 3. The data sets and software used in the analysis, and modes of event reconstruction, are given in Ch. 4. The analysis strategy, covering the regions probed by the analysis, and the trigger and offline selections that define these regions and signal event categorisation, are covered in Ch. 5. The likelihood model for this analysis, and a discussion on the statistical and systematic uncertainties that enter into the fitting procedure, are found in Ch. 6. An in-depth discussion of the fit results for the $t\bar{t}H$ /resolved VH analysis are in Ch. 7, and for the combination with other Run 1 and Run 2 analyses in Ch. 8. The results are finally summarised in Ch. 9.

The $t\bar{t}H$ /resolved VH analysis has been a group effort over a number of years, with

studies important to the narrative of $H \rightarrow \text{inv}$ searches performed by various people over this time. The areas in which I have specialised or been a key contributor include: use and development of all software featured in Sec. 4.1, the addition of some signal and background samples mentioned in Sec. 4.2, and optimising the offline event selection using various metrics in Sec. 5.2 and Sec. 5.3, namely in the $t\bar{t} + \text{jets}$ background suppression study, developing the categorisation scheme to the current version featured in Table 5.2, and some minor studies with the QCD multijet suppression selection. I have also contributed heavily to the implementation of the fit model established in Ch. 6, and to various systematic uncertainty measurements, mainly for the photon normalisation, signal normalisation, boosted object tagging, jet energies, and general bookkeeping of all systematic uncertainties to produce the results in Table 6.1. The bulk of my work for the analysis was in understanding the statistical interpretation of our results and use of the HIGGS COMBINE software. This incurred conducting the maximum-likelihood fit to obtain exclusion limits as presented in Ch. 7, alongside profile likelihood scans, goodness-of-fit (GF) tests, and measuring the impacts and pulls of systematic uncertainties on the signal strength. Similarly, much of my time has been dedicated to the combination of the $t\bar{t}H$ /resolved VH analysis with other Higgs boson production channels, presented in Ch. 8, where I have contributed to all aspects and specialised in establishing the set of Higgs portal models and motivating literature, and their comparisons to select results from direct-detection experiments.

Work such as object tagging, corrections and corresponding uncertainties, and trigger selections for analysers are provided centrally by the CMS Collaboration, and are often the recommended or established method. Some of these lack publicly-available literature, which is highlighted in the text where applicable.

The AutoDQM project presented in Sec. 3.4 follows the work of the AutoDQM group, of which I have been a member since the beginning. Amongst other smaller contributions, I have been the primary contributor to the implementation of the layer 2 calorimeter trigger and subsequent performance evaluation.

Chapter 2

The Standard Model and the Dark Sector

The focus of this thesis is on Higgs boson decays to the invisible by analysing signal channels using a decade of data from the CMS detector. However, to understand the importance of the Higgs boson and its hypothesised connection to the DM sector, a grounding in the theory of the SM and the Brout-Englert-Higgs (BEH) mechanism [2] is essential. Phenomenological arguments for using the LHC and DM-dedicated detectors to probe the substance that 27% of the universe is believed to be made up of is also necessary. Throughout this chapter, several matrix expressions are used that are listed in Appendix A.

2.1 The Standard Model

The Standard Model of particle physics is the most accurate description of the universe to date, and intends to model the large scales of the universe from the fewest principles using a set of particles at the quantum level. Although not reconciled with the gravitational force, the other major forces governing the universe are well-described, those being the strong nuclear force, the weak nuclear force and electromagnetism (EM). Each of these forces has a corresponding *carrier* particle whose interaction with other SM particles is responsible for its propagation. In the context of the LHC which operates at temperatures similar to that of the early universe, i.e. energies of $\mathcal{O}(\text{TeV})$, the weak force and EM are comfortably unified into a single *electroweak* (EW) force [17–19], with their carriers effectively coupled in EW theory. The threshold for EW unification to break down is for energies below 100 GeV, although the study here is entirely focused on the high energy SM.

In the simplest form, the SM is comprised of two sets of particles, distinguishable based on their intrinsic *spin*, a property that arises within quantum mechanics. The two categories of SM particle are fermions, defined by their half-integer spin, and include quarks and leptons which constitute normal matter, and bosons, which bear whole integer spin and couple to fermions to mediate the interactions (or forces) between them.

Fermions are again divided into two subsets. The first of these are quarks, of which there are six *flavours*. The six flavours pair together in three *generations*, with one of the quarks bearing an electric charge of $+\frac{2}{3}$, and the other bearing an electric charge of $-\frac{1}{3}$. These flavour pairs are: in the first generation, up u and down d ; in the second, charm c and strange s ; and in the third, top t and bottom (or beauty) b . Quarks have a vast range of masses, from the lightest quark $m_u = 2.2$ MeV to the heaviest quark (and in fact the heaviest particle in the SM) $m_t = 173$ GeV [20]. As well as having electric charge, mass and spin, quarks carry colour charges, for which there are three colours: red, green and blue (with antiquarks carrying the complementary anti-colour of each). Owing to *colour confinement*, a property of the strong force, quarks cannot be directly observed. Instead, their properties are inferred from *hadrons* such as protons, composed of three quarks with distinct colours, or *mesons* such as pions, e.g. $\pi^\pm$, composed of a quark-antiquark pair with same colour-anticolour charge. Further groupings of quarks into tetraquark and pentaquark states have been discovered [21, 22], although these are short-lived. All quark groupings must be colour-neutral, and are referred to as *baryons*, from which the term *baryonic* matter, to mean normal matter, stems. Quarks interact via all of the fundamental forces.

The other category of fermions contains leptons. These are again grouped into generations, this time with a particle with an electric charge of -1, and an electrically neutral, near-massless *neutrino* counterpart [23]. The lepton pairs are: in the first generation, electron e and electron neutrino, ν_e ; in the second, muon μ and muon neutrino, ν_μ ; and in the third, tau lepton τ and tau neutrino, ν_τ . The charged leptons can interact via the EM and weak forces, while the neutrino only interacts via the weak force, although neutrino interactions are rare owing to their extremely low mass. Leptons do not carry colour charge, and therefore are inert to the strong force.

An additional property of SM particles is weak isospin, T , and describes the chirality, doublet nature and matter or antimatter nature of particles. A particle is said to have a left-handed (LH) chirality if the direction of its spin and momentum are anti-parallel; a particle with aligned spin and momentum direction is right-handed

(RH), which is the result of a parity transformation applied to the LH particle. This is a conserved quantity under the weak force, where in the LH fermion doublets, u -type quarks and neutrinos have weak isospin $+\frac{1}{2}$, and their RH antifermion counterparts have weak isospin $-\frac{1}{2}$. On the other hand, d -type quarks and charged leptons have weak isospin $-\frac{1}{2}$, while their RH antifermion analogues have weak isospin $+\frac{1}{2}$. Due to preferential chirality whereby RH fermions and LH antifermions do not couple to the weak force in the SM, their weak isospin is 0. Based on observation [24], this coupling is heavily suppressed.

A summary of the SM fermions is provided in Table 2.1. A vast zoo of particles exist thanks to work over the last 120 years, from the discovery of the electron in J.J. Thompson's cathode ray experiment [25], Ernest Rutherford's work on probing the atomic nucleus and discovering the proton [26], the discovery of the muon in 1937 [27] amidst the advent of antimatter and meson discoveries in the 1930s-50s, the neutrino discovery in 1956 by Reines and Cowan [28], to the birth of linear and circular collider technology for discovering heavier particles, such as at the Stanford Linear Accelerator Center (SLAC) used in the discoveries of the u , d , and s quarks in deep inelastic electron-proton scattering experiments in 1969 [29, 30], the c quark (jointly at Brookhaven National Laboratory) in 1974 [31, 32], tau lepton in 1975 [33], followed by the discovery of the b quark at Fermilab in 1977 [34], and the t quark with the Tevatron at Fermilab in 1995 [35, 36]. Today, more distinct or exotic particles or particle excitations are being produced using the LHC, the successor to the Tevatron, as higher and higher energies are achieved. However all of these new particles can be broken down into the elementary particles above.

The four carrier particles in the SM are referred to as *gauge bosons*, or *vector bosons*. Gauge bosons are spin-1 (vector) particles which obey a given gauge symmetry within their quantum field. The strong force is mediated by the gluon, g , which is a massless, electrically neutral particle that exists in eight flavours, each of which correspond to compositions of two distinct colour charges. The EM force is carried by the photon, γ , which is also massless and electrically neutral but interacts with electrically charged particles. The weak force is mediated by the weak bosons, $W^\pm$ and Z^0 , which are massive vector bosons that interact with all fermions in the SM. Given these are weak force mediators, the weak vector bosons carry weak isospin of ± 1 for $W^\pm$ and 0 for Z^0 . The Higgs boson H is not a gauge boson: it is not a force carrier and is a scalar spin-0 particle. The Higgs boson arises from the excitation of the Higgs field [2], which is the mechanism by which elementary particles possess mass. The Higgs boson is electrically neutral, with 0 spin, but has weak isospin $-\frac{1}{2}$ due to the doublet nature of the Higgs

Table 2.1: The complete list of fermions that are accounted for in the SM, with quark and charged lepton masses taken from Ref. [20], and neutrino mass limit taken from Ref. [23]. Gen. denotes the fermion generation, and the weak isospin T is for LH particles only (where for RH particles, $T = 0$).

Type	Gen.	Particle	Spin	T	Electric charge	Mass
Quark	1	Up (u)	1/2	+1/2	+2/3 e	$2.16^{+0.49}_{-0.26}$ MeV
		Down (d)		-1/2	-1/3 e	$4.67^{+0.48}_{-0.17}$ MeV
	2	Charm (c)		+1/2	+2/3 e	1.27 ± 0.02 GeV
		Strange (s)		-1/2	-1/3 e	$93.4^{+8.6}_{-3.4}$ MeV
	3	Top (t)		+1/2	+2/3 e	172.69 ± 0.30 GeV
		Bottom (b)		-1/2	-1/3 e	$4.18^{+0.03}_{-0.02}$ GeV
Lepton	1	Electron (e)	1/2	-1/2	-1 e	0.511 MeV
		Electron neutrino (ν_e)		+1/2	0	< 0.09 eV
	2	Muon (μ)		-1/2	-1 e	105.66 MeV
		Muon neutrino (ν_μ)		+1/2	0	< 0.09 eV
	3	Tau (τ)		-1/2	-1 e	1776.86 ± 0.12 MeV
		Tau neutrino (ν_τ)		+1/2	0	< 0.09 eV

field, described in Sec. 2.3. A summary of the SM bosons is provided in Table 2.2.

The timeline for the discovery of SM bosons begins in the 19th Century through various light detection experiments at various wavelengths corresponding to the photon, followed much later by the gluon in the Positron–Electron Tandem Ring Accelerator (PETRA) at DESY in 1979 [37], the $W^\pm$ and Z^0 bosons with the CERN Super Proton Synchrotron (SPS) in 1983 [38, 39], and most recently the Higgs boson discovery with the CERN LHC in 2012 [40, 41].

All known (i.e. baryonic) matter in the universe can be constructed from these particles, although stable matter is only formed from electrons, protons and neutrons, with a proton composed of the tri-quark udu and a neutron dud . A lone neutron decays to a proton with a mean lifetime of about 15 minutes, while a proton is assumed to be stable in the SM, with studies of proton decay a topic of interest in New Physics (NP) searches [42].

2.2 Gauge Theories in the SM

In physics conservation laws are derived from symmetries in an underlying field theory. Field theories representative of a physical system are often described by a *Lagrangian*

Table 2.2: The complete list of known bosons in the SM, accounting for three of the four fundamental forces of the universe. Boson masses are taken from Ref. [20].

Force	Carrier	Spin	T	Electric charge	Mass
Strong	Gluon (g)	1	0	0	0
EM	Photon (γ)	1	0	0	0
Weak	W bosons ($W^\pm$)	1	± 1	$\pm 1 e$	$80.377 \pm 0.012 \text{ GeV}$
Weak	Z boson (Z^0)	1	0	0	$91.1876 \pm 0.0021 \text{ GeV}$
–	Higgs boson (H)	0	$-1/2$	0	$125.25 \pm 0.17 \text{ GeV}$

density $\mathcal{L}$, for which there exists a conserved charge or current associated with a particular symmetry transformation if $\mathcal{L}$ is invariant under the transformation. This is the principle of Noether’s theorem [43].

The SM is a *gauge* quantum field theory (QFT), for which there exist multiple symmetries. A gauge theory is one for which the physical system is invariant under transformations between a group of gauges, called the symmetry group. Under QFT, particles are considered excitations of their underlying quantum fields, these fields being the fundamental element of the theory, and not the particles, given particle number is not conserved in the quantum universe. *Interactions*, or transformations, between these fields are described by continuous symmetry groups. Consequently, an interaction due to a given group must conserve the charges that are borne from the symmetries of that group. Charge conservation is mediated by gauge boson fields in the SM, which are *generators* of the symmetry group. Equivalently, for each generator of a symmetry group there is necessarily an associated gauge field. Gauge fields are given as terms in $\mathcal{L}$ to ensure invariance under gauge transformations, from which SM particles can be derived. Three of the four fundamental forces have been described by gauge theories that have been successful in determining the extent to which these forces interact within the SM.

2.2.1 Electromagnetism

The simplest gauge theory in the SM is EM, whose Lagrangian $\mathcal{L}_{EM}$ ¹ is invariant under unitary transformations according to the symmetry group $\mathbb{U}(1)_{EM}$. Starting

¹The choice here is to use EM to denote electromagnetism, although formally this is quantum electrodynamics (QED).

with the Dirac equation for a field $\Phi(x)$ of free particles of mass m ,

$$(i\gamma^\mu\partial_\mu - m)\Phi = 0, \quad (2.1)$$

where γ^μ are the set of Dirac matrices, the corresponding formula for $\mathcal{L}_{EM}$ is given by

$$\mathcal{L}_{EM} = \bar{\Phi}(i\gamma^\mu\partial_\mu - m)\Phi. \quad (2.2)$$

A *local* symmetry under $\mathbb{U}(1)_{EM}$ transformations is crudely given by

$$\Phi(x) \rightarrow \Phi'(x) = \exp[ie\theta(x)]\Phi(x), \quad (2.3)$$

for which $\theta(x)$ is some gauge transformation that necessarily depends on space-time coordinates (locality), and e is introduced as the electric charge for the sake of the EM picture. Applying the transformation to Eq. 2.2 leads to a $\partial_\mu\theta$ term that is non-invariant. To restore invariance, the derivative in Eq. 2.2 is modified to include a cancelling term, so giving the covariant derivative

$$\partial_\mu \rightarrow D_\mu = \partial_\mu - ieA_\mu, \quad (2.4)$$

where A_μ is a vector potential that is introduced to absorb the remainder term, i.e. the gauge is set such that

$$A_\mu \rightarrow A'_\mu \hat{=} A_\mu + \frac{1}{e}\partial_\mu\theta,^2 \quad (2.5)$$

and $\mathcal{L}_{EM}$ invariance is restored. The final $\mathcal{L}_{EM}$ form is similar to Eq. 2.2 but is naturally extended to accommodate the vector potential terms, so giving

$$\mathcal{L}_{EM} = \bar{\Phi}(i\gamma^\mu\partial_\mu - m)\Phi + eA_\mu\bar{\Phi}\gamma^\mu\Phi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu},^3 \quad (2.6)$$

where the term $eA_\mu\bar{\Phi}\gamma^\mu\Phi = A_\mu J^\mu$ is the charge density current (assuming this is no longer a field of free particles), and the kinetic term $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$ contains the EM field strength tensor,

$$F_{\mu\nu} \hat{=} \partial_\mu A_\nu - \partial_\nu A_\mu. \quad (2.7)$$

The physical consequence of this is that A_μ is a photon field in the group $\mathbb{U}(1)_{EM}$, and as such the electron-photon interaction is assured in the EM mechanism via the $A_\mu J^\mu$, while $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$ is the kinetic term that permits a propagating photon field. Notably

²The use of the symbol $\hat{=}$ denotes that the equality is by design.

³In taking the total Lagrangian of $\mathcal{L}_{EM}$, the set of Maxwell's equations of EM can be derived.

there is no mass term for the photon field here, as this would violate gauge invariance; the concept of the photon being massless is in agreement with observations that set the upper mass limit $m_\gamma < 1 \times 10^{-18}$ eV [20].

This now established EM theory, where the photon is massless and can interact with electrons, successfully describing interactions such as pair production, electron-positron annihilation, and Compton scattering. The addition of a mass term or additional non-gauge-invariant terms in $\mathcal{L}_{EM}$ would mean charge density current conservation would break down, violating experimental observation. The introduction of gauge invariance to EM as shown above is the simplest demonstration of the covariant formalism, and is expanded on with the more elaborate theories to follow.

2.2.2 Quantum Chromodynamics

The strong interaction is modelled by *quantum chromodynamics* (QCD), for which there are eight unique colour states of the gluon, which couple only to quarks in the SM owing to their colour-charged nature. The $\text{SU}(3)_C$ symmetry group describes the colour nature of QCD, for which the corresponding Lagrangian summed over all quark flavours f is given by

$$\mathcal{L}_{QCD} = \sum_f \bar{\Phi}_i^{(f)} (i\gamma_\mu D_{ij}^\mu - m_f \delta_{ij}) \Phi_j^{(f)} - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}, \quad (2.8)$$

where analogously to the covariant formalism adopted in the EM Lagrangian,

$$D_{ij}^\mu = \partial^\mu \delta_{ij} + ig A_\mu^a \lambda_{ij}^a, \quad (2.9)$$

where A_μ^a are the gluon fields with colour index a , and λ^a are the set of Gell-Mann matrices which are generators of the $\text{SU}(3)_C$ group, and

$$G_{\mu\nu}^a \hat{=} \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g_s f^{abc} A_\mu^b A_\nu^c \quad (2.10)$$

is the gluon field strength tensor of coupling strength g_s and structure constant f^{abc} of the $\text{SU}(3)_C$ group. Note that this produces a single QCD parameter g_s , where if $g_s = 0$ quarks and gluons would propagate freely. The quark masses are EW in origin.

Quarks are held by the strong force in a bound state at low energies by gluons: in a proton, this is three quarks held together by three gluons. At high energies, this regular structure breaks down while still obeying the principle of colour confinement. This is the exhibition of *asymptotic freedom*, whereby the interaction strength between

particles diminishes as the energy scale is increased. In practice, a strongly-bound quark-antiquark pair would require a higher energy to be separated proportional to their increasing separation, to the point that the force of separation plateaus with increasing energy scale until it is sufficient to form two newly-bound quark-antiquark states that are colour-confined. This guarantees that the observation of an isolated quark is impossible in nature.

Given SM quarks cannot be observed as free particles, a potential whereby the strong force increases linearly with distance to be consistent with asymptotic freedom and colour confinement can be written as

$$V(r) = -\frac{4\alpha_s}{3r} + kr, \quad (2.11)$$

where α_s is the strong interaction running coupling constant, given by

$$\alpha_s = \frac{g_s^2}{4\pi}, \quad (2.12)$$

and k is a constant parametrising the linear potential term. The coupling strength α_s is dependent on the energy scale⁴, due to the effect of screening, whereby the colour charge of a lone quark can be shielded by the creation of virtual quark-antiquark loops in gluon splitting, and the effect of antiscreening, whereby virtual gluon loops strengthen the colour charge of a lone quark. This occurs in QCD and not EM because, unlike the electrically-neutral photon, the massless gluon carries colour charge. This leads to gluon self-interaction, which means the strong force is range-limited due to screening but increases linearly with distance at a certain energy scale. Both screening effects are embedded in the one-loop coupling formula parametrised as a function of the energy scale μ^2 , given by

$$\alpha_s(\mu^2) = \frac{\alpha_s(\mu_R^2)}{1 - \alpha_s(\mu_R^2) \cdot \beta \cdot \ln \frac{\mu^2}{\mu_R^2}}, \quad (2.13)$$

where μ_R^2 is the *renormalisation* energy scale, and

$$\beta = \frac{1}{12\pi} (11N_c - 2N_f), \quad (2.14)$$

for which N_c is the number of colour charges, and N_f is the number of quark flavours.

⁴This energy scale dependence is also true for the QED coupling strength where perturbation theory is valid for $\mu \ll 10^{90}$ GeV. The energy scales of the LHC and for pp collisions are more sensitive in QCD given perturbation theory breaks down for much lower energies.

In the SM, these values are $N_c = 3$ and $N_f = 6$. From this formula, for all models for which there are at most 16 quark flavours, the antiscreening effect is greater than the screening, and the asymptotic freedom of quarks is maintained.

The screening effect follows from EM, where the coupling constant, or *fine-structure constant*, at rest is $\alpha_{EM} \approx \frac{1}{137}$ and increases at higher energies, for example at $\mu = 90$ GeV, $\alpha_{EM} \approx \frac{1}{127}$. To visualise this, a real electron exists in a *polarised* space in the vacuum by being *dressed* by a cloud of virtual photons and positively-charged particles attracted from the vacuum, which effectively *screens* the electric charge for a second real electron at rest a large distance from the first, which reduces the size of the coupling α_{EM} . Colliding the two electrons at high energy reduces their separation, with higher energy collisions increasing the effective coupling between them as each charge penetrates the virtual cloud of the other electron. As the electrons are separated at smaller distances at higher energies, the ability for the positively-charged virtual clouds to screen the electrons weakens, and their coupling strength increases.

In QCD, the antiscreening effect from virtual gluon loops dominates, and from Eq. 2.13 it can be seen that α_s decreases at higher energies and are asymptotically free and the strong force no longer manifests. At energies below $\mu = 300$ MeV, the coupling α_s is too strong and quark states too confined to probe the QCD structure, relying instead on experimental measurements and simulations in lattice QCD theories.

Renormalisation in the SM is a technique that allows for diverging terms arising from computation of certain quantities such as coupling amplitudes in Feynman diagrams containing loops. This accounts for the (anti-)screening effects in EM and QCD, by replacing the bare masses and couplings with their renormalised values, which are energy-scale dependent. The consequence of this is the bare masses and coupling strengths predicted by theory do not necessarily agree with the observed mass and coupling measurements from experiment, and that the removal of divergences means masses and couplings are dependent on the renormalisation scale. In QCD, this treatment is derived under perturbation theory, which is valid for energy scales $\mu > 300$ MeV.

In summary, by operating colliders above the renormalisation energy scale, QCD processes are a key component of proton-proton collision events as quarks and gluons are in free states directly after the collision, and offers a physical interpretation of the antiscreening effect which is a direct consequence of gluons carrying the colour charge conserved by the $\text{SU}(3)_C$ group.

2.2.3 Electroweak Unification

Electroweak unification is the combination of EM with the weak force. The gauge formalisms for the weak and EW forces are more elaborate than EM, as additional properties of the weak force such as electric charge, fermion flavour, weak isospin, field *chirality* and massive $W^\pm$ and Z bosons have to be accounted for. This is best represented by the $\mathbb{SU}(2)_L$ group, with L denoting LH fields. Given this also needs to encompass EM, the resulting symmetry group is given by the product $\mathbb{SU}(2)_L \otimes \mathbb{U}(1)_Y$, with the weak hypercharge denoted Y for which the symmetry group $\mathbb{U}(1)_Y$ is agnostic to $\mathbb{U}(1)_{EM}$.

Consider quarks in the SM⁵, for which there is an LH doublet, $\phi_1(x)$ that interacts with the weak force and EM, and RH singlets, $\phi_{2,3}(x)$ that interact only via EM. The quark fields are denoted

$$\phi_1(x) = \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad \phi_2(x) = u_R, \quad \phi_3(x) = d_R, \quad (2.15)$$

for which the Lagrangian for a free particle field, given by

$$\mathcal{L}_{free} = \sum_{j=1}^3 i \bar{\phi}_j(x) \gamma^\mu \partial_\mu \phi_j(x), \quad (2.16)$$

must be invariant under the following $\mathbb{SU}(2)_L \otimes \mathbb{U}(1)_Y$ gauge transformations:

$$\phi_1(x) \rightarrow \phi'_1(x) \equiv \exp [ig' y_1 \theta(x)] U_L \phi_1(x), \quad (2.17)$$

$$\phi_{2,3}(x) \rightarrow \phi'_{2,3}(x) \equiv \exp [ig' y_{2,3} \theta(x)] \phi_{2,3}(x), \quad (2.18)$$

where g' is the coupling constant for $\mathbb{U}(1)_Y$, and $U_L(x)$ is the $\mathbb{SU}(2)_L$ gauge transformation

$$U_L(x) \hat{=} \exp \left[ig \frac{\sigma_i}{2} \omega_i(x) \right] \quad (2.19)$$

acting on $\phi_1(x)$, where σ_i denotes the set of Pauli matrices, g is the coupling constant for $\mathbb{SU}(2)_L$, and $\omega_i(x)$ is the set of gauge transformations distinct from $\theta(x)$.

This leaves four gauge parameters, θ and ω_i from which four different gauge fields

⁵The equivalent for leptons is valid in this derivation.

are derived according to the covariant formalism introduced in Eq. 2.4, so giving

$$\begin{aligned} D_\mu \phi_1(x) &= \left[\partial_\mu - ig \frac{\sigma_i}{2} W_\mu^i(x) - ig' y_1 B_\mu(x) \right] \phi_1(x), \\ D_\mu \phi_{2,3}(x) &= \left[\partial_\mu - ig' y_{2,3} B_\mu(x) \right] \phi_{2,3}(x). \end{aligned} \quad (2.20)$$

For the covariant derivative $D_\mu \phi_j(x)$ to transform identically to the field $\phi_j(x)$, the gauge fields must satisfy the conditions

$$B_\mu(x) \rightarrow B'_\mu(x) \equiv B_\mu(x) + \frac{1}{g'} \partial_\mu \theta(x), \quad (2.21)$$

$$\tilde{W}_\mu(x) \rightarrow \tilde{W}'_\mu(x) \equiv U_L(x) \tilde{W}_\mu U_L^\dagger(x) - \frac{i}{g'} \partial_\mu U_L(x) U_L^\dagger(x), \quad (2.22)$$

where $\tilde{W}_\mu(x) \hat{=} \frac{\sigma_i}{2} W_\mu^i(x)$. Here, Eq. 2.21 and Eq. 2.5 are identical, and the analogous ϕ_j - B_μ coupling is free, i.e. the weak hypercharges y_j are arbitrary. This freedom does not apply to Eq. 2.22, whose extra terms are borne directly from the non-linear $\text{SU}(2)$ group, and permit a single coupling g .

Both the strong and EW gauge groups are referred to as *Yang-Mills* theories, whereby the theory is described by an N -dimensional special unitary group, $\text{SU}(N)$, for which there are N^2-1 massless gauge boson fields associated with that group. The strong force is represented by a $\text{SU}(3)$ gauge group for which there are eight dimensions, giving rise to the eight gluon colour states. Colour charge is the preserved quantity here. The EW force is described by an $\text{SU}(2) \otimes \text{U}(1)$ gauge theory, for which the photon and three weak vector bosons $W^\pm$ and Z^0 are the observed proponents of the boson fields B and W_i . In EW theory, weak isospin is conserved within the $\text{SU}(2)_L$ group, while weak hypercharge is conserved within the $\text{U}(1)_Y$ group. The $\text{U}(1)_{EM}$ field theory is linear, and not self-interacting i.e. there exist no photon-photon interactions, and photons will only mediate fermion interactions. Yang-Mills theories are non-linear and contain Lagrangian terms that imply self-interacting gauge bosons. This self-interaction is observed in the gauge-invariant kinetic Lagrangian terms, which are constructed from the field strength tensors

$$B_{\mu\nu} \hat{=} \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (2.23)$$

$$\tilde{W}_{\mu\nu} \hat{=} \partial_\mu \tilde{W}_\nu - \partial_\nu \tilde{W}_\mu - ig [W_\mu, W_\nu], \quad (2.24)$$

where

$$\tilde{W}_{\mu\nu} \hat{=} \frac{\sigma_i}{2} W_{\mu\nu}^i, \quad W_{\mu\nu}^i = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i + g \epsilon^{ijk} W_\mu^j W_\nu^k, \quad (2.25)$$

for which ϵ^{ijk} is the Levi-Civita tensor. In this case, the gauge transformation of $B_{\mu\nu}$ is invariant, while $\tilde{W}_{\mu\nu}$ transforms covariantly, i.e.

$$\tilde{W}_{\mu\nu} \rightarrow U_L \tilde{W}_{\mu\nu} U_L^\dagger, \quad (2.26)$$

and thus the normalised kinetic Lagrangian term is given by

$$\mathcal{L}_{kin} = -\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^i W_i^{\mu\nu}, \quad (2.27)$$

which contains quartic and quadratic terms for which the EW gauge bosons permissibly interact, which can be shown to correspond to the EW self-interaction diagrams in Fig. 2.1.




Figure 2.1: The set of EW self-interaction vertices as is permissible in the SM according to the EW kinetic Lagrangian term. Note that $\mathbf{SU}(2)_L$ algebra fails to generate a neutral vertex purely formed from photons and Z bosons.

So far the EW Lagrangian has yielded fermion-EW boson interaction and EW boson self-interaction terms, but is forbidden from yielding mass terms for bosons or fermions to preserve gauge invariance; therefore $\mathbf{SU}(2)_L \otimes \mathbf{U}(1)_Y$ symmetry only yields massless fields. This is because their masses emerge from the symmetry breaking of the EW gauge group.

2.3 Electroweak Symmetry Breaking

The process of *spontaneous symmetry breaking* (SSB) is the mechanism by which SM particles inherit their masses. The masses of weak gauge bosons derive from interaction with the BEH mechanism [2–7], which is considered the cornerstone of particle physics. The mechanism permits massive gauge fields to manifest while preserving

gauge invariance in the EW Lagrangian. The most elementary formalism by which the BEH mechanism exists in the SM begins with a complex scalar doublet,

$$\Phi = \begin{pmatrix} \Phi^+ \\ \Phi_0 \end{pmatrix} \hat{=} \frac{1}{\sqrt{2}} \begin{pmatrix} \Phi_1 + i\Phi_2 \\ \Phi_3 - i\Phi_4 \end{pmatrix}, \quad (2.28)$$

for which there are four unconstrained variables, or degrees of freedom, before SSB that correspond to the four massless gauge fields B and W^i . As established in Sec. 2.2.3, bare mass terms for gauge bosons violate gauge invariance here, so each boson has two degrees of freedom corresponding to two polarisation states. There is a total of twelve degrees of freedom from the combined massless gauge fields and complex scalar doublet.

The most generic scalar potential $V(\Phi)$ compatible with a gauge invariant $\mathbb{SU}(2)_L \otimes \mathbb{U}(1)_Y$ symmetry is given by

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (2.29)$$

for which the quadratic and quartic coefficients, μ^2 and λ , relate to the Higgs boson mass and the Higgs boson self-interaction coupling strength, respectively. Were these both positive in value, the potential would bear a single, stable minimum, which would not comply with a system with a broken symmetry and so would still have massless gauge fields⁶ This minimum corresponds to the physical vacuum, where the expected value of Φ is zero and therefore gauge invariant, visualised in 1D in Fig. 2.2 (red).

Instead, it is required that $\lambda > 0$ and $\mu^2 < 0$, for which the previously stable minimum is now an unstable maximum. The new minimum of the Higgs potential from SSB now forms a circle of degenerate points in 4D space, demonstrated in Fig. 2.3. The minimum, or *vacuum expectation value* (VEV), of $V(\Phi)$ occurs at a point v , shown in 1D in Fig. 2.2 (purple), where

$$v^2 = -\frac{\mu^2}{\lambda}. \quad (2.30)$$

The choice of vacuum state can be taken from any of the degenerate minima, as the symmetry applies at all points.

⁶Although it is true that EW bosons take their mass from the BEH mechanism, it is not strictly true that without this mechanism the EW bosons would be massless. Assuming a positive VEV (Eq. 2.30), the masses can be taken from QCD theory in quark-antiquark condensates, which possess the same quantum numbers as the Higgs boson. Much of this extends from the QCD model of composite bosons, which would have been a possible resolution had the Higgs boson not been discovered.




Figure 2.2: Demonstration of SSB for different solutions of the Higgs potential described in Eq. 2.29, where $\mu^2 > 0$ (red) has no non-zero solutions for the VEV v , while $\mu^2 < 0$ (purple) has a single local maximum and two minima along the ϕ_1 axis at v .

Expanding Φ about the VEV gives the form

$$\Phi = \frac{1}{\sqrt{2}} \exp \left[\frac{i\sigma_i \theta^i(x)}{v} \right] \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (2.31)$$

where $H(x)$ is the real Higgs component of the scalar field. Given the exponent can be gauged under the unitary transformation

$$U(x) \hat{=} \exp \left[-\frac{i\sigma_i \theta^i(x)}{v} \right], \quad (2.32)$$

which rotates away the fields $\theta^i(x)$, the resulting vacuum state can be inserted in the expansion of the scalar potential in Eq. 2.3, so giving

$$V(\Phi)|_{VEV} = \frac{1}{2} (2\lambda v^2) H^2 + \lambda v H^3 + \frac{\lambda}{4} H^4 + \text{const}, \quad (2.33)$$

for which there are cubic and quartic Higgs boson self-coupling terms preceded by a mass term. The expansion therefore gives the Higgs boson mass,

$$m_H = \sqrt{2\lambda v^2} = \sqrt{-2\mu^2}. \quad (2.34)$$




Figure 2.3: The Higgs potential in 4D space following SSB, which is consistent with a universe in which EW gauge bosons have mass. The minimum of $V(\Phi)$ occurs at v , where v^2 is defined in Eq. 2.30.

As such, excitations about the VEV of the Higgs field correspond to the presence of the Higgs boson. This highlights that the significance of the discovery of the Higgs boson particle by the ATLAS and CMS Collaborations [40, 41], with a mass of 126.0 ± 0.4 (stat.) ± 0.4 (syst.) GeV and 125.3 ± 0.4 (stat.) ± 0.5 (syst.) GeV respectively. Not only does this prove the presence of the BEH mechanism, but also provides a constraint on SM parameters dependent on the Higgs mass.

The gauge boson mass terms are sourced from the the matter coupling component of the EW Lagrangian,

$$\mathcal{L}_{matter} = (D^\mu \Phi)^\dagger D_\mu \Phi - \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2, \quad (2.35)$$

using the covariant derivative defined in Eq. 2.20 to ensure invariance under $\text{SU}(2)_L \otimes \text{U}(1)_Y$ transformations. Substituting the Higgs field translated according to the VEV as shown in Eq. 2.31 directly into Eq. 2.35 yields a series of terms quadratic in the gauge boson fields,

$$(D_\mu \Phi)^\dagger D^\mu \Phi = \frac{1}{2} \partial_\mu H \partial^\mu H + \left[\left(\frac{gv}{2} \right)^2 W_\mu^+ W^{\mu-} + \frac{v^2}{8} (g' B_\mu - g W_\mu^3)^2 \right] \left(1 + \frac{H}{v} \right)^2, \quad (2.36)$$

which uses a substitution for the $W^\pm$ boson term given by

$$W_\mu^\pm \triangleq \frac{W_\mu^1 \mp iW_\mu^2}{\sqrt{2}} \quad (2.37)$$

to parametrise the masses for the $W^\pm$ bosons,

$$m_W = \frac{vg}{2}. \quad (2.38)$$

The Z boson and photon fields are mixed within the neutral gauge bosons B_μ and W_μ^3 by the *weak mixing* or Weinberg angle, θ_W , such that

$$\begin{aligned} Z_\mu &= B_\mu \cos \theta_W - W_\mu^3 \sin \theta_W, \\ A_\mu &= B_\mu \sin \theta_W + W_\mu^3 \cos \theta_W, \end{aligned} \quad (2.39)$$

where the hypercharge-weak field coupling relation $\tan \theta_W = g'/g$. The resulting mass values are computed from the eigenstates of the above relations, as

$$\begin{aligned} m_Z &= \frac{v(g^2 + g'^2)^{\frac{1}{2}}}{2} = \frac{m_W}{\cos \theta_W}, \\ m_A &= 0. \end{aligned} \quad (2.40)$$

This demonstrates how the BEH mechanism breaks the $\text{SU}(2)_L \otimes \text{U}(1)_Y$ symmetry, yielding three massive weak bosons and a massless gauge boson consistent with the photon given the $\text{U}(1)_{EM}$ symmetry is unbroken after SSB. Considering the twelve degrees of freedom, after SSB three of these are taken by the massless gauge fields W^i to become massive, and mixed by the BEH mechanism to manifest as $W^\pm$ and Z . Massive gauge fields have three degrees of freedom corresponding to their spin, which uses a further three degrees of freedom. The photon term A remains massless and so only has two polarisation states. The complex scalar doublet becomes a real scalar singlet in the process, which is the Higgs boson, and occupies the last degree of freedom.

Conservation of the electric charge e is maintained after SSB, manifesting through Y and T .

Using Eq. 2.38 and Eq. 2.40, the Weinberg angle can be calculated purely from the weak boson masses

$$\cos \theta_W = \frac{m_W}{m_Z}, \quad (2.41)$$

which is a free parameter in the SM. From observation, $\theta_W = 28.21^\circ$ and $v = 246 \text{ GeV}$ [20].

2.3.1 SM masses and couplings

From the BEH mechanism, the Higgs coupling to vector bosons and fermions is well-described. Taking Eq. 2.36 which describes the EW boson masses and couplings to the Higgs field, and using the definitions in Eq. 2.40, the Higgs-vector boson (HVV) coupling Lagrangian can be derived to first order in H/v , given by

$$\mathcal{L}_{HVV} = \frac{2m_W^2}{v} W_\mu^+ W^{\mu-} H + \frac{m_Z^2}{v} Z_\mu Z^\mu H = g m_W W_\mu^+ W^{\mu-} H + \frac{g m_Z}{2 \cos \theta_W} Z_\mu Z^\mu H. \quad (2.42)$$

From this, the HWW and HZZ coupling exists with strength proportional to the square of the vector boson mass, i.e. $g_{HVV} \propto m_V^2/v$. Additionally, the Higgs boson does not couple to photons under the EW mechanism, given it is massless, and therefore can only be loop-induced via quarks or vector bosons. Quartic couplings $HHWW$ and $HHZZ$ are also predicted when extracting terms to second order in H/v , although these vertices are yet to be observed in nature.

The Higgs coupling to fermions is also the mechanism by which fermions have mass, where direct mass terms of the form $m_f \phi \phi$, for some fermion field ϕ and mass m_f , are not invariant under $\text{SU}(2)_L \otimes \text{U}(1)_Y$ gauge transformations. Instead, the fermion mass must arise from the Higgs mechanism, and is illustrated in the *Yukawa* interaction terms of the SM Lagrangian, so-called due to its HFF structure (where F denotes some fermion), given by

$$\mathcal{L}_{Yukawa} = -Y_d^{ij} \bar{Q}_L^i \Phi D_R^j - Y_u^{ij} \bar{Q}_L^i \Phi_c U_R^j - Y_e^{ij} \bar{L}_L^i \Phi E_R^j + \text{h.c.}, \quad (2.43)$$

where Q , U and D are the quark fields defined in Eq. 2.15, i.e. $Q = \phi_1(x)$, $U = \phi_2(x)$, and $D = \phi_3(x)$ which are composed of a linear combination of the mass eigenstates u and d , L and E denote the lepton fields, $Y_{d,u,e}$ denote 3×3 complex matrices, and $\Phi_c = i\sigma_2 \Phi^*$ for the scalar field Φ as given in Eq. 2.31. Taking the case of the first lepton generation, when the Higgs field is at its VEV, the Yukawa Lagrangian becomes

$$\mathcal{L}_{Yukawa,e} = -\frac{y}{2} (v + H) (\bar{e}_L e_R + \bar{e}_R e_L), \quad (2.44)$$

where y is the Yukawa coupling strength, and is the corresponding component in the Y_e^{ij} matrix for the electron. This result demonstrates that the mass term in the Yukawa Lagrangian is of the form $m_e = yv/2$, and fermions couple to the Higgs boson with HFF vertices. A similar derivation exists for quarks, however requires a change of basis to show the same relation. Therefore, fermions couple to the Higgs boson with a strength proportional to their mass, i.e. $y_f \propto m_f/v$. The Higgs boson couplings in

the SM and their coupling strengths are presented in Fig. 2.4.




Figure 2.4: The set of Higgs boson interactions with SM fermions and bosons, with their accompanying coupling strengths in proportions of m_V or m_f , and the VEV, v . The quartic coupling term has not been observed experimentally.

2.4 The parton model

As aforementioned, a proton consists of a udu quark configuration bound by gluons. At higher energies these gluons can produce virtual quark-antiquark pairs. To distinguish these from the real quarks, the udu quarks are referred to as *valence* quarks, and those from virtual processes are referred to as *sea* quarks. All quarks and gluons within the proton are referred to as *partons* in HEP, and form the basis of the parton model.

In high energy pp collisions, each parton carries some fraction of the momentum of the proton. The distribution of parton momenta follows a probability distribution called a *parton density function* (PDF), $f_i^{(P_j)}(x_i)$, akin to measuring the probability of observing a parton i within proton j with momentum x_i normalised by the total proton momentum. The total cross section σ_{pp} of the collision is the incoherent sum of all partonic cross sections $\hat{\sigma}$, i.e. the cross section for two partons colliding between two protons is given by

$$\sigma_{pp} = \int dx_1 dx_2 f_1^{(P_1)}(x_1) f_2^{(P_2)}(x_2) \hat{\sigma}(x_1 x_2 s), \quad (2.45)$$

where s is the centre-of-mass energy.

The PDFs and $\hat{\sigma}$ undergo a more elaborate prescription than the above, to account for other phenomena. Examples that have to be considered include the phenomenon of Bjorken scaling [44], which under certain assumptions show that PDFs vary with

momentum fraction x_i and independently of Q^2 , which illustrates how valence quarks behave as point-like particles in the proton. Corrections to this phenomenon are required due to scaling violation as the initial assumptions break down, and are explained hereafter. The sea quark momentum distribution must also be considered, where an infinite number of sufficiently low momentum quark-antiquark pairs can be radiated from gluons. From measurement, the valence and sea quarks carry approximately half of the momentum of the proton, with the rest carried by the gluons. Radiative corrections to $\hat{\sigma}$ due to emission of a single real or virtual gluon are also applied, which affects the initial energy entering the scattering. The resulting calculations introduce a *factorisation scale*, which similarly to the QCD renormalisation scale is required to negate diverging terms from initial gluon radiation collinear to the parton direction. Such collinear divergences are therefore absorbed into the PDF so long as their transverse momentum is below the factorisation scale. The new collision cross section following radiative corrections becomes

$$\sigma = \int dx_1 dx_2 f_1^{(P_1)}(x_1, \mu^2) f_2^{(P_2)}(x_2, \mu^2) \hat{\sigma}(x_1 x_2 s, \mu^2). \quad (2.46)$$

PDFs can only be extracted experimentally for a given proton energy, although can be transformed to different energy scales using the *Dokshitzer-Gribov-Lipatov-Atarelli-Parisi* (DGLAP) equations [45–47]. The DGLAP equations are required due to the evolution of PDFs in a collision, which can occur when the incoming parton is produced from the splitting of a parent parton, or the incoming parton itself splits. The evolution equation depends on the momentum fraction z split from the parton with momentum fraction x and PDF $f(x, \mu^2)$, and is given by

$$\mu^2 \frac{\partial f(x, \mu^2)}{\partial \mu^2} = \int_z^1 \frac{dz}{z} \frac{\alpha_s}{2\pi} P(z) f\left(\frac{x}{z}, \mu^2\right), \quad (2.47)$$

which is applicable to any splitting process, and allows for the prediction of the evolution of PDFs universally. Examples of splitting processes at leading order (LO) are presented in Fig. 2.5. Splitting functions have been computed in QCD to LO, next-to-LO (NLO) [48], and next-to-NLO (NNLO) [49], and are essential to PDF computation in the LHC.

In summary, in perturbative QCD values for masses, couplings, and PDFs cannot be predicted, however their evolution with the energy scale can be. The PDF can be measured for some process such as deep inelastic scattering, from which the quark and gluon PDFs can be accessed via the DGLAP equations which allow their PDF to




Figure 2.5: Examples of splitting functions at LO that are accounted for by the DGLAP equation, which can predict how the incoming parton PDF evolves due to these processes. The incoming parton, of momentum fraction x , splits to produce a new parton of momentum fraction z of x , with the splitting parton carrying momentum fraction $(1 - z)$ of x .

evolve with the kinematics of the system. From Run 1 and Run 2 data at the LHC, PDF measurements have been made for various collider objects, such as leptons and photons, with a good consensus of the PDF uncertainties by various object groups. The parton model is a good basis for visualising proton-proton collisions at the LHC, while the renormalisation and factorisation scales, and PDFs, for various processes form an important discussion for the analysis. Several processes of importance to $H \rightarrow \text{inv}$ analyses such as this are generated at various orders of precision in QCD, and contribute to the overall uncertainty of the results.

Chapter 3

The CMS Detector

The purpose of this chapter is to describe the operations at the CERN LHC and the setup of the CMS detector. CERN, or the *organisation européenne pour la recherche nucléaire*, hosts and coordinates the largest collider experiment, the LHC, amongst many others on the Franco-Swiss border, near Geneva. The LHC is primarily a pp collider that operates at 13.6 TeV as of the start of Run 3 in July 2022. This is by far the highest energy particle collider in the world, and provides an ideal environment for measuring SM parameters and producing high mass particles given its high operating energy and high luminosity (i.e. high data-taking rate). An in-depth discussion of the LHC is presented in Sec. 3.1. The resulting data is recorded in the CMS detector, which post-filtering of unwanted events provides data sets for use in analyses. The CMS detector setup and general performance is detailed in Sec. 3.2, while the modes of data acquisition (DAQ) and the state of data from Run 1 and Run 2 are explained in Sec. 3.3. The potential for data quality monitoring (DQM) using conventional statistics and machine-learning (ML) in future runs is covered with the introduction of AutoDQM, a statistics-orientated program introduced in Sec. 3.4.

3.1 The Large Hadron Collider

The LHC is a circular pp collider situated in a tunnel 27 km in circumference and 175 m underground at its deepest point. The main LHC operations site lies on the Franco-Swiss border, 8 km north-west of Geneva, with the ring largely lying inside France and extending as far west as the Jura mountains. A circumference of this size is necessary for accelerating protons in beams of energy $\mathcal{O}(\text{TeV})$, a prerequisite in the search for the Higgs boson for which the LHC was largely purpose-built. The technical design report for the LHC is detailed in Ref. [50].

The LHC has four lead experiments that lie at the four beam *interaction points* (IPs) about the ring: ATLAS, CMS, LHCb (LHC *beauty*, specialising in heavy flavour physics via studies involving B hadrons) and ALICE (A Large Ion Collider Experiment, specialising in studies of quark-gluon plasmas produced in heavy ion collisions). ATLAS and CMS [51] are the largest detectors, and are operated by the largest collaborations. These are general-purpose high energy detectors, with a range of specialist physics groups that use the resulting pp data for various purposes.

ATLAS and CMS are placed at opposite sides of the LHC ring, and are independent collaborations that both first verified the production of a particle consistent with the 125 GeV Higgs boson in 2012 [40, 41]. The discovery in the case of CMS used 5.1 fb^{-1} data at $\sqrt{s} = 7 \text{ TeV}$ and 5.3 fb^{-1} data at $\sqrt{s} = 8 \text{ TeV}$. Following the discovery in Run 1, the LHC has had two long shutdowns for upgrades to the ring and detector subsystems, either side of the $\sqrt{s} = 13 \text{ TeV}$ operating period in Run 2. With the CMS detector, a total of 177 fb^{-1} of data has been collected in Run 1 and Run 2. From July 2022 Run 3 has commenced, with pp collisions taking place at $\sqrt{s} = 13.6 \text{ TeV}$, and a target data intake of 280 fb^{-1} .

Before protons reach the LHC ring, they have to be boosted in smaller feeder accelerators first. A schematic of the LHC complex is shown in Fig. 3.1. The operation until the end of Run 2 was as so: hydrogen gas was injected into the 36 m long linear accelerator LINAC2, where a beam containing 3×10^{14} protons formed from ionisation of the gas. LINAC2 accelerated the protons to 50 MeV (as of Run 3, the 86 m long LINAC4 is used, boosting the protons to 160 MeV), from which the single injection provided the next link in the LHC chain, the Proton Synchrotron Booster (PSB), with a proton beam energy of 1.4 GeV (2 GeV for Run 3). Consistent between Run 2 and Run 3, the beam then reaches the Proton Synchrotron (PS), where it achieves a boost to 26 GeV before entering the SPS, to reach an energy of 450 GeV.

Two beams are finally accelerated into the 27 km LHC ring, antiparallel to one another, where after approximately 20 minutes each achieved (achieves) an energy of 6.5 TeV (6.8 TeV) for Run 2 (Run 3). The protons in the beam are synchronised into *bunches* via radio frequency cavities placed in isolated gaps along the ring. The synchronisation process in fact begins in the PS and SPS, and continues in the LHC. The bunches are finely separated into 25 ns packets by providing a longitudinal oscillating frequency for the protons in the beam, which controls the proton speeds within a bunch. This ensures any protons lagging behind the bunch are kicked with a precise energy to stay with the main bunch, while the speed of protons ahead of the bunch is reduced. The resulting distribution of individual proton energies per bunch is thereby




Figure 3.1: The Run 3 accelerator complex showing the boosting stages from LINAC4 to the PS and SPS required to generate a 6.8 GeV proton beam from hydrogen gas through a series of accelerators. Various LHC detectors and experiments are also shown here, including ATLAS and CMS.

minimised. The beam direction is maintained mainly with dipole and quadrupole electromagnets. The former generate 8.3 T magnetic fields and steer the beam around the ring, while the latter help to focus the beam. This is preferable given a tighter beam cross section increases the probability of collisions between the protons in contraflow beams at the IPs.

During Run 2, a single beam contained a maximum of 2,808 bunches, and these collided at discrete intervals. The advantage to this versus a continuous beam is that detectors can distinguish individual beam crossings, therefore can better filter out pileup interactions to more accurately calculate kinematic variables in a given beam crossing. The pileup is the number of additional interactions that occur within the same or nearby bunch crossings, which for Run 2 reached an average of 32 per crossing.

Crossings occurred at a rate of 40 MHz, therefore per-second 1 billion pp collisions took place at the IP.

3.1.1 Luminosity

The *luminosity*, $\mathcal{L}_{\text{inst}}$, of a particle collider is a measure of the ability to produce the desired number of interactions in the accelerator, and relates the rate of events, $\frac{dN}{dt}$, to the production cross section, σ_p , of a collider, via

$$\frac{dN}{dt} = \mathcal{L}_{\text{inst}} \cdot \sigma_p. \quad (3.1)$$

The quantity $\mathcal{L}_{\text{inst}}$ here is referred to as the *instantaneous* luminosity, and is an effective measure of the performance of the accelerator. The rate is in terms of N , the number of events. The production cross section σ_p is the measure of the number of final state products from the collision, although this is hard to estimate from theory given there are many possible collision final states. Instead, this is measured in the LHC as the total cross section of inelastic interactions, which has been extrapolated from LHCb data to be measured at 75.4 ± 5.4 mb [52]. A *barn* is a unit of area equal to 10^{-28} m², and is a useful measure of cross section on the subatomic scale.

The instantaneous luminosity can be increased in a circular collider by increasing the number of bunches per beam and tightening the geometric cross section of the beam through increasing the magnetic field strength of the quadrupole magnets. These improvements were made between Run 1 and Run 2 where the peak instantaneous luminosity increased from $\mathcal{L}_{\text{inst}} = 1.2 \times 10^{34}$ cm⁻² s⁻¹ to 2×10^{34} cm⁻² s⁻¹. In the LHC, when integrating $\mathcal{L}_{\text{inst}}$ over the time for which collisions take place, a measure of the amount of data recorded can be obtained. Integrating this value is more nuanced than the above would suggest, given $\mathcal{L}_{\text{inst}}$ varies during daily operations of the LHC. An example of change to $\mathcal{L}_{\text{inst}}$ is the number of protons in each beam decreasing as collisions take place, thereby reducing the $\mathcal{L}_{\text{inst}}$ from the peak value. The *integrated* luminosity, $\mathcal{L}_I$, is therefore given as

$$\mathcal{L}_I = \int dt \mathcal{L}_{\text{inst}} = \frac{N}{\sigma_p}. \quad (3.2)$$

The amount of data recorded in the LHC and by the detectors is measured as the $\mathcal{L}_I$, and is quoted typically in units of *inverse femtobarns*, fb⁻¹. Across Run 2, a total of 163.6 fb⁻¹ of data was delivered, of which 150.9 fb⁻¹ was recorded by the CMS detector. Using the inelastic interaction cross section above [52], this corresponds to

10^{16} pp collision events recorded by the CMS detector in the years 2015–2018. The values broken down by year for Run 1 and Run 2 are given in Table 3.1 [53, 54].

Table 3.1: The integrated luminosity (in fb^{-1}) and centre-of-mass energies associated with each year of the LHC operating and recording pp collisions with the CMS detector.

Year(s)	Run 1	2015	2016	2017	2018	Run 2	Run 1 and Run 2
$\sqrt{s}$	7 TeV + 8 TeV	13 TeV	13 TeV	13 TeV	13 TeV	13 TeV	13 TeV
LHC delivered	29.4	4.3	41.6	49.8	67.9	163.6	193.0
CMS recorded	27.4	3.9	38.3	45.0	63.7	150.9	178.3
CMS verified	24.8	2.3	36.3	41.5	59.7	139.9	164.7

The true number of collisions does not, however, correspond to the number of events recorded; that is to say, only 10^{16} events were recorded for Run 2, about one-thirtieth of the total number of collisions. Events are selected by *triggering* on a set of criteria, known as a *menu*, which is explained in Sec. 3.3. The number of proton interactions per bunch crossing was larger by on average a factor 20 during Run 1 and a factor 34 during Run 2.

3.1.2 Pileup

A single pp collision event is identified by finding the primary interaction vertex (PV) corresponding to the *hardest scattering* process as evaluated using tracker information alone, as described in Ref. [55], and then calculating which *soft parton showers* or jets, where gluons and quarks hadronise to visible particles that further decay and interact with the detector until too low energy to be detected, belong to that event. During the bunch crossing that produced the event, additional soft *pileup* interactions take place in the LHC: specifically, during Run 2 an average of 34 pileup interactions per bunch crossing were recorded by the CMS detector, with a peak pileup of more than 50 interactions. The mean number of pileup interactions per year is presented in Fig. 3.2.

In the detector, these interactions overlap with the event of interest, making reconstruction difficult. Potential new physics at low-to-mid range energies is hard to observe as a result, and as $\mathcal{L}_{\text{inst}}$ is increased in the collider, the contamination becomes a greater obstacle. Since the start of the LHC, however, pileup mitigation techniques have been employed, and have grown more sophisticated with time [56–58].




Figure 3.2: The average number of interactions per year, effectively a measure of the pileup, of Run 1 and Run 2, presented as luminosity as recorded by the CMS detector per unit of pileup.

The CMS tracking detector can distinguish the PV and interaction vertices of charged particles due to the event of interest from vertices due to pileup interactions. Neutral particles make up 40% of a pp collision on average. In the target event, neutral particles are typically aligned with the charged particles from the same event, while neutral and charged particles from pileup vertices are more likely to shower isotropically in the detector. The technique of filtering the pileup interactions from the event of interest like this is referred to as pileup per particle identification (PUPPI) [59, 60], and is illustrated in Fig. 3.3. Here, the charged particles from pileup vertices are identified and rejected. Then, neutral particles are weighted according to the probability that they come from a pileup vertex based on the charged particle distribution surrounding them. An in-depth explanation at the reconstruction level is detailed in Sec. 4.3.

An example of pileup mitigation using the accelerator itself is by increasing the angle at which the beams cross. Given the LHC is circular, the crossing points are designed such that the angle at which the beams collide can be manipulated slightly. A larger angle decreases $\mathcal{L}_{inst}$ and the interaction cross section σ_p but also has the effect of reducing the number of pileup interactions. In conjunction with the vertex tracking and identification algorithms in the ATLAS and CMS detectors, pileup is




Figure 3.3: The PUPPI technique of identifying vertices based on the charged particle distribution, rejecting those coming from pileup vertices. The neutral particles are weighted based on the distribution of charged particles in its vicinity.

reasonably mitigated in Run 1 and Run 2 data.

3.2 Compact Muon Solenoid

CMS is a general-purpose particle detector extending 29 m in length, 15 m in diameter and weighing 14 kilotonnes. It was constructed in 15 layers above ground at its location, Point 5, in Cessy, France, before being lowered 100 m into the cavern in individual stages for assembly. The layers of the detector are built around a large solenoid magnet, giving CMS its namesake. The solenoid is a superconducting magnet that measures 13 m in length, 6 m in diameter and is cooled to 1 K. Operating at 4 T, it is the most powerful magnet of its kind ever built, and bends the trajectories of charged particles as they travel out of the detector. Equating the magnetic force to the centripetal force of particles traversing the detector in the transverse plane, the magnitude of a particle's momentum in the plane transverse to the beam, $\vec{p}_T$, is given by

$$\vec{p}_T = q \vec{r} \times \vec{B}, \quad (3.3)$$

where q is the charge of the particle, while $\vec{r}$ is the radial distance of the particle and $\vec{B}$ is the strength of the magnetic field flowing longitudinally, both measured with respect to the beampipe. The solenoid extends radially to 3 m then is confined within a steel *yoke* [51], as illustrated in Fig. 3.4.

To accurately reconstruct an event backwards from the final state particles in the detector subsystems to the pp collision IP, CMS uses a coordinate system to which every subsystem module is sensitive. For a cylindrical detector, a cylindrical coordinate




Figure 3.4: A graphic of the CMS detector, illustrating the steel yoke around which the silicon tracker, calorimeters and muon detector is situated.

system is employed, with the polar (azimuthal) angle θ (ϕ) measured with respect to the anticlockwise beam axis. The pseudorapidity, η , is an effective measure of θ that is determined by

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right], \quad (3.4)$$

Thus, the p_T can be calculated as a function of η given

$$p_T = \frac{|\vec{p}|}{\cosh \eta}. \quad (3.5)$$

To measure the angular distance between two particles in the topological space of the detector, the variable ΔR is defined as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \quad (3.6)$$

This parameter serves to represent the size of certain objects in clustering, and is a good measure of overlap between reconstructed objects.

The detector reaches a coverage of $|\eta| < 3.0$ for all SM particles except neutrinos, and up to $|\eta| = 5.2$ for all except muons and neutrinos. Neutrinos are the only SM particle to escape the CMS detector undetected, given they have no charge and only very weakly interacts with the material in the detector.

Each subsystem is designed to probe different particles produced from collisions. Subsystems are modelled as a cylinder and are divided into *barrel* and *end cap* (EC) components, corresponding to the curved surface and the circular planes at each end of a cylinder, respectively. The beampipe runs down the axis of the cylinder, with the IP focused in the longitudinal centre.

3.2.1 Tracker

The tracker is situated about the IP, and can detect charged particles from between 4 cm to 11 cm from the beampipe [61, 62]. For most of Run 2, $100 \times 150 \mu\text{m}$ silicon pixel sensors equating to 124 million readout channels lay closest to the beam, structured into a barrel and ECs providing coverage over a surface area of 1.9 m^2 . The pixel detector consists of four concentric barrel layers up to 0.16 m in the radial direction, and three disk layers up to 0.516 m along the z axis, shown in green in Fig. 3.5. Silicon strip detectors make up the outer layers of the tracker, indicated in red and blue segments for single- and double-sided strips in Fig. 3.5, and cover up to 1.1m radially, and 2.7m in the z axis. The tracker used in 2016, consisting of 1440 silicon pixel and 15,148 silicon strip detector modules, collectively measured charged particles within the range $|\eta| < 2.5$, with the barrel-to-EC boundary at $|\eta| = 1.6$. This was upgraded at the start of 2017, with the new tracker [63] consisting of 1856 silicon pixel and 15,148 silicon strip detector modules, and measured particles up to $|\eta| < 3.0$ during the 2017 and 2018 data-taking periods [64]. The strip and pixel sensors, and electronic readout equipment, are designed to withstand the high fluence conditions due to constant bombardment from highly radioactive collision products.

Charged particles interact with the tracker, where sophisticated track-finding algorithms are used to reconstruct the trajectories of several particles simultaneously from the hits recorded in the tracker using pattern recognition software [58]. These algorithms determine whether a sequence of 10 to 14 hits corresponds to a single trajectory with position resolution within $10 \mu\text{m}$. Particle charge and $\vec{p}_T$ can be calculated based on the track curvature inside the 4 T magnetic field. In this context, $\vec{p}_T$ is the vectorial sum of the momenta, as measured in the plane transverse to the beampipe, of particles




Figure 3.5: A graphic of the CMS tracker in the r - z plane used for 2017 and 2018, demonstrating the layering of silicon strip detectors in the barrel and EC. The barrel-to-EC boundary is situated at $|\eta| < 1.6$, and the maximum acceptance of the tracker is at $|\eta| < 3.0$. The pixel detector, shown in green, is situated entirely within the inner barrel, with the red and blue segments representing the single- and double-sided strip detector modules, respectively.

in an event. Formally, for each object i with transverse momentum $\vec{p}_{T,i}$,

$$\vec{p}_T = \sum_i \vec{p}_{T,i}. \quad (3.7)$$

For the proton-proton collision, this is zero as the partons are accelerated longitudinally, therefore the contribution owing to objects that have failed reconstruction, have fallen outside of the acceptance of the detector, or have not been detected can be computed. The *transverse energy* E_T is the magnitude of $\vec{p}_T$. Vertex reconstruction algorithms then determine the location of the PV, and some *secondary vertices*, for example from t quark and tau lepton decays to b quarks. Secondary tracks can be determined to belong to the PV by measuring the *impact parameter*. This is the closest distance between the PV and a particle track from a secondary vertex. Vertex reconstruction is an important aspect of event reconstruction as features in this analysis, and leads to efficient identification of final states containing b quarks, differentiating between prompt top quark decay, $t \rightarrow bW$, and tau lepton decay from a displaced vertex, $\tau \rightarrow \nu_\tau W$.

Track reconstruction efficiency is affected by the precision of $\vec{p}_T$ measurements. For muons, the efficiency is practically 100% within the tracker acceptance, whereas for electrons and hadrons this falls to 95% in the central region and as low as 80% in the forward regions. In the 2016 tracker, nonisolated particles in the range $1 < p_T < 10$ GeV and $|\eta| < 1.4$ had a typical p_T resolution of 1.5% and impact parameter resolution of 25–90 (45–150) μm in the transverse (longitudinal) direction [58]. The upgraded

tracker for 2017 and 2018 produced typical p_T resolution of 1.5% in p_T and impact parameter resolution of 20–75 μm in the transverse plane [65] for nonisolated particles of $1 < p_T < 10$ GeV. The $\vec{p}_T$ resolution for reconstructed tracks in the central region is 1% to 2% for muons with $p_T = 100$ GeV, and 5% for muons with $p_T = 1$ TeV; the resolution generally degrades in the forward region, with the fall off a lot faster for muons with lower momenta in the range $1 < p_T < 10$ GeV due to multiple scattering. The variation of the track reconstruction efficiency of muons using 2017 data with p_T and η is shown in Fig. 3.6, following an internal study by the CMS Collaboration.




Figure 3.6: Track reconstruction efficiency for muons as a function of muon p_T (left) and η (right), which is practically 100% efficient for low p_T and low $|\eta|$.

3.2.2 Electromagnetic Calorimeter

Two calorimeters surround the tracker. The inner of these is the *EM calorimeter* (ECAL), a hermetic, homogeneous cylinder housing 75,848 fine-grained lead tungstate crystals, each weighing 1.5 kg [66]. The ECAL is designed to detect mainly photons and electrons leaving the tracker, although charged particles can interact with the dense crystals. The ECAL central barrel (EB) covers a range $|\eta| < 1.48$, with the ECAL ECs (EE) extending up to $|\eta| = 3.0$, exceeding tracker acceptance. The ECAL crystals collectively cover a volume of 11 m^3 and weigh 92 tonnes. A preshower detector (ES), populated by planes of 140,000 silicon strips and 4,300 sensors, is placed in front of the ECs, with an acceptance range of $1.65 < |\eta| < 2.6$ and a thickness of 26 mm used to improve spatial resolution between photons and π^0 .

The crystals scintillate on interaction with charged particles, with two $5 \times 5 \text{ mm}^2$

avalanche photodiodes (APDs) per crystal used to read out the signal in the EB, and single 280 mm² vacuum phototriodes (VPTs) used in the EE, which are more tolerant to the higher concentration of radiation in the longitudinal direction. Given these crystals are not segmented longitudinally, measurements of photon angles to the IP depend on vertex reconstruction in the tracker.

ECAL readout is sensitive to energy deposits between 50 MeV and 2 TeV, and is optimised to minimise event pileup. Electronics are linked in clusters of 5×5 crystals, referred to as a trigger tower in the EB and a super crystal in the EE, and transfer readout through a chain of digitisers and amplifiers to the trigger system. The EM energy resolution, σ_E^{ECAL}/E , can be generalised to

$$\frac{\sigma_E^{\text{ECAL}}}{E} = \sqrt{\left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2}, \quad (3.8)$$

where E is the energy (in GeV), S is a stochastic term, N is the contribution due to noise in the ECAL electronics, and C is a constant. Using Run 1 data, these terms were calculated to be $S = 0.028 \text{ GeV}^{\frac{1}{2}}$, $N = 0.12 \text{ GeV}$ and $C = 0.003$, with a resulting σ_E^{ECAL}/E of 0.50% [67] for a 100 GeV electron. By comparison, σ_E^{ECAL}/E has been measured to between 1% and 3.4%, depending on $|\eta|$ region, using electrons from $Z \rightarrow ee$ decays throughout Run 2 [68].

In the Higgs sector, the processes $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$, the discovery channels used for the Higgs boson in Run 1, as well as $H \rightarrow WW^*$, $H \rightarrow b\bar{b}$ and $H \rightarrow t\bar{t}$ decay processes, are reliant on ECAL information. The key particles probed by the ECAL are electrons and photons, which deposit their energy in showers before stopping completely inside the calorimeter. Typically, charged hadrons, muons and tau leptons are not stopped in the ECAL, instead only depositing a small amount of energy, therefore other subsystems are needed to reconstruct processes involving these particles.

For high precision ECAL measurements, correction factors need to be applied to readout. Namely, as electrons and photons pass through the crystals, they produce EM showers which appear as clusters of hits in the readout. A correction is required to account for hit clusters belonging to one object. Another correction is the degradation of relative response in the laser monitoring system for measuring crystal transparency. Over the lifetime of the ECAL crystals radiation damage causes them to become less responsive during data taking insofar as, for 2018, the response fell by 10% in the EB and as much as 99% in the EE. For Run 1 and Run 2, the degradation in response over time is shown in Fig. 3.7 [69], demonstrating the importance of response corrections.




Figure 3.7: The degradation of ECAL response over time during data-taking periods throughout Run 1 and Run 2. The response is seen to fall 10% in the EB and as much as 99% in the EE by 2018. Response corrections are essential in precisely measuring energy deposits in the ECAL.

3.2.3 Hadronic Calorimeter

The outer of the two calorimeters surrounding the tracker is the *hadronic calorimeter* (HCAL), which measures energy deposits of composite particles that escape the ECAL. The HCAL comprises of brass plate absorbers and plastic scintillators in the HCAL barrel (HB) and HCAL ECs (HE), which stop hadrons corresponding to quark and gluon signatures from the IP [70, 71]. The HB covers the region $|\eta| < 1.30$ and is composed of absorber-scintillator wedges, while the HE covers range $1.30 < |\eta| < 2.96$ and is arranged into discs. These sections act as a sampling calorimeter, with light from the plastic scintillator collected by wavelength-shifting (WLS) and optic fibres. The WLS light is transported to a series of silicon photomultipliers (SiPM), a multipixel APD offering signal gains between 10^4 and 10^6 . The charge in each 25 ns interval is integrated, and the SiPM output is then received by a digitiser that can filter event information through back-end electronics.

The forward hadron (HF) calorimeter is a Cherenkov calorimeter made from steel absorbers and quartz fibres which run along the beampipe axis and are bundled into

wedges to capture Cherenkov radiation within acceptance $3.0 < |\eta| < 5.2$, covering the very forward regions of the detector. Each of the HF calorimeter halves consists of 432 readout towers composed of long and short quartz fibres lying parallel to the beampipe. The long fibres extend along the full 165 cm depth of the HF calorimeter while the short fibres are laid from a depth of 22 cm from the front of the detector. The fibres are spaced 5 mm apart, and optical-to-electrical signal conversion is performed using photomultiplier tubes. After transport to the phototransducers, the signal is fed to the same electronics chain used by the HB and HE for readout. The signal is read out from the two sets of fibres individually, in order to distinguish showers produced by electrons and photons, which deposit a large fraction of their energy in the long fibres, from showers produced by hadrons, which typically deposit their energy equally between the long and short fibres. The HF calorimeter is essential for forward jet identification, and determining the missing energy by providing more information on p_T measurements than the calorimeters in the central region alone, and additional luminosity measurements. Higgs boson production in the VBF channel relies on forward jets to complement Higgs boson coupling and decay width measurements.

The HCAL is segmented in η - ϕ coordinates which map to *calorimeter towers* projected radially from the IP, with each tower containing a given number of HCAL scintillator modules and ECAL crystals depending on location in the detector. The granularity, or density of towers, is greatest in the central region about $\eta = 0$, and decreases with increasing $|\eta|$. Combined with the ECAL, the calorimeters measure energy deposits corresponding to most particles produced in the detector, which are used to identify and reconstruct namely jets, charged leptons and photons.

Corrections to the HCAL response and clustering, much like in the ECAL, are applied during data taking, and an intercalibration correction between the two subdetectors is necessary to coordinate and maintain precision measurements. The hadronic energy resolution, σ_E^h/E applicable to the combined ECAL and HCAL system, is given by

$$\frac{\sigma_E^h}{E} = \sqrt{\left(\frac{S}{\sqrt{E}}\right)^2 + C^2}, \quad (3.9)$$

where E is the energy (in GeV), S is a stochastic term and C is a constant [71]. For Run 1, these value were estimated for both the HB and HE as $S = 0.847 \text{ GeV}^{\frac{1}{2}}$ and $C = 0.074$ [72], and for the HF calorimeter as $S = 1.98 \text{ GeV}^{\frac{1}{2}}$ and $C = 0.09$ [73]. Given forward jets are typically of very high energy, S is expected to be higher for the HF calorimeter than the inner calorimeters.

3.2.4 Muon Chamber and Yoke

The outermost subsystem consists of the muon detector embedded inside the steel return yoke, extending from 3 to 7.5 m. The yoke is a 12.5 kilotonne frame within which the muon subsystems are concentrically layered. It operates much like the solenoid, however at only 2 T and returns a magnetic field in the opposite direction to that of the inner detector. Similarly, the $\vec{p}_T$ of a particle traversing the outer detector can be measured with Eq. 3.3.

Muons are produced in pp collisions with very high energy and $\vec{p}_T$. Their rest mass of 0.106 GeV is $\mathcal{O}(200)$ times larger than that of the electron, and they are minimum ionising particles, meaning their energy loss rate is close to minimal such that they can escape the detector. The signatures for the core set of particles produced in pp collisions is presented in a transverse slice of the CMS detector in Fig. 3.8, showing how the muon escapes the detector and therefore only its tracking can be performed. Muons rarely deposit significant amounts of energy in the calorimeters, hence to perform precision studies of muons outside the tracker, a dedicated subsystem has been installed within the return yoke, comprising of gas-filled drift tubes, cathode strip planes, and resistive plates, all of which rely principally on gaseous ionisation by passing muons for precise position and timing measurements [74, 75]. For tracking performed in a subsystem on the scale of the muon chambers, it was necessary to opt for drift tracking over a system akin to that of the silicon tracker, although there is typically no loss of position and time resolution by comparison. Muon data is a key component in several studies including Higgs physics, given the muon is often a byproduct of signatures of interest that can be easily resolved over background processes in the LHC.

Within the barrel, covering $|\eta| < 1.2$, resides 480 resistive plate chambers (RPCs) and 250 drift tubes (DTs), while in the ECs, covering $0.9 < |\eta| < 2.4$, is situated 576 RPCs, reaching $|\eta| = 1.9$, and 540 cathode strip chambers (CSCs). An overlap region of barrel and EC modules exists in the range $0.9 < |\eta| < 1.2$, to ensure no gaps in acceptance exist. The muon chambers are arranged into four stations layered radially outwards, embedded within the return yoke.

The muon chambers are designed to perform three tasks: triggering on events producing single and multiple muons, muon identification and high resolution momentum measurements. The latter is provided by the solenoid and return yoke, which offer a high magnetic field strength and high spatial resolution far from the IP. The momentum and charge of charged objects in the muon detector is determined for a large kinematic range.




Figure 3.8: A transverse slice of the CMS detector showing how different particles produce signatures as they traverse the four main subsystems. The silicon tracker shows tracks of charged particles which bend under the influence of the superconducting solenoid. Electron and photon energy deposits are shown as showers in the ECAL, with heavier hadrons such as π^0 showering likewise in the HCAL. Muons traverse the detector without depositing much of their energy, but appear in the muon chambers as a track leaving the detector. The bend in their trajectory is in the opposite direction to that of their path in the tracker given the direction of the magnetic field of the return yoke is antiparallel to that of the solenoid.

The DTs are arranged into long aluminium cells with an anode wire along its centre. Muon position is determined in the DTs by calculating the drift time of the electron liberated by the traversing muon to the anode wire, under the influence of an optimised electric field. Each DT chamber contains 8 drift cells, providing a spatial resolution of $100\text{ }\mu\text{m}$. In the CSCs the cathode strips run radially outward, and operate much like the DTs, although with additional resolving power given by the charge induced on each strip. Advantages to using CSCs include precise position and time resolution of muon hits, and ability to operate under high exposure and in a non-uniform magnetic field, giving for a single CSC a position and time resolution respectively of around 50 to $140\text{ }\mu\text{m}$ and 3 ns . The RPCs are dual-layer chambers each containing gas under a high potential using a parallel plate capacitor made with graphite-Bakelite electrodes. The result is an optimal time resolution of 1.5 ns per RPC for fast trigger readout.

The momentum resolution in the CMS detector uses information from both the tracker and muon detector subsystems where available. For Run 2 muons with momenta up to $p_T = 100\text{ GeV}$, the scale- and momentum-corrected resolution is 1% in

the barrel and 3% in the ECs [74].

3.3 Triggering and Data Acquisition

At the design luminosity of the LHC, the CMS detector operates under a pp collision rate exceeding 1 GHz, which would require a bandwidth $\mathcal{O}(1 \text{ PB/s})$ to record all collisions. To filter this to a rate which can be handled by the detector's electronic readout, a trigger is used to select events of potential interest. CMS uses a two tier trigger system: the first is the level 1 trigger (L1T), consisting of processors that reduce data output to 100 kHz; the second is the high level trigger (HLT), a software farm that reduces event readout to a manageable $\mathcal{O}(1 \text{ kHz})$. The trigger is part of the DAQ system, which stores events of interest as raw data in ROOT file containers for final processing before being released as analysis object data (AOD) for use by analysers.

3.3.1 Level 1 Trigger

The L1T is the first layer of identifying candidate objects, such as ionisation deposits matching to a muon signature, or energy clusters consistent with a photon, electron, tau lepton or jet. To do this, a trigger *menu*, or series of event-filtering algorithms, is programmed into custom hardware via FPGA firmware chips [76, 77]. The target is to reduce the pileup-subtracted event rate of 40 MHz down to 100 kHz. Under a fixed latency synchronised to the LHC clock, the L1T decides whether to accept or reject an event based on calorimeter or muon trigger information within 4 μs of the collision producing that event.

Trigger primitives (TPs) are crude information sent by the calorimeter and muon subsystems to be processed by the L1T. TPs are processed at 40 MHz separately for calorimeter and muon chamber readout before being combined and evaluated in the global trigger (GT) that ultimately determines if an event is L1 accepted (L1A).

The L1 calorimeter trigger utilises a two-layer operation. Calorimeter TPs from the ECAL and HCAL containing p_T and quality flags are transmitted via $\mathcal{O}(10^3)$ multi-gigabit fibre optic links to the Layer 1 calorimeter trigger which then processes the information in parallel. The TP energy deposits are summed into calorimeter towers, with local corrections applied to p_T and kinematic quantities with η -, energy- and cluster shape-dependent calibrations before being passed to the Layer 2 calorimeter trigger via high-bandwidth links. Layer 2 clusters towers together to form calorimeter objects. Algorithms find and reconstruct these objects, e.g. electrons, photons, tau

leptons and jets, alongside additional calibrations and pileup subtraction. Global energy sums are also defined at this stage. Across the L1T vectorial energy sums are calculated using tracker information, such as the missing transverse momentum, $\vec{p}_T^{\text{miss}}$, which is the negative sum of transverse momenta $\vec{p}_T$ of objects defined in the L1T in an event. Formally, for each trigger object i with transverse momentum $\vec{p}_{T,i}$,

$$\vec{p}_T^{\text{miss}} = - \sum_i^{\text{HLT}} \vec{p}_{T,i}, \quad (3.10)$$

for which the *missing transverse energy* p_T^{miss} is the magnitude of $\vec{p}_T^{\text{miss}}$. The Layer 2 calorimeter trigger then sends this information to the GT.

Each of the DT, CSC and RPC systems in the muon detector are combined in the L1 muon trigger to optimise coverage and redundancy. For the DT and CSC systems, front-end electronics register hit information to identify muon track segments in each muon station. Track segments are processed by regional muon track finders (MTFs), which apply pattern recognition algorithms to identify muon candidates and compute their momenta. Track segment information is shared between the CSC and DT track finders for improved coverage in the overlap region. RPC hits are transmitted directly from front-end electronics to a pattern comparator trigger, a series of logic boards that determine muon candidates. Up to 4 muon candidates from the CSCs and DTs and up to 8 from the RPCs are received by the global muon trigger (GMT) per bunch crossing, with each candidate assigned $\vec{p}_T$, quality factor and position coordinates. The kinematic information is sufficient to eliminate candidate degeneracy, and quality checks are required to produce good muon reconstruction.

The final L1T stage is the GT, which discriminates over the final set of event objects from both the Layer 2 calorimeter trigger and the GMT for L1A. Often the menu requirements are based on e.g. p_T threshold of a given single object or topological matching based on coincidence of several objects from an event. This usually involves taking the highest $\vec{p}_T$ jets or clusters, determines if these objects pass the menu criteria, then for L1A a signal ID is transferred to the subsystems to readout the event data for DAQ. This selection implements between 350 and 400 trigger criteria before passing information to the HLT algorithms at a rate of 100 kHz. A summary of the L1T from the detector subsystems to the GT is presented in Fig. 3.9.




Figure 3.9: Schematic of the L1T from the ECAL, HCAL and muon chambers, to the GT, with the subsystem-specific algorithms applied to TPs via a series of trigger menus. This reduces the rate of event data taking from 40 MHz to 100 kHz, which is delivered to the HLT.

3.3.2 Layer 2 Calorimeter and DQM Overview

The layer 2 calorimeter in the L1T is physically a micro telecommunications computing architecture (microTCA) crate containing nine master processor (MP7) cards, referred to as *nodes* [78, 79]. Each node reads a bunch crossing at a time, all of which are time-multiplexed. The concept of a time-multiplexed trigger (TMT) was introduced for Run 2, where from Run 1 the bunch crossing window decreased by half to 25 ns, pileup increased to a maximum average value of 50, and electron/photon event rates increased by as much as a factor 10 to 50 kHz. For analysis to remain efficient at low energies and enhanced signature recognition for L1A, data from a single event is passed to 18 calorimeter trigger processor (CTP7) cards that format and preprocess data from various regions in the ECAL, HCAL, and HF calorimeter in Layer 1, and

is then passed to the 9 MP7 nodes of Layer 2. TMT means to feed data to the nodes based on timestamp over processing triaged by detector segmentation (as done in Run 1), which ensures each MP7 node has access to full calorimeter information from the Layer 1 cards per bunch crossing, hence full calorimeter granularity down to the tower geometry can be processed, with the necessary algorithms to cluster towers into calorimeter objects held on single FPGAs in the MP7 card.

A demultiplexer board (DEMUX), also an MP7 card, formats data processed by each node for the GT. An AMC13 card is used in the timing and control distribution system (TCDS) to apply clock and signal timestamps. The trigger, timing and control (TTC) readout is then transferred by optic fibre for the GT and DAQ. The microTCA crate is configured as shown in Fig. 3.10.




Figure 3.10: MicroTCA crate setup as used in the layer 2 calorimeter trigger. A series of TMT nodes receive full calorimeter tower granularity from the layer 1 CTP7 pre-processors, and layer 2 algorithms are applied by single FPGAs on a single MP7 node per bunch crossing. Readout is TTC-stamped by the TCDS, and is passed to the GT and DAQ.

In the process of DAQ, DQM is performed to ensure the data readout from the detector is accurate, all subsystems are functioning normally and there's no obvious interference with the readout when a run is being recorded. DQM is performed both online and offline. Online DQM involves real-time diagnostics by *shifters*. These are people who study the output from each subsystem and trigger in CMS while a run is active. Detector experts who are on-call during runs for extended periods are referred to as DOCs (detector on-call), and are responsible for making final decisions for DQM

and detector operations. CMSSW modules process data as it becomes available, and histograms of readout are generated to aid the job of shifters. Offline DQM uses runs for which the full data set is available, some time after the run has ended. This tends to be for efficiency or module performance studies. Specifically for the layer 2 calorimeter trigger, this can be anything from comparing calorimeter object kinematics between the data and emulator-generated results to studying diagnostic plots such as bunch crossing timing and object isolation.

In layer 2 DQM, object variable distributions are a good means of verifying readout and performance. Studying central and forward jet histograms can help trace issues relating to a faulty calorimeter module by localising spikes or dead zones in the p_T or occupancy distributions in η - ϕ plots. 1D plots for kinematic variable distributions may show anomalous peaks corresponding to an unusually high rate from the calorimeters. Diagnostic plots such as bunch crossing occupancy and timing are better suited for debugging firmware or checking for spikes in data-taking rates that are due to faulty trigger links or from trigger tower geometry. These issues can be diagnosed by tracing back from layer 2 through the L1T to output directly from the ECAL or HCAL.

Spikes in the rate of readout from the ECAL or HCAL subsystems are referred to as *hot towers*, and are the lead cause of high L1T calorimeter rates. If the layer 1 trigger readout rate is high for both the jet and the electron and photon collections, **SingleJet** and **SingleEG** respectively, then the hot tower is from the ECAL; if only from **SingleJet**, the hot tower is from the HCAL.

Towers are masked on-the-fly by ECAL or HCAL DOCs to bring rates down to normal, although if this persists, the cause of high rates may be a faulty trigger link. These occur between the ECAL/HCAL readout and layer 1, and typically produce a characteristic set of spikes. For Run 3, with the increased rate of object monitoring, a record amount of data passes through the DEMUX in the layer 2 trigger, therefore early diagnosis of hot towers or faults in firmware is essential.

A typical study done in layer 2 is for asymmetries in each η hemisphere of the detector across a series of runs, checking for hot towers, *holes* where towers are dead in the detector or have been masked already (leaving zero occupancy), or large dead zones due to subsystem malfunction or inactivity. A simple check of the layer 1 histograms can identify if a mask has been applied or not. These checks are performed for all layer 2 calorimeter objects, which include: central and forward jets, isolated and non-isolated EM showers, and isolated and non-isolated tau leptons. Typical distributions are produced in 1D and 2D at this level for each object, including: occupancy, η - ϕ occupancy, η - ϕ - p_T , η , ϕ and isolation. Energy sum distributions are also available in

layer 2, such as p_T and similar kinematic variables for both the central and forward regions. For shifters, checking every histogram for these patterns can be exhaustive, and discerning a good reference run from a bad one is difficult without the use of software. One machine-learning (ML) orientated program is AutoDQM, a statistics-orientated toolkit introduced in Sec. 3.4.

3.3.3 High Level Trigger

The HLT hardware contains a single processor farm housing 22,000 CPU cores to process L1A readout [80, 81]. The HLT menu consists of more than 400 paths dedicated to reconstructing objects and identifying events of interest to analysts. Each path functions to sequentially trigger on events using filter and reconstruction modules sorted by increasing complexity. HLT software uses a CMSSW framework, with the menu containing complex and in some cases bespoke triggers designed to target topologies of interest to certain analysis groups. An example of this is the use of two triggers targeting VBF topologies. The algorithms in the path are triaged such that those fastest at processing are applied first, and the offline-style reconstruction algorithms run last. Full-precision collision data that was initially input into the L1T and identified as L1A is accessible by the HLT, as well as tracker information. The HLT acceptance must be determined within 220 ms per node to handle the 100 kHz input rate from the L1T.

HLT algorithms of note are the anti- k_T jet clustering algorithm [82, 83] to classify light and heavy jets based on individual p_T deposits. A crude description of this is: considering a set of large, well-spaced energy deposits corresponding to a set of hard jets, and a series of soft particles surrounding each energy deposit, then it is reasonable to assume the deposits are a signature resulting from large p_T jet showering in the calorimeters. Soft deposits within a radius ΔR , defined in Eq. 3.6, of the largest jet p_T are attributed to that jet, giving it a circular conical shape; the hard jet with the next largest p_T is attributed the soft showering energy of the remaining deposits inside radius ΔR . An illustration of this is provided in Fig. 3.11. Examples of this are the AK4 jet, a light jet defined by a radius $R = 0.4$, and AK8 jets, termed *fat* jets, with radius $R = 0.8$ often attributed to objects with boosted topologies. The HLT also recalculates global quantities including $\vec{p}_T$ and $\vec{p}_T^{\text{miss}}$, to greater accuracy than the L1T measurement.

Classification algorithms are also employed in the HLT, an example being the deepCSV neural network which is used to tag b jets, with an efficiency of correctly identifying a b jet of 68%, and a probability of misidentifying a light-flavoured jet




Figure 3.11: The anti- k_T algorithm for jet clustering into a single jet of radius ΔR and momentum p_T .

of 1% [84]. Particle flow (PF) reconstruction is a global event reconstruction algorithm [85] that is designed to identify and reconstruct all final state particles from an event. In these cases, the algorithms can be computationally exhaustive, therefore for use in the HLT typically these are approximations of the intended-scale algorithm, which can be run offline outside of the LHC cycle on the raw event data that is retained. By the end of the paths, the event rate is reduced by a factor 100 to 1 kHz for DAQ, with data saved to hard drives worldwide at a streaming rate of $\mathcal{O}(6 \text{ GB/s})$. The data is then accessible for analysts, although some further central processing and reduction may be necessary.

3.3.4 Data Reduction and Storage

Full event information from prompt online reconstruction is taken directly from the HLT, containing *raw* detector details which is stored in ROOT file containers around 1 MB in size per event. Offline corrections and comprehensive object reconstruction is performed afterwards, and stored in the *RECO* tier. The size of RECO data per event can reach as much as 2-3 MB, therefore only a subset of this data is retained for detector performance and low-level development studies. Typically, AOD files are produced from the postprocessing of RECO tier files to contain only information

required for analyses. These are nominally available, although again these are very information-dense for Run 2, at around 500 kB per event.

To reduce the footprint of AOD files, a compression algorithm further removes unnecessary information via *slimming* and *skimming* ROOT files. The former omits entire branches from event trees, while the latter applies a selection to event trees, with those failing omitted. This yields miniAOD files at 50 kB per event, and further still nanoAOD files at $\mathcal{O}(1 \text{ kb})$ per event. NanoAOD uses the LZMA package to remove more information and convert types e.g. from float64 to float32 precision. There may be too little residue from the nanoAOD format for some analyses, however for analyses orientated on invisible Higgs boson decays this is appropriate.

3.4 AutoDQM

AutoDQM is a statistical tool adapted to read from the CMS DQM stores and facilitate the role of shifters in fast analysis of hundreds of histograms from a range of runs to qualify readout data. An arsenal of statistical tests are available, both conventional and ML-based, with its primary focus to spot hard-to-find issues faster than shifters. Tests began with Run 2 data and have continued with Run 3 data using online and offline resources. Full documentation for AutoDQM and its ML-orientated analogue is available in Ref. [86, 87].

AutoDQM conventionally takes histograms from two different runs and compares them. The histograms are taken from any part of the trigger, for example the layer 2 calorimeter trigger, and will analyse both 1D and 2D plots. For 2D histograms such as η - ϕ occupancy or η - ϕ - p_T plots, a pull value is calculated for each corresponding bin between two histograms, given by

$$\text{pull value} = \frac{(x_1 - x_2)^2}{\epsilon_1^2 + \epsilon_2^2}, \quad (3.11)$$

where x_i is the bin population in histogram i , and ϵ_i is Poissonian error on the bin. The result is a new 2D histogram that maps the pull value for each coordinate, with some bins appearing empty (masks or dead towers) and some with a high value (hot towers).

Another statistical test is the Kolmogorov-Smirnov (KS) test, applicable to 1D distributions. This test works by finding the distance between each corresponding point in two comparison histograms, and qualifying the equality of the two distributions. This ensures the test is sensitive to both location and shape of each histogram. For

both the KS and pull-value test, plots are declared anomalous if some pull-value, KS or χ^2 threshold is exceeded. In other words, for a good reference run and an untested run, the run would be declared *bad* if their histograms poorly match, due to too many hot, masked or dead towers or zones based on the layer 2 calorimeter trigger. The abnormality threshold is set at a value according to input from the AutoDQM strategists and by detector performance groups (DPGs) and detector experts.

AutoDQM has a series of subsystem or trigger level `json` files, containing the names of and paths to the histograms of interest, and sets the statistical comparator(s) for each histogram. Normalisation in 1D and 2D is also available, which is required when comparing two runs with vastly different numbers of events. In η - ϕ plots, normalising in ϕ helps to remove occupancy features due to the calorimeter substructure, such as the barrel-to-EC boundary or barrel-to-EC occupancy variations. This approach still retains all the holes and dead zone information in the comparison plots, and avoids weighting the pull values by the error on the more populated distribution. The choice to write files in this manner allows users to implement their own custom subsystems and add them locally.

A demonstration of the AutoDQM GUI is shown in Fig. 3.12, displaying the output for offline layer 2 histograms when comparing two 2022 runs. The reference run, 356902, is a run known to have been marked internally by CMS shift experts as a *golden*, or certified data set. The data run, 356709, has been marked as a bad run by shifters due to significant L1 jet *pre-firing*. This phenomenon occurs when the ECAL trigger timing gradually shifts during data taking, and can vary depending on the ECAL region. The GUI shows plots corresponding to the results of the KS and pull-value tests for **CenJets** and **EGamma** (EG, or electron and photon objects) distributions.

The corresponding run histograms from online DQM demonstrate the AutoDQM process in 2D, as shown in Fig. 3.13 for the **CenJets** η - ϕ - p_T distribution. The first histogram is the distribution from the reference run, while the second is the bad run due to pre-firing. During online DQM, this would have been flagged in the ECAL output where particularly high p_T is recorded in the central η - ϕ region. The third histogram, taken from the pull-value test performed in offline AutoDQM, shows how this region of high p_T appears when compared to the reference run, and is quantified by the pull value exceeding the normal threshold, therefore an anomaly has been found. This illustrates that the use of conventional statistical tools in AutoDQM can perform as well as shifters and DOCs.




Figure 3.12: The AutoDQM GUI for conventional statistical tests such as the KS test for 1D and pull-value tests for 2D. Taken from layer 2 output for two 2022 runs, the bad data run, 356709, and a golden reference run, 356902, the histograms shown are those marked as anomalous by AutoDQM from KS and pull-value tests in JetMET samples.




Figure 3.13: Input layer 2 CenJets η - ϕ - p_T distributions from online DQM runs, with the reference run 356968 (top) and the bad run 356709 (middle), where the resulting pull-value test performed by AutoDQM offline shows the anomalous region of the bad run in the central η - ϕ region, in red (bottom).

3.4.1 AutoDQM ML

Further to the conventional statistical tests in AutoDQM, the implementation of principal component analysis (PCA) and autoencoder (AE) ML methods is at the forefront of AutoDQM development. These both use ML to visually judge anomalies in 1D and 2D based on a reference set of runs, rather than a single particular run, over which the AE or PCA is trained. The objective is for the trained data set to learn the typical structures of histograms, e.g. η and ϕ plots from good runs, and to use these to identify bad runs based on anomalies potentially not visible by-eye to DQM shifters.

An AE is a type of neural network (NN) designed to learn an efficient representation of some data without supervision. Crudely, it works as follows: the AE takes images as input, then it encodes or *compresses* information about these images into latent space. The compression imposes a bottleneck in the NN, that is to say a representation is produced by forcibly reducing information in many dimensions into just a few key descriptors. Without the bottleneck, no representation will be learnt and the input will be memorised as-is, which is inefficient. However, for the representation to be accurate, the input has to be correlated, i.e. there is some set of features or patterns that are common across every bit of data in the input. After the bottleneck, the representation is decoded or *reconstructed* and compared to the original input using a *loss function* to test the AE performance. A graphic illustrating this chain is given in Fig. 3.14. Fundamentally, an ideal AE will be sensitive enough to input that it produces a good reconstruction, but insensitive enough to input to not memorise or overfit during the training.

The AE trains the NN on the input data to recognise the structures in the image. Anomalies are easier to flag when comparing an example histogram from a test run to reconstructed version produced by the trained NN, and can be done unsupervised. The point is that anomalies are intrinsically rare in runs: the rarer the better, as they will more likely be distinguishable from the reconstruction that relies on many reference good runs.

A rectified linear unit (RELU) activation function is required to ensure the reconstruction does not produce negative event counts, given the AE is not inherently informed that negative counts is non-physical. This function is enacted on the last layer of the decoding, and assigns any NN output with a negative number to a bin with zero counts. This does not affect the tuning of weights in the functioning of the AE, nor is the performance affected by applying the RELU function, by design: the reconstruction is near-identical with and without this function.

The PCA is run in parallel with the AE, and is not as abstract a tool. The NN




Figure 3.14: The steps of an AE to form a representation of the input data by compressing it into the fewest dimensions possible using a bottleneck before reconstructing the representation to compare to the original input using a loss function.

in PCA training is linear, whereas the AE can exploit learning non-linear behaviour. This does not guarantee a better performance for the AE. The PCA workflow, outlined in Fig. 3.15, has so far yielded a 0.25% false-positive score based on training with 1D histograms from the CSC muon chamber subsystem.

The AE and PCA are trained on all 2022 CMS JetMET data sets within Run block C, spanning runs 355862 to 357482 inclusive, that are verified as having good data taking to produce **GenJets** ϕ and η - ϕ - p_T distributions, according to CMS internal data certification reports. The choice of a single run block for training ensures that computation time is much shorter than multiple blocks, and ensures that the training data set more precisely models the runs within that block, without biasing by blocks produced under different running conditions. The exclusion of bad runs from the training is to ensure that the NN neither learns any anomalous features nor biases what the reconstruction of a good run should look like.

The reconstruction of the ϕ and η - ϕ - p_T distributions in the AE and PCA are compared against the original offline histogram for the bad run 356709, and is presented in Fig. 3.16. Note that 2D histograms are projected into 1D histograms before being run through the AE or PCA, which both saves on computation time and better illustrates discrepancies in the output evaluation stage. A ratio of the event fraction in each bin between the reconstruction and original histogram is also presented below each plot. For this run, it is visually clear that the reconstruction is heavily biased by the trained data set, and as such is a poor description of the original histogram. The sum




Figure 3.15: The steps for a PCA-trained workflow, from input to output and flagging. This operates much the same as the AE but only uses a linear function for dimension reduction in the compression stage.

of squared errors score (SSE) is a performance metric that calculates the difference between the original and reconstructed distributions. The SSE anomaly threshold is yet to be determined for layer 2 for the AE and PCA.

For comparison, the reconstruction of the ϕ and η - ϕ - p_T distributions taken from a certified good run are presented in Fig. 3.17. Here, the SSE scores for the PCA are ≈ 1000 times smaller than for the bad run, while the AE is only ≈ 1.5 -2 times smaller. This would suggest the PCA is performing far better in this scenario, although further studies are necessary before drawing conclusions. So far, studies like this have been performed across all runs marked as bad in 2022 Run block C, and the SSE scores compared to test the AE and PCA performances in correctly marking runs as good or bad. This also allows a performance evaluation of these NN on a cause-by-cause basis, where multiple issues are flagged in data certification reports that have caused a run to be bad, including large trigger dead times, and HCAL exclusion due to faulty readout.

Currently, the AE and PCA are completely abstract, and have no initial bias towards any input histogram features. It is possible to introduce a pre-programming that biases towards symmetry in η hemispheres or anomalously high peaks, which would be better-informed to identify bad runs without compromising the NN's objec-




Figure 3.16: Reconstructions of CenJets ϕ and η - ϕ - p_T distributions from the AE (orange) and PCA (green), taken from bad run 356709 in which a significant amount of L1 pre-firing has occurred, and overlaid on the original histogram for that run. The SSE is an effective χ^2 measurement between the original and reconstructed histograms.

tives. Larger-scale performance of the AE and PCA is also envisioned, with performance plots underway that would depend on the average number of histograms flagged as anomalous that are certified good runs versus the average number of histograms flagged that are from bad runs, and measures of the runs with at least N runs flagged




Figure 3.17: Reconstructions of CenJets ϕ and η - ϕ - p_T distributions from the AE (orange) and PCA (green), taken from the good run 356968, and overlaid on the original histogram for that run. The SSE is an effective χ^2 measurement between the original and reconstructed histograms.

as anomalous, where N is expected to be a small number < 10 .

Chapter 4

Event reconstruction and simulation

The main signal processes studied in this analysis are $t\bar{t}H$ and VH (illustrated in Fig. 1.1), where the final states are fully hadronic. Additional signal contributions from VBF and ggH topologies are also accounted for. The typical signature for $H \rightarrow \text{inv}$ decays is a high p_T^{miss} in the event. Background signatures that mimic $H \rightarrow \text{inv}$ signal include the *lost lepton* processes, ℓ_{lost} (where ℓ is either a μ or e), where a W boson decays semi-leptonically ($W \rightarrow \ell\nu$) with the charged lepton failing reconstruction or falling outside of the detector acceptance, and cannot be reduced by simply removing a charged lepton. The resulting signature is a large p_T^{miss} as the charged lepton is *lost* in the detector. Processes falling under the ℓ_{lost} umbrella include $W(\ell\nu) + \text{jets}$, $t\bar{t}$, $t\bar{t}X$ where the associated X is any boson W , Z , γ , or a visibly decaying H , and single t quark production. Similarly, processes where a Z boson decays invisibly ($Z \rightarrow \text{inv}$) mimic the $H \rightarrow \text{inv}$ signal, which involves a Z boson decaying to two neutrinos ($Z \rightarrow \nu\bar{\nu}$). The $Z \rightarrow \text{inv}$ contribution can be estimated based on kinematically similar processes, including $Z(\ell^+\ell^-) + \text{jets}$, Drell-Yan (DY, or $Z/\gamma^*(\ell^+\ell^-) + \text{jets}$), $\gamma + \text{jets}$, and diboson (VV) and triboson (VVV) production. The background due to QCD multijet processes typically does not have a large genuine p_T^{miss} , however due to their large production cross section, even a small fraction of events where a high p_T jet is mismeasured, is affected by detector noise, or falls outside of the detector acceptance has the potential to contribute to the hadronic $H \rightarrow \text{inv}$ background. These tend to dominate at low p_T^{miss} , hence can be mitigated by selecting signal from a high p_T^{miss} region. Representative diagrams for ℓ_{lost} , $Z \rightarrow \text{inv}$, and QCD multijet processes are provided in Fig. 4.1.

The physics objects considered in this analysis, including muons, electrons, pho-




Figure 4.1: The representative Feynman diagrams for the ℓ_{lost} , $Z \rightarrow \text{inv}$, and QCD multijet background processes.

tons, and jets, comply with specific identification *working points* (WPs) as recommended by the various Physics Object Groups (POGs) of CMS. The object samples have varying degrees of efficiency and purity depending on the tightness of the WP criteria. A tighter (looser) criteria will produce a sample of greater purity (efficiency). Tight, loose, and medium WPs are considered for different objects, for both selecting and vetoing events containing certain objects.

The following chapter serves as a preamble to the analysis, introducing the software used for the analysis in Sec. 4.1, followed by the data and simulated samples in Sec. 4.2 that are used, including their generators, their simulated precision in QCD and EW production, and details of higher order corrections applied. The methods by which events are reconstructed in the detector, including the definitions of objects found in the detector that are relevant to the analysis, are described in Sec. 4.3.

4.1 Software

This analysis makes use of the NANO AOD-TOOLS [88] toolkit for the purposes of skimming large input NanoAOD [89] to more manageable sizes. This is necessary when acquiring data and simulated sample sets from NanoAOD in a format ready for use by analysers. Having lightweight, flat ROOT trees as input provides the option of having additional branches containing analysis specific variables without suffering a huge loss in storage capacity. Their creation is done by building on top of modules that are officially provided by the CMS experiment using NANO AOD-TOOLS. This streamlined nature of the NanoAOD samples and the flexibility of the accompanying NANO AOD-TOOLS software also provided a good basis to reduce the processing time of samples for the analysis from $\mathcal{O}(\text{weeks})$ to $\mathcal{O}(\text{days})$. Analysis-specific modules enable actions such as: preselection by taking only the relevant triggers into account, creation

of new object collections while applying all POG-recommended definitions, applying triggers and filters to available events, and many more.

The FASTER ANALYSIS SOFTWARE TASKFORCE (F.A.S.T.) [90] data processing package is also used extensively, which allows users to perform analysis on dataframes synthesised from flattened NanoAOD files, bypassing direct interaction with ROOT. The F.A.S.T. package is primarily used for applying trigger filters and offline selection. The toolkit provides a series of packages that can be tailored for use in analyses. First, skimmed trees are referenced with the F.A.S.T.-CURATOR[91] package. The F.A.S.T.-CARPENTER [92] package is the primary tool used to process data in the analysis, and streamlines the production of dataframes by bypassing direct user interaction with ROOT. Instead, the user feeds a file configuring, for example, definitions of new variables, systematic weights, event selections, and variables to store, that is then interpreted by F.A.S.T.-FLOW [93] at run time to F.A.S.T.-CARPENTER. F.A.S.T.-CARPENTER then utilises vectorisation and batch computing to process events quickly and produce binned PANDAS dataframes. The dataframes can be multidimensional and include systematic weights, ideal for visualisation with F.A.S.T.-PLOTTER [94] to produce histogram-style figures. The F.A.S.T.-CALIPER package is also used to optimise signal selections for given variables in multiple dimensions, by computing the cumulative total of yields bin-wise to obtain significance maxima in 1D and 2D corresponding to the boundary on which to make the selection. Dataframes are available as input for other tools made use of in the subsequent steps of the analysis, such as F.A.S.T.-DATACARD [95], a tool which produces datacards for the CMS HIGGS COMBINE statistics package [96] that performs the final fit and returns the limit on $\mathcal{B}(H \rightarrow \text{inv})$.

4.2 Simulated samples

The analysis of the $t\bar{t}H$ and resolved VH channels in fully hadronic final states makes use of CMS data collected during the data-taking period of 2016–2018, corresponding to an integrated luminosity of 138 fb^{-1} . The data recorded by the CMS detector is supplemented by simulated events generated typically under SM assumptions. Simulations such as these helped in the development of the CMS detector for the higher luminosity Run 3 conditions, as well as for testing performance of proposed detectors for future colliders. Simulated samples are also useful in studying the kinematics of signal processes in searches beyond the SM (BSM) and background processes kinematically similar to the signal. The samples in this analysis are generated using *Monte*

Carlo (MC) sampling. MC-generated results can be parametrised according to the SM or any BSM or DM models that provide signal and background samples on which to perform analyses alongside data. Large numbers of simulated events are generated to minimise statistical uncertainty, where the generation relies on random Gaussian-based sampling.

To simulate an event, first the hard scattering process is modelled, typically to LO, and in some cases where the theory is better-understood to greater precision, in either MADGRAPH5_AMC@NLO version 2.4.2 or higher [97] (which can generate processes at LO or NLO) or POWHEG v.1.0 or higher [98] matrix element (ME) generators. The ME is encoded with the maximum amount of information available for a hard scattering event, with which soft processes such as parton showering, hadronisation, and other underlying event properties can be simulated. Showering from the hard process incurs the simulation of free quarks and gluons fragmenting and radiating at high energies. The showered partons then form visible particles via hadronisation in the simulation. Information underlying the event such as softer collisions with the incident partons not involved in the hard scattering are also simulated. All three stages are referred to as parton showering simulation, and are modelled by the FxFx [99] matching scheme for MADGRAPH5_AMC@NLO-generated samples, and by PYTHIA8 [100] for POWHEG-generated samples. The PYTHIA8 event tune CUETP8M1 (CP5) is used for the 2016 (2017–18) data-taking periods [101]. Samples for 2016 (2017 and 2018) make use of the NNPDF3.0 LO or NLO (NNPDF3.1 NNLO) PDFs [102]. Finally, the simulated samples are postprocessed similarly to real data, ensuring that simulation more accurately represents the data. Postprocessing is performed by propagating all final state particles through a simulation of the CMS detector using the GEANT4 [103] toolkit, which recreates the conditions and geometry of the detector and reconstructs physics objects from simulated collisions. With GEANT4, the response of the CMS subsystems to simulated events is mimicked in the event reconstruction. The resulting samples are then ready for use in the analysis.

The signal processes VBF, VH, $t\bar{t}H$, and ggH are simulated with POWHEG v.2.0 [104–107] at NLO accuracy in QCD. These samples require the 125 GeV SM Higgs boson to decay to four neutrinos ($H \rightarrow ZZ^* \rightarrow 4\nu$) with a branching fraction of unity. The cross sections are appropriately normalised to the corresponding SM predictions computed at NNLO (VBF and VH), NLO ($t\bar{t}H$), and next-to-NNLO (N3LO) (ggH) accuracy in QCD corrections, and to NLO accuracy in EW corrections [108]. Exceptionally, the $gg \rightarrow Z(q\bar{q})H$ sample is generated at NNLO accuracy in QCD corrections. The signal sample specifications are summarised in Table 4.1.

Table 4.1: List of signal processes used in the analysis, with their production cross sections and precision of cross section corrections in QCD and EW production. In all cases these samples are generated in POWHEG v.2.0.

Process	Subprocess	Cross section / pb ⁻¹	QCD precision	EW precision
VBF		3.77×10^0	NNLO	NLO
V(q $\bar{q}$)H		2.24×10^0	NNLO	NLO
	W ⁺ (q $\bar{q}$)H	8.31×10^{-1}	NNLO	NLO
	W ⁻ (q $\bar{q}$)H	5.27×10^{-1}	NNLO	NLO
	Z(q $\bar{q}$)H	8.80×10^{-1}	NNLO	NLO
t $\bar{t}$ H		5.07×10^{-1}	NLO	NLO
ggH		4.86×10^1	N3LO	NLO
gg $\rightarrow$ Z(q $\bar{q}$)H		1.23×10^{-1}	NNLO	NLO

Simulated samples of background processes are typically generated to the highest precision available, however in some cases a lower-order accuracy is corrected to match that of a higher order. This is performed for samples where the higher-order sample is not available centrally or the agreement between data and simulation is poorer.

The largest background contribution is from V + jets processes, where samples are generated at LO accuracy using MADGRAPH5_AMC@NLO v.2.6.5 with up to four partons in the final state using the MLM [109] matching scheme. The V + jets samples cover the processes: Z($\nu\bar{\nu}$) + jets, for which two or more jets are produced from QCD vertices, or two or more jets are produced from EW vertices; DY, or Z/ $\gamma^*(\ell^+\ell^-)$ + jets, where at least two jets are produced from QCD vertices only, and a requirement that the dilepton mass $m_{\ell\ell} > 50$ GeV is applied to ensure DY production is from the Z boson peak; and W($\ell\nu$) + jets, where at least two jets are produced from QCD or EW vertices. The V + jets samples are generated in bins of the *hadronic transverse energy*, H_T , which is the scalar sum of jet p_T reconstructed at the generator level. Formally, for each jet j in an event with transverse momentum $p_{T,j}$,

$$H_T = \sum_j^{\text{jet}} p_{T,j}. \quad (4.1)$$

The simulation at LO is corrected to account for missing higher-order diagrams using *K-factors*, equal to the ratio of the NLO-to-LO cross section, that are derived from the analogous MADGRAPH5_AMC@NLO-generated NLO QCD $V + \text{jets}$ sample [110], with up to two partons in the final state. The K-factors are extracted differentially as a function of the p_T of the V boson, p_T^V , and the p_T of the *leading* jet, the jet with the largest p_T in an event, $p_{T,1}^j$. Additionally, $\gamma + \text{jets}$ processes, which are kinematically similar to V boson production in VH topologies, are generated at NLO in bins of the p_T of the photon, p_T^γ . The binning scheme is defined at the ME level to increase the statistical precision in the region containing a single high-energy photon that is probed by the analysis.

The signal-like ZH production, where the Higgs boson decays as $H \rightarrow b\bar{b}$, is a relatively small background, given the small cross section for $H(b\bar{b})$ decay from ZH production, and is generated at NLO using MADGRAPH5_AMC@NLO v.2.6.5 for 2016 samples, and POWHEG v.2.0 [111] for 2017 and 2018 samples.

Background $t\bar{t} + \text{jets}$, $t\bar{t}X + \text{jets}$ (where X is any boson γ , V , or a visibly-decaying Higgs boson), and single t quark processes can contribute significantly to the ℓ_{lost} background in $t\bar{t}H$ searches, where $W(\ell\nu)$ boson decay can result in genuine, large p_T^{miss} in the event following the decay of the t quark(s). The largest t quark background, QCD-induced $t\bar{t} + \text{jets}$ production, is generated at NLO with POWHEG v.2.0 [112, 113] with up to two additional partons in ME computation. This includes the scenarios for which the $t\bar{t}$ system decays to a hadronic, semi-leptonic and dileptonic final state. The t quark p_T spectrum is corrected to match the spectrum obtained from NNLO QCD + NLO EW simulation, following the method in Ref. [114].

Rare $t\bar{t}X + \text{jets}$ backgrounds include $t\bar{t}\gamma + \text{jets}$, $t\bar{t}W + \text{jets}$, and $t\bar{t}X + \text{jets}$ processes, and are generated using MADGRAPH5_AMC@NLO v.2.6.5, with subsequent decays generated using MADSPIN [115] to account for spin correlations in the former two cases. The $t\bar{t}H + \text{jets}$ sample, for which the H decays to visible states, is generated using POWHEG v.2.0. Single t quark generation is EW-induced in s , t and tW channels using MADGRAPH5_AMC@NLO v.2.6.5, POWHEG v.2.0, and POWHEG v.1.0 [116], respectively. For the t -channel, decays of the W boson exchanged between the initial state quark and b quark are exceptionally modelled using MADSPIN [115].

The multiboson processes VV and VVV are subdominant backgrounds in which at least one V boson decays leptonically, thereby producing jets and a large p_T^{miss} in the final state. Diboson processes WZ and ZZ are generated at NLO using PYTHIA8, while the WW diboson process is generated at NNLO in QCD using POWHEG v.2.0 [117].

Triboson events are generated at NNLO using MADGRAPH5_AMC@NLO v.2.6.5.

The QCD multijet samples are generated at LO using MADGRAPH5_AMC@NLO v.2.6.5 in exclusive ranges of H_T , which increases the statistical precision in the non-leptonic, non-photonic regions studied by the analysis. A summary of the background processes, their precision in EW and QCD generation, and their relative cross sections, are presented in Table 4.2.

Table 4.2: List of background processes used in the analysis, with their production cross sections above $H_T > 200$ GeV, and their precision in QCD and EW (where applicable). In some cases, corrections at higher orders are applied to the samples via method of cross section reweighting, as stated below.

Process	QCD	EW	Cross section (pb)	Corrections
$Z(\nu\bar{\nu}) + \text{jets}$	LO	LO	1.25×10^2	NLO QCD and EW reweighting
$DY + \text{jets}$	LO	LO	5.84×10^1	NLO QCD and EW reweighting
$W(\ell\nu) + \text{jets}$	LO	LO	5.41×10^2	NLO QCD and EW reweighting
$\gamma + \text{jets}$	NLO	NLO	1.21×10^3	
$ZH(b\bar{b})$	NLO	NLO	7.91×10^{-1}	
$t\bar{t} + \text{jets}$	NLO	—	8.32×10^2	NNLO QCD and NLO EW reweighting
$t\bar{t}X + \text{jets}$	NLO	—	5.48×10^0	
Single t (s -channel)	—	NLO	1.03×10^1	
Single t (t -channel)	—	NLO	2.20×10^2	
Single t (tW -channel)	—	NLO	7.17×10^1	
WW	NNLO	—	1.19×10^2	
WZ, ZZ	NLO	NLO	6.37×10^1	
VVV	NNLO	NNLO	4.30×10^{-1}	
QCD multijet	LO	LO	1.92×10^6	

4.3 Event reconstruction

Final-state particles of interest to the analysis include light-quark jets, leptons, photons, b-quark jets, and heavy Lorentz-boosted objects such as t quarks and W bosons. Before categorising events based on the particles in the final state, a discussion on the techniques used to identify and reconstruct PF candidates must first be addressed. For some physics objects, different identification and requirement definitions are used depending on whether to select a hadronic or leptonic/photonic final state, which affects the purity of the corresponding region in desired events. An increase in the purity of the reconstruction of an object is generally at the expense of the efficiency with which that object is reconstructed. Such a consideration is necessary when defining a region

based on the identification of an object in an event, or the removal (or veto) of events containing that object from a region.

Candidate leptons and photons are required to have a sufficiently high p_T to be confidently identified and reconstructed above detector noise and background contributions $\mathcal{O}(\text{GeV})$, and be detected within the acceptance of the subsystem that targets them. Additional criteria depend on the tightness of isolation of a given particle (lepton or photon), or the relative energy of that particle within a cone of small (tight) or large (loose) radius, which are used in muon, electron, and photon event reconstruction. Candidate muons require $p_T > 10 \text{ GeV}$ and $|\eta| < 2.4$ [74]. For electrons and photons, reconstruction is performed such that the EM shower that they produce in the ECAL is contained entirely in the EB or in the EE, so candidates are not considered for reconstruction within the ECAL transition region of $1.44 < |\eta| < 1.57$. Candidate electrons require $p_T > 10 \text{ GeV}$ and either $|\eta| < 1.44$ or $1.57 < |\eta| < 2.5$ [68], and candidate photons $p_T > 15 \text{ GeV}$ and either $|\eta| < 1.44$ or $1.57 < |\eta| < 2.5$ [68]. Hadronically-decaying candidate tau leptons require $p_T > 20 \text{ GeV}$ and $|\eta| < 2.3$ [118]. To mitigate any overlap between muon, electron, and photon collections, either due to misreconstruction or genuine physics processes producing such particles in close proximity, a particle identification hierarchy is established within each event. Electrons are removed from events containing loose muons if they are located within a cone of $\Delta R < 0.3$ of the muon direction. Thereafter, photons are removed from events containing loose muons and loose electrons if they are located within a cone of $\Delta R < 0.3$ of the muon or electron direction. Finally, tau leptons are removed from all loose corrections of muons, electrons, and photons if they are located within a cone of $\Delta R < 0.3$ of the muon, electron, or photon direction. Loose identification and isolation criteria are used to veto on events in the hadronic regions that contain leptons or photons, which have efficiencies of > 99 , $\simeq 95$, and $\simeq 90\%$ for loosely-isolated muons, electrons, and photons, respectively. Tight and loose identification and isolation criteria are used to select and count muons, electrons, and photons for non-hadronic events, which are found to enhance the purity at little expense to the efficiency. These achieve typical efficiencies of $\simeq 95$, 70 , and 70 ($\simeq 98$, 95 and 90)%, for tight (loose) muons, electrons, and photons, respectively.

Jets are reconstructed by clustering all PF candidates originating from the PV with the anti- k_T jet clustering algorithm [82, 83], using a distance parameter $R = 0.4$ (AK4). This requires a separation of $\Delta R > 0.4$ from any muons, electrons, or photons previously identified. The jet momentum is calculated as the vectorial sum of all particle momenta within the jet cluster, which in simulation is on average within 5

and 10% of the true momentum across the full p_T spectrum and detector acceptance. Charged hadron subtraction (CHS) [119] is then applied to remove charged particles from pileup vertices [120]. To ensure the measured jet energy matches that of the particle level jets, jet energy corrections (JEC) derived from simulation as functions of p_T and η are applied for jets with $p_T > 15$ GeV. Further corrections are performed for residual discrepancies between the jet energy scale (JES) in data and simulated samples [119]. Additionally, each jet must pass selection criteria to filter jets adversely affected by instrumentation or reconstruction failure. The jet energy resolution (JER) in simulated samples is smeared to match that of the data, given the JER in the data is poorer, typically corresponding to between 15 and 20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV [119].

An additional series of tight identification criteria are applied to jets as recommended by the Jet POG, which vary depending on their location in the detector and the data-taking year, and conditionally for the veto of leptons in the data. These depend primarily on the *energy fraction*, f_E , defined as the fraction of energy attributed to a given reconstructed object within a PF jet, and the multiplicity of various PF objects, N . The criteria for jets in 2016 data comply with the requirements of the tight WP for the pileup identification algorithm [121] that ensures jets are not reconstructed from pileup vertices. The algorithm trains a boosted decision tree on 2016 samples, factoring in tracker information and object multiplicity. For 2017 and 2018 data, tight jet identification conditions are required only for jets with $p_T < 50$ GeV, with no requirements necessary for higher p_T jets. Also for 2018 data, the inner η region is increased from $|\eta| < 2.4$ to $|\eta| < 2.6$ following an extension of the tracker subsystem. The requirements for all years are applied to the energy fractions attributed to muons (f_E^μ), and neutral and charged hadrons (f_E^{h0} and $f_E^{h\pm}$) and EM showers ($f_E^{\text{EM}0}$ and $f_E^{\text{EM}\pm}$), for jets with $p_T > 15$ GeV, consistent with the recommendations from the Jet POG presented in Table 4.3, and taken from internal studies performed by the JetMET POG. The AK4 jets require $p_T > 30$ GeV and $|\eta| < 5.0$, that they are tightly identified, and that those associated with loose leptons and photons located within a cone of $\Delta R < 0.4$ of the jet direction are removed. These jet identification criteria ensure a high AK4 jet purity while achieving an identification and background rejection efficiency of $> 98\%$ in all detector regions.

A series of filter selections are applied to remove events that are poorly measured or misreconstructed either in data or simulation, and are applied across all years. The filters discriminate primarily according to the PF jet energy fractions for various particles, and vary according to the region of the detector in which the objects are

Table 4.3: The requirements applied to data for jet candidates with $p_T > 15$ GeV to pass the tight jet identification conditionally with a lepton veto (denoted *), for 2016, 2017, and 2018, respectively. Note that requirements are applied to data with an inner region of $|\eta| < 2.4$ for 2016 and 2017, and $|\eta| < 2.6$ for 2018. The Jet POG recommendations are taken from internal studies performed by the JetMET POG.

Criterion	$ \eta < 2.4$ (2.6)	2.4 (2.6) $< \eta < 2.7$	$2.7 < \eta < 3.0$	$ \eta > 3.0$
f_E^{h0}	< 0.9	< 0.9	< 0.98	> 0.02
f_E^{EM0}	< 0.9	< 0.9	> 0.01	< 0.9
$N_{\text{constituents}}$	> 1	> 1	—	—
N_{neutral}	—	—	> 2	> 10
$f_E^{h\pm}$	> 0	—	—	—
$f_E^{EM\pm}$	< 0.99 (0.90*)	—	—	—
N_{charged}	> 0	—	—	—
f_E^μ	$< 0.80^*$	$< 0.80^*$	—	—
f_E^{h0}	< 0.9	< 0.9	—	> 0.02
f_E^{EM0}	< 0.9	< 0.9	[0.02,0.99]	< 0.9
$N_{\text{constituents}}$	> 1	> 1	—	—
N_{neutral}	—	—	> 2	> 10
$f_E^{h\pm}$	> 0	—	—	—
$f_E^{EM\pm}$	$< 0.80^*$	—	—	—
N_{charged}	> 0	—	—	—
f_E^μ	$< 0.80^*$	$< 0.80^*$	—	—
f_E^{h0}	< 0.9	< 0.9	—	> 0.02
f_E^{EM0}	< 0.9	< 0.9	[0.02,0.99]	< 0.9
$N_{\text{constituents}}$	> 1	—	—	—
N_{neutral}	—	—	> 2	> 10
$f_E^{h\pm}$	> 0	—	—	—
$f_E^{EM\pm}$	$< 0.80^*$	$< 0.80^*$	—	—
N_{charged}	> 0	> 0	—	—
f_E^μ	$< 0.80^*$	$< 0.80^*$	—	—

located. These include a filter that rejects events containing forward jets (> 2.4) with $p_T > 50$ GeV, as these are more susceptible to mismeasurement during reconstruction.

This filter also reduces background contributions, given the signal topologies probed are more likely to be located in the central regions of the detector. More stringent neutral and charged hadron fraction requirements are applied than those imposed by the tight jet identification conditions, focused on the leading (subleading) jets, with events retained if $f_{E, 1(2)}^{h0, j} < 0.8$ and $f_{E, 1(2)}^{h\pm, j} > 0.1$. A muon jet filter is applied that requires all jets with $p_T > 200$ GeV to have $f_E^\mu < 0.5$, or the azimuthal separation between the $\vec{p}_T$ of a jet and $\vec{p}_T^{\text{miss}}$, $|\Delta\phi(\vec{p}_T^{\text{miss}}, \vec{p}_{T,1234})| < \pi - 0.4$, in order to reject events with misreconstructed muons.

Heavy-flavoured AK4 jets can be identified as b jets using the DEEPCSV deep neural network (DNN) algorithm [84]. The medium WP is adopted for the analysis, where b jets with $p_T > 20$ GeV are correctly identified with a probability of 80%, and where a charm or light jet is misidentified as a b jet with 1% probability. For b jets to be considered, a requirement of $p_T > 30$ GeV and $|\eta| < 2.4$ ($|\eta| < 2.5$) is applied to the AK4 jet collections produced for 2016 (2017 and 2018) samples that are the input for the DNN. The medium WP corresponds to a DEEPCSV b + bb identification discriminator threshold of between 0.4184 and 0.6321 depending on the year of data taking. Simulated events containing b jets are corrected to be in agreement with the data by deriving efficiency corrections from data control samples. This is performed separately for various final states, such as in events where $t\bar{t}$ decays into an opposite-sign $e\mu$ pair, or in QCD multijet events.

The pileup per particle identification algorithm (PUPPI) [59, 60] is employed to mitigate pileup effects at the reconstructed particle level by defining a local shape variable that can discriminate between particles originating from the PV and from pileup vertices. The reduction of the effect of pileup relies on mitigation techniques [122] that filter energy deposits associated with pileup vertices and remove objects not associated with the PV. Charged particles originating from pileup are discarded. For neutral particles, a local shape variable is computed based on the information from charged particles in its vicinity that originate from the PV within the tracker acceptance, and information from both charged and neutral particles outside this acceptance. The momenta of neutral particles are then rescaled based on the probability that they originated from the PV as deduced from the local shape variable [59].

When a high- p_T t quark or V boson decays hadronically, a large set of collimated particles crosses the detector. These can be clustered within a single jet of radius $R = 0.8$ (AK8) using the anti- k_T algorithm. In order to reduce pileup effects, PUPPI PF candidates are used to seed the AK8 jet finder. The main feature that distinguishes hadronically decaying t quarks or V bosons from the fragmentation of a single

quark or gluon is the jet mass. To improve the resolution, the modified mass-drop algorithm [123–125] (also known as the soft-drop, SD) with angular exponent $\beta = 0$, soft cutoff threshold $z_{\text{cut}} < 0.1$, and characteristic radius $R_0 = 0.8$ [126], is applied to each AK8 jet to remove soft wide-angle radiation. In addition, a DNN classifier called the DEEPAK8 [127] algorithm is employed, which takes as input up to 100 constituent PF candidates, as well as information from up to seven secondary vertices, and assigns a set of numerical scores to each reconstructed AK8 jet corresponding to the probabilities that it originates from particular heavy resonance decays, for example $Z \rightarrow b\bar{b}$, $Z \rightarrow q\bar{q}$ (where q is a light or c quark), $W \rightarrow cs$. Reconstructed AK8 jets originating from t quarks (V bosons) are selected by requiring $p_T > 400$ (200) GeV, SD mass m_{SD} between 120 and 210 (65 and 120) GeV, and a DeepAK8 probability score for t quarks versus QCD (W bosons versus QCD) larger than between 72.5 and 83.4 (91.8 and 92.5)% depending on the year of data taking. The resulting t quark (W boson) identification efficiency at the $p_T = 400$ (200) GeV threshold is estimated in simulation as 28 (25)%, with a 1% misidentification rate from QCD jets. Simulated events containing AK8 jets are corrected to agree with the data using data-derived efficiency correction factors, and dedicated JEC are also applied [127]. The corrections and uncertainties to the JER and JES are also propagated in AK8 jets. Overlap mitigation between AK4 and AK8 jets is only required in the AK4 jet collection, given AK8 jets are clustered from the same hadronic energy deposits as AK4 jets.

A series of filters for p_T^{miss} calculation follows the prescription recommended internally by the Missing Transverse Energy (MET) POG, which targets mismeasured $\vec{p}_T^{\text{miss}}$ from various sources. The filters are applicable to data, and generally serve to negate the effects of EE noise and crystal deterioration in the ECAL. These include a PV filter to remove events that fail vertex quality criteria, a filter to address the *beam halo* arising from undesired interactions with the LHC apparatus, a filter to reduce HB and HE noise, a filter for dead trigger towers in the ECAL during reconstruction, and a filter for low-quality PF muon events.

The $t\bar{t}H$ /resolved VH has been outlined, along with the most relevant background processes highlighted, and the software used to perform the analysis described. The considerations in this chapter have established the signal and backgrounds relevant to the $t\bar{t}H$ and resolved VH channels, their relative abundance as determined from the production cross sections, and how various PF candidates, such as leptons, photons, jets and p_T^{miss} are defined. Pileup mitigation techniques, erroneous event filters, and corrections to data and simulation are also touched upon, largely following the methods recommended by the CMS Collaboration and associated POGs. The event

selection bespoke to the analysis begins in Sec. 5. The discussions on aspects of simulated samples, and on object reconstruction efficiency and WPs (where applicable) are important for the discussions on the systematic uncertainties described in Sec. 6.

Chapter 5

Analysis strategy

Having established the data sets and software employed by the analysis, the strategy by which $\mathcal{B}(H \rightarrow \text{inv})$ is extracted from data, and the methods by which the SM background contributions are estimated, are outlined. The $t\bar{t}H$ and resolved VH final states are required to be fully hadronic, for which a signal region (SR) is defined using various kinematic considerations and lepton/photon vetoes to remove any leptonic final states. The ℓ_{lost} and $Z \rightarrow \text{inv}$ background contributions in the SR can be estimated using well-defined lepton and photon control regions (CRs); single-muon and single-electron CRs are effective at estimating corrections to the expected simulated ℓ_{lost} contributions, while dimuon, dielectron, and single-photon CRs are effective at estimating corrections to the expected simulated $Z \rightarrow \text{inv}$ contributions. Where possible, the CRs have identical kinematic requirements to the SR, and the lepton or photon considerations are used in the region definition but otherwise ignored in the calculation of event observables. To handle hadronic background contributions in the SR, angular variables are introduced to target primarily QCD multijet and $t\bar{t} + \text{jets}$ backgrounds. For QCD multijet specifically, a dedicated QCD-enriched hadronic sideband (HS) is defined by inverting and constricting the QCD suppression selection applied to the SR. An illustration of the SR and CR setup is provided by Fig. 5.1. As is standard practice, studies are performed with data in the SR *blinded* to the analysers, to prevent biasing of the final results¹. In order to validate the analysis techniques in the SR, an additional region adjacent to the SR and HS with identical lepton and photon vetoes, referred to as the *validation region* (VR), is studied. This acts as a mirror for the SR, with an intermediate level of QCD multijet background and signal presence versus the SR and HS. And given this is not blinded, the VR allows for a direct comparison to

¹Blind analysis can consist of using simulation-only results, data sampling, hiding parameters, and performing the analysis with added random numbers, to verify selections and data sets.

the effects on the SR caused by background suppression, selection optimisation, and agreement between data and simulation.




Figure 5.1: An illustration of the SR and CRs considered for the analysis, including lepton/photon CRs used for the estimation of the ℓ_{lost} and $Z \rightarrow \text{inv}$ contributions to the SR, and the QCD-enriched HS used to estimate QCD multijet background in the QCD-suppressed SR.

The signal is extracted from a fit to the distribution of *hadronic recoil* (vector $\vec{p}_{\text{T}}^{\text{miss}}$), which is defined as the vectorial sum of the $\vec{p}_{\text{T}}^{\text{miss}}$ and the p_{T} of any selected charged lepton(s) or photon in an event. In the SR, the hadronic recoil corresponds to the $p_{\text{T}}^{\text{miss}}$, while in the CRs it effectively measures the p_{T}^{V} or p_{T}^{γ} in the leptonic CRs or photonic CR, respectively. This p_{T} subtraction ensures good correspondence between the SR and CRs. As with $\vec{p}_{\text{T}}^{\text{miss}}$, and $\vec{H}_{\text{T}}^{\text{miss}}$, the hadronic recoil is calculated using the AK4 jet collection, therefore JECs are propagated through the use of the $\vec{p}_{\text{T}}$ -corrected jets.

The initial selection collects events of interest via a suite of triggers whose require-

ments are applied to variables calculated using PF candidates reconstructed at the HLT. This trigger selection, along with a discussion on their selection efficiencies, is detailed in Sec. 5.1. Following the trigger suite, a series of requirements are applied that enhance the signal and optimise event quality for the hadronic recoil search, which are described in Sec. 5.2. Event categorisation, namely in how the $t\bar{t}H$ and resolved VH signals are defined, and the SM background estimations using the CRs, is explained in Sec. 5.3.

5.1 Trigger selections and efficiencies

The trigger selections vary between regions and data-taking periods. Events in the SR, HS, and the single-muon and dimuon CRs are collected using HLT selection criteria on p_T^{miss} and the missing hadronic transverse energy, H_T^{miss} , which is the magnitude of the negative $\vec{p}_T$ sum of jets reconstructed at the HLT level with a p_T^{miss} threshold of 30 GeV applied. Muonic contributions are subtracted from the computation of PF p_T^{miss} and PF H_T^{miss} to allow the same trigger to be used in the SR and the muon CRs, which produces a typical trigger efficiency of $> 90\%$ for $p_T^{\text{miss}} > 250$ GeV. The use of the combined p_T^{miss} and H_T^{miss} triggers in the muon CRs instead of single-muon triggers improves correspondence with the selection in the SR and minimises any selection bias. Both the p_T^{miss} and H_T^{miss} thresholds in 2016 (2017 and 2018) at the HLT level vary between 90 and 120 (are equal to 120) GeV in these regions. A subset of the data collected in 2017 has an additional trigger requirement of PF $H_T > 60$ GeV given this was not present in the trigger menu for the entirety of that year. During data taking in 2017, EE noise at high $|\eta|$ affected data quality, therefore events from 2017 data and simulation are corrected to reduce the effect on PF jet and p_T^{miss} measurements [128]. Additionally, for 2016 and 2017 data-taking periods, there was an inefficiency due to the ECAL trigger pre-firing in the region $|\eta| > 2.0$ [129]. Consequently, a significant proportion of forward jets were mistakenly associated as belonging to the previous bunch crossing, and many events were rejected due to failed reconstruction. This resulted in events containing an electron or photon (jet) with $p_T > 50$ (100) GeV having an efficiency loss of up to 20%, depending on p_T and η , as estimated internally by the CMS Collaboration. Correction factors for this trigger inefficiency are obtained from 2016 and 2017 data and applied to simulation samples as a function of η , calculated in Sec. 6.2.2.

Events in the single-electron and dielectron CRs from the years 2016, 2017, and 2018 are required to pass a tight single-electron trigger with a p_T threshold of 27, 35,

and 32 GeV respectively, or pass a loose single-electron trigger with a p_T threshold of 105, 115, and 115 GeV respectively, or pass a single-photon trigger with a p_T threshold of 175, 200, and 200 GeV, respectively. The single-photon trigger for 2016 is lowered from 175 GeV to 165 GeV if the fraction of the HCAL-to-ECAL energy deposit is below 10%. For 2016 and 2017 data in the electron CR, events passing the single-photon trigger must also fail both the tight and loose single-electron triggers. The low-threshold single-electron triggers require the electron candidate to pass a tight isolation condition, while the high-threshold trigger imposes a looser selection on the isolation to improve the efficiency at high p_T . Photon events are required to pass the single-photon trigger as described above. In simulation, the efficiency of electron and photon event selection is corrected with data-derived efficiency factors.

During significant periods of data taking in 2018, the HCAL sector covering $-1.57 < \phi < -0.87$ and $-3.0 < \eta < -1.39$ was not functional. This is referred to as the hadron EC module (HEM) issue, for which approximately 65% of certified CMS data was affected. A HEM filter is therefore used to remove event data from the hadronic regions with $-1.8 < \phi(\overrightarrow{\text{recoil}}) < -0.6$ if they contain jets within the affected region, and similarly for the electron CRs. To ensure good correspondence between data and simulation, the latter is reweighted to account for the efficiency loss.

5.1.1 Trigger efficiency

The trigger selection efficiency is estimated for both data and simulation in the single-muon and dimuon regions as obtained using the combined PF p_T^{miss} and PF H_T^{miss} HLT paths, with the correction applied to the simulation. The efficiency in these trigger selections is obtained from an orthogonal single-muon (dimuon) data set when compared to the simulated $W(\ell\nu) + \text{jets}$ (DY + jets) samples, with the muon $\vec{p}_T$ added to event recoil to emulate the genuine p_T^{miss} in the single-muon (dimuon) CR. This is to avoid biasing the uncertainty estimation. A curated suite of requirements imposed offline on the single-muon data and $W(\ell\nu) + \text{jets}$ samples (dimuon data and DY + jets samples) is designed to recreate the region probed in the analysis, by requiring one (two) tight muon(s) with $p_T > 30$ (22, 11) GeV, two leading jets with $p_T > 80, 40$ GeV, $H_T^{\text{miss}} > 200$ GeV, $H_T^{\text{miss}}/p_T^{\text{miss}} < 1.2$, and no other lepton vetoes. The efficiency, ϵ is computed for each p_T^{miss} interval, given by

$$\epsilon(p_T^{\text{miss}}) = \frac{N(\text{offline} + \text{trigger})}{N(\text{offline})}, \quad (5.1)$$

where $N(\text{offline} + \text{trigger})$ is the number of events on application of the combined

PF p_T^{miss} and PF H_T^{miss} trigger and offline requirements to the single-muon (dimuon) data set and simulated $W(\ell\nu) + \text{jets}$ (DY + jets) sample, and $N(\text{offline})$ is the number of events imposed by the offline selection alone. The combined PF p_T^{miss} and PF H_T^{miss} trigger efficiency is measured for individual years, presented for the single-muon data and $W(\ell\nu) + \text{jets}$ samples (dimuon data and DY + jets samples) in Fig. 5.2 5.3). Trigger efficiency correction factors are defined as the ratio of ϵ between data and simulation. The factors are applied to simulation, and presented in Fig. 5.4 in intervals of p_T^{miss} , with larger p_T^{miss} intervals summed together to ensure a sufficient population of events. An uncertainty of 2% is identified for each year to account for the uncertainty in the trigger efficiency, extracted from Fig. 5.4 in the range $p_T^{\text{miss}} < 800$ GeV to cover the bulk of the events and avoid effects due to low numbers of events in higher p_T^{miss} intervals.




Figure 5.2: Combined PF p_T^{miss} and PF H_T^{miss} trigger efficiency in single-muon data and $W(\ell\nu) + \text{jets}$ MC samples for 2016 (top left), 2017 (top right), and 2018 (bottom).




Figure 5.3: Combined PF p_T^{miss} and PF H_T^{miss} trigger efficiency in dimuon data and DY + jets MC samples for 2016 (top left), 2017 (top right), and 2018 (bottom).

The efficiency of the trigger event selection in the single-electron, dielectron, and single-photon CRs is measured following the methods described in Ref. [11] for the combined boosted VH and ggH channels, and Ref. [8] for the VBF channel. In this case, both the correction factors and their uncertainties are expected to be similar between the analyses, where any effects on the corrections due to the larger jet activity in $t\bar{t}H$ events are subdominant overall.

5.2 Offline event selection

The offline event selection is designed to refine the quality of events by selecting those with a large amount of activity from jets and a sizeable hadronic recoil within all regions. Signal purity is enhanced by selecting events with hadronic recoil > 200 GeV




Figure 5.4: Trigger efficiency correction factors, or scale factors ($SF = \text{Data}/\text{Monte Carlo}$) for 2016 (top left), 2017 (top right), and 2018 (bottom). A green band demarcates the 2% uncertainty attributed to the trigger efficiency as measured in the ratio across all p_T^{miss} intervals.

and $H_T^{\text{miss}} > 200$ GeV, as imposed by the HLT, and requiring $p_{T,1}^j > 80$ GeV.

The quality of an event is a measure of how well reconstructed it is. In such events, the hadronic recoil and H_T^{miss} are in good agreement, which is assured via the requirement that $H_T^{\text{miss}} / \text{recoil} < 1.2$, and azimuthal separation $|\Delta\phi(\vec{\text{recoil}}, \vec{H}_T^{\text{miss}})| < 0.5$. For a well-defined hadronic region, an additional selection on p_T^{miss} is applied as calculated using only charged PF particles, of $p_{T,\text{track}}^{\text{miss}} > 60$ GeV, and azimuthal separation $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_T^{\text{miss}_{\text{track}}})| < 1$. These ensure good correspondence in the hadronic recoil distributions of the SR and HS as measured when charged lepton PF candidates are included and excluded. In $t\bar{t}H$ events, this is further tightened by requiring $\sqrt{(\Delta\phi(\vec{\text{recoil}}, \vec{p}_T^{\text{miss}_{\text{track}}})^2 + 4\Delta\phi(\vec{\text{recoil}}, \vec{H}_T^{\text{miss}})^2)} < 1.0$, due to a small number of poorly-described events. The consequence of this selection is minimal: the targeted events are

removed, and there is neither a discernible effect on the shape of the hadronic recoil distribution nor a bias introduced to the result.

Given the SR is similar to that of the other $H \rightarrow \text{inv}$ searches in the VBF channel [8] and the ggH/boosted VH channels [11], there is a high potential for overlap across the regions. A particularly severe contamination would impede the combination across all channels, therefore event overlap is mitigated through additional selections in this analysis. To facilitate the combination with the VBF analysis, a veto is implemented via inversion of the kinematic selection of VBF topologies. This is performed by removing events that fail any one of the following criteria: hadronic recoil > 250 GeV; the p_T of the *subleading* AK4 jet (i.e. the AK4 jet with the second largest p_T in an event), $p_{T,2}^j > 40$ GeV; the invariant mass of the two leading AK4 jets, or the *dijet mass*, $m_{jj} > 200$ GeV; the jets are well separated in η , such that $|\Delta\eta(\vec{p}_{T,1}^j, \vec{p}_{T,2}^j)| > 1.0$; their azimuthal separation is sufficiently small, such that $|\Delta\phi(\vec{p}_{T,1}^j, \vec{p}_{T,2}^j)| < 1.5$; and the jets are in opposite η -hemispheres, such that $\eta_1^j \cdot \eta_2^j < 0$. Additionally, events containing highly forward AK4 jets are vetoed, by requiring $|\eta_1^j| < 2.4$ and $|\eta_2^j| < 2.4$.

Orthogonality with respect to leptonic $t\bar{t}H$ decays [12, 13] is accounted for with requirements on the *transverse mass*, m_T , which is the component of the invariant mass in the transverse plane of a particle or set of particles. This is commonly used at the LHC in the context of a set of decay products, where one of the daughter particles is invisible, and contributes to the hadronic recoil instead. For a pair of daughter particles, with p_T of $\vec{p}_{T,1}$ and $\vec{p}_{T,2}$, and transverse energy $E_{T,1}$ and $E_{T,2}$, the m_T of the system is given by

$$\begin{aligned} m_T^2 &= (E_{T,1} + E_{T,2})^2 - (\vec{p}_{T,1} + \vec{p}_{T,2})^2 \\ &= m_1^2 + m_2^2 + 2(E_{T,1}E_{T,2} - \vec{p}_{T,1} \cdot \vec{p}_{T,2}), \end{aligned} \quad (5.2)$$

which in the event that one daughter is invisible, and in the limit where the daughter particles are highly energetic such that the rest mass $m_0 \ll p_T$, becomes

$$m_T^2 = 2p_T \cdot (|\overrightarrow{\text{recoil}}|) \cdot (1 - \cos(\Delta\phi)), \quad (5.3)$$

where $\Delta\phi = \phi - \phi(\overrightarrow{\text{recoil}})$, and p_T and ϕ pertain to the visible daughter particle. In this case, the m_T of the lepton and hadronic recoil combined (i.e. the component of the invariant mass of the lepton and hadronic recoil in the transverse plane) requires $m_T^\ell < 110$ GeV to be applied in the single-lepton CRs, to be consistent with W boson decay. A requirement on the dilepton mass $m_{\ell+\ell-} < 120$ GeV is implemented in the dilepton CRs. The selection on $m_{\ell+\ell-}$ is also subtly effective at suppressing $t\bar{t}H$ signal

contamination in the CRs. The requirements implemented for the offline selection are summarised in Table 5.1.

Table 5.1: The offline selections applied in the various regions probed by this analysis in order to improve signal purity, event quality, and mitigate overlap with the regions of other $H \rightarrow \text{inv}$ searches.

Variable	Selection	Purpose	Region
Hadronic recoil	$> 200 \text{ GeV}$	Signal purity	All
H_T^{miss}	$> 200 \text{ GeV}$		All
$p_{T,1}^j$	$> 80 \text{ GeV}$		All
$H_T^{\text{miss}} / \text{recoil}$	< 1.2	Event quality	All
$ \Delta\phi(\text{recoil}, \vec{H}_T^{\text{miss}}) $	< 0.5		All
$p_{T,\text{track}}^{\text{miss}}$	$< 60 \text{ GeV}$		SR, HS
$ \Delta\phi(\text{recoil}, \vec{p}_T^{\text{misstrack}}) $	< 1.0		SR, HS
$ \eta_1 , \eta_2 $	< 2.4	Analysis orthogonalisation	All
VBF signal	Veto by inversion		All
m_T^ℓ	$< 110 \text{ GeV}$		Single-lepton CRs
$m_{\ell^+\ell^-}$	$< 120 \text{ GeV}$		Dilepton CRs

5.3 Event categorisation

Having established the region definitions, and refined the purity and quality of the selected events, the next stage is to establish how the final states are categorised into the $t\bar{t}H$ and resolved VH topologies, and how the CRs are used to estimate the SM background contributions within each channel. There are three categories of hadronic final state sought, the first requiring at least one boosted t quark or W boson, belonging to boosted $t\bar{t}H$ topologies. Events without boosted objects require at least one b jet, targeting the resolved $t\bar{t}H$ topology. Both topologies in the $t\bar{t}H$ category require at least 5 AK4 jets. Finally, the remaining events are allocated to the resolved VH topology if they contain two resolved AK4 jets with a dijet mass of $65 < m_{jj} < 120 \text{ GeV}$, compatible with that of a V boson. Within each category, events are subcategorised according to the AK4 jet (n_j), b jet (n_b), boosted t quark (n_t), and boosted W boson (n_W) multiplicities. Subcategories bear a $p_{T,2}^j$ restriction to improve event quality, with a requirement of $p_{T,2}^j > 80 \text{ GeV}$ in the $t\bar{t}H$ category to suppress

background processes, and $p_{T,2}^j > 30$ GeV in the VH category to ensure the presence of a second AK4 jet but with a sufficiently high p_T threshold to be significant against background contributions.

The expected number of events categorised in the SR according to this prescription can be much smaller for some subcategories than others, especially at high hadronic recoil. This is addressed by modifying the width of the hadronic recoil intervals for some subcategories, to improve the statistical precision of those with only $\mathcal{O}(1-10)$ events. The lower recoil bound of 200 GeV is imposed by the trigger selection, and increases in intervals of 100 GeV to 500 GeV, after which the recoil interval has no upper bound. This is imposed to ensure the population of the highest interval in the dilepton CRs is sufficiently large. The binning regime is common between the SR and CRs to ensure good correspondence between the events in each region and hadronic recoil interval. By the same argument, the highest interval is combined with lower recoil intervals for $t\bar{t}H$ 2Boosted1b and 2Boosted2b, and VH 2j1b and 2b2b subcategories. The definitions and binning regime in hadronic recoil for each subcategory are optimally selected to achieve a set of event samples with high purity for a given production mode, and minimal background contamination or signal cross contamination. Optimisation is determined using *figures of merit*, such as the expected significance, Z_{exp} , between signal S and background B yields, equal to

$$Z_{\text{exp}} = \frac{S}{\sqrt{B + \sigma_B^2}}, \quad (5.4)$$

in the Poissonian regime, for a systematic uncertainty of σ_B . This is an effective measure of a signal count being statistically significant above a background fluctuation, where one standard deviation in the background B is equal to $\sqrt{B}$. The peak in the figure of merit for the distribution of a given variable corresponds to its selection threshold. Another figure of merit is the *Asimov significance*, given by

$$Z_A = \sqrt{2 \left((S + B) \ln \left(1 + \frac{S}{B} \right) - S \right)}, \quad (5.5)$$

when not accounting for the systematic uncertainty, and

$$Z_A = \sqrt{2 \left((S + B) \ln \left(\frac{(S + B)(B + \sigma_B^2)}{(B^2 + (S + B)\sigma_B^2)} \right) - \left(\frac{B}{\sigma_B} \right)^2 \ln \left(1 + \frac{S\sigma_B^2}{B(B + \sigma_B^2)} \right) \right)}, \quad (5.6)$$

when σ_B is assumed to be 5%. In the limit $S \ll B$, $Z_A \approx Z_{\text{exp}}$ is retrieved [130]. The Asimov significance is typically adopted in LHC and other HEP analyses, where the convention is to estimate a likelihood ratio for hypothesis testing. In the asymptotic limit of such large sample sizes as those available at the LHC in data and simulation, Z_A is an appropriate measure of the significance [130, 131].

In the resolved $t\bar{t}H$ category, one of the largest backgrounds is due to $t\bar{t} + \text{jets}$ events. The topologies of the $t\bar{t}$ and $t\bar{t}H$ systems in the transverse plane are depicted in Fig. 5.5 (upper) along with the $\vec{\text{recoil}}$, where t quarks are highly boosted and emitted back-to-back from the IP, while in the $t\bar{t}H$ system the $t\bar{t}$ pair tends to be boosted into the opposite azimuthal hemisphere to that of the invisibly-decaying Higgs boson. The $\vec{\text{recoil}}$ is therefore often correlated with the $\vec{p}_T$ of one of the two b jets produced in the $t\bar{t}$ system, however is typically aligned with the Higgs boson direction in the $t\bar{t}H$ system. By considering pairs of selections according to the azimuthal separation between $\vec{p}_{T,1}^j$, the leading b jet $\vec{p}_T$, $\vec{p}_{T,1}^b$, or the subleading b jet $\vec{p}_T$, $\vec{p}_{T,2}^b$, and $\vec{\text{recoil}}$, $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^j)|$, $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^b)|$, or $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,2}^b)|$ respectively, it is possible to mitigate the presence of $t\bar{t} + \text{jets}$ background in $t\bar{t}H$ events with fully resolved hadronic t quark decays. An additional requirement ensures that the leading AK4 jet is not the leading b jet when measuring $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^j)|$. These angular variables prove to be sensitive to the $t\bar{t} + \text{jets}$ background, and are an effective discriminator between this and $t\bar{t}H$ signal, as demonstrated in Fig. 5.5 (lower). Here, events containing resolved $t\bar{t}H$ topologies are used to measure the correlation between $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^b)|$ and $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^j)|$ ($|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,2}^b)|$) for resolved $t\bar{t}H$ events containing one b jet (at least two b jets).

The optimisation of the selection for $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^b)|$ and $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^j)|$, and $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^b)|$ and $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,2}^b)|$, is performed on Z_A in 2D using the F.A.S.T.-CALIPER package [132]. The cumulative total of S and B is computed bin-wise in 2D to find a maximal regional value for Z_A for $\sigma_B = 5\%$. The significance distribution and the resulting F.A.S.T.-CALIPER output is illustrated in Fig. 5.6 in the case of the two b jet resolved $t\bar{t}H$ subcategories, with the optimal requirements $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,1}^b)| > 1.0$ and $|\Delta\phi(\vec{\text{recoil}}, \vec{p}_{T,2}^b)| > \pi/2$ indicated.

The category definitions and the hadronic recoil binning regime optimised for signal selection, along with the $t\bar{t} + \text{jets}$ suppression requirements, are presented in Table 5.2.

5.3.1 ℓ_{lost} and $Z \rightarrow \text{inv}$ background estimation

The single-lepton CRs are used to calculate the background contributions due to ℓ_{lost} processes in the SR, which mainly appear in the $t\bar{t}H$ and VH 2j2b categories. In




Figure 5.5: Upper: The topology of the $t\bar{t}$ and $t\bar{t}H$ system in the transverse plane of the detector, where the $\vec{p}_T^{\text{recoil}}$, denoted by a solid arrow, is typically aligned with the leading or subleading b-jet direction in $t\bar{t}$ events, and aligned with the invisibly decaying Higgs boson in $t\bar{t}H$ events. Lower: The distribution of $t\bar{t}$ and $t\bar{t}H$ processes in $|\Delta\phi(\vec{p}_T^{\text{recoil}}, \vec{p}_{T,1}^b)|$ and either $|\Delta\phi(\vec{p}_T^{\text{recoil}}, \vec{p}_{T,1}^j)|$ (for the 5/6j1b subcategories) or $|\Delta\phi(\vec{p}_T^{\text{recoil}}, \vec{p}_{T,2}^b)|$ (for the 5/6j2b subcategories), demonstrating the alignment or anti-alignment of $\vec{p}_T^{\text{recoil}}$ commonly with the lead b jet for $t\bar{t}$ processes, and the lack of correlation with the presence of the invisibly decaying Higgs boson in $t\bar{t}H$ processes.




Figure 5.6: Left: The Z_A distribution in the two b jet resolved $t\bar{t}H$ subcategories in intervals of $|\Delta\phi(\text{recoil}, \vec{p}_{T,1}^b)|$ and $|\Delta\phi(\text{recoil}, \vec{p}_{T,2}^b)|$, assuming a background systematic uncertainty $\sigma_B = 5\%$. The selections on each variable are indicated as white lines, and correspond to the optimal values of Z_A in each region. Right: The distributions of $|\Delta\phi(\text{recoil}, \vec{p}_{T,1}^b)|$ (top) and $|\Delta\phi(\text{recoil}, \vec{p}_{T,2}^b)|$ (bottom) in Z_A for each variable as computed during optimisation, with the optimum thresholds indicated in each case.

the $t\bar{t}H$ category, ℓ_{lost} is mainly from $W + \text{jets}$ and t quark-related processes, while in the VH category, this is mainly from $W + \text{jets}$ events. The single-muon (single-electron) CR is defined by requiring exactly one tight muon (electron) with $p_T > 20$ (40) GeV [68, 74]. The m_T of the charged lepton p_T and the hadronic recoil in the single-lepton CRs is given by

$$m_T^\ell = \sqrt{2p_T^\ell \cdot (|\vec{\text{recoil}}|) \cdot (1 - \cos(\phi(\vec{p}_T^\ell) - \phi(\vec{\text{recoil}})))}, \quad (5.7)$$

where a mass window of $50 < m_T^\ell < 110$ GeV is required in both the single-muon and single-electron CRs.

The dilepton CRs are used to calculate the background contribution due to $Z \rightarrow \text{inv}$ processes in the SR, which dominates the VH category and the $t\bar{t}H$ category at high hadronic recoil. Given the processes $Z \rightarrow \nu\bar{\nu}$ and $Z \rightarrow \ell^+\ell^-$ are kinematically identical as a direct result of lepton universality, the dilepton CRs can accurately constrain the $Z \rightarrow \text{inv}$ contribution and minimise any theoretical uncertainty mapping between the CR and SR process. The dimuon (dielectron) CR is defined by requiring one tight muon (electron) with $p_T > 20$ (40) GeV, and one loosely-isolated muon (electron) with the opposite charge and $p_T > 10$ (10) GeV [68, 74]. The invariant mass of the dimuon system, $m_{\mu\mu}$ (dielectron system, m_{ee}) must be compatible with that of a

Table 5.2: The categorisation of SR events into the boosted $t\bar{t}H$, resolved $t\bar{t}H$ and VH production modes sought in the analysis. The intervals of hadronic recoil for each subcategory are designed to ensure the bins are sufficiently populated in all regions.

Category	Subcategory	n_j	n_b	n_t	n_W	$p_{T,2}^j$ (GeV)	Other	Hadronic recoil intervals (GeV)
$t\bar{t}H$ boosted	2Boosted1b	≥ 5	1	2				
	2Boosted2b	≥ 5	≥ 2	2				[200, 300, $+\infty$]
	1t1b	≥ 5	1	1	0	> 80	-	
	1t2b	≥ 5	≥ 2	1	0			
	1W1b	≥ 5	1	0	1			[200, 300, 400, 500, $+\infty$]
	1W2b	≥ 5	≥ 2	0	1			
$t\bar{t}H$ resolved	5j1b	5	1	0	0	> 80	$ \Delta\phi(\overrightarrow{\text{recoil}}, \vec{p}_{T,1}^b) > 1.0,$	[200, 300, 400, 500, $+\infty$]
	6j1b	≥ 6	1	0	0		$ \Delta\phi(\overrightarrow{\text{recoil}}, \vec{p}_{T,1}^j) > \pi/2$	
	5j2b	5	≥ 2	0	0		$ \Delta\phi(\overrightarrow{\text{recoil}}, \vec{p}_{T,1}^b) > 1.0,$	[200, 300, 400, 500, $+\infty$]
	6j2b	≥ 6	≥ 2	0	0		$ \Delta\phi(\overrightarrow{\text{recoil}}, \vec{p}_{T,2}^b) > \pi/2$	
VH	2j0b	2	0	0	0	> 30	$65 < m_{jj} < 120$ GeV	[200, 300, 400, 500, $+\infty$]
	2j1b	2	1	0	0			[200, 300, 400, $+\infty$]
	2j2b	2	2	0	0			[200, 300, $+\infty$]

Z boson, therefore the mass is between 75 and 105 GeV in the $t\bar{t}H$ category, and between 60 and 120 GeV in the VH category. The mass constraints are optimised to suppress additional background about the Z boson production resonance, with the tighter invariant mass constraint for the $t\bar{t}H$ category designed to reduce the larger combinatorial background in dileptonic $t\bar{t}$ decays. The $t\bar{t}H$ category has an additional selection of $|\Delta\phi(\overrightarrow{\text{recoil}}, \vec{p}_T^{\text{miss}_{\text{track}}})| < \pi/2$, which further reduces the $t\bar{t} + \text{jets}$ background and favours DY production in the dilepton CRs. In the $t\bar{t}H$ category, the subcategories can suffer from low numbers of events, and as such are dominated by high statistical uncertainty. This is addressed by summing the events in each subcategory together into the $t\bar{t}H$ boosted and $t\bar{t}H$ resolved topologies when extracting $\mathcal{B}(H \rightarrow \text{inv})$. In the $t\bar{t}H$ boosted category, the dimuon and dielectron CRs are summed together to form a single dilepton CR. The hadronic recoil intervals are preserved when summing over subcategories.

The single-photon CR is used for $Z \rightarrow \text{inv}$ background estimation in the VH category, where the event kinematics and topologies in $Z + \text{jets}$ and $\gamma + \text{jets}$ processes are similar. This improves the sensitivity to VH signal at high recoil because of the larger number of events than in the dilepton CRs. The single-photon CR is not used for the $t\bar{t}H$ categories as the theory behind the mapping of $\gamma + \text{jets}$ to $Z + \text{jets}$ processes has sizeable uncertainties in events containing more than two b jets. The

single-photon CR, based on the associated trigger, is defined by requiring exactly one loose photon with $p_T > 230$ GeV [68]. During identification, photons can typically be discriminated from other sources of ECAL deposits using the properties of the deposits themselves, such as isolation in ECAL and HCAL, or the shape of the EM showers. However, other particles can potentially be incorrectly identified as photons, for example in QCD multijet events where a jet is misidentified as a photon. In order to estimate the contribution from misidentified photons in the single-photon CR, a purity measurement is performed. The purity is defined as the fraction of reconstructed photon candidates that correspond to genuine isolated photons originating from the PV in the event. The photon purity is measured in data based on the lateral width $\sigma_{i\eta i\eta}$ [133], which parametrises the shape of the energy deposit associated with the photon in the ECAL. The characteristic $\sigma_{i\eta i\eta}$ distribution from genuine photons peaks at $\sigma_{i\eta i\eta} < 0.01$, while the distribution due to misidentified photons possesses a less pronounced peak with a much broader decline for $\sigma_{i\eta i\eta} > 0.01$. A template fit to the $\sigma_{i\eta i\eta}$ distribution is performed, where for genuine photons simulated $\gamma + \text{jets}$ events are used, while for misidentified photons a sample enriched in misidentified photon events is obtained by inverting the isolation requirements of the photon identification criteria in the single-photon CR. Two examples of the template fits are presented in Fig. 5.7 for the lowest and highest p_T^γ intervals, distributed in $\sigma_{i\eta i\eta}$, considered in the 2017 data set. The purity is defined as the fraction of genuine photons extracted from the fit that pass the $\sigma_{i\eta i\eta}$ selection. The photon purity is measured separately in bins of p_T^γ for each data-taking period, with the purity distributed in p_T^γ presented in Fig. 5.8. The contamination is the fraction of misidentified photons in the single-photon CR, and is estimated at around 4% for $p_T^\gamma > 200$ GeV. The QCD multijet contribution in the single-photon CR is then estimated by weighting events in data for each p_T^γ bin by the corresponding contamination. A systematic uncertainty of 25% is estimated by performing the procedure for different $\sigma_{i\eta i\eta}$ binning in the template fit, which accounts for any mismodelling of the simulated $\sigma_{i\eta i\eta}$ distribution. In comparison, the statistical uncertainty in the photon purity estimate in each p_T^γ bin is found to be much smaller than the systematic one. The full requirements for the CRs used in this analysis are presented in Table 5.3.

5.3.2 QCD multijet background estimation

QCD multijet events typically feature several jets in the final state, which if all are well measured will yield a recoil of zero. However, jet mismeasurement from calorimeter response or during reconstruction, or QCD background due to semileptonic decays




Figure 5.7: Fits of genuine (real) and misidentified (fake) photon templates to data in the $\sigma_{i\eta i\eta}$ distribution. This is performed for the lowest (left) and highest (right) p_T^γ interval considered in the 2017 data set.

of heavy-flavour objects such as B mesons, can result in a non-negligible recoil in an event. Such events will have a reasonably high p_T jet closely aligned with the $\vec{\text{recoil}}$, such that their angular separation $|\Delta\phi_{\min}(\vec{\text{recoil}}, \vec{p}_{T,1234})|$ is small. Both effects are especially pertinent to QCD multijet processes, due to the prominence of the QCD production cross section. Consequently the causes of mismeasurement, including fluctuations in calorimeter response, are more likely to appear in QCD multijet events. Mismeasurement is not an exclusive phenomenon to QCD multijet events, however background processes, such as $Z(\nu\bar{\nu}) + \text{jets}$ processes with a large genuine recoil, are less affected by jet mismeasurement due to the smaller production cross section.

Suppression of QCD multijet processes can therefore be carried out in the SR, removing events where the $\vec{\text{recoil}}$ is aligned with a mismeasured jet. Two angular variables are introduced to perform the suppression. The first requirement is on the minimum azimuthal separation between $\vec{\text{recoil}}$ and the direction of any of the four highest- p_T jets, $p_{T,1234}$. Any SR events require $|\Delta\phi_{\min}(\vec{\text{recoil}}, \vec{p}_{T,1234})| > 0.5$, while the inversion, $|\Delta\phi_{\min}(\vec{\text{recoil}}, \vec{p}_{T,1234})| < 0.5$, applies to the HS. A selection is then applied on a parameter $\tilde{\omega}_{\min}$, which is a variable designed to suppress events where the recoil is caused by a jet p_T mismeasurement, as defined in Ref. [134]. This variable is motivated by considering how the H_T^{miss} can be minimised via rescaling the p_T of a single jet in the event. The factors leading to the jet mismeasurement typically affect the p_T of the jet, but not the jet direction, hence a simple rescaling of the p_T to which H_T^{miss} is highly sensitive is indicative of the jet having been truly mismeasured. The distributions in $\tilde{\omega}_{\min}$ for the $t\bar{t}H$ category in the *representative* SR, with no requirement on $|\Delta\phi_{\min}(\vec{\text{recoil}}, \vec{p}_{T,1234})|$, is presented in Fig. 5.9 (left), illustrating the contribution of QCD multijet predominantly in the low $\tilde{\omega}_{\min}$ range. The choice for the $\tilde{\omega}_{\min}$ threshold




Figure 5.8: The fraction of genuine photons extracted from the fit distributed in p_T^γ for each data-taking year. An exponential fit yields an uncertainty of 25% when the binning regime in $\sigma_{i\eta i\eta}$ is varied.

in the SR was chosen by optimising on Z_{exp} and Z_A for $\sigma_B = 0\%$ or 5% , which leads to the change in QCD multijet contribution pivoted about the range $0.1 < \tilde{\omega}_{\text{min}} < 0.4$ for both categories, illustrated in Fig. 5.9 (right). The optimisation is performed based on the $t\bar{t}H$ category where QCD multijet background is more prevalent compared to the VH channel, although the selections that define the HS and SR are applied to both categories.

A selection of $\tilde{\omega}_{\text{min}} > 0.3$ is applied to the SR, which is effective at QCD suppression particularly in categories containing no b jets following the $|\Delta\phi_{\text{min}}(\overrightarrow{\text{recoil}}, \vec{p}_{T,1234})| > 0.5$ restriction. The inversion is more stringent in the HS, by requiring $\tilde{\omega}_{\text{min}} < 0.2$. The $\tilde{\omega}_{\text{min}}$ selection is determined by optimising the HS to be as QCD multijet-enriched as possible while ensuring the SR has negligible QCD multijet background. All other kinematic selections in the HS are identical to those applied to the SR. The VR is defined identically to the SR, but for the selection $0.2 < \tilde{\omega}_{\text{min}} < 0.3$ adjacent to it.

While the event selection in the SR aims to reduce QCD multijet contributions as

Table 5.3: Summary of the selection defining the lepton and photon CRs used to estimate the ℓ_{lost} and $Z \rightarrow \text{inv}$ backgrounds in the analysis, excluding the QCD multijet background suppression selection discussed in Sec. 5.3.2.

Control region	Category	Mass reqs. (GeV)	p_T reqs. (GeV)
Single-muon CR	$t\bar{t}H$ VH	$50 < m_T^\mu < 110$	$p_{T,1}^\mu > 20$
Single-electron CR	$t\bar{t}H$ VH	$50 < m_T^e < 110$	$p_{T,1}^e > 40$
Dimuon CR	$t\bar{t}H$ VH	$75 < m_{\mu\mu} < 105$ $60 < m_{\mu\mu} < 120$	$p_{T,1}^\mu > 20, p_{T,2}^\mu > 10$ $p_{T,1}^\mu > 20, p_{T,2}^\mu > 10, \Delta\phi(\text{recoil}, \vec{p}_T^{\text{misstrack}}) > \pi/2$
Dielectron	$t\bar{t}H$ VH	$75 < m_{ee} < 105$ $60 < m_{ee} < 120$	$p_{T,1}^e > 40, p_{T,2}^e > 10$ $p_{T,1}^e > 40, p_{T,2}^e > 10, \Delta\phi(\text{recoil}, \vec{p}_T^{\text{misstrack}}) > \pi/2$
Single-photon CR	VH	—	$p_T^\gamma > 230$




Figure 5.9: Left: The $\tilde{\omega}_{\min}$ distribution in the representative $t\bar{t}H$ SR, with no requirement on $|\Delta\phi_{\min}(\text{recoil}, \vec{p}_{T,1234})|$, illustrating the decrease in the QCD multijet yield as $\tilde{\omega}_{\min}$ is increased. Right: The significance distributions in the representative region, where the $\tilde{\omega}_{\min}$ selection closely follows the optimal values of all figures of merit for the $t\bar{t}H$ category. The shape of the significance distribution changes very little for each figure of merit, with their curves closely overlapping. The resulting HS, SR, and VR is indicated in red, green, and yellow areas, respectively.

much as possible, a QCD-enriched HS is necessary to estimate the remaining QCD multijet contribution in the SR. A data-driven approach is employed that applies two

selections to define the HS, and then uses a transfer factor equal to the ratio of the simulated number of events in the SR and HS in order to estimate the QCD multijet contribution in the SR. In both the $t\bar{t}H$ and VH categories, there are too few QCD multijet events to reliably define a transfer factor for each subcategory and hadronic recoil interval. Therefore, the selection in $|\Delta\phi_{\min}(\overrightarrow{\text{recoil}}, \vec{p}_{T,1234})|$ and $\tilde{\omega}_{\min}$ is applied to all regions in the VH category to ensure good correspondence between SR and CR events, but only to the SR for the $t\bar{t}H$ category. This is to preserve the simulated event yields in the CRs of the $t\bar{t}H$ category, particularly in the boosted $t\bar{t}H$ subcategories. Given the shape of the hadronic recoil and relative population of the $t\bar{t}H$ subcategories are observed not to depend on $|\Delta\phi_{\min}(\overrightarrow{\text{recoil}}, \vec{p}_{T,1234})|$ and $\tilde{\omega}_{\min}$ within the statistical precision of simulated QCD multijet samples, the expected QCD yields in the HS are integrated over all $t\bar{t}H$ subcategories and hadronic recoil intervals. Similarly, the hadronic recoil intervals are summed together in each VH subcategory. In neither category does the selection have a measurable effect on the shape of the hadronic recoil distribution.

The estimated QCD multijet background yield, $N_{i,j}^{\text{QCD,SR}_{t\bar{t}H}}$, in the $t\bar{t}H$ SR for a given subcategory I and hadronic recoil interval J is defined as

$$N_{I,J}^{\text{QCD,SR}_{t\bar{t}H}} = \sum_p \sum_q (N_{p,q}^{\text{data,HS}_{t\bar{t}H}} - N_{p,q}^{\text{QCD,HS}_{t\bar{t}H}}) \tau_{\text{QCD}}^{t\bar{t}H} f_I^{\text{SR}_{t\bar{t}H}} f_J^{\text{SR}_{t\bar{t}H}}, \quad (5.8)$$

where QCD refers to all processes that are not QCD multijet, the summation indices p and q correspond to each subcategory and hadronic recoil interval, respectively, $f_I^{\text{SR}_{t\bar{t}H}}$ and $f_J^{\text{SR}_{t\bar{t}H}}$ are the fractions of simulated QCD multijet events in each SR subcategory and hadronic recoil interval, respectively, and $\tau_{\text{QCD}}^{t\bar{t}H}$ is the ratio of the total expected number of simulated QCD multijet events in the SR and HS, given by

$$\tau_{\text{QCD}}^{t\bar{t}H} = \frac{N_{\text{SR}}^{\text{QCD}}}{N_{\text{HS}}^{\text{QCD}}}. \quad (5.9)$$

Here, the statement $N^{\text{data,HS}} - N^{\text{QCD,HS}} = N^{\text{QCD,HS}}$ assumes that any excess of data in the HS arises solely from QCD multijet processes.

For the VH category, the individual subcategories are sufficiently populated in the HS that a sum over the subcategories is not required when calculating the expected QCD multijet background in the HS. However, due to few QCD multijet events in the SR, an efficiency factor, $\epsilon_{m_{jj}}$, is introduced as the ratio between selecting QCD multijet events with the pre-existing $65 < m_{jj} < 120$ GeV selection and the inverted requirement. The efficiency is defined for each subcategory and hadronic recoil interval,

and is given by

$$\epsilon_{m_{jj}} = \sum_p \sum_q \frac{N_{p,q}^{\text{QCD,HS}_{\text{VH}}}(65 < m_{jj} < 120)}{N_{p,q}^{\text{QCD,HS}_{\text{VH}}}(m_{jj} < 65, m_{jj} > 120)}. \quad (5.10)$$

The estimated QCD multijet background yield is therefore given by

$$N_{I,J}^{\text{QCD,SR}_{\text{VH}}} = \sum_q (N_q^{\text{data,HS}_{\text{VH}}} - N_q^{\text{QCD,HS}_{\text{VH}}}) \tau_{\text{QCD}}^{\text{VH}} \epsilon_{m_{jj}} f_I^{\text{SR}_{\text{VH}}} f_J^{\text{SR}_{\text{VH}}}. \quad (5.11)$$

In general, the estimated QCD multijet contribution in the SR is found to be small in comparison to background from ℓ_{lost} and $Z \rightarrow \text{inv}$ processes. In addition to the statistical uncertainties, a 100% systematic uncertainty is assigned to the QCD multijet estimate. The actual uncertainty in the QCD prediction is measured at around 50% across the entire $t\bar{t}H$ category, and is applied to the SR using the method described. It is inflated to 100% to be more conservative when handling the individual $t\bar{t}H$ subcategories, which are limited by event counts especially at larger hadronic recoil. Using the more conservative 100% uncertainty compared to the 50% uncertainty was found to have negligible impact on the final fit.

5.3.3 Signal region composition

Following the event selection, signal processes are enhanced against background in the SR for the $t\bar{t}H$ and VH categories, as well as enriched in the category topology that is targeted over other signal topologies. This is to enhance the presence of $H \rightarrow \text{inv}$ events as much as possible, however a substantial amount of background remains, along with a small amount of signal contamination despite the orthogonality requirements to the topologies as tagged in other searches. The boosted $t\bar{t}H$ categories are typically highly pure in $t\bar{t}H$ events as these are well defined via the boosted t quark and W boson tagging. For the resolved $t\bar{t}H$ categories, defined by high jet multiplicity, there is some contamination from resolved VH topologies, where the V boson mass is not reconstructed and additional jets appear in the final state, and ggH processes, which have a dominant production cross section and are typically associated with final states with large numbers of jets. The VH category suffers similarly from ggH contamination, although similarly to the $t\bar{t}H$ category predominantly contains events from the topology of interest.

The relative background composition in the $t\bar{t}H$ category are predominantly from $t\bar{t} + \text{jets}$ processes, with some large $Z(\nu\bar{\nu}) + \text{jets}$ contributions in the resolved $t\bar{t}H$

subcategories. In the VH categories, $Z(\nu\bar{\nu}) + \text{jets}$ and $W(\ell\nu) + \text{jets}$ are the dominant background contributions, along with the subdominant $t\bar{t} + \text{jets}$ processes. The composition of the SR for each subcategory, using simulated samples from data taking in 2017, is illustrated in Fig. 5.10 under the assumption that $\mathcal{B}(H \rightarrow \text{inv}) = 1$.





Figure 5.10: The fractional SR composition of signal (left) and background (right) processes for each subcategory using simulated 2017 CMS samples.

The figures of merit using the signal and background yields in the SR for each hadronic recoil interval of the $t\bar{t}H$ and VH categories are presented in Fig. 5.11, alongside the distribution of signal-to-background for reference.

In summary, the hadronic recoil has been defined as the variable from which signal is extracted simultaneously across the now-established SR and CRs. The event selection for the analysis, with region definitions and definitions for the $t\bar{t}H$ and resolved VH categories, have been established, and background mitigation and estimation techniques are introduced to produce the relative yield compositions in Fig. 5.10. The next stage is to establish how to extract the signal, which involves defining the likelihood model for the analysis. The statistical interpretation of the results of a maximum-likelihood fit first requires a discussion of the statistical uncertainties in simulated event yields, and theoretical and experimental systematic uncertainties, or NPs, that determine the confidence in a discovery or exclusion limit on the observation of $H \rightarrow \text{inv}$ events, as presented in Ch. 6.



Chapter 6

Statistical interpretation

The upper limit on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ is obtained by performing a binned maximum-likelihood fit simultaneously for each year, category, region, and hadronic recoil interval. The systematic uncertainties are estimated and encoded into the likelihood model as nuisance parameters (NPs), and are correlated to varying degrees across years and categories depending on their nature. The likelihood model can be written as

$$\mathcal{L} = \mathcal{L}_{\text{SR}} \cdot \mathcal{L}_{\mu} \cdot \mathcal{L}_{\text{e}} \cdot \mathcal{L}_{\mu\mu} \cdot \mathcal{L}_{\text{ee}} \cdot \mathcal{L}_{\gamma}, \quad (6.1)$$

where $\mathcal{L}_{\text{SR}}$ is the likelihood function for any of the SRs in the boosted $\text{t}\bar{\text{t}}\text{H}$, resolved $\text{t}\bar{\text{t}}\text{H}$, and VH category, and $\mathcal{L}_{\mu}$, $\mathcal{L}_{\text{e}}$, $\mathcal{L}_{\mu\mu}$, $\mathcal{L}_{\text{ee}}$, and $\mathcal{L}_{\gamma}$ designate the likelihood functions for the single-muon, single-electron, dimuon, dielectron and single-photon CRs, respectively. The likelihood function for the SR over all n_{I} subcategories and n_{J}^i recoil intervals in each subcategory i is defined as

$$\mathcal{L}_{\text{SR}} = \prod_i^{n_{\text{I}}} \prod_{j(i)}^{n_{\text{J}}^i} \text{Poisson}(n_{\text{obs}}^{i,j} \mid n_{\text{pred}}^{i,j}), \quad (6.2)$$

where

$$n_{\text{pred}}^{i,j} = \mu s^{i,j} \rho_s^{i,j} + b_{\ell_{\text{lost}}}^{i,j} A_{\ell_{\text{lost}}}^{i,j} \rho_{\ell_{\text{lost}}}^{i,j} + b_{\text{Z} \rightarrow \text{inv}}^{i,j} A_{\text{Z} \rightarrow \text{inv}}^{i,j} \rho_{\text{Z} \rightarrow \text{inv}}^{i,j} + b_{\text{QCD}}^{i,j} \rho_{\text{QCD}}^{i,j}. \quad (6.3)$$

Here, μ is the signal strength and is interpreted as an estimator for $\mathcal{B}(\text{H} \rightarrow \text{inv})$, where the signal prediction assumes that $\mathcal{B}(\text{H} \rightarrow \text{inv}) = 1$. Roman indices i and j denote the subcategory and hadronic recoil interval, respectively, within the SR. In the first term, $s^{i,j}$ is the predicted number of simulated signal events, and $\rho_s^{i,j}$ is the estimate of the contribution due to systematic uncertainties affecting the prediction $s^{i,j}$. In the second

(third) term, $b_{\ell_{\text{lost}}}^{i,j}$ ($b_{Z \rightarrow \text{inv}}^{i,j}$) denotes the predicted number of simulated ℓ_{lost} ($Z \rightarrow \text{inv}$) events, $\rho_{\ell_{\text{lost}}}^{i,j}$ ($\rho_{Z \rightarrow \text{inv}}^{i,j}$) is the contribution due to systematic uncertainties affecting $b_{\ell_{\text{lost}}}^{i,j}$ ($b_{Z \rightarrow \text{inv}}^{i,j}$), and $A_{\ell_{\text{lost}}}^{i,j}$ ($A_{Z \rightarrow \text{inv}}^{i,j}$) is a free normalisation parameter that simultaneously scales the ℓ_{lost} ($Z \rightarrow \text{inv}$) contributions in the SR and the sum of $W + \text{jets}$, $t\bar{t} + \text{jets}$, and single t quark ($DY + \text{jets}$, $\gamma + \text{jets}$, $t\bar{t}Z$, and multiboson) background contributions in the single-lepton (dilepton and photon) CRs. The role of these parameters is described in Sec. 6.3. In the fourth term, $b_{\text{QCD}}^{i,j}$ denotes the predicted number of QCD events in the SR, and $\rho_{\text{QCD}}^{i,j}$ is the contribution due to systematic uncertainties affecting $b_{\text{QCD}}^{i,j}$.

The likelihood function for the single-muon and single-electron CR is defined as

$$\mathcal{L}_{\mu,e} = \prod_i^{n_I} \prod_{j(i)}^{n_j^i} \text{Poisson}(n_{\text{obs}}^{i,j} \mid B_{t/W}^{i,j} A_{t/W}^{i,j} \rho_{t/W}^{i,j} + B_{\cancel{t/W}}^{i,j} \rho_{\cancel{t/W}}^{i,j}), \quad (6.4)$$

where $B_{t/W}^{i,j}$ is the sum of simulated $W + \text{jets}$, $t\bar{t} + \text{jets}$, and single t quark yields, which are scaled by the free parameter $A_{t/W}^{i,j}$ simultaneously with $A_{\ell_{\text{lost}}}^{i,j}$ in the SR, and $B_{\cancel{t/W}}^{i,j}$ is the sum of all other SM background processes in the CR. Similarly, the likelihood function for the dimuon, dielectron, and single-photon CR is defined as

$$\mathcal{L}_{\mu\mu,ee,\gamma} = \prod_i^{n_I} \prod_{j(i)}^{n_j^i} \text{Poisson}(n_{\text{obs}}^{i,j} \mid B_{Z/\gamma}^{i,j} A_{Z/\gamma}^{i,j} \rho_{Z/\gamma}^{i,j} + B_{\cancel{Z/\gamma}}^{i,j} \rho_{\cancel{Z/\gamma}}^{i,j}), \quad (6.5)$$

where $B_{Z/\gamma}^{i,j}$ is the sum of simulated $DY + \text{jets}$, $\gamma + \text{jets}$, $t\bar{t}Z$, and multiboson yields, which are scaled by the free parameter $A_{Z/\gamma}^{i,j}$ simultaneously with $A_{Z \rightarrow \text{inv}}^{i,j}$ in the SR, and $B_{\cancel{Z/\gamma}}^{i,j}$ is the sum of all other SM background processes in the CR. In all CR likelihood functions, $\rho^{i,j}$ is the contribution due to systematic uncertainties affecting $B^{i,j}$.

The yields in the dimuon and dielectron CRs for the $t\bar{t}H$ category are summed into the boosted and resolved $t\bar{t}H$ classes due to low event counts in the individual subcategories. Additionally, yields for the boosted $t\bar{t}H$ class are summed to form a single dilepton CR. Consequently, the index i corresponds to either the boosted or resolved $t\bar{t}H$ class in Eq. 6.5, and the free parameters $A^{i,j}$ are shared between subcategories.

To perform the maximum-likelihood fit, the sources of systematic uncertainty are measured. These are treated as NPs, and are both theoretical and experimental in origin, as presented in Table 6.1 with the pre-fit effect on the yields of each category and region post-event selection, across year of data taking, subcategory, hadronic recoil interval, and all SM background processes, when the respective systematic uncertainty

is changed within ± 1 standard deviation. Some systematic uncertainties are grouped to improve presentation, and muon and electron CRs are combined into single-lepton and dilepton CRs, with the largest effects in the systematic uncertainty in that group, and in the given region, presented. The systematic uncertainties are described in detail in the following sections.

6.1 Theoretical systematic uncertainties

In simulated signal samples, theoretical uncertainties associated with the uncertainties in PDF parametrisation and missing higher order corrections in the QCD and EW perturbative expansions are estimated following the procedure outlined in Ref. [135]. The uncertainties in QCD scale, $\delta(\text{QCD scale})$, and PDF and α_s , $\delta(\text{PDF}+\alpha_s)$, for the $t\bar{t}H$, various VH processes, VBF, and ggH production cross sections at $\sqrt{s} = 13$ TeV are correlated across years. The values are correlated across ggH and $q\bar{q}H$ processes for $\delta(\text{PDF}+\alpha_s)$, with the processes grouped respectively as $t\bar{t}H$ and ggH, and VH and VBF, processes, and their values and correlations summarised in Table 6.2. These are log-normal constrained parameters in the fit.

For uncertainties in background samples, simulations of $V + \text{jets}$ processes are sourced from LO samples with NLO corrections applied using K-factors. The corrections are applied separately in QCD and EW accuracy to the QCD production of $V + \text{jets}$, with corrections derived from NLO $W + \text{jets}$ samples applied to the LO $W(\ell\nu) + \text{jets}$ processes, and corrections derived from NLO $DY + \text{jets}$ samples applied to the LO $DY + \text{jets}$ and $Z(\nu\bar{\nu}) + \text{jets}$ processes. The ME calculations for NLO samples contain up to two additional jets, and require a selection kinematically similar to that of the SR, such that $p_{T,1}^j > 80$ GeV, $p_T^V > 200$ GeV, $H_T > 200$ GeV, and $H_T/p_T^V < 1.2$, where all variables are defined according to the vector sum of jet or V-boson p_T at generator level. Given leptons appear in the calculation of these variables at generator level, an initial overlap removal is performed. Any generator-level jet within a cone of $\Delta R < 0.4$ of a selected lepton is removed from the calculation of these variables. The NLO K-factor distributions are presented in Fig. 6.1 as applied to QCD-produced $V + \text{jets}$ samples irrespective of the year of data taking. The uncertainty associated with the QCD renormalisation scale, factorisation scale, and PDF of the K-factors are treated as uncorrelated sources individually, although their respective values are correlated across each year.

The EW corrections to LO $V + \text{jets}$ samples require a different K-factor to reweight events to NLO in QCD production. The corrections are provided in Ref. [110] for

Table 6.1: The pre-fit maximum and minimum effects of systematic uncertainties on the event yields in each region and category, calculated by year, subcategory, recoil interval, and process, when each uncertainty is changed to ± 1 standard deviation.

Systematic uncertainties on background yields (pre-fit)							
	SR		Single-lepton		Dilepton		Single-photon
	tH cat.	VH cat.	tH cat.	VH cat.	tH cat.	VH cat.	VH cat.
Theoretical uncertainties							
Fact. scale V + jets (QCD)	<1.0-7.7 %	<1.0-19 %	<1.0-2.6 %	<1.0-11 %	<1.0-20 %	<1.0-22 %	6.0 %
Ren. scale V + jets (QCD)	<1.0-7.2 %	<1.0-8.6 %	<1.0-3.6 %	<1.0-10 %	<1.0-14 %	2.0-11 %	12 %
PDF V + jets	<1.0-9.1 %	2.0-23 %	<1.0-3.1 %	<1.0-15 %	<1.0-23 %	<1.0-26 %	8.0 %
Ren. & Fact. scale tH (QCD)	<1.0-1.7 %	<1.0 %	<1.0-1.4 %	<1.0-1.4 %	<1.0 %	<1.0 %	—
Ren. & Fact. scale t and tX (QCD)	7.8-15 %	2.5-9.3 %	6.4-17 %	<1.0-6.3 %	<1.0-5.8 %	<1.0-5.8 %	—
p _T ^b reweighting and PDF	<1.0-3.1 %	<1.0-1.2 %	<1.0-4.0 %	<1.0-3.9 %	<1.0 %	<1.0 %	—
Ren. & Fact. scale VV (QCD)	<1.0 %	<1.0 %	<1.0 %	<1.0 %	<1.0 %	<1.0 %	<1.0 %
tH and VH cat. cross section (QCD)	5.8-9.2 %	<1.0-3.8 %	—	—	—	—	—
tH and VH cat. cross section (PDF & α _s)	3.6 %	1.6-1.8 %	—	—	—	—	—
ISR	2.0 %	3.0-6.0 %	2.0 %	<1.0-4.2 %	2.0 %	6.0 %	<1.0-4.0 %
FSR	5.0 %	3.0-5.0 %	2.0-2.2 %	<1.0-3.1 %	4.6-5.0 %	5.0 %	2.0-3.0 %
Photon normalisation	—	—	—	—	—	—	40 %
Experimental uncertainties							
Integrated luminosity	1.2-2.5 %	1.2-2.5 %	1.2-2.5 %	1.2-2.5 %	1.2-2.5 %	1.2-2.5 %	1.2-2.5 %
t quark tagging	3.2-6.5 %	—	2.1-5.7 %	—	—	—	—
W boson tagging	7.8-18 %	—	7.1-18 %	—	—	—	—
b jet tagging	8.2-12 %	8.2-22 %	6.5-11 %	2.4-11 %	5.6-8.7 %	1.6-9.6 %	6.6-9.0 %
Electron identification and isolation	—	—	3.7-11 %	4.7-9.6 %	<1.0-15 %	<1.0-20 %	—
Electron reconstruction	—	—	<1.0-1.8 %	<1.0 %	1.0-1.5 %	<1.0-1.4 %	—
Muon identification	—	—	<1.0-1.0 %	<1.0-1.0 %	<1.0-1.8 %	<1.0-1.9 %	—
Muon isolation	—	—	<1.0 %	<1.0 %	<1.0 %	<1.0 %	—
Lepton veto	<1.0 %	<1.0 %	—	—	—	—	—
Photon identification and isolation	—	—	—	—	—	—	2.4-12 %
Photon reconstruction	—	—	—	—	—	—	<1.0 %
Pileup	1.4-8.8 %	<1.0-4.5 %	<1.0-4.8 %	<1.0-4.7 %	<1.0-2.1 %	<1.0-7.9 %	<1.0-3.3 %
Trigger inefficiency (pre-fire)	<1.0-12 %	<1.0-1.4 %	<1.0-3.4 %	<1.0-2.4 %	<1.0-1.6 %	<1.0-1.5 %	<1.0-0.3 %
Trigger selection efficiency	2.0 %	2.0 %	2.0 %	2.0 %	2.0 %	2.0 %	2.0 %
Tau lepton veto	<1.0 %	<1.0 %	<1.0-1.0 %	<1.0-2.4 %	<1.0 %	<1.0 %	<1.0 %
JER	2.4-3.6 %	<1.0-1.1 %	1.7-3.0 %	<1.0-1.5 %	<1.0-3.5 %	<1.0-1.4 %	<1.0-2.9 %
JES	<1.0-6.3 %	<1.0-2.9 %	<1.0-5.0 %	<1.0-2.2 %	<1.0-6.7 %	<1.0-2.8 %	<1.0-3.8 %
QCD prediction	100 %	100 %	—	—	—	—	—

Table 6.2: Uncertainty values for the QCD scale, $\delta(\text{QCD scale})$, and PDF and α_s , $\delta(\text{PDF}+\alpha_s)$, in the signal process production cross sections, $t\bar{t}H$, VH , VBF , and ggH , and the correlation scheme for $\delta(\text{PDF}+\alpha_s)$ across the processes. The values are taken from Ref. [135], and are treated as log-normal constrained in the fit.

Signal	$\delta(\text{QCD scale})$ (%)	$\delta(\text{PDF}+\alpha_s)$ (%)	$\delta(\text{PDF}+\alpha_s)$ correlation
VBF	+0.4%	$\pm 2.1\%$	$q\bar{q}H$
$W^\pm H$	-0.3%	$\pm 1.8\%$	$q\bar{q}H$
ZH	+0.5%	$\pm 1.6\%$	$q\bar{q}H$
$t\bar{t}H$	-0.7%	$\pm 3.6\%$	ggH
ggH	+3.8%	$\pm 2.4\%$	ggH
	-3.1%		
	+5.8%		
	-9.2%		
	+4.6%		
	-6.7%		

$V + \text{jets}$ and $\gamma + \text{jets}$ processes as functions of p_T^V and p_T^γ defined at generator level, respectively. The EW-derived K-factors are applied to the LO samples, with NLO DY + jets samples used to apply corrections to LO DY + jets and $Z(\nu\bar{\nu}) + \text{jets}$ processes. The values are presented in Fig. 6.2, where the drop-off at $p_T^V > 1250$ GeV has negligible impact as there are few-to-no events in this p_T^V range. The uncertainty in the PDF, renormalisation scale, and factorisation scale in simulated $\gamma + \text{jets}$ samples is accounted for by the K-factors, and set at flat values of 8%, 12%, and 6%, respectively, across hadronic recoil intervals and correlated across all years of data taking. A combined QCD renormalisation and factorisation scale uncertainty is applied to the shape of the hadronic recoil spectrum for multiboson processes, and is treated as correlated across years.

The dilepton and single-photon CRs are used in the VH category to estimate the $Z \rightarrow \text{inv}$ contribution in the SR. As the $Z(\ell^+\ell^-) + \text{jets}$ processes in the CRs are kinematically identical to the $Z(\nu\bar{\nu}) + \text{jets}$ backgrounds to which they map in the SR, there is no consideration for systematic uncertainty in this mapping. In the single-photon CR, $\gamma + \text{jets}$ processes are not kinematically identical to $Z(\nu\bar{\nu}) + \text{jets}$, and therefore there is an uncertainty in the mapping between their p_T spectra. Although there is good agreement between the hadronic recoil distributions in these CRs, there are differences in the normalisation between the dilepton CRs, which agree well amongst themselves, and the photon CR. This is accounted for by measuring the fraction of data-to-simulation ratios in each hadronic recoil interval between the dimuon and single-photon CRs, and between the dielectron and single-photon CRs, and adopting the largest fraction as the normalisation uncertainty. Due to low statistics in the high-




Figure 6.1: The QCD-derived K-factors applied to LO QCD simulations of DY + jets (top left), Z($\nu\bar{\nu}$) + jets (top right), and W($\ell\nu$) + jets (bottom) derived differentially as functions of p_T^V and $p_{T,1}^j$ from NLO QCD DY + jets (top) and W + jets (bottom) samples.

est hadronic recoil intervals, disagreement measurements are only taken in the lower recoil ranges. A systematic uncertainty of 40% is estimated and applied as a flat uncertainty across all hadronic recoil intervals in the single-photon CR. The uncertainty is uncorrelated between the individual VH subcategories, and is only correlated across 2017 and 2018, where the CP5 tune is used to generate γ + jets samples, and the CUETP8M1 tune is used for the 2016 sample. The systematic uncertainty obtained from template fits to genuine and misidentified photon events for different $\sigma_{i\eta i\eta}$ distribution binning regimes to account for mismodelling in the simulated $\sigma_{i\eta i\eta}$ distribution is measured to be 25%, and is applied to individual years of data taking.

More minor NP contributions to the overall theoretical uncertainty include the effects of initial and final state radiation (ISR and FSR, respectively), which are derived from PSWEIGHTS at the NanoAOD level by varying the renormalisation scale in the parton shower. As these variations were introduced only for 2018 samples, the ISR




Figure 6.2: The EW-derived K-factors applied to LO QCD simulations of $DY + \text{jets}$, $Z(\nu\bar{\nu}) + \text{jets}$, $W(\ell\nu) + \text{jets}$, and $\gamma + \text{jets}$ processes derived as functions of p_T^V or p_T^γ from NLO QCD $DY + \text{jets}$, $W + \text{jets}$, and $\gamma + \text{jets}$ samples, where values are obtained from Ref. [110].

and FSR uncertainties are derived here and applied to all years. Their values are estimated in the $t\bar{t}H$ and VH categories for the signal processes, and for ℓ_{lost} , $Z \rightarrow \text{inv}$, and $\gamma + \text{jets}$ separately. Due to the difference in PYTHIA8 event tune for samples from 2016 and samples from 2017 and 2018, the systematic uncertainty is correlated across 2017 and 2018, and independent of 2016. The average value is assigned to the uncertainty across region and hadronic recoil interval, and is presented in Table 6.3.

6.1.1 Systematic uncertainties in t quark samples

The $t\bar{t} + \text{jets}$ and $t\bar{t}X + \text{jets}$ backgrounds, and $t\bar{t}H$ signal, need to be correctly modelled given their prominence in the $t\bar{t}H$ category. The default renormalisation μ_R and factorisation scale μ_F used in generating samples of $t\bar{t} + \text{jets}$ and $t\bar{t}X + \text{jets}$ are not necessarily accurate, and so are a poor description of the nature of $t\bar{t}$ production. Instead, event weights are derived from combinations of $\mu_R - \mu_F$ pairs dependent on the nominal (μ_{nom}) value and the variation ($\mu_\uparrow, \mu_\downarrow$) in these scales, yielding eight independent weight variants (excluding the $\mu_{R,\text{nom}} + \mu_{F,\text{nom}}$ pair). The weight variants

Table 6.3: The ISR and FSR uncertainties derived from 2018 samples by varying PSWeights at the NanoAOD level for individual processes and categories, which are correlated across process and category. The values are applied to 2016 separately from 2017 and 2018 due to the difference in sample tunes.

Radiation	Category	ℓ_{lost}	Process			
			$Z \rightarrow \text{inv}$	$\gamma + \text{jets}$	$t\bar{t}H$	VH
ISR	$t\bar{t}H$	0.02	0.02	—	0.01	0.03
	VH	0.04	0.06	0.08	0.01	0.03
FSR	$t\bar{t}H$	0.02	0.05	—	0.01	0.03
	VH	0.03	0.04	0.03	0.01	0.03

are multiplied by an additional factor f_{var} , given by

$$f_{\text{var}} = \frac{\sum_{\text{Events}} w^{\mu_{\text{R,nom}} + \mu_{\text{F,nom}}}}{\sum_{\text{Events}} w_{\text{weightvariant}}}, \quad (6.6)$$

as calculated over all events pre-skimming the NanoAOD data sets. This ensures the dependence on the hadronic recoil is extracted while the normalisation remains fixed. The normalisation dependence in background processes is measured later in this section. The uncertainty for signal $t\bar{t}H$ samples is applied as shown in Table 6.2. The weight variants $\mu_{\text{R},\uparrow} + \mu_{\text{F},\uparrow}$ and $\mu_{\text{R},\downarrow} + \mu_{\text{F},\downarrow}$ characterise the envelope in each sample, and are distributed similarly between regions, where regional effects cancel, shown in Fig. 6.3 for the boosted $t\bar{t}H$ class, and Fig. 6.4 for the resolved $t\bar{t}H$ class, using 2017 data sets. The analysis uses hadronic recoil intervals up to and including 500 GeV, while the envelope obscures around 1000 GeV where there are few $t\bar{t} + \text{jets}$ events. The combined systematic QCD renormalisation and factorisation scale uncertainties are applied to simulation. These QCD scale uncertainties are independent of the year of data taking, and applicable only to the shape of the hadronic recoil spectrum for each process. The uncertainty is uncorrelated between the samples themselves, considering that these samples are produced using different tunes, although for $t\bar{t}H + \text{jets}$ processes the uncertainty is correlated between the $t\bar{t}H$ signal (with an invisibly-decaying Higgs boson) and $t\bar{t}H$ background (with a visibly-decaying Higgs boson), given the generators and kinematic parameters are identical.

The distribution of the generator-level t quark p_{T} , p_{T}^t , in $t\bar{t} + \text{jets}$ samples is




Figure 6.3: The weight variants for QCD renormalisation and factorisation scales for $t\bar{t} + \text{jets}$ samples in the SR of the boosted $t\bar{t}H$ class as a function of hadronic recoil, where the corresponding QCD scale uncertainty is equal to the envelope of the distribution characterised by the variants $\mu_{R,\uparrow} + \mu_{F,\uparrow}$ and $\mu_{R,\downarrow} + \mu_{F,\downarrow}$. The analysis has hadronic recoil intervals up to the region around 500 GeV, below where the envelope obscures at higher hadronic recoil.

reweighted to achieve NNLO accuracy in QCD and NLO accuracy in EW production, following the example of Ref. [136]. This is performed using an analytic function $f(p_T^t)$ that fits to the p_T^t distribution to determine the weights. The function is independent of the year of data taking, and takes the form

$$f(p_T^t) = \exp a + b \cdot p_T^t + c \cdot (p_T^t)^2, \quad (6.7)$$

where the coefficients are estimated at $a = 1.614 \times 10^{-2}$, $b = -1.966 \times 10^{-4}$, and $c = -1.454 \times 10^{-8}$. The value $f(p_T^t)$ is equal to the correction factor for a particular




Figure 6.4: The weight variants for QCD renormalisation and factorisation scales for $t\bar{t} + \text{jets}$ samples in the SR of the resolved $t\bar{t}H$ class as a function of hadronic recoil, where the corresponding QCD scale uncertainty is equal to the envelope of the distribution characterised by the variants $\mu_{R,\uparrow} + \mu_{F,\uparrow}$ and $\mu_{R,\downarrow} + \mu_{F,\downarrow}$. The analysis has hadronic recoil intervals up to the region around 500 GeV, below where the envelope obscures at higher hadronic recoil.

t quark, with the event weight $w_{t\bar{t}}$ defined as

$$w_{t\bar{t}} = \sqrt{\prod_{i=t,\bar{t}} f_i(p_{T,i}^t)}. \quad (6.8)$$

Uncertainties from the fit to $f(p_T^t)$ are used instead of the theoretical uncertainties in the samples themselves, to avoid overlapping with the hadronic recoil dependence accounted for in the renormalisation and factorisation scale uncertainties. The variations in the coefficients a , b , and c are used as input in each iteration of the fit, in order to determine the resulting systematic uncertainties in the event weight. Each co-

efficient has a corresponding, independent uncertainty, with the variations presented in Fig. 6.5. The dependence on normalisation was similarly negated to ensure the uncertainty is only applicable to the shape of the hadronic recoil distribution. The systematic uncertainty in the PDF for $t\bar{t}$ + jets samples is also extracted in the p_T^t reweighting.




Figure 6.5: The ratio of the generator level p_T^t spectrum in NNLO QCD and NLO EW simulation of $t\bar{t}$ + jets samples to that of the POWHEG generated NLO QCD sample used in the analysis. The systematic uncertainty in each coefficient a , b , and c is equal to the deviation in par0, par1, and par2, respectively, with an additional systematic uncertainty for the PDF extracted in the reweighting.

6.2 Experimental systematic uncertainties

In object tagging and event selection, efficiency correction factors are commonly applied to simulation equal to the weight or efficiency ϵ attributed to the object of interest. At the event level, the selection efficiency w_{sel} , or the tagging efficiency for

selecting events containing the objects of interest, is defined as

$$w_{\text{sel}} = \prod_i^{n_{\text{obj}}} \epsilon_i, \quad (6.9)$$

for an event containing n_{obj} objects. The uncertainty in the selection efficiency is then equal to w_{sel} when each of ϵ_i is taken to its extrema. These are typically decorrelated between years as w_{sel} varies between years from changes in running conditions.

6.2.1 Lepton and photon identification and isolation

The systematic uncertainties in lepton and photon identification and isolation criteria are propagated via efficiency correction factors used to weight simulated events to compensate for discrepancies with the data. For electrons and photons, w_{sel} is derived from for the identification, isolation, or reconstruction efficiency differentially as a function of ECAL supercluster p_T and η of the leading electrons and photons in the event. In this case, the identification and isolation uncertainty are grouped together, independent of the reconstruction uncertainty. The average uncertainty in w_{sel} is between 1 and 5%. The efficiency correction factors and the associated uncertainties are provided internally for the CMS Collaboration courtesy of the EGamma POG, where the weights vary between the sample and the year of data taking. As such, the electron and photon identification and isolation, and reconstruction uncertainties are uncorrelated for each year and for either the $t\bar{t}H$ or VH category, and are separated between electrons and photons. Similarly, the identification and isolation efficiency for muons is calculated in p_T and η , although are not combined like electrons and photons. Typically these have uncertainties of around 1%. The muon reconstruction uncertainty is negligible compared to other uncertainties, and therefore is not considered in the analysis [74]. The muon identification and isolation uncertainties are uncorrelated for each year and each category. The loose lepton and photon veto efficiency in the SR has an associated uncertainty that is set at 0.75%, uncorrelated between years. The uncertainty associated with the tau lepton veto efficiency is similarly applied to the SR, and is decorrelated between years.

6.2.2 Trigger efficiency and pre-firing

The uncertainty in the combined PF p_T^{miss} and H_T^{miss} trigger efficiency is computed using the single-muon and dimuon CRs, following the study outlined in Sec. 5.1.1. The uncertainty in the trigger efficiency correction factors is measured at 2%. The

uncertainty in the trigger is applied separately for p_T^{miss} and H_T^{miss} , electrons, and photons, all at 2%, and for individual years. An additional uncertainty associated with the pre-fire trigger inefficiency is applied to the data-taking years 2016 and 2017 individually, following the event reweighting technique recommended internally by the CMS Collaboration. This is achieved using simulations that emulate the effects of the ECAL trigger issue, with event weight $w_{\text{pre-fire}}$ defined by

$$w_{\text{pre-fire}} = \prod_{i=\gamma, \text{jet}} 1 - \epsilon_i(p_T^{\text{EM}}, \eta), \quad (6.10)$$

where ϵ_i is the probability for an ECAL object to trigger prematurely as a function of η and the EM p_T , p_T^{EM} , equal to p_T^γ for the photon, and to the p_T multiplied by the sum of $f_E^{\text{EM}\pm}$ and $f_E^{\text{EM}0}$ for jets. The probabilities are assumed to be uncorrelated between photons and jets. The systematic uncertainty in $w_{\text{pre-fire}}$ is estimated by using the maximum value of either $0.2\epsilon_i$ or the statistical uncertainty in the particular $p_T^{\text{EM}} - \eta$ region from which ϵ_i is extracted. The trigger inefficiency is estimated at $< 1\%$.

6.2.3 Object tagging

The efficiency corrections applied to simulated events containing b jets as tagged using the DEEPCSV algorithm are defined as the ratio of the the tagging efficiency in data to that of the simulated sample for a given WP, as a function of jet p_T . The correction factors are calculated across multiple topologies, and a weighted average is obtained with a corresponding systematic uncertainty of typically 5% [84]. Efficiency corrections are similarly obtained using the DEEPAK8 algorithm as a function of AK8 jet p_T . This is performed separately for reconstructed AK8 jets tagged as t quarks and W bosons, with the associated systematic uncertainties typically between 1% and 10% in each case [127]. The b-jet, t-quark, and W-boson tagging uncertainties are applied independently in the fit, but are each correlated across years.

6.2.4 Luminosity and pileup

The uncertainty in integrated luminosity varies between 1.2 and 2.5% depending on the data-taking year [53, 137, 138], and is composed of both correlated and uncorrelated components. The uncertainty in 2016 is 1.0% correlated across 2016–2018 and 0.6% uncorrelated, in 2017 is 2.0% correlated across 2016–2018, 0.6% correlated across 2017 and 2018, and 0.9% uncorrelated, and in 2018 is 2.0% correlated across 2016–2018, 0.2% correlated across 2017 and 2018, and 1.5% uncorrelated.

Simulated events are reweighted at the NanoAOD level to adjust for pileup effects from changing beam conditions over each year of data taking, and have an associated uncertainty. Samples are generated with a certain distribution of the number of pileup interactions, which are reweighted by drawing random values from distributions provided in short LHC reference runs available for use by analysers. The recommended runs assume an inelastic pp cross section of 69.2 mb and an associated uncertainty of 4.6% for all Run 2 data sets. In this analysis, the reweighting produces a pileup uncertainty between 1 and 2% in all regions. As the pileup distributions, and therefore the weights, vary between years, the uncertainty is not correlated.

6.2.5 QCD prediction

The QCD prediction uncertainty is assigned at 100%, as detailed in Sec. 5.3.2, and is correlated separately across the boosted $t\bar{t}H$ and resolved $t\bar{t}H$ classes, but uncorrelated between the VH subcategories. Differences in the QCD multijet background distributions in the $t\bar{t}H$ and VH categories, driven by the difference in jet multiplicities, mean that the uncertainty is not correlated across category. The prediction is performed independently of the year of data taking, and includes the statistical uncertainty in the simulated samples of each hadronic recoil interval.

6.2.6 Jet energy uncertainties

The uncertainties associated with the JER and JES vary according to the detector region or the nature of the JEC. The analysis is mildly sensitive to both JER and JES effects, and therefore requires a particular treatment for Run 2 JEC uncertainties. For JES uncertainties, a prototypical grouping of constituent sources is adopted, which are roughly organised according to detector region, and are assumed to be 100% correlated if the source is common across all of 2016–2018, and 0% correlated if unique to a particular year of data taking. The prototype JES sources are listed in Table 6.4, with uncertainty values provided for internal use by the JetMET POG.

Variations due to JEC are estimated by applying correction factors separately for JER and each JES source, which are extracted differentially as functions of p_T and η , to the p_T of AK4 jets. As the corrections affect the hadronic recoil in each event, the weighted event yields are obtained for each region, subcategory, and hadronic recoil interval by applying identical selections and treatments as those applied to the nominal event yields. However, due to irregularities in symmetry between the variations in the different regions, categories, and hadronic recoil intervals, a *smoothing* technique is

Table 6.4: The list of prototype JES sources for Run 2, applied to AK4 jets with correlations approximated as 100% for sources common to all years, and 0% for those specific to a given year, obtained internally for analysers within the CMS Collaboration.

JES source	Correlation	Description
Absolute	100%	Flat absolute scale uncertainties, mainly from combined γ , $Z \rightarrow ee$, and $Z \rightarrow \mu\mu$ reference scales, and ISR and FSR corrections
Absolute 2016/17/18	0%	
BBEC1	100%	Grouped source referring to the central barrel (BB, $ \eta < 1.3$) and inner EC (EC1, $1.3 < \eta < 2.5$)
BBEC1 2016/17/18	0%	
EC2	100%	Source referring to the EC2 ($2.5 < \eta < 3.0$) detector region
EC2 2016/17/18	0%	
FlavorQCD	100%	Jet flavour uncertainty due to Pythia6 Z2/Herwig++2.3 differences in light/c/b quark and gluon responses
HF	100%	Source referring to the HF detector region, $ \eta > 3.0$
HF 2016/17/18	0%	
RelativeBal	100%	Full difference between log-linear fits of MPF and p_T balance methods (significant for Run 2, and not fully understood)
RelativeSample 2016/17/18	0%	η -dependent uncertainty due to difference between relative residuals observed with dijets, $Z + \text{jets}$ and $\gamma + \text{jets}$

applied by weighting the variations by the average variation normalised by the nominal yield in each category and region. The event selection is largely identical between SR and CRs, and therefore the effects of JER and JES variations are expected to cancel between the sets of regions post-fit. Residual discrepancies between regions remain after smoothing post-fit, leading to an average systematic uncertainty of 2% for JER and 3% for all JES sources.

6.3 Statistical uncertainty

The statistical uncertainty in simulated samples for each subcategory and hadronic recoil interval is introduced for the maximum-likelihood fit via individual, uncorrelated NPs. A single NP is assigned to each subcategory and hadronic recoil interval that scales the sum of simulated events across all processes in that interval, which is constrained by the total statistical uncertainty rather than individual statistical uncertainties for each process. This is the Barlow-Beeston-light (BBL) approach [139], designed to minimise the number of parameters required in the fit. The method requires an input threshold per bin (corresponding to a particular subcategory and hadronic recoil interval), which is the number of effective, unweighted events, $n^{\text{threshold}}$. The total simulated event yield n_{tot} and statistical uncertainty σ_{tot} for each bin is calculated, where

$$n_{\text{tot}} = \sum_{i \in \text{bkg}} n_i, \quad e_{\text{tot}} = \sqrt{\sum_{i \in \text{bkg}} e_i^2}, \quad (6.11)$$

for each process i in that bin. The effective number of unweighted events is defined as $n_{\text{tot}}^{\text{eff}} = n_{\text{tot}}^2 / e_{\text{tot}}^2$, rounded to the nearest integer. If $n_{\text{tot}}^{\text{eff}} \leq n^{\text{threshold}}$, each process will be assigned its respective statistical uncertainty within the bin. In this event, each process i for which $n_i^{\text{eff}} \leq n^{\text{threshold}}$, where $n_i^{\text{eff}} = n_i^2 / e_i^2$, a Poisson-constrained NP is created, or otherwise a Gaussian-constrained NP is assigned. If $n_{\text{tot}}^{\text{eff}} > n^{\text{threshold}}$, a single Gaussian-constrained BBL NP is created that scales the total yield in the bin. The Gaussian-constrained NP x_{Gauss} has a nominal value of 0, and scales the yield as $n_{\text{tot}} + x_{\text{Gauss}} \cdot e_{\text{tot}}$, while in the Poissonian case, the NP x_{Po} has a nominal value of 1, and scale the yield as $n_{\text{tot}} \cdot x_{\text{Po}}$. An upper threshold of $n^{\text{threshold}} = 10$ is adopted for the analysis. The advantage of the BBL approach is that NPs have a simple analytic form dependent only on the total simulated event yields and the observed number of data events in each bin. This means that on maximising the likelihood fit of the entire model, obtaining the best-fit values for statistical NPs is not required, as these can be acquired from bin-wise fits, which is of particular importance to this analysis given the large number of subcategories and hadronic recoil intervals in the model.

Another consideration is the NPs that scale, or normalise, the expected rate of a particular process in a given bin. Such rate parameters (RPs) are used to bring simulated yields into agreement with data in the CRs, and apply such a correction to simulated yields in the SR. In this analysis, these are unconstrained NPs that behave as rate multipliers for each bin, and simultaneously scale the ℓ_{lost} and $Z \rightarrow \text{inv}$ backgrounds in the SR and the corresponding background processes in the lepton and photon CRs. The only constraint on the RPs is amongst themselves, where a single

RP mediates the rate of the ℓ_{lost} background in the SR, the single-muon CR, and the single-electron CR in each subcategory and hadronic recoil interval simultaneously in the fit, and another RP mediates the $Z \rightarrow \text{inv}$ background in the SR, the dilepton CR (for the $t\bar{t}H$ boosted class), the dimuon and dielectron CRs (for the $t\bar{t}H$ resolved class and individual VH subcategories), and the photon CR (for the VH subcategories). As aforementioned, the yields in the dimuon and dielectron CRs are summed across $t\bar{t}H$ subcategories into $t\bar{t}H$ boosted and $t\bar{t}H$ resolved classes, and additionally summed into a single dilepton CR for $t\bar{t}H$ boosted class, motivated by the limited number of simulated events in individual subcategories. Consequently, each $t\bar{t}H$ class and CR (dimuon and dielectron, or dilepton) has a corresponding RP.

6.4 Statistical interpretation

The following is an overview of the approach to general $H \rightarrow \text{inv}$ searches and limit setting using the CL_S procedure. An in-depth discussion of the method of maximum-likelihood fitting, and fit evaluation, is presented in Appendix B.

All $H \rightarrow \text{inv}$ searches aim to discover new signal processes that are statistically distinct from background, by estimating the signal strength μ observed in data relative to the expected signal strength according to the SM prediction,

$$\mu = \frac{\sigma_{\text{obs}}(H \rightarrow \text{inv})}{\sigma_{\text{SM}}(H \rightarrow \text{inv})}. \quad (6.12)$$

An estimate of the upper limit on μ parametrises the sensitivity of a given channel to $H \rightarrow \text{inv}$ events. Exclusion limits are set in the absence of a statistically significant excess of events in the SR. The expected upper limit on the signal strength is constrained by the B -only fit, while the $S + B$ fit constrains the observed upper limit.

The CL_S method is adopted for limit setting in LHC analyses [131, 140–142]. This procedure estimates the p -value, or probability under a given hypothesis of obtaining a result as extreme as the observed data, in both the B -only (p_B) and $S + B$ (p_{S+B}) fits, and computes the CL_S solution

$$\text{CL}_S = \frac{p_{S+B}}{1 - p_B} \geq \alpha, \quad (6.13)$$

where $1 - \alpha$ is the *confidence level* (CL), or probability that the series of estimates scanned for various values of μ covers the true value of the signal strength. Conventionally, the CL is set at 95% for LHC analyses. The asymptotic limit for large sample

sizes is also adopted [131], and assumes that the set of simulated samples can be replaced by a single, representative Asimov data set using approximate formulae in the maximum-likelihood fit. The Asimov data set is a set of pseudo data generated based on simulated samples only, assuming the SM expectation that $\mu = 0$, such that the maximum-likelihood best-fit value of all NPs in the model are equal to their generated value. This is done to determine the expected upper limit on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ at the 95% CL, while the observed upper limit is estimated using the true data.

Having established the likelihood model in Eq. 6.1, and described the inputs μ as an estimator for $\mathcal{B}(\text{H} \rightarrow \text{inv})$, along with all the statistical and systematic uncertainties that behave as NPs in the fit, the maximum-likelihood fit is computed. The results are presented in Ch. 7 for the $t\bar{t}\text{H}$ /resolved VH channels, as well as the fit performance evaluation following the techniques detailed in Appendix B.

Chapter 7

Fit Results

The signal and background yields for the hadronic recoil distributions are obtained from a combined fit across all $t\bar{t}H$ and VH categories according to the likelihood function in Eq. 6.1. Results are presented according to a fit simultaneously to all regions (CR+SR), in both the scenarios with $\mathcal{B}(H \rightarrow \text{inv}) = 0$ (B -only fit) and with both signal and background contributions ($S + B$ fit). The former is a means of ensuring the SM backgrounds alone can accurately model the data, while in the $S + B$ fit, the signal contribution S is weighted by the best-fit value for $\mathcal{B}(H \rightarrow \text{inv})$, $\hat{\mu}$. A fit assuming $\mathcal{B}(H \rightarrow \text{inv}) = 0$ using only the CRs (CR only) is also performed to estimate the SM background in the SR independently of the observed data in the SR. In all cases, uncertainties are inclusive of statistical and systematic contributions, and yields correspond to all 2016–2018 results, although each year is treated separately in the fit. All statistical tools used to evaluate the results are introduced in Appendix B.

The predicted background yield from the CR-only and CR+SR fit in the single-muon (single-electron) CR is presented in Fig. 7.1 (Fig. 7.2) for the $t\bar{t}H$ and VH categories. In these CRs, $t\bar{t} + \text{jets}$, $W(\ell\nu) + \text{jets}$, and single t quark production processes dominate, with smaller contributions from multiboson and $t\bar{t}X$ processes. The CR-only and $S+B$ fit profiles are consistent with the data within uncertainty, with discrepancies between prediction and data appearing mainly in the highest intervals of hadronic recoil, where event counts are small. The effect of the RPs in the single-muon and single-electron CRs is also highlighted in the data-to-simulation ratios of Fig. 7.1 and Fig. 7.2, where the post-fit yields are situated equidistant from unity.

The dimuon, dielectron, dilepton (only for $t\bar{t}H$), and single-photon (only for VH) CR distributions used for the prediction of backgrounds stemming from $Z \rightarrow \text{inv}$ decays are presented in Fig. 7.3. The largest SM contributions in the lepton CRs are from $DY + \text{jets}$, with subdominant contributions from $t\bar{t} + \text{jets}$, $t\bar{t}Z + \text{jets}$, and single




Figure 7.1: Distributions of hadronic recoil in the $t\bar{t}H$ (upper) and VH (lower) categories for the single-muon CR. The black histogram shows the total background (bkg.) prediction from a CR only, B -only fit, while the red histogram shows the yields from a CR+SR $S + B$ fit.




Figure 7.2: Distributions of hadronic recoil in the $t\bar{t}H$ (upper) and VH (lower) categories for the single-electron CR. The black histogram shows the total background (bkg.) prediction from a CR only, B -only fit, while the red histogram shows the yields from a CR+SR $S+B$ fit.

t quark production in the $t\bar{t}H$ category, and from multiboson processes in the VH category. For these CRs, prediction is consistent with data throughout the $t\bar{t}H$ category, while in the VH category there are greater outliers in the highest recoil intervals from both the CR-only and $S + B$ fits. These outliers can be attributed to disagreement between differences in the data-to-prediction ratios in the dilepton CRs and the single-photon CR, which are constrained by the same RP. The single-photon CR is dominated by $\gamma + \text{jets}$ processes, with data well represented by the prediction, although this can be attributed to the conservative photon-to- $Z \rightarrow \text{inv}$ mapping systematic uncertainty, which allows the normalisation to vary by 40%.

The total SM background prediction in the SR is shown for the $t\bar{t}H$ and VH category in Fig. 7.4. The observed best-fit estimate of $\mathcal{B}(H \rightarrow \text{inv})$ for the $t\bar{t}H$ and VH categories using 2016–2018 data is $\hat{\mu} = 0.07_{-0.10}^{+0.10}(\text{stat.})_{-0.17}^{+0.18}(\text{syst.})$, under the pre-fit assumption $\mathcal{B}(H \rightarrow \text{inv}) = 1$, and the best-fit SM expectation of $\mathcal{B}(H \rightarrow \text{inv})$ estimated at $0.00_{-0.10}^{+0.10}(\text{stat.})_{-0.16}^{+0.17}(\text{syst.})$. The fit also considers the contribution from the other signal processes, including ggH and VBF. The non-negligible presence of ggH signal is largely from its predominant production cross section, as shown in Table 4.1. The overall analysis is demonstrably dominated by systematic uncertainties, however in some subcategories the results are sensitive to the low numbers of events, especially in the case of the $t\bar{t}H$ category.

The largest SR background contributions in the $t\bar{t}H$ category are from $t\bar{t} + \text{jets}$ and single t quark processes, mainly appearing in the boosted $t\bar{t}H$ class. The $Z(\nu\bar{\nu}) + \text{jets}$ and $W(\ell\nu) + \text{jets}$ background processes are subdominant in the $t\bar{t}H$ category. The predominant QCD production cross section means contributions from QCD multijet background are non-negligible, contributing up to $\approx 8\%$ of the total yield in boosted $t\bar{t}H$ subcategories at low hadronic recoil intervals. The prediction agrees with observation in all cases, with some discrepancies in the highest recoil intervals of the $t\bar{t}H$ boosted subcategories, which have the lowest event yields. The best-fit estimate for this category is $\hat{\mu} = -0.16_{-0.26}^{+0.26}$ ($0.00_{-0.25}^{+0.26}$ expected).

The SR background in the VH category is dominated by $Z(\nu\bar{\nu}) + \text{jets}$ and $W(\ell\nu) + \text{jets}$ processes, followed by $t\bar{t} + \text{jets}$, single t quark, and multiboson production. QCD production is a similarly small contribution in the VH category as in the $t\bar{t}H$ category. The best-fit estimate for this category is $\hat{\mu} = 0.28_{-0.27}^{+0.27}$ ($0.00_{-0.26}^{+0.27}$ expected).

The ℓ_{lost} , $Z \rightarrow \text{inv}$, QCD multijet, and signal event yields following a CR only, B-only fit and a CR+SR, $S + B$ fit for each subcategory and hadronic recoil interval summed across 2016–2018 are tabulated along with observation in Tables 7.1 and 7.2, respectively. The uncertainties are inclusive of statistical and systematic contributions.




Figure 7.3: Distributions of hadronic recoil in the $t\bar{t}H$ category for the dimuon, dielectron, and dilepton CRs (upper), and the VH category for the dimuon, dielectron, and single-photon CRs (lower). The black histogram shows the total background (bkg.) prediction from a CR only, B -only fit, while the red histogram shows the yields from a CR+SR $S + B$ fit.




Figure 7.4: Distributions of hadronic recoil in the $t\bar{t}H$ (upper) and VH (lower) categories for the SR, showing the signal contributions from $t\bar{t}H$, VH , ggH , and VBF weighted by $B(H \rightarrow \text{inv}) = 0.07$. The black histogram shows the total background (bkg.) prediction from a CR only, B -only fit, while the red histogram shows the yields from a CR+SR $S+B$ fit.

Table 7.1: Total post-fit SR yields for each subcategory and recoil interval across 2016–2018, obtained from a B -only fit using CR only predictions inclusive of statistical and systematic uncertainties, and the observed data for reference.

Subcategory	Hadronic recoil	ℓ_{test}	$Z \rightarrow \text{inv}$	QCD	Total SM background	Data
ttH 1tlb	[200, 300)	251.1 ± 9.5	35.2 ± 4.1	23.1 ± 16.8	309.4 ± 19.8	288.0 ± 17.0
	[300, 400)	235.2 ± 9.5	35.7 ± 5.0	5.2 ± 4.2	276.1 ± 11.5	257.0 ± 16.0
	[400, 500)	97.5 ± 5.3	27.6 ± 4.9	0.9 ± 0.6	126.1 ± 7.2	145.0 ± 12.0
	[500, ∞)	37.5 ± 2.9	26.1 ± 4.9	0.3 ± 0.3	63.9 ± 5.7	66.0 ± 8.1
	[200, 300)	312.5 ± 12.0	19.0 ± 2.2	10.9 ± 8.6	342.4 ± 14.9	298.0 ± 17.3
ttH 1t2b	[300, 400)	265.9 ± 10.7	20.2 ± 2.7	2.5 ± 1.7	288.6 ± 11.2	299.0 ± 17.3
	[400, 500)	93.6 ± 5.1	15.4 ± 2.6	0.4 ± 0.3	109.5 ± 5.7	136.0 ± 11.7
	[500, ∞)	35.4 ± 2.9	13.8 ± 2.5	0.2 ± 0.0	49.4 ± 3.9	53.0 ± 7.3
	[200, 300)	1704.6 ± 49.9	190.7 ± 21.2	18.8 ± 16.8	1914.1 ± 56.8	1819.0 ± 42.6
	[300, 400)	395.6 ± 15.1	90.2 ± 12.7	4.3 ± 2.9	490.0 ± 19.9	486.0 ± 22.0
ttH 1W1b	[400, 500)	56.2 ± 3.9	35.8 ± 6.5	0.8 ± 0.5	92.7 ± 7.7	111.0 ± 10.5
	[500, ∞)	9.9 ± 1.3	13.9 ± 2.9	0.3 ± 0.0	24.1 ± 3.2	37.0 ± 6.1
	[200, 300)	1295.8 ± 40.7	53.1 ± 5.7	5.6 ± 3.8	1354.5 ± 41.3	1276.0 ± 35.7
	[300, 400)	266.2 ± 11.8	27.2 ± 3.8	1.3 ± 0.9	294.7 ± 12.4	298.0 ± 17.3
	[400, 500)	38.3 ± 3.3	8.1 ± 1.5	0.2 ± 0.0	46.6 ± 3.7	47.0 ± 6.9
ttH 1W2b	[500, ∞)	6.0 ± 1.0	3.7 ± 0.7	0.1 ± 0.0	9.9 ± 1.2	17.0 ± 4.1
	[200, 300)	20.2 ± 3.6	3.8 ± 0.4	0.3 ± 0.3	24.3 ± 3.6	14.0 ± 3.7
	[300, ∞)	6.3 ± 1.4	6.1 ± 0.9	0.1 ± 0.0	12.5 ± 1.7	15.0 ± 3.9
	[200, 300)	15.8 ± 2.9	3.9 ± 0.9	0.3 ± 0.0	20.0 ± 3.1	15.0 ± 3.9
	[300, ∞)	5.4 ± 1.3	3.8 ± 0.5	0.1 ± 0.0	9.3 ± 1.4	6.0 ± 2.4
ttH 5j1b	[200, 300)	5279.7 ± 114.4	1703.7 ± 82.8	99.1 ± 78.5	7082.4 ± 161.6	7207.0 ± 84.9
	[300, 400)	1135.0 ± 31.8	836.4 ± 50.0	22.5 ± 17.3	1994.0 ± 61.7	1907.0 ± 43.7
	[400, 500)	182.2 ± 9.0	267.5 ± 24.9	4.0 ± 2.8	453.6 ± 26.6	427.0 ± 20.7
	[500, ∞)	54.2 ± 3.7	146.0 ± 20.3	1.5 ± 1.0	201.7 ± 20.6	221.0 ± 14.9
	[200, 300)	1317.8 ± 47.3	350.0 ± 16.6	11.8 ± 8.5	1679.6 ± 50.9	1602.0 ± 40.0
ttH 5j2b	[300, 400)	188.7 ± 9.2	174.1 ± 10.4	2.7 ± 2.0	365.5 ± 14.1	367.0 ± 19.2
	[400, 500)	33.6 ± 3.5	53.8 ± 5.1	0.5 ± 0.3	87.9 ± 6.2	91.0 ± 9.5
	[500, ∞)	8.2 ± 1.4	24.6 ± 3.5	0.2 ± 0.0	33.0 ± 3.8	36.0 ± 6.0
	[200, 300)	3851.5 ± 87.9	805.5 ± 38.8	85.9 ± 66.3	4742.9 ± 116.7	4632.0 ± 68.1
	[300, 400)	876.0 ± 27.5	438.8 ± 26.1	19.5 ± 13.4	1334.2 ± 40.2	1371.0 ± 37.0
ttH 6j1b	[400, 500)	179.6 ± 8.5	162.8 ± 15.4	3.4 ± 2.5	345.9 ± 17.8	312.0 ± 17.7
	[500, ∞)	61.0 ± 4.0	98.2 ± 13.6	1.3 ± 1.0	160.5 ± 14.3	197.0 ± 14.0
	[200, 300)	1214.0 ± 38.7	237.2 ± 11.4	15.6 ± 12.0	1466.8 ± 42.1	1404.0 ± 37.5
	[300, 400)	237.9 ± 12.0	118.8 ± 7.1	3.6 ± 2.9	360.3 ± 14.2	341.0 ± 18.5
	[400, 500)	38.8 ± 3.8	40.9 ± 4.0	0.6 ± 0.4	80.3 ± 5.6	91.0 ± 9.5
VH 2j0b	[500, ∞)	12.9 ± 1.7	21.6 ± 3.0	0.2 ± 0.0	34.7 ± 3.5	41.0 ± 6.4
	[200, 300)	17753.9 ± 373.6	29102.3 ± 655.5	105.8 ± 68.3	46962.1 ± 757.6	47559.0 ± 218.1
	[300, 400)	2535.2 ± 69.4	5505.3 ± 155.0	16.8 ± 12.0	8057.3 ± 170.3	8106.0 ± 90.0
	[400, 500)	278.9 ± 16.1	684.1 ± 34.7	2.8 ± 1.8	965.8 ± 38.3	938.0 ± 30.6
	[500, ∞)	19.2 ± 3.1	76.9 ± 8.1	0.9 ± 0.5	97.1 ± 8.7	98.0 ± 9.9
VH 2j1b	[200, 300)	3020.1 ± 84.0	2490.4 ± 114.7	26.2 ± 24.5	5536.8 ± 144.3	5883.0 ± 76.7
	[300, 400)	360.1 ± 17.3	609.0 ± 44.1	3.6 ± 3.0	972.7 ± 47.5	949.0 ± 30.8
	[400, ∞)	36.3 ± 4.5	66.7 ± 7.3	0.6 ± 0.5	103.7 ± 8.6	120.0 ± 11.0
	[200, 300)	209.4 ± 14.0	422.3 ± 46.6	2.0 ± 1.2	633.7 ± 48.6	617.0 ± 24.8
	[300, ∞)	30.7 ± 3.5	102.6 ± 15.4	0.2 ± 0.0	133.6 ± 15.8	128.0 ± 11.3

Table 7.2: Total post-fit SR yields for each subcategory and recoil interval across 2016–2018, obtained from an $S + B$ fit using CR+SR predictions inclusive of statistical and systematic uncertainties, and the observed data for reference. The extracted S yields are weighted by the best-fit value for $\mathcal{B}(\text{H} \rightarrow \text{inv})$, $\hat{\mu} = 0.07$.

Subcategory	Hadronic recoil	ℓ_{lost}	$Z \rightarrow \text{inv}$	QCD	Total SM background	Data	Signal ($\mathcal{B}(\text{H} \rightarrow \text{inv}) = 0.07$)
ttH 1t1b	[200, 300)	246.3 ± 8.3	31.8 ± 3.6	16.4 ± 8.0	294.5 ± 12.1	288.0 ± 17.0	1.0 ± 0.8
	[300, 400)	227.3 ± 8.6	35.8 ± 4.2	3.8 ± 1.6	266.9 ± 9.7	257.0 ± 16.0	1.3 ± 1.0
	[400, 500)	98.7 ± 5.3	35.3 ± 4.8	0.7 ± 0.3	134.7 ± 7.1	145.0 ± 12.0	1.0 ± 0.8
	[500, ∞)	36.5 ± 2.6	32.6 ± 4.4	0.3 ± 0.0	69.4 ± 5.1	66.0 ± 8.1	0.9 ± 0.7
	[500, ∞)	302.8 ± 9.6	16.6 ± 1.9	7.7 ± 3.4	327.1 ± 10.4	298.0 ± 17.3	1.4 ± 1.2
ttH 1t2b	[200, 300)	264.3 ± 9.3	19.8 ± 2.4	1.9 ± 0.8	286.0 ± 9.6	299.0 ± 17.3	1.6 ± 1.3
	[300, 400)	95.8 ± 5.0	19.6 ± 2.7	0.3 ± 0.0	115.8 ± 5.7	136.0 ± 11.7	1.2 ± 0.9
	[400, 500)	35.1 ± 2.6	17.1 ± 2.3	0.1 ± 0.0	52.4 ± 3.5	53.0 ± 7.3	1.0 ± 0.8
	[500, ∞)	1659.7 ± 36.1	176.7 ± 18.9	13.7 ± 6.2	1850.0 ± 41.2	1819.0 ± 42.6	5.7 ± 4.0
	[500, ∞)	386.2 ± 12.7	93.1 ± 10.8	3.2 ± 1.4	482.5 ± 16.7	486.0 ± 22.0	2.9 ± 1.9
ttH 1W1b	[200, 300)	55.5 ± 3.8	46.7 ± 6.4	0.6 ± 0.3	102.8 ± 7.4	111.0 ± 10.5	0.9 ± 0.6
	[300, 400)	10.0 ± 1.2	18.8 ± 3.0	0.2 ± 0.0	29.1 ± 3.2	37.0 ± 6.1	0.4 ± 0.3
	[400, 500)	1256.2 ± 28.7	47.7 ± 5.0	4.0 ± 1.8	1307.9 ± 29.2	1276.0 ± 35.7	3.9 ± 3.2
	[500, ∞)	261.8 ± 9.6	27.3 ± 3.1	0.9 ± 0.4	290.0 ± 10.1	298.0 ± 17.3	1.9 ± 1.6
	[500, ∞)	36.9 ± 3.0	10.1 ± 1.4	0.2 ± 0.0	47.2 ± 3.3	47.0 ± 6.9	0.6 ± 0.4
ttH 1W2b	[200, 300)	6.4 ± 1.0	4.7 ± 0.7	0.1 ± 0.0	11.1 ± 1.2	17.0 ± 4.1	0.2 ± 0.0
	[300, 400)	16.6 ± 2.4	3.5 ± 0.4	0.2 ± 0.0	20.3 ± 2.5	14.0 ± 3.7	0.5 ± 0.3
	[400, 500)	6.5 ± 1.4	6.2 ± 0.8	0.1 ± 0.0	12.7 ± 1.6	15.0 ± 3.9	0.5 ± 0.4
	[500, ∞)	14.4 ± 2.4	3.3 ± 0.7	0.2 ± 0.0	17.9 ± 2.5	15.0 ± 3.9	0.4 ± 0.3
	[500, ∞)	4.8 ± 1.2	3.7 ± 0.4	0.1 ± 0.0	8.5 ± 1.3	6.0 ± 2.4	0.5 ± 0.4
ttH 2Boosted1b	[200, 300)	5247.6 ± 78.2	1751.5 ± 72.2	104.9 ± 66.7	7104.1 ± 125.6	7207.0 ± 84.9	14.4 ± 7.7
	[300, 400)	1101.4 ± 24.5	829.5 ± 34.2	21.8 ± 11.5	1952.6 ± 43.6	1907.0 ± 43.7	7.4 ± 3.8
	[400, 500)	176.7 ± 7.6	255.4 ± 14.4	3.9 ± 2.3	436.1 ± 16.4	427.0 ± 20.7	2.7 ± 1.4
	[500, ∞)	52.8 ± 3.1	170.1 ± 11.0	1.5 ± 0.8	224.4 ± 11.5	221.0 ± 14.9	1.5 ± 0.8
	[500, ∞)	1263.4 ± 29.9	355.1 ± 14.5	11.6 ± 5.8	1630.1 ± 33.8	1602.0 ± 40.0	6.3 ± 4.2
ttH 2Boosted2b	[200, 300)	186.3 ± 8.0	172.4 ± 7.8	2.7 ± 1.4	361.4 ± 11.2	367.0 ± 19.2	2.9 ± 1.8
	[300, 400)	33.6 ± 3.1	51.6 ± 3.2	0.5 ± 0.3	85.7 ± 4.4	91.0 ± 9.5	0.9 ± 0.5
	[400, 500)	7.8 ± 1.3	28.4 ± 2.1	0.2 ± 0.0	36.5 ± 2.5	36.0 ± 6.0	0.5 ± 0.3
	[500, ∞)	3752.9 ± 57.9	821.3 ± 34.1	85.2 ± 47.0	4659.4 ± 82.0	4632.0 ± 68.1	12.3 ± 8.1
	[500, ∞)	867.8 ± 20.3	438.8 ± 18.5	20.1 ± 10.6	1326.7 ± 29.5	1371.0 ± 37.0	6.7 ± 4.0
ttH 6j1b	[200, 300)	170.6 ± 7.2	154.9 ± 8.9	3.4 ± 1.7	328.9 ± 11.6	312.0 ± 17.7	2.4 ± 1.4
	[300, 400)	61.0 ± 3.8	115.4 ± 7.7	1.3 ± 0.8	177.8 ± 8.6	197.0 ± 14.0	1.6 ± 0.8
	[400, 500)	1174.2 ± 26.7	238.1 ± 10.0	15.5 ± 8.9	1427.7 ± 29.9	1404.0 ± 37.5	7.8 ± 6.1
	[500, ∞)	230.0 ± 9.1	116.3 ± 5.3	3.6 ± 1.9	349.9 ± 10.7	341.0 ± 18.5	3.8 ± 2.9
	[500, ∞)	39.9 ± 3.5	38.8 ± 2.4	0.6 ± 0.3	79.3 ± 4.3	91.0 ± 9.5	1.4 ± 1.0
ttH 6j2b	[200, 300)	12.5 ± 1.6	25.0 ± 1.8	0.2 ± 0.0	37.8 ± 2.4	41.0 ± 6.4	0.7 ± 0.4
	[300, 400)	17794.5 ± 282.9	29422.1 ± 391.1	103.2 ± 55.3	47319.8 ± 485.9	47559.0 ± 218.1	185.6 ± 92.5
	[400, 500)	2535.3 ± 62.1	5484.6 ± 92.6	16.5 ± 8.7	8036.4 ± 111.8	8106.0 ± 90.0	44.3 ± 23.0
	[500, ∞)	277.3 ± 13.8	659.3 ± 23.7	2.7 ± 1.4	939.4 ± 27.5	938.0 ± 30.6	6.6 ± 3.4
	[500, ∞)	19.2 ± 3.0	75.1 ± 6.0	0.9 ± 0.4	95.2 ± 6.7	98.0 ± 9.9	0.6 ± 0.3
VH 2j1b	[200, 300)	3051.7 ± 63.0	2707.6 ± 84.3	29.4 ± 22.6	5788.7 ± 107.6	5883.0 ± 76.7	20.3 ± 10.0
	[300, 400)	348.7 ± 14.0	605.2 ± 27.2	3.9 ± 2.6	957.8 ± 30.8	949.0 ± 30.8	5.2 ± 2.8
	[400, 500)	37.6 ± 4.2	72.5 ± 6.2	0.7 ± 0.5	110.8 ± 7.5	120.0 ± 11.0	0.7 ± 0.4
	[500, ∞)	201.4 ± 12.3	407.8 ± 26.0	2.0 ± 1.6	611.2 ± 28.8	617.0 ± 24.8	10.8 ± 7.9
	[500, ∞)	29.8 ± 3.3	98.6 ± 9.7	0.2 ± 0.0	128.6 ± 10.3	128.0 ± 11.3	3.5 ± 2.5

No significant excess in the data is reported that can be attributed to signal contributions alone, and therefore the null hypothesis representing contributions from SM background alone cannot be rejected in favour of a discovery. Therefore, exclusion limits are computed on $\mathcal{B}(H \rightarrow \text{inv})$ as defined in Eq. 6.12, adhering to the CL_S criterion [140, 141] under the asymptotic approximation [131], and are found to be 0.43 (0.52 expected) and 0.74 (0.53 expected) at 95% CL for the $t\bar{t}H$ and VH categories, respectively, with a combined upper limit of 0.54 (0.39 expected). These results are shown in Fig. 7.5 together with the observed and expected profile likelihood distribution, where the corresponding 68 and 95% CL intervals are extracted following the procedure outlined in Ref. [143] and Ref. [144]. The observed limits obtained are compatible with the SM background expectation. The breakdown of 95% CL upper limits by year and category is presented in Table 7.3. The expected limits improve as anticipated with the increasing luminosity associated with each data-taking year for the VH category.




Figure 7.5: Left: Observed and expected limits at 95% CL for the $t\bar{t}H$ and VH categories using 2016–2018 data. Right: The profile likelihood scan corresponding to observed and expected (where $\mathcal{B}(H \rightarrow \text{inv}) = 0$) limits in the fit to the $t\bar{t}H$ and VH categories.

A GF test is performed to determine the compatibility of the model with the data. The expected distribution of χ^2_{sat} is produced by generating 1000 toy experiments, with the p -value extracted as the fraction of those larger than the test statistic obtained in the fit to data, as explained in Appendix B.1. In brief, to generate these toy experiments, NPs are fixed to their best-fit values from a CR+SR fit in light of the data, while the constrained component of these NPs are randomised in the evaluation of the likelihood fit. A p -value of 0.2% is observed, below what is expected for a well-

Table 7.3: The breakdown of the observed and expected (exp.) 95% CL upper limits on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ in each year and category.

Data	$t\bar{t}\text{H}$	VH	Combined
2016	0.77 (0.78 exp.)	0.79 (1.04 exp.)	0.52 (0.62 exp.)
2017	0.93 (0.86 exp.)	1.67 (0.94 exp.)	1.04 (0.64 exp.)
2018	0.59 (0.87 exp.)	0.90 (0.77 exp.)	0.54 (0.60 exp.)
2016–2018	0.43 (0.52 exp.)	0.74 (0.53 exp.)	0.54 (0.39 exp.)

modelled analysis. Statistical constraints in certain subcategory and hadronic recoil intervals can lead to a poor GF as the expected probability distribution may not be well fitted by a χ^2_{sat} distribution. An additional study was performed by removing a subset of subcategory and hadronic recoil intervals in the single-lepton CRs, which share a RP. Where the single-muon CR and single-electron CR data fluctuates in opposite directions relative to the simulated event yields, as visualised in comparing the ratio plots of Fig. 7.1 to Fig. 7.2, the seven bins with the greatest disagreement between data and simulation were removed. This resulted in an improved p -value of 6% across the combined $t\bar{t}\text{H}$ and VH categories with negligible effect on the sensitivity. Corroborating this statement is the higher p -values obtained from GF measurements for 2016 and 2018 data sets, presented in Table 7.4. Taking this into consideration, the low p -value across all years is attributed to statistical fluctuations in $\approx 3\%$ of all subcategories and hadronic recoil intervals studied, with a visualisation of the GF distributions provided in Appendix C.

This is the first limit on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ in the $t\bar{t}\text{H}$ channel in a fully hadronic final state reported by CMS using post-2016 data. A reinterpretation of a stop SUSY search using 2016 data corresponding to 35.9 fb^{-1} analyses the $t\bar{t}\text{H}$ channel with up to two leptons, with the zero-lepton channel 95% CL upper limit measured at $\mathcal{B}(\text{H} \rightarrow \text{inv}) = 0.85$ (0.73 expected) [145], consistent with the upper limit $\mathcal{B}(\text{H} \rightarrow \text{inv}) = 0.77$ (0.78 expected) measured here. Differences between the two values stem from the additional systematic uncertainties considered for this analysis, including those from the p_{T}^{t} reweighting, and differences resulting from the different NP correlation schemes in the fit, as performed in this analysis across years, and in Ref. [145] across lepton multiplicity channels. For direct comparison, a search by the ATLAS experiment in the same channel [146] reports a 95% CL upper limit on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ of 0.95 (0.52 expected), using full Run

Table 7.4: The breakdown of p -values obtained from GF studies in each year and category.

Data	$t\bar{t}H$	VH	Combined
2016	55.3%	5.6%	10.7%
2017	6.1%	31.5%	3.5%
2018	33.5%	30.4%	20.0%
2016–2018	7.4%	1.6%	0.2%

2 data. The boosted $t\bar{t}H$ subcategories have a greater signal purity than those of the resolved $t\bar{t}H$ class, evidenced by the S/B measurements in Fig. 5.11.

No exclusion limits on $\mathcal{B}(H \rightarrow \text{inv})$ in the resolved hadronic VH channel have been reported by the ATLAS or CMS experiments using post-2015 data. A 95% CL upper limit of $\mathcal{B}(H \rightarrow \text{inv}) = 0.37$ (0.31 expected) is found in the boosted VH channel of Ref. [11], using 2016–2018 data corresponding to 137 fb^{-1} . The VH 2j0b and VH 2j2b subcategories provide the greatest sensitivity to $H \rightarrow \text{inv}$, given these are pure in the $V(q\bar{q})H$ and $Z(b\bar{b})H$ background processes that they target, respectively, while the subdominant VH 2j1b subcategory relies on events with a falsely-tagged or untagged b jet (under the DEEPCSV algorithm). This subcategory nevertheless is retained in the analysis as it improves the sensitivity of the overall limit on $\mathcal{B}(H \rightarrow \text{inv})$ as measured in the VH category.

The systematic uncertainty with the largest impact on the $\mathcal{B}(H \rightarrow \text{inv})$ measurement is associated with the JES, while the statistical uncertainty contributes significantly to the overall uncertainty on $\mathcal{B}(H \rightarrow \text{inv})$. The breakdown of the impacts into uncertainty groups are presented in Table 7.5, together with the expectation values. The likelihood scans performed to obtain these impacts are provided as supporting material in Appendix C, which are also presented with the largest pulls on systematic and statistical uncertainties in the fit.

In summary, falling short of a discovery, upper limits on $\mathcal{B}(H \rightarrow \text{inv})$ are established in hadronic final states of the $t\bar{t}H$ and resolved VH channels at the 95% CL. These show a stronger-than-expected limit in the $t\bar{t}H$ channel, of 0.43 (0.52 expected), and a weaker-than-expected limit in the VH channel of 0.74 (0.53 expected). Their combined limit of 0.54 (0.39) shows an excess of data of 1 standard deviation. The model yields a p -value of 0.2%, although this is considered to be biased by a small number

Table 7.5: The observed and expected impacts on $\mathcal{B}(\text{H} \rightarrow \text{inv})$ for different groups of uncertainties, where the expected results are produced with $\mathcal{B}(\text{H} \rightarrow \text{inv}) = 0$.

Uncertainty group	Impact on $\mathcal{B}(\text{H} \rightarrow \text{inv})$	
	Observed	Expected
Jet energy calibration	± 0.11	± 0.11
Lepton veto	± 0.05	$+0.05$ -0.04
Lepton/photon identification	± 0.06	± 0.06
Theory	$+0.07$ -0.06	$+0.06$ -0.05
Integrated luminosity/pileup	± 0.02	$+0.02$ -0.03
QCD prediction	± 0.02	± 0.02
Boosted object/b-jet tagging	± 0.02	± 0.02
Triggers	± 0.04	± 0.03
Stat. uncertainty of simulation	± 0.08	± 0.08
Stat. uncertainty in data	± 0.10	± 0.10

of hadronic recoil intervals where the correlation single-muon and single-electron CR data-to-simulation agreement is poorest, accounting for $\mathcal{O}(1\%)$ of all subcategories and hadronic recoil intervals. These are the first $\text{H} \rightarrow \text{inv}$ limits measured in the hadronic $t\bar{t}\text{H}$ channel using CMS data from 2017 and 2018, and the first in the resolved VH channel using post-2015 data by either of the ATLAS or CMS Collaborations.

Chapter 8

Combination with other Higgs boson production channels

Invisible Higgs boson decays have been analysed in multiple channels dating from the start of LHC operations onward. Each search has a bespoke set of search regions, kinematic properties, particle identification and tagging requirements, and systematic uncertainties. A combined result across Higgs boson production modes requires these to be well understood. This includes synchronising trigger requirements, mitigating overlapping phase space, and handling the correlation between the sources of systematic uncertainty that arise from each analysis. A combination of the results of this analysis, analyses also covering the years 2016–2018, and earlier published CMS combination results using Run 1 and 2015 data [16] collected at $\sqrt{s} = 7, 8$, and 13 TeV corresponding to integrated luminosities of 4.9, 19.7 and 140 fb⁻¹, respectively, is performed. Run results from 2015 correspond to 2.3 fb⁻¹ data collected at $\sqrt{s} = 13$ TeV, and are described separately as the analyses using this data were combined historically with Run 1 data, and use a separate fit modelling and uncertainty correlation scheme to later searches. A combined likelihood fit, accounting for appropriate event overlap reduction and the correlation of systematic uncertainties across regions, is outlined. The Higgs boson production modes VBF, VH, $t\bar{t}H$, and ggH from various CMS analyses are included, with their respective data sets detailed in Table 8.1. The Run 1 and Run 2 analysis combines results from the VBF [8] channel, the $Z(\ell^+\ell^-)H$ [15] channel, the combined ggH and boosted VH channels [11], the semi- and dileptonic $t\bar{t}H$ channels [12–14] channels, and the Run 1 and 2015 combinations that include VBF, $Z(\ell^+\ell^-)H$, $Z(b\bar{b})H$, $V(jj)H$, and ggH searches [16, 147, 148]. The parameters of the individual likelihood functions are treated as independent for each analysis, unless explicitly specified otherwise. As several analyses are discussed, the $t\bar{t}H$ and resolved

VH search described previously is labelled as the $t\bar{t}H$ /resolved VH analysis.

Table 8.1: The data sets used in the combination, with the respective integrated luminosities for each production mode across Run 1 and Run 2.

Analysis tag	Production mode	Integrated luminosity (fb^{-1})		
		7 TeV	8 TeV	13 TeV
VBF-tagged	VBF	—	19.2 [147]	140 [16][8]
VH-tagged	$Z(\ell^+\ell^-)H$	4.9 [147]	19.7 [147]	140 [16][15]
	$Z(b\bar{b})H$	—	18.9 [147]	—
	$V(jj)H$	—	19.7 [148]	140 [16][10]
	Boosted VH	—	—	138 [11]
$t\bar{t}H$ -tagged	$t\bar{t}H$ (hadronic)	—	—	138 [10]
	$t\bar{t}H$ (leptonic)	—	—	138 [12–14]
ggH-tagged	ggH	—	19.7 [148]	140 [16][11]

The VBF channel provides the most stringent limit on $\mathcal{B}(H \rightarrow \text{inv})$, which produces a 95% CL upper limit of $\mathcal{B}(H \rightarrow \text{inv}) = 0.18$ (0.10 expected) across Run 1 and Run 2. In this channel, the signal is identified with a large hadronic recoil and two high- p_T forward jets located in opposite η hemispheres with small azimuthal separation. The VBF signal is enhanced by requirements on the dijet invariant mass, and their angular separation, which suppress DY contributions in the $Z \rightarrow \text{inv}$ background. Single-muon, single-electron, dimuon, di-electron, and single-photon CRs are used to measure the dominant $V + \text{jets}$ background in the SR by constraining the yield fractions $Z(\ell^+\ell^-) + \text{jets}/W(\ell\nu) + \text{jets}$ and $\gamma + \text{jets}/Z(\ell^+\ell^-) + \text{jets}$. This accounts for the low statistical precision of the dilepton CRs, where the $Z + \text{jets}$ production cross section is ≈ 100 times smaller than that of $\gamma + \text{jets}$. The analysis is described in Ref. [8], which defines a dominant *missing momentum triggered region* (MTR) and a subdominant *VBF jets triggered region* (VTR) category, where events are selected for the former preferentially, and orthogonality between the two is ensured by requiring a large hadronic recoil > 250 GeV in the MTR, and $160 < \text{recoil} < 250$ GeV in the VTR. The categories are defined for 2017 and 2018 data sets, and use the same statistical approach to combine with the 2012, 2015, and 2016 data sets as that which is taken here. The dominant uncertainties in the analysis are the theoretical uncertainties in

the yield fractions between the $Z + \text{jets}$, $W + \text{jets}$, and $\gamma + \text{jets}$ processes. While the ratios can lead to partial cancellations in the theoretical uncertainties, no dedicated theory treatment exists for the highly-forward p_T nature of the VBF phase space, hence manual, conservative calculations are performed.

The next most sensitive channel to $\mathcal{B}(H \rightarrow \text{inv})$ is the associated ZH production, or *MonoZ*, channel with a fully leptonic final state, $Z(\ell^+\ell^-)$. As the analysis uses lepton triggers, the p_T^{miss} threshold is lower than the hadronic search analogues, at 100 GeV. The dominant background contributions are from diboson processes such as ZZ and WZ production, where one Z boson decays to electrons or muons, and the other decays to at least one neutrino. To constrain the diboson background, CRs with one or two additional charged leptons are defined. The leading systematic uncertainties relate to the theoretical uncertainties in the absolute normalisation and transfer factors between WZ and ZZ contributions in the SR and CRs, and the experimental uncertainty in the reconstruction of electrons and muons. The search is also statistically limited by the data sets produced in the Run 2 era. The Run 2 MonoZ limit is described in Ref. [15], where the likelihood function is entirely independent of the data-taking year. A 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.29 (0.25 expected) is reported.

The MonoJet/MonoV analysis is similarly sensitive to $\mathcal{B}(H \rightarrow \text{inv})$, and is composed of a search for ggH and VH production, where the latter identifies boosted V boson decays with merged hadronic decay products, complementarily to the resolved $V(\text{jj})H$ topology in the $t\bar{t}H$ /resolved VH analysis. Events are selected in fully hadronic final states, requiring a well-identified leading jet and large hadronic recoil > 250 GeV. Events are categorised according to substructure properties of the leading jet, firstly by requiring a jet $p_T > 250$ GeV and dividing boosted VH events into low or high purity according to the DEEPAK8 score, referred to as the loose and tight MonoV category, respectively, followed by selecting ggH events for the MonoJet category with $p_{T,1}^j > 100$ GeV having failed the boosted VH category requirements. The categories also apply a veto on b-tagged jets, and are split according to 2017 and 2018 data sets, with the 2016 data set analysed previously [149]. The same approach to background estimation as the VBF analysis uses is adopted, where again the dominant background is from $V + \text{jets}$ processes. The yields for $W + \text{jets}$ processes in the SR, $\gamma + \text{jets}$ in the single-photon CR, and $DY + \text{jets}$ in the dilepton CRs are defined as transfer factors to normalise the $Z \rightarrow \text{inv}$ contribution in the SR. Similarly, the $W(\ell\nu) + \text{jets}$ backgrounds in the single-lepton CRs are defined as transfer factors relative to the ℓ_{lost} contribution in the SR. Subdominant backgrounds such as t quark and diboson processes are estimated directly from simulation, while a dedicated estimation is used for QCD multijet

contributions. Consequently, the largest uncertainties in the MonoJet/MonoV search stem from the WZ normalisation uncertainty, mainly affecting the MonoV channel where t quark and diboson processes contribute to $\approx 25\%$ of the overall background. A 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.28 (0.27 expected) is reported in Ref. [15] using 2016–2018 data sets.

Finally, $t\bar{t}H$ analyses with leptonic final states are included, which require a reinterpretation of the supersymmetry (SUSY) searches in the semi- and dileptonic $t\bar{t}$ decay channels in Refs. [12, 13] in the context of the $t\bar{t} + \text{DM}$ model studied in Ref. [14]. The reinterpretation is valid for p_T^{miss} searches, where the SUSY study is intended to be sensitive to the production of SUSY partners of the t quark, stop quarks, which decay to a t quark and the lightest SUSY particle (LSP). The final state will contain t quark decay products and a large p_T^{miss} corresponding to the undetected LSP, which is kinematically similar to a $H \rightarrow \text{inv}$ signature. The semi- and dileptonic $t\bar{t}H$ final state searches both use 2016–2018 data sets, and produce a 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.32 (0.38 expected), and is the only analysis to report a slightly stronger than expected limit. The limit set by the $t\bar{t}H$ /resolved VH analysis is therefore similar in sensitivity to the leptonic $t\bar{t}H$ channel, at 0.54 (0.39 expected).

8.1 Selection overlap

A statistical combination between the different analysis channels necessitates removal of any significant overlap in selected events if the results are to be statistically valid. Analyses with hadronic final states have partially overlapping SRs amongst themselves. The overlap is systematically verified between each analysis by generating the list of events in each region and identifying the fraction of events that are common to each relative to the total number of selected events. Aside from the $t\bar{t}H$ /resolved VH analysis, the following selections are made to reduce the overlap in each analysis after their respective results have been published, and hence the results from the fit, including exclusion limits, are expected to differ following the combination. However, the selections are chosen such that these changes are small, with changes in the exclusion limits by analysis $\mathcal{O}(0.1\text{--}1\%)$.

The most significant overlap is observed between the MonoJet and VBF channels. As the MonoJet category requires the presence of a hard leading jet with $p_{T,1}^j > 100 \text{ GeV}$ and a large hadronic recoil $> 250 \text{ GeV}$, which is also characteristic of the VBF selection but without the more refined jet selection, a subset of the MonoJet events is selected in the VBF analysis. The relative overlap is found to be

similar between each SR and each flavour of CR. As the VBF analysis is dominant in sensitivity to $\mathcal{B}(H \rightarrow \text{inv})$, events that would otherwise pass the VBF selection are removed from the MonoJet selection. Events containing at least two jets that pass the tight jet identification criteria, with the leading and subleading jet requiring $p_{T,1}^j > 80$ GeV and $p_{T,2}^j > 40$ GeV, respectively, and neither being highly forward such that $|\eta_{1,2}^j| < 4.7$, are considered for rejection. To invert the VBF selection, discrimination is performed according to the dijet system requirements, where events are rejected if the jets have a combined mass $m_{jj} > 200$ GeV, are located in opposite hemispheres, $\eta_1^j \cdot \eta_2^j < 0$, and are well separated in η and not in ϕ , $|\Delta\eta(\vec{p}_{T,1}^j, \vec{p}_{T,2}^j)| > 1.0$ and $|\Delta\phi(\vec{p}_{T,1}^j, \vec{p}_{T,2}^j)| < 1.5$. The overlap mitigation has a greater effect at higher hadronic recoil given events are more likely to have additional high- p_T jets than the leading one. An average of between 10% and 20% of background events are removed depending on the hadronic recoil, as demonstrated before and after applying the overlap filter in Fig. 8.1 in the MonoJet category (upper), loose MonoV category (middle), and tight MonoV category (lower). The plots for this overlap study are produced by colleagues on the combination. The ggH signal is more strongly affected by this selection, being reduced by between 13 and 24%, than background processes, likely due to the initial state of the gluon, where the presence of additional jets is enhanced. The VBF signal contribution is reduced the most among all processes in the MonoJet category, by up to 40%, although the effect of this is much less pronounced considering the far smaller production cross section compared to ggH. The effects of the removal are generally weaker in the MonoV categories, although as MonoV events require a boosted object tag according to the DEEPAK8 algorithm, the contamination by VBF-like topologies is inherently minimal. The expected 95% CL upper limit for either of 2017 or 2018 worsens in the MonoJet (MonoV) category by approximately 15 (2.5)%, although as the overall sensitivity is dominated by the MonoV channel, the exclusion limit changes negligibly. The overlap removal in the 2016 data set is introduced in Ref. [150] and included in Ref. [11].

The $t\bar{t}H$ /resolved VH analysis enforces orthogonality with VBF-tagged events by inverting the kinematic selection and removing events containing forward leading and subleading AK4 jets with $|\eta^j| > 2.4$, as detailed in Sec. 5.2 and identically for the MonoJet category above. A non-negligible level of overlap is measured between the SR and CRs of the $t\bar{t}H$ /resolved VH analysis and the ggH/boosted VH channels of the MonoJet/MonoV analysis. Observations show that 1% of ggH events account for between 30 and 40% of events in the VH 2j0b subcategory, however as any correction would change the ggH yields $\mathcal{O}(1\%)$, no correction is applied to the resolved VH




Figure 8.1: The effect of reducing VBF signal overlap within the MonoJet (upper), loose MonoV (middle), and tight MonoV (lower) categories for each data-taking year. The relative change in yields for each signal process due to the overlap mitigation selection is presented as a function of the hadronic recoil, with the average change given in parentheses. This study was performed by colleagues on the combination.

category. In the loose MonoV category, 10% of events account for 20% of events in the same subcategory, due to the difference in choice of the DEEPAK8 V boson tagging WP, where a loose WP is used the loose MonoV category, and a medium veto WP is employed in the resolved VH channel. To mitigate this, events in the loose MonoV category are removed if exactly two AK4 jets are present with a dijet mass $65 < m_{jj} < 120$ GeV. The effect of this selection is presented in Table 8.2, where the before-and-after fraction of overlap of loose MonoV events in the VH 2j0b subcategory is given first, and vice versa second, for each region of the respective analyses. The resulting overlap of events in the loose MonoV category is rendered to $\mathcal{O}(1\%)$, and no further selection is necessary for the resolved VH channel. The sensitivity of the loose MonoV category to $\mathcal{B}(H \rightarrow \text{inv})$ changes negligibly, while the overall MonoV result is dominated by the tight MonoV category.

Table 8.2: The effect of overlap mitigation between the corresponding SR and CRs in the loose MonoV category and the VH 2j0b subcategory. The values represent the fraction of overlap before and after the veto is applied (loose MonoV events in VH 2j0b/vice versa).

Region	Change in overlap (Loose MonoV / VH 2j0b)
SR	0.13 / 0.15 $\rightarrow$ 0.04 / 0.03
Single-lepton CR	0.06 / 0.17 $\rightarrow$ 0.02 / 0.04
Dilepton CR	0.12 / 0.19 $\rightarrow$ 0.03 / 0.04
Single-photon CR	0.12 / 0.25 $\rightarrow$ 0.03 / 0.04

8.2 Systematic uncertainties

A correlation scheme for the systematic uncertainties between analyses is designed such that, generally, uncertainties with the same values, such as on the integrated luminosity, or those associated with the same sources, such as common trigger paths or centrally-provided JECs, are correlated. All systematic uncertainties between data sets with different centre-of-mass energies are uncorrelated. An overview of the uncertainty correlation scheme adopted for the analyses using data recorded from 2016 onward is presented in Table 8.4, excluding the signal theory uncertainties which depend on production mode and initial state rather than analysis channel. The uncertainty in

the integrated luminosity is calculated independently for Run 1 analyses. For analyses pre-dating 2016, the correlation scheme established in Ref. [16] is used.

8.2.1 Theoretical uncertainties

The most important theoretical uncertainties are those associated with the modelling of leading background processes in each analysis. In the hadronic VBF, MonoJet/MonoV, and $t\bar{t}H$ /resolved VH analyses, the leading background processes are from $V + \text{jets}$ and $t\bar{t} + \text{jets}$ events. Each analysis performs a maximum-likelihood fit in which the SR and CR $V + \text{jets}$ yields are estimated using RPs. In the case of the VBF and MonoJet/MonoV analyses, these RPs propagate the background theory uncertainties. In the VBF channel, focused in a highly-forward p_T phase space, there is no dedicated theory treatment of the $Z(\ell^+\ell^-) + \text{jets}/W(\ell\nu) + \text{jets}$ and $\gamma + \text{jets}/Z(\ell^+\ell^-) + \text{jets}$ yield fractions. Therefore, the uncertainty in these ratios is calculated via conservative manual calculations, where it is assumed there is no strong cancellation between the different processes, which results in a typical uncertainty of $\mathcal{O}(10\%)$ [8]. In the MonoJet/MonoV analysis, the uncertainties in the yield fractions are estimated following the prescription described in Ref. [110], which employs a splitting scheme for theory nuisances that accounts for realistic uncertainty cancellations between the component processes. This leads to a typical uncertainty of a few percent. The $t\bar{t}H$ /resolved VH analysis applies the same corrections and theory uncertainties to the background samples, however the estimation of yields between SR and CRs for these processes is calculated using unconstrained RPs during the fit. Therefore, given the differences in region definitions, assumptions in the uncertainty modelling, and estimation methods, the background theory uncertainties are completely uncorrelated between the analyses.

The leading background contributions in the MonoZ analysis are from WZ and ZZ production, which are estimated in tri- and quad-lepton CRs. Background $V + \text{jets}$ processes are normalised using data alone, with shape uncertainties providing additional flexibility to the p_T^{miss} spectra. The diboson processes contributing here are also applicable to the VBF and MonoJet/MonoV channels, although are limited to the overall normalisation of VV and VVV processes and not the p_T^{miss} spectrum, unlike the MonoZ analysis. Given the difference in dependence of this uncertainty, independent NPs are assigned to the MonoZ analysis and the hadronic analyses. No correlation is introduced between the hadronic analyses either, for the same reasons discussed for the $V + \text{jets}$ uncertainties, and likewise for the leptonic $t\bar{t}H$ channel.

The PDF uncertainties on these background processes are minor contributions in all analyses. The extent of cancellation of PDF uncertainties between various processes

in the different analysis regions varies between channels, and therefore the residual PDF uncertainty is uncorrelated between each analysis.

The uncertainties in signal modelling address higher-order QCD corrections and PDFs. The former are estimated by varying the QCD renormalisation and factorisation scales, with the resulting cross section uncertainty, determined for each signal production mode in VBF, WH, ZH, $t\bar{t}H$, ggH, and ggZH, applied independently per channel. The PDF uncertainties are dependent on the initial state of each signal production mode, and are correlated into $q\bar{q}H$ and ggH states. The values for the QCD scale and PDF uncertainties are obtained from Ref. [135], and presented in Table 8.3, and are log-normal constrained. In some channels, additional uncertainty contributions relating to signal acceptance modelling are considered, which are treated as uncorrelated between analyses.

Table 8.3: The log-normal-constrained uncertainties in higher-order QCD corrections, applied independently to each process, and the PDF uncertainties associated with each initial state $q\bar{q}H$ and ggH, correlated across the corresponding final states.

Source	Process	Magnitude
QCD scale uncorrelated by process	VBF	0.997 / 1.004
	WH	0.993 / 1.005
	ZH	0.969 / 1.038
	$t\bar{t}H$	0.918 / 1.058
	ggH	0.933 / 1.046
	ggZH	0.811 / 1.251
Correlated $q\bar{q}H$ PDF	VBF	1.016
	WH	1.021
	ZH	1.018
Correlated ggH PDF	$t\bar{t}H$	1.036
	ggH	1.032
	ggZH	1.024

8.2.2 Experimental uncertainties

The largest experimental uncertainties are typically related to the identification, isolation, and reconstruction criteria of physics objects, namely muons, electrons, and

photons. In the MonoZ and leptonic $t\bar{t}H$ analyses, muons and electrons are identified for the SR within a lower p_T spectrum than the leptonic CRs in the hadronic analyses, and so are uncorrelated between analyses with hadronic and leptonic final states. Additionally, the lepton identification and reconstruction criteria, as well as the parametrisation of their associated uncertainties, differ sufficiently between the MonoZ and leptonic $t\bar{t}H$ searches that the uncertainties are not correlated. Within the hadronic analyses, a p_T^{miss} threshold of 250 GeV is applied to all regions in the VBF and MonoJet/MonoV channels, meaning that the lepton and photon p_T spectra are similar. However, VBF signal extraction is performed in the m_{jj} distribution, which is only weakly correlated with the hadronic recoil spectrum, and so is mostly sensitive to events just above the recoil threshold. Signal extraction in the MonoJet/MonoV analysis is performed in the hadronic recoil distribution directly, and given the average lepton and photon p_T is dependent on the hadronic recoil interval, the dependence extends to the uncertainty in lepton and photon identification. Therefore, the systematic uncertainties in lepton and photon identification, isolation, and reconstruction are not correlated between the VBF and MonoJet/MonoV analyses. In the $t\bar{t}H$ /resolved VH search, signal extraction is performed across the hadronic recoil distribution from a lower 200 GeV threshold, and so the corresponding systematic uncertainties are not correlated with the other analyses.

The same trigger paths are used in the VBF, MonoJet/MonoV, and $t\bar{t}H$ /resolved VH analyses, and so the uncertainties associated with the PF p_T^{miss} , electron, and photon triggers are correlated between analyses for each trigger object. The effect of these uncertainties is dominated by events just above their respective thresholds, which has little dependence on the average PF object p_T scale, and so these uncertainties are log-normal constrained. In the MonoZ and leptonic $t\bar{t}H$ analyses, the trigger criteria are different, namely that a muon trigger is used in place of a PF p_T^{miss} HLT path. Both analyses use single and dilepton triggers, and so the trigger efficiency selection uncertainties are correlated, but are uncorrelated with the uncertainties of the hadronic analyses.

The uncertainties in boosted object tagging using the DEEPAK8 algorithm are only applicable to the MonoJet/MonoV, $t\bar{t}H$ /resolved VH, and leptonic $t\bar{t}H$ searches. For V boson tagging in the MonoJet/MonoV and $t\bar{t}H$ /resolved VH analyses, these sources of uncertainty are not correlated across analyses given different WPs are adopted in each case. Similarly, as the WPs differ for t quark tagging in the $t\bar{t}H$ /resolved VH, and leptonic $t\bar{t}H$ analyses, the uncertainties are not correlated. In the VBF, MonoJet/MonoV, and MonoZ analyses, the uncertainty in b jet tagging is separated into

the efficiency uncertainty associated with tagging a b jet that is truly a b jet, and the uncertainty in the rate at which light jets are misidentified as a b jet. A veto is applied on the presence of b-tagged jets in these analyses, in order to reduce background contributions from processes involving t quarks. As the same b jet identification WPs are adopted in these analyses, the uncertainties are correlated. In the $t\bar{t}H$ /resolved VH analysis, a single b-tagging NP combines the effect of the efficiency selection and misidentification, and hence it is not correlated with the other analyses. There is also no b-jet veto applied in this analysis, given hadronically decaying t quarks are targeted. The leptonic $t\bar{t}H$ analysis uses different modes of b tagging, therefore is not correlated with the other analyses.

The muon, electron, and photon vetoes used in the applicable analyses do not have a significant effect on the overall uncertainty, and are therefore left decorrelated given their implementations vary and are non-trivial to combine into single NPs. All searches apply a veto on well-identified hadronically decaying tau leptons in the SR, specifically removing $W(\tau\nu)$ events in $W + \text{jets}$ samples used in the hadronic analyses, and WZ samples in the MonoZ case. Identical tau lepton identification criteria are applied in the MonoJet/MonoV and MonoZ searches, and so the tau lepton veto uncertainty is correlated, while the VBF analysis uses the DEEPTAU identification algorithm [151], so is not correlated. The tau veto uncertainty is also independent in the $t\bar{t}H$ /resolved VH and the leptonic $t\bar{t}H$ analyses, given their respective identification criteria vary again.

All analyses incur some level of calibration of the JES and JER, which in the hadronic analyses are provided by the JME POG and split into the same prototype sources of jet energy uncertainties. Consequently, all sources are correlated between the hadronic analyses, including the uncertainty components correlated and uncorrelated between the years 2016, 2017, and 2018. The JECs in the MonoZ analysis are represented by a single uncertainty source, which accounts for effects in the low p_T^{miss} region due to experimental effects, such as jet mismeasurement in DY events, as well as for SM events with real p_T^{miss} , namely $ZZ^* \rightarrow 4\nu$ events. The systematic uncertainty in the MonoZ channel is therefore independent of those measured in the hadronic analyses. Similarly, the leptonic $t\bar{t}H$ analysis applies NPs for JER and JES alone, therefore it is not correlated with the uncertainties of other analyses.

The determination of the integrated luminosity is affected by a number of sources of uncertainty, which are assumed to be correlated amongst all channels, and correlated amongst data sets irrespective of the inclusion of those from 2015. The uncertainty on the luminosity is largely insignificant, and therefore neither accounting for the correlation between the 2015 and 2016 data sets, nor the 2015 data set with other Run

2 data sets, has no measurable effect on the total uncertainty.

Table 8.4: Overview of uncertainty correlation scheme adopted for the combination of analyses using data from 2016 onward. The correlation pattern is defined by a symbol or colour per analysis. Analyses with the same letter and colour in a given row have correlated NP effects for the given uncertainty source.

Uncertainty group	Source	Analysis channel					
		MonoJet	MonoV	VBF	MonoZ	t $\bar{t}$ H/res. VH	Lep. t $\bar{t}$ H
Background theory	V / γ theory	A	A	B	—	C	D
	VV norm.	A	B	C	—	D	E
Object identification, isolation, and reconstruction	Electron	A	A	B	C	D	E
	Muon	A	A	B	C	D	E
	Photon	A	A	B	—	C	D
	V tagging	A	A	—	—	B	—
	t tagging	—	—	—	—	A	B
	b tagging	A	A	—	—	B	C
Trigger	p_T^{miss}	A	A	A	—	A	—
	Muon	—	—	—	A	—	A
	Electron	A	A	A	B	C	B
	Photon	A	A	A	—	A	—
Veto	b tag	A	A	A	A	—	—
	Hadronic tau	A	A	B	A	C	D
	Muon, electron	A	A	B	C	D	E
JECs	JER, JES	A	A	A	B	A	C
Luminosity		A	A	A	A	A	A

8.3 Limits and evaluation of results

The results from the statistical combination of all channels and data sets featured in Table 8.1 are presented, along with a breakdown according to the respective analyses,

the run period, and by final state. The GF test is performed by generating 1000 toy experiments, and a p -value of 12.1% is extracted for combined Run 1 and Run 2 data corresponding to the probability that the combined $S + B$ model is consistent with data from the CMS experiment.

All results are obtained from maximum-likelihood fits to the individual analyses or Higgs boson production channels, as well as their combination. The fit accounts for both statistical and systematic uncertainties, and therefore all results are inclusive of both effects unless otherwise stated. The exclusion limits on $\mathcal{B}(H \rightarrow \text{inv})$ are calculated relative to SM production cross sections, with a 95% CL upper limit of 0.16 (0.09 expected) obtained for 2016–2018 data, and 0.15 (0.08 expected) obtained on inclusion of Run 1 and 2015 data sets. Therefore, an excess in the data of two standard deviations is reported by the CMS combination using Run 1 and Run 2 data. The Run 1 and 2015 data sets have corresponding exclusion limits of 0.33 (0.27) and 0.43 (0.44), respectively, as shown in Fig. 8.2 (left), so have sensitivities subdominant to some more recent analyses. The sensitivity in the Run 2 result is dominated outright by the VBF analysis, which for 2016–2018 data sets yields a 95% CL upper limit of 0.18 (0.11 expected). The MonoZ and MonoJet/MonoV channels have subdominant sensitivities to $\mathcal{B}(H \rightarrow \text{inv})$, with minor contributions in the leptonic $t\bar{t}H$ and the $t\bar{t}H$ /resolved VH analyses, demonstrated in Fig. 8.2 (right). The values obtained are equivalent to the published results quoted by the respective analyses within the precision of the combination, reflecting the negligible effect the overlap removal has on the overall results.

The combination contains all final states of the VBF, VH, $t\bar{t}H$, and ggH channels, therefore enables a breakdown into each Higgs boson production mode for the first time using Run 1 and Run 2 data. The VBF, VH, $t\bar{t}H$, and ggH channels are determined according to the category defined in each analysis, which themselves are composed of varying levels of signal cross contamination as is input to each fit, as discussed in Appendix D. The 95% CL upper limits are presented in Fig. 8.3 (left) for all Run 1 and Run 2 data sets. The final combination represents an improvement in sensitivity of approximately 20% relative to the most sensitive single channel (VBF), which demonstrates the predominant sensitivity of the VBF channel to $\mathcal{B}(H \rightarrow \text{inv})$, as well as shows that the two standard-deviation data excess in the combined Run 1 and Run 2 limit is driven by the VBF result. All search channels show weaker than expected limits, with the exception of $t\bar{t}H$. A finer breakdown into the hadronic and leptonic final states in each production mode, following the listing in Table 8.1, is presented in Appendix D. The dependence of the profile likelihood functions on $\hat{\mu}$ is




Figure 8.2: Left: Observed and expected limits at 95% CL for the combination, broken down by data-taking period. Right: Observed and expected limits at 95% CL for the combination, broken down by analysis.

shown in Fig. 8.3 (right), with the breakdown of the best-fit $\mathcal{B}(H \rightarrow \text{inv})$ values for each channel, and corresponding 95% CL exclusion limits for reference, presented in Table 8.5. The best-fit values of $\mathcal{B}(H \rightarrow \text{inv})$ for the individual production channels are compatible with one another and with the combined value of $\hat{\mu} = 0.08^{+0.04}_{-0.04}$, and the observed signal strength is compatible within two standard deviations with the SM prediction.

Table 8.5: The observed best-fit estimates of $\mathcal{B}(H \rightarrow \text{inv})$, for each analysis channel in the combination, and the 95% CL upper limits on $\mathcal{B}(H \rightarrow \text{inv})$.

Channel	Best-fit $\mathcal{B}(H \rightarrow \text{inv})$	95% CL upper limit
Combined	$0.08^{+0.04}_{-0.04}$	0.15 (0.08 exp.)
VBF-tag	$0.09^{+0.05}_{-0.05}$	0.18 (0.10 exp.)
VH-tag	$0.07^{+0.09}_{-0.09}$	0.24 (0.18 exp.)
ttH-tag	$-0.11^{+0.15}_{-0.15}$	0.25 (0.30 exp.)
ggH-tag	$0.22^{+0.16}_{-0.16}$	0.49 (0.32 exp.)

The CMS combination result is comparable to the recently-published ATLAS analogue [152], which reports a combined 95% CL upper limit of 0.107 (0.077 expected) using Run 1 and Run 2 data. In the ATLAS analysis, the channels VBF, ttH, $Z(\ell^+\ell^-)H$,




Figure 8.3: Left: Observed and expected limits at 95% CL for the combination, broken down by the Higgs boson production mode. Right: The profile likelihood scan corresponding to the observed limit in the fit to the individual production channels.

VBF + γ , and ggH are included, resulting in an excess of only one standard deviation from the SM expectation.

8.4 Dark matter interpretation

The upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ is interpreted in the context of a set of Higgs portal models governing DM interactions, where a stable WIMP, such as a scalar, fermionic, or vectorial singlet state has a substantial coupling to the 125 GeV Higgs boson. For sufficiently light DM particles, with mass $m_{\text{DM}} < \frac{m_H}{2}$, these WIMPs will appear as invisible decay products of the Higgs boson. In direct-detection experiments, the interaction of a WIMP with an atomic nucleus can occur via the exchange of a Higgs boson, and the resulting nuclear recoil is measured to obtain an upper bound on the spin-independent DM-nucleon scattering cross section, $\sigma_{\text{DM-nucleon}}^{\text{SI}}$. Therefore, the interaction between WIMPs and nucleons can be probed at low energies in direct-detection apparatus, as well as in colliders at high energies via $H \rightarrow \text{inv}$ decays.

An effective field theory (EFT) approach is considered for WIMPs of scalar and Majorana-like fermionic nature, introduced in Ref. [153], while in the vector case two UV-complete DM models are considered, given the EFT analogue is known to violate unitarity. The vectorial WIMP model (Vector DM^{UV-comp}) is introduced in Refs. [153, 154]. The UV-complete solution necessitates the existence of a VEV borne from some new SM singlet scalar, Φ , required to guarantee the invariance of a new gauge sym-

metry $U(1)'$. The corresponding Lagrangian is given a generic scalar Higgs portal coupling, permitting mixing between the SM Higgs boson and Φ . As this is not a unique UV completion, a radiative portal approach (Vector $DM_{m_2}^{\text{radiative}}$) is introduced in Ref. [155] for dark Higgs boson masses $m_2 = 65$ and 100 GeV, and with a mixing angle between the SM and dark Higgs bosons $\theta = 0.2$. The choice of m_2 values is derived from the best and worst limits obtained in the UV-complete model when varying m_2 in the range $[65, 1000]$ GeV [156]. The radiative UV-complete model introduces the Higgs portal via additional heavy fermions that interact under both $U(1)'$ and SM gauge symmetries. These fermions mediate Higgs boson and vectorial WIMP interactions analogously to t quarks mediating Higgs boson production in ggH at the one-loop level, illustrated in Fig. 1.1. The effect of the radiative portal model is a trivial quartic correction to $\sigma_{\text{DM-nucleon}}^{\text{SI}}$ as a function of θ and m_2 . The resulting formulae for $\sigma_{\text{DM-nucleon}}^{\text{SI}}$ in the EFT spin-zero and spin- $\frac{1}{2}$ models, $\sigma_{S\text{-nucleon}}$ and $\sigma_{f\text{-nucleon}}$, respectively, and the UV-complete spin-one models with and without the radiative portal correction, $\sigma_{V\text{-nucleon}}^{\text{rad}}$ and $\sigma_{V\text{-nucleon}}^{\text{UV-comp}}$, respectively, are taken from Ref. [156] and given by

$$\sigma_{S\text{-nucleon}} = \frac{8\Gamma_{H \rightarrow SS}^{\text{inv}} m_N^4 f_N^2}{m_H^3 v^2 \beta_S (m_S + m_N)^2}, \quad (8.1)$$

$$\sigma_{f\text{-nucleon}} = \frac{16\Gamma_{H \rightarrow ff}^{\text{inv}} m_f^2 m_N^4 f_N^2}{m_H^5 v^2 \beta_f^3 (m_f + m_N)^2}, \quad (8.2)$$

$$\sigma_{V\text{-nucleon}}^{\text{UV-comp}} = \frac{32\Gamma_{H \rightarrow VV}^{\text{inv}} m_V^4 m_N^4 f_N^2}{m_H^7 v^2 \beta_V (m_V + m_N)^2} \cdot \left(1 - 4\frac{m_V^2}{m_H^2} + 12\frac{m_V^4}{m_H^4}\right)^{-1}, \quad (8.3)$$

$$\sigma_{V\text{-nucleon}}^{\text{rad}} = \sigma_{V\text{-nucleon}}^{\text{UV-comp}} \cdot \cos^4 \theta \cdot \left(1 - \frac{m_H^2}{m_2^2}\right)^2, \quad (8.4)$$

where $\Gamma_{H \rightarrow \chi\chi}^{\text{inv}}$ is the Higgs boson partial decay width to WIMP χ with mass m_χ and β_χ function given by

$$\beta_\chi = \sqrt{1 - 4\frac{m_\chi^2}{m_H^2}}, \quad (8.5)$$

m_N is the nucleon mass, and f_N , equal to 0.308 ± 0.018 , is the Higgs-nucleon coupling parametrisation factor¹. The value for $\Gamma_{H \rightarrow \chi\chi}^{\text{inv}}$ is estimated using the relation to the

¹The Higgs VEV v has value 246 GeV throughout this thesis, however some Higgs portal literature use the convention $v = 173$ GeV, which is a factor $\sqrt{2}$ smaller.

branching fraction

$$\mathcal{B}(H \rightarrow \text{inv}) = \frac{\Gamma_{H \rightarrow \chi\chi}^{\text{inv}}}{\Gamma_{\text{SM}}^{\text{inv}} + \Gamma_{H \rightarrow \chi\chi}^{\text{inv}}} \Leftrightarrow \Gamma_{H \rightarrow \chi\chi}^{\text{inv}} = \frac{\Gamma_{\text{SM}}^{\text{inv}} \cdot \mathcal{B}(H \rightarrow \text{inv})}{1 - \mathcal{B}(H \rightarrow \text{inv})}, \quad (8.6)$$

where $\Gamma_{\text{SM}}^{\text{inv}} = 0.00407$ GeV is the SM Higgs boson decay width to invisible particles [20], and $\Gamma_{H \rightarrow \chi\chi}^{\text{inv}} = \mathcal{B}(H \rightarrow \text{inv}) \cdot \Gamma_H^{\text{tot}}$, for Γ_H^{tot} the total Higgs decay width, following Ref. [156].

Uncertainties in $\sigma_{\text{DM-nucleon}}^{\text{SI}}$ are obtained from extrema of f_N that are derived from lattice theory [157], given this is the dominant effect on $\sigma_{\text{DM-nucleon}}^{\text{SI}}$. The results are compared with those from direct-detection searches, which are designed around the principle that DM particles interact with the atomic nuclei of the active materials in which they specialise. The limits featured are the most recent reported by the XENON1T-Migdal [158], DarkSide-50 [159], Panda-X 4T [160], and LUX-ZEPLIN [161] experiments. The 90% CL upper limits on $\sigma_{\text{DM-nucleon}}^{\text{SI}}$ as a function of m_{DM} in the range [0.1,1000] GeV are presented in Fig. 8.4. The observed exclusion threshold using all Run 1 and Run 2 data is estimated at 0.14 at the 90% CL for an equivalent comparison to the direct-detection results. From comparison to direct detection, the Higgs portal limits are better for light DM masses but are highly model-dependent.

The sensitivity of the Run 1 and Run 2 channels to $\mathcal{B}(H \rightarrow \text{inv})$ depends on the cross sections for the different Higgs boson production modes. These can be parametrised by the coupling strength of the Higgs boson to V bosons and fermions. The coupling strength modifiers that scale these cross sections, κ_V and κ_F , are assumed to be equal to 1. However, BSM scenarios can be tested by varying the parameters κ_V and κ_F [162]. Using CMS analysis data, the observed 95% CL upper limits on $\mathcal{B}(H \rightarrow \text{inv})$ are evaluated as functions of κ_V and κ_F , and presented in Fig. 8.5 alongside best-fit estimates $\hat{\kappa}_V$ and $\hat{\kappa}_F$ and the respective one and two standard-deviation contours, as taken from Ref. [163]. The 95% CL limit on $\mathcal{B}(H \rightarrow \text{inv})$ is found to be 0.15 at the best-fit value for κ_V and κ_F from Ref. [163], and varies between 0.13 and 0.17 inside the 95% CL contour.




Figure 8.4: 90% CL upper limits on $\sigma_{\text{DM-nucleon}}^{\text{SI}}$ as a function of m_{DM} . Results are presented for a fermion (red) and scalar (yellow) DM candidate derived from EFT models, and two cases of a vector DM candidate derived from UV-complete approaches, the first denoted Vector DM $^{\text{UV-comp}}$ [154] (burgundy), and the second a radiative portal version denoted Vector DM $^{\text{radiative}}_{m_2}$ [155] (orange), with dark Higgs boson masses of $m_2 = 65$ and 100 GeV and a dark-SM Higgs boson mixing angle of $\theta = 0.2$. Uncertainties are derived from Refs. [157]. Results are compared to direct-detection searches from XENON1T-Migdal [158], DarkSide-50 [159], PandaX-4T [160] and LUX-ZEPLIN [161].




Figure 8.5: Observed 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ as a function of coupling strength modifiers, κ_V and κ_F , for a Higgs boson of mass 125 GeV. Values for $\hat{\kappa}_V$ and $\hat{\kappa}_F$ from Ref. [163] are shown as a black cross, together with 68 and 95% CL contours.

Chapter 9

Summary

As outlined at the start of this thesis, searches for invisible decays of the Higgs boson are a means of putting our understanding of the standard model (SM) to the test, and probing the universe beyond the SM (BSM). The target sector outside of the SM is the dark sector, a hypothetical set of frameworks that are designed to describe very real phenomena like dark matter (DM), investigating how it exists in our universe, and how it interacts with the SM. Data from particle colliders, such as those recorded at the CERN LHC using the CMS detector, can be analysed for $H \rightarrow \text{inv}$ events, where the SM predicted branching fraction $\mathcal{B}(H \rightarrow \text{inv}) = 0.12\%$.

Using 2016–2018 CMS data corresponding to 138 fb^{-1} , the $t\bar{t}H$ and VH production modes with fully hadronic final states are analysed. The former targets final states rich in jet multiplicity, and includes final states containing b jets, or boosted t quarks or W bosons. The latter focuses on resolving a dijet pair with an invariant mass compatible with the decay of a V boson. No significant excess of events is observed above those from SM background processes, hence a 95% confidence level (CL) upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.54 (0.39 expected) is set in the combined $t\bar{t}H$ and resolved VH channels, assuming SM production cross sections.

In the $t\bar{t}H$ channel, the same exclusion threshold is set at 0.43 (0.52 expected), which is a clear improvement on the previous CMS limit obtained using 2016 data corresponding to 35.9 fb^{-1} of $\mathcal{B}(H \rightarrow \text{inv}) = 0.85$ (0.73 expected) [145] in the zero-lepton channel analysed by Ref. [145]. The year-for-year comparison shows that this result is consistent with the 95% CL upper limit of 0.81 (0.76 expected) obtained in this analysis using the 2016 data set alone. A search by the ATLAS experiment in the same channel measured a 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.49 (0.64 expected) using full Run 2 data [146]. Therefore, the CMS result represents an improvement in sensitivity to $\mathcal{B}(H \rightarrow \text{inv})$ of $\approx 20\%$, however a small data deficit is measured in

the CMS case. The 95% CL upper limit in the resolved VH channel is measured at 0.74 (0.53 expected), for which there is no published comparison using post-2015 data aside from the orthogonally-defined boosted VH topology studied in Ref. [11], which reports a limit of 0.37 (0.31 expected), and so the excess of data in the combined VH channel with a fully hadronic final state was anticipated.

The results are combined with previous $\mathcal{B}(H \rightarrow \text{inv})$ searches carried out at $\sqrt{s} = 7, 8$, and 13 TeV in complementary production modes with the CMS detector. Combining the channels VBF, VH, $t\bar{t}H$, and ggH in all analysed final states, a combined 95% CL upper limit on $\mathcal{B}(H \rightarrow \text{inv})$ of 0.15 (0.08 expected) is obtained using Run 1 (2011–2012) and Run 2 (2015–2018) data. This is a relative improvement in sensitivity of $\approx 20\%$ compared to the most sensitive single channel, VBF [8], and a relative improvement of $\approx 65\%$ relative to the previous combination result using Run 1 and 2015 data [16], although in this combination no $t\bar{t}H$ channel had been established. The ATLAS combined 95% CL upper limit using similar channels is reported at 0.107 (0.077 expected) using data from the same periods [152], and therefore both ATLAS and CMS experiments report consistent upper limits. The CMS combination results are interpreted in the context of Higgs portal models governing DM interactions with the SM. Exclusion limits for scalar-, fermionic-, and vector-spin DM particle models are produced and compared to contemporary direct-detection limits. Limits from the Higgs portal models are found to be superior for light DM masses, although are highly model-dependent. All results of the $t\bar{t}H$ and resolved VH analysis, and the accompanying combination, are reported in Ref. [10].

The outlook for future $H \rightarrow \text{inv}$ searches lies in pursuing these further at the end of Run 3, which has already completed data taking for 2022 and is scheduled to continue through to the end of 2025, collecting double the data taken during Run 2. However, to achieve the precision required to truly probe a branching fraction $\mathcal{O}(0.1\%)$, an improvement on the order of $100\times$ is required. In the Future Circular Collider under the hadron-collider regime (FCC-hh), 30 ab^{-1} of data is anticipated, therefore $\mathcal{B}(H \rightarrow \text{inv})$ can be probed to the level of 10^{-4} [164] with well-constrained theoretical and experimental uncertainties and an extremely precise estimation of the $Z(\nu\bar{\nu}) + \text{jets}$ spectrum. Alternatively, the less expensive e^+e^- International Linear Collider (ILC) is also in contention for construction, operating as a Higgs boson production factory via VBF and VH mechanisms, the most sensitive channels to $\mathcal{B}(H \rightarrow \text{inv})$, although such a machine would have fewer future prospects than the FCC.

Appendix A

Look-up values

The Pauli matrices, σ_i , $i \in [1, 2, 3]$, and the identity matrix $\mathbb{I}$, are given by

$$\sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \mathbb{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (\text{A.1})$$

The Dirac matrices, γ^μ , $\mu \in [0, 1, 2, 3]$, are given by

$$\gamma_0 = \begin{pmatrix} \mathbb{I} & \mathbf{0} \\ \mathbf{0} & -\mathbb{I} \end{pmatrix}, \quad \gamma_1 = \begin{pmatrix} \mathbf{0} & \sigma_1 \\ -\sigma_1 & \mathbf{0} \end{pmatrix}, \quad \gamma_2 = \begin{pmatrix} \mathbf{0} & \sigma_2 \\ -\sigma_2 & \mathbf{0} \end{pmatrix}, \quad \gamma_3 = \begin{pmatrix} \mathbf{0} & \sigma_3 \\ -\sigma_3 & \mathbf{0} \end{pmatrix}, \quad (\text{A.2})$$

where the 2×2 zero matrix is given by

$$\mathbf{0}_{[2,2]} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}. \quad (\text{A.3})$$

The Gell-Mann matrices, λ_i , $i \in [1, 2, 3, 4, 5, 6, 7, 8]$, are given by

$$\lambda_1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \lambda_2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \lambda_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (\text{A.4})$$

$$\lambda_4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \quad \lambda_5 = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \quad (\text{A.5})$$

$$\lambda_6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \lambda_7 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \quad \lambda_8 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}. \quad (\text{A.6})$$

Appendix B

Statistical analysis

The following serves as a primer to the statistical approach to general $H \rightarrow \text{inv}$ searches, introducing the maximum-likelihood estimation and methods testing the strength of the fit model, as means of quantifying a significance on any discovery of $H \rightarrow \text{inv}$ events or measurements of upper limits on $\mathcal{B}(H \rightarrow \text{inv})$ that are obtained in these analyses. These largely adhere to methods outlined in Refs. [131, 142] and references therein, as recognised by the LHC experiments as the correct approach adopted in analyses. All $H \rightarrow \text{inv}$ searches aim to discover new signal processes that are statistically distinct from the background. In such an experiment, the total number of observed events $N = \mu S(H \rightarrow \text{inv}) + B$, where μ is the *signal strength* relative to the expected strength of the signal in the SM, given by

$$\mu = \frac{\sigma_{\text{obs}}(H \rightarrow \text{inv})}{\sigma_{\text{SM}}(H \rightarrow \text{inv})}. \quad (\text{B.1})$$

The probability of some $H \rightarrow \text{inv}$ event x based on several event observations in LHC data is encoded in a *hypothesis*, H , which statistical tests are designed to validate in light of observation. The *null* hypothesis H_0 corresponds to $\mu = 0$, where only SM background processes are observed. The *alternative* hypothesis, H_1 , is where $\mu > 0$ and there is some non-zero signal contribution. A hypothesis H bears a *probability distribution* $f(x|H)$, and the probability that an observation x is compatible with H given H is true is defined as the *likelihood*, denoted $\mathcal{L}(x|H) \equiv P(x|H)$.

A test of H_0 is defined by specifying a critical region W , the boundary of which corresponds to a probability α that the data observed in W , assuming H_0 is true, is correct. This is equivalent to $P(x \in W|H_0) < \alpha$, where if $x \in W$, H_0 is rejected. The probability α denotes the *significance level* (SL) of the test. In general, W is not a unique region for a given SL, thus H_1 must be included in the test of H_0 . A rule of

thumb is that W is a region of low probability if H_0 is true, and high probability if H_1 is true. An example of this is shown in Fig. B.1.




Figure B.1: A Gaussian distributed x with critical region W indicated either side of the hard border corresponding to the SL α (with an x -value x_c), along with the respective $f(x|H)$ given the hypotheses H_0 (with mean μ_0) and H_1 (with mean μ).

In the frequentist approach adopted here, it is only the observation x that determines the probability that H is accepted or rejected assuming it (or the alternative) is true. Rejecting H_0 if proven-true is referred to as a *Type-I error*. The maximum probability for this is the SL, α . Similarly, accepting H_0 if proven-false (and H_1 rejected when proven-true) is referred to as a *Type-II error*. This occurs with probability $P(x \notin W|H_1) = \beta$, where the *power* of the test with respect to H_1 is given by $1 - \beta$. A test for H_0 ought to be designed such that the power is maximised with respect to H_1 , meaning to maximise the probability H_0 is rejected when H_1 is true. However, it is possible that there is no single H_1 for which the power is maximal over all possible alternative H , for example in some Gaussian distributed observations x with mean μ , uncertainty σ , and probability distribution

$$f(x; \mu, \sigma) \begin{cases} H_0 & \mu = \mu_0 \\ H_1 & \mu > \mu_0. \end{cases} \quad (\text{B.2})$$

The first step is to find the critical boundary x_c for which the power is maximised. This is determined by the SL, where $\alpha = P(x \geq x_c|\mu_0)$, and is referred to as a *one-tailed*

test, illustrated in Fig. B.2 (left). The test is not necessarily exclusive of the region for which $\mu < \mu_0$, which requires a *two-tailed* test, demonstrated in Fig. B.2 (right), where the probability of a result in each region assuming H_0 is true now has value $\alpha/2$. The result of a two-tailed test gives a reasonable power in the region for which $\mu < \mu_0$, but weakens the power in the region $\mu > \mu_0$ versus the one-tailed test.




Figure B.2: Visual representations of the possible hypothesis tests, demonstrating the one-tailed (left) and two-tailed tests (right) distributed in μ . The two-tailed test improves on the statistical test of an alternative hypothesis for $\mu < \mu_0$, but loses power versus the one-tailed test in the region for which $\mu > \mu_0$. This is a case where no single test can achieve maximal power.

For $H \rightarrow \text{inv}$ searches, tests are constructed with a well-understood false rate α that in the SM, $\mathcal{B}(H \rightarrow \text{inv}) = \mathcal{B}(H \rightarrow ZZ^* \rightarrow 4\nu)$, and a high power with respect to H_1 , $\mathcal{B}(H \rightarrow \text{inv}) > \mathcal{B}(H \rightarrow ZZ^* \rightarrow 4\nu)$. This will have a form

$$\begin{cases} H_0(\text{SM}) & \mathcal{B}(H \rightarrow \text{inv}) = \mathcal{B}(H \rightarrow ZZ^* \rightarrow 4\nu) \\ H_1(\text{BSM}) & \mathcal{B}(H \rightarrow \text{inv}) > \mathcal{B}(H \rightarrow ZZ^* \rightarrow 4\nu). \end{cases} \quad (\text{B.3})$$

The event selection divides the sample space into regions based on parameter acceptances dictated by H . The highest power for a given SL in a test of the SM $\mathcal{B}(H \rightarrow \text{inv})$ versus a BSM $\mathcal{B}(H \rightarrow \text{inv})$ will give a critical region W , controlled by constant c that gives a test of a desired size, with $\alpha = P(t(x) \leq c)$. This is the Neyman-Pearson

lemma, where the optimal test statistic $t(x)$ is given by the likelihood ratio

$$t(x) = \frac{\mathcal{L}(x|H_1)}{\mathcal{L}(x|H_0)} \begin{cases} > c & \text{inside } W \\ \leq c & \text{outside } W, \end{cases} \quad (\text{B.4})$$

The optimal boundary on W is not always calculable, given there may not be enough information about the likelihoods $\mathcal{L}(x|H_0(\text{SM}))$ or $\mathcal{L}(x|H_1(\text{BSM}))$. Note that the BSM case will always be categorised under some model, as there is no model-independent test given an alternative model with a higher or lower power will always exist.

B.1 Goodness-of-fit tests

Let a prediction $f(\mathbf{x}|H)$ be based on a set of observations $\mathbf{x}$ under H . On observing a single point, x_{obs} , the validity of H can be tested. The validity of this result is quantified using a p -value, which is the probability that any sample of data is of equal or lesser compatibility with H versus the observed data, assuming H is true. The p -value is dependent on the data, and is extracted from a *goodness-of-fit* (GF) test. If this is found from some test statistic $t(\mathbf{x})$ on H , then the p -value is related via

$$p_H = \int_t^\infty f(t'|H) dt', \quad (\text{B.5})$$

where the probability that p_0 , corresponding to H_0 , is less than α is given by

$$P(p_0 \leq \alpha|H_0) = \alpha. \quad (\text{B.6})$$

Therefore, reformulating x_{obs} as the scalar $t(\mathbf{x})$ defines the critical region W in H_0 with size α as the set of data for which $p_0 < \alpha$. The p -value quantifies signal-like fluctuations only, assuming H_0 is true, and neglects other statistical phenomena such as effects sensitive to the hadronic recoil binning and the look elsewhere effect (LEE)¹.

Even though here, p_0 is dependent on H_0 , the resulting test will have a power with respect to a given H_1 . Obtaining the p -value determines the significance Z , which is the number of standard deviations σ a Gaussian distributed set of the observations $\mathbf{x}$ a statistical fluctuation is in one direction that will give the same p -value as that of

¹The LEE is a measure of the probability to find a signal peak at least as significant as the one observed anywhere else in the full distribution in $\mathbf{x}$.

the experiment. This is illustrated in Fig. B.3, so-giving the relation

$$p = \frac{1}{\sqrt{2\pi}} \int_Z^\infty e^{-\frac{x^2}{2}} dx = 1 - \Phi(Z) \Rightarrow Z = \Phi^{-1}(1 - p). \quad (\text{B.7})$$




Figure B.3: The relation between Z and p -value, demonstrating visually the definition of the significance Z as the fluctuation in a Gaussian distributed set of observations $\mathbf{x}$ in a single direction with the same p -value as that produced in the actual experiment.

The typical test statistic used to measure the GF based on observation $\mathbf{x}$ composed of $\mathcal{N}_{\text{dof}}$ degrees of freedom to the predicted $\boldsymbol{\mu}$ is Pearson's χ^2 statistic, defined as

$$\chi^2 = \sum_{i=1}^{\mathcal{N}_{\text{dof}}} \frac{(x_i - \mu_i)^2}{\sigma_i^2} \xrightarrow{x_i \sim \text{Po}(\mu_i)} \sum_{i=1}^{\mathcal{N}_{\text{dof}}} \frac{(x_i - \mu_i)^2}{\mu_i}, \quad (\text{B.8})$$

where the first expression represents a Gaussian distributed x_i with mean μ_i and standard deviation σ_i , and the second expression arises for Poisson distributed x_i . If in the Poissonian case the means $\mu_i \gg 1$, then their distribution can be appropriately Gaussian modelled, and the χ^2 test is invariant of the distribution. The p -value is generally soluble from the χ^2 probability distribution,

$$p = \int_{\chi^2}^\infty f_{\chi^2}(\chi^2; N) d\chi^2, \quad (\text{B.9})$$

however in the event that some hadronic recoil intervals contain few events, the χ^2 will not follow the expected probability distribution, and the GF is poorly measured. The

saturated model method is employed, outlined in Refs. [165, 166], which estimates a quantity χ_{sat}^2 that is similar to χ^2 but is calculable for arbitrary combinations of hadronic recoil intervals with arbitrary associated constraints. The relation $\chi_{\text{sat}}^2 = \chi^2$ is asymptotically valid, in the large sample limit in all bins, and only under certain conditions. Nevertheless, the test is still a useful measure of GF, and is indicative of the level of agreement between data and prediction.

B.2 Parameter fitting and maximum likelihood

For some probability distribution $f(\mathbf{x}; \Theta)$ describing data based on observations $\mathbf{x}$, the parameters Θ are constants that characterise the shape of $f(\mathbf{x}; \Theta)$, for which a test statistic $t(\mathbf{x})$ is a function of $\mathbf{x}$ for which all Θ are known. The general *estimator* $\hat{\theta}_i(\mathbf{x})$ of a particular $\theta_i \in \Theta$ is a test statistic that can be used to estimate some property of $f(\mathbf{x}; \theta_i)$. A good estimator will converge to the real θ_i in the limit of a large set of observations. The procedure to estimate $\hat{\theta}(\mathbf{x})$ given $\mathbf{x}$ is called *parameter fitting*, where repeating the procedure yields estimates that follow a probability distribution $g(\hat{\theta}_i; \theta_i)$.

Formally, the parameters Θ contain a parameter of interest (POI), the signal strength μ , and many NPs, θ , such that $\mathcal{L}(\theta) = f(\mathbf{x}; \theta)$. The *best-fit* or *maximum-likelihood* estimators (MLEs), $\hat{\theta}_{\text{MLE}}$, correspond to the parameter values for which $\mathcal{L}(\theta)$ is maximised, and the probability of obtaining further observations like $\mathbf{x}$ is greatest. For Gaussian distributed $\mathbf{x}$ with true mean μ and variance σ^2 , the likelihood is given by

$$\mathcal{L}(\mu, \sigma^2) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} \exp -\frac{(x_i - \mu)^2}{2\sigma^2}, \quad (\text{B.10})$$

where the test statistics $\hat{\mu}$ and $\hat{\sigma}^2$ that maximise $\mathcal{L}(\mu, \sigma^2)$ also maximise the logarithmic likelihood (LL) of $\mathcal{L}(\mu, \sigma^2)$,

$$\ln(\mathcal{L}(\mu, \sigma^2)) = \sum_{i=1}^n \ln(f(x_i; \mu, \sigma^2)) = \sum_{i=1}^n \left(\ln\left(\frac{1}{\sqrt{2\pi}}\right) + \frac{1}{2} \ln\left(\frac{1}{\sigma^2}\right) - \frac{(x_i - \mu)^2}{2\sigma^2} \right), \quad (\text{B.11})$$

where the convention to use the LL is purely given likelihood functions can be numerically very large.

Once an estimate of some $\hat{\theta} \in \theta$ is obtained, the statistical error, or width of the distribution of $\hat{\theta}$ on many repeats, can be quantified. Estimates of $\hat{\theta}$ are almost always Gaussian distributed in the large sample limit as they become statistically unbiased.

An information inequality sets a lower bound on the variance of a general estimator, which in the large sample limit expanded about $\hat{\theta}$ yields

$$\ln(\mathcal{L}(\theta)) \approx \ln(\mathcal{L}_{\max}) - \frac{(\theta - \hat{\theta})^2}{2\hat{\sigma}_{\hat{\theta}}^2}, \quad (\text{B.12})$$

where $\ln(\mathcal{L}(\hat{\theta})) = \ln(\mathcal{L}_{\max})$ on assuming $\hat{\theta} = \hat{\theta}_{i,\text{MLE}}$. If θ is changed by $\pm\hat{\sigma}$ about the MLE, the LL becomes

$$\ln(\mathcal{L}(\theta)) \approx \ln(\mathcal{L}_{\max}) - \frac{1}{2} \quad (\text{B.13})$$

and the LL decreases by half from its maximum value.

For two correlated parameters, α and β , the maximum of the corresponding LL, $\ln(\mathcal{L}(\alpha, \beta))$, is computed numerically using simulated event generation. A maximum-likelihood fit between simulation and data using some pre-fit values for α and β results in MLEs $\hat{\alpha}$ and $\hat{\beta}$, and their covariance $\text{cov}[\hat{\alpha}, \hat{\beta}]$, which is a measure of the correlation between the two parameters. Repeated maximum-likelihood fits will produce MLE averages close to their true values. In the large sample limit, $\ln(\mathcal{L}(\alpha, \beta))$ tends to a quadratic about the maximum.

B.3 Interval estimates and limit setting

An interval for a statistical uncertainty is associated with an estimate $\hat{\theta}$ of the true value θ , and an estimate of the standard deviation $\sigma_{\hat{\theta}}$. Therefore, across repeat experiments, estimates of $\hat{\theta}$ will follow a *sampling distribution* $g(\hat{\theta})$ centred about $\theta \pm \sigma_{\hat{\theta}}$, which is Gaussian distributed in the large sample limit². However, in general $g(\hat{\theta})$ is not Gaussian distributed, and instead a *confidence interval* (CI) is reported. The upper and lower tail probabilities of $g(\hat{\theta}; \theta)$ are denoted as α and β , respectively. The aim is to find the boundaries of α and β , $\mu_{\alpha}(\theta)$ and $\nu_{\beta}(\theta)$, respectively, where

$$\alpha = P(\hat{\theta} \geq u_{\alpha}(\theta)) = \int_{u_{\alpha} \rightarrow \hat{\theta}_{\text{obs}}}^{\infty} g(\hat{\theta}; \theta) d\hat{\theta}, \quad (\text{B.14})$$

$$\beta = P(\hat{\theta} \leq v_{\beta}(\theta)) = \int_{-\infty}^{v_{\beta} \rightarrow \hat{\theta}_{\text{obs}}} g(\hat{\theta}; \theta) d\hat{\theta}, \quad (\text{B.15})$$

an illustration for which is provided in Fig. B.4 (left) about $\hat{\theta}$. The distribution of $\hat{\theta}$ as a function of θ will produce a region between boundaries $\mu_{\alpha}(\theta)$ and $\nu_{\beta}(\theta)$ known as

²Over several parameters, $g(\hat{\theta})$ is a multidimensional Gaussian distribution with parameters characterised by covariance matrix V .

the *confidence belt*, illustrated in Fig. B.4 (right). The points where $\hat{\theta}_{\text{obs}}$ intersect the confidence belt define the CI $[a, b]$. The CL is the probability to cover the true value of the parameter within the CI, and is equal to $1 - \alpha - \beta$.




Figure B.4: Left: The upper and lower tail probabilities, α and β , respectively (left), are illustrated in the sampling distribution $g(\hat{\theta}; \theta)$. Right: The confidence belt constructed from the region between u_α and v_β , where the intersection with θ_{obs} gives the boundaries a and b of the CI in θ .

In the context of $H \rightarrow \text{inv}$ searches, when seeking the signal S in a region W , a histogram depicting Poisson distributed yields N_i with mean $\mu S_i + B_i$ across $\mathcal{N}_{\text{bins}}$ subcategory and hadronic recoil intervals is modelled. Further subsidiary measurements, denoted $\mathbf{m}$, that are applicable to $\mathcal{M}_{\text{bins}}$ subcategory or hadronic recoil intervals, that help constrain the background or hadronic recoil shape parameters are typically Poisson distributed with mean $u_i(\boldsymbol{\theta})$ depending on $\boldsymbol{\theta}$. The likelihood is then given by

$$\mathcal{L}(\mu, \boldsymbol{\theta}) = \prod_{i=1}^{\mathcal{N}_{\text{bins}}} \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-(\mu s_i + b_i)} \prod_{j=1}^M \frac{u_j^{m_j}}{m_j!} e^{-u_j}. \quad (\text{B.16})$$

The profile likelihood ratio (PLR) is used to test H , given by

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}})}{\mathcal{L}(\hat{\mu}, \hat{\hat{\boldsymbol{\theta}}})}, \quad (\text{B.17})$$

where $\mathcal{L}(\hat{\mu}, \hat{\hat{\boldsymbol{\theta}}})$ is the maximised likelihood function, based on MLEs $\hat{\mu}$ and $\hat{\hat{\boldsymbol{\theta}}}$, and $\hat{\hat{\boldsymbol{\theta}}}$ is the value of $\boldsymbol{\theta}$ that maximises $\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}})$ for a given value of μ . The presence of the NPs leads to the broadening of $\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}})$ as a function of μ compared to the maximised likelihood, and is indicative of the loss of sensitivity to μ . Conventionally, the PLR test statistic t_μ is defined as

$$t_\mu = -2 \ln(\lambda(\mu)), \quad (\text{B.18})$$

where higher values of t_μ represent an increasing incompatibility between data and μ . A general assumption for this is that μ , corresponding to the POI $\mathcal{B}(\text{H} \rightarrow \text{inv})$, is non-negative given a negative signal contribution is non-physical in $\text{H} \rightarrow \text{inv}$ searches. However, it is convenient to let $\hat{\mu}$ be an *effective estimator* of $\mathcal{B}(\text{H} \rightarrow \text{inv})$ that is allowed to be negative in order to maximise $\mathcal{L}(\hat{\mu}, \hat{\theta})$.

In discovery experiments, a special case of t_μ , denoted q_μ , is used, which in the B -only H_0 ($\mu = 0$) is given by

$$q_0 = \begin{cases} -2 \ln(\lambda(0)) & \hat{\mu} \geq 0 \\ 0 & \hat{\mu} < 0, \end{cases} \quad (\text{B.19})$$

where the data is interpreted to lack agreement with H_0 only if $\hat{\mu} > 0$, as for any $\hat{\mu} < 0$ there is only evidence against H_0 due to some systematic discrepancy. In computing exclusion limits, the test statistic is given by

$$q_\mu = \begin{cases} -2 \ln(\lambda) & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu, \end{cases} \quad (\text{B.20})$$

where an upward fluctuation of the data, with $\hat{\mu} > \mu$, is not treated as having less compatibility with H_1 than the data obtained and is not reason to reject it.

The Wald approximation of the PLR demonstrates that $\hat{\mu}$ is approximately Gaussian distributed with mean μ' and variance σ^2 , and as such in the large sample limit p -values can be derived from $f(q_0|0)$ and $f(q_\mu|\mu')$, and hence their discovery or exclusion significances, Z_0 or Z_μ , respectively, given by

$$p_0 = 1 - F(q_0|0) \Leftrightarrow Z_0 = \Phi^{-1}(1 - p_0) = \sqrt{q_0} \quad (\text{B.21})$$

$$p_\mu = 1 - F(q_\mu|\mu) \Leftrightarrow Z_\mu = \Phi^{-1}(1 - p_\mu) = \sqrt{q_\mu}, \quad (\text{B.22})$$

where $F(q_0|0)$ and $F(q_\mu|\mu)$ are the cumulative distributions of $f(q_0|0)$ and $f(q_\mu|\mu)$, respectively.

Therefore, to find which values are consistent with the data, a test of size α is performed for all μ . The values not rejected constitute a CI for μ at the $1 - \alpha$ CL. If formulated in terms of a p -value p_μ , then the CI represents the μ values for which $p_\mu > \alpha$. To set an upper limit, the test statistic q_μ defined in Eq. B.20 is used, with p -value

$$p_\mu = \int_{q_{\mu,\text{obs}}}^{\infty} f(q_\mu|\mu) dq_\mu \approx 1 - \Phi(\sqrt{q_\mu}), \quad (\text{B.23})$$

where the approximation is valid in the large sample limit. Therefore, the 95% CL on μ is the highest μ value possible for which $p_\mu \geq 0.05$.

It is possible that the effect of the hypothesised μ is very small in comparison with the B -only hypothesis. As such, $f(q_\mu|0)$ and $f(q_\mu|\mu)$ will appear very similar, illustrated in Fig. B.5 (left), where the probability distributions closely overlap. In contrast, sensitivity to μ would imply that the distributions are highly separated, and the power (probability to reject μ if $\mu = 0$) is substantially higher than α , as shown by Fig. B.5 (right). By construction, for low sensitivity the probability to reject μ assuming μ is true is α , while the probability to reject μ for $\mu = 0$, the power, is slightly greater than α . This means, for probability $\alpha = 5\%$, any hypotheses with no sensitivity are excluded. This is referred to as ‘spurious exclusion’.




Figure B.5: Left: The probability distributions for $f(q_\mu|0)$ and $f(q_\mu|\mu)$ closely overlapping, indicative of the hypothesised μ having a small impact relative to H_0 ($\mu = 0$). Right: The probability distributions for $f(q_\mu|0)$ and $f(q_\mu|\mu)$ when the power is much higher than α and the distributions are highly separated.

A technique designed to handle this is the CL_S procedure, where H_0 ($\mu = 0$) and H_1 ($\mu > 0$) are tested together with the same test statistic

$$Q = -2 \ln \left(\frac{\mathcal{L}_{S+B}}{\mathcal{L}_B} \right), \quad (\text{B.24})$$

where distributions are illustrated in Fig. B.6 for reasonable (left) and low sensitivities (right). The CL_s solution divides the p -value by CL_B , equal to $1 - p_0$, such that

$$CL_S = \frac{CL_{S+B}}{CL_B} = \frac{p_{S+B}}{1 - p_B} \equiv \frac{p_\mu}{1 - p_0}, \quad (\text{B.25})$$

where $S + B$ is rejected if $CL_S \leq \alpha$. This increases the effective p -value when the distributions are close, so prevents exclusion in the event of low sensitivity. The CL_S procedure can be performed on the parameter $\mu = \frac{\sigma_{\text{obs}}(H \rightarrow \text{inv})}{\sigma_{\text{SM}}(H \rightarrow \text{inv})}$, yielding an upper limit

on $\mathcal{B}(H \rightarrow \text{inv})$.




Figure B.6: Left: The CL_S technique applied to H_0 and H_1 distributed in the test statistic Q , where a reasonable sensitivity under Q gives well separated distributions. Right: The CL_S technique for low sensitivity under Q , leading to a large overlap and hence a small p_0 in the B -only hypothesis.

B.4 Nuisance parameters and systematic uncertainty

Uncertainties propagate in an experiment in two ways: randomly (statistically) and systematically. In the former, large statistical fluctuation is due to a limited number of observations or limited accuracy in each observation, resulting in a large spread of estimates about the true value. Systematic errors are harder to handle, and result in imprecise end values that can require corrections to account for the source of the errors. These can appear during observations, or when comparing data from different experimental sources. In some cases parameters are correlated across data sets, for which it is common practice to combine these variables to minimise their uncertainties. In this case, μ has a common value across all channels, and the likelihood abstractly takes the form

$$\mathcal{L}(\mathbf{x}|\mu, \boldsymbol{\theta}) = \prod_i \mathcal{L}_i(\mathbf{x}_i|\mu, \boldsymbol{\theta}_{\text{corr}}, \boldsymbol{\theta}_i), \quad (\text{B.26})$$

where $\boldsymbol{\theta}$ is the complete set of NPs, for which $\boldsymbol{\theta}_{\text{corr}}$ are NPs correlated across certain regions, and $\boldsymbol{\theta}_i$ are calculated independent of the region i .

Sometimes the best-fit estimates of parameters can lie outside the range of individual parameter values. This is more likely to occur if the uncertainty estimate depends on the parameter being estimated, which motivates the use of the likelihood method. This procedure estimates the probability distribution as a function of the data for fixed parameters, $P(\mathbf{x}|\boldsymbol{\theta})$, and constructs the likelihood in the opposite sense, so that $\mathcal{L}(\boldsymbol{\theta}|\mathbf{x})$ is a function of the parameters for given data. The LL is parabolic in the large

sample limit, for which the standard deviation is determined by

$$\ln(\mathcal{L}(\mu_0 \pm \sigma)) = \ln(\mathcal{L}(\mu_0)) - \frac{1}{2}. \quad (\text{B.27})$$

In general, a model of the data is not perfect. Models can be improved by including additional adjustable parameters in the model, with $\mathcal{L}(x|\theta) \rightarrow \mathcal{L}(x|\theta, \nu)$. However, this incorporates additional systematic uncertainty that increases the parameter space, which decreases the sensitivity to $\mathcal{B}(\text{H} \rightarrow \text{inv})$, and increases the variance of the estimate. For statistic q_θ used to test a hypothetical value of NP θ , such that its p -value

$$p_\theta = \int_{q_{\theta, \text{obs}}}^{\infty} f(q_\theta|\theta, \nu) dq_\theta, \quad (\text{B.28})$$

the value ν must satisfy that θ is rejected only if $p_\theta < \alpha \forall \nu$ if the exact CI is to be obtained. For test statistics based on the PLR, $f(q_\theta|\theta, \nu)$ becomes independent of any NPs in the large sample limit. In general, for finite data samples, this is not true, and some θ cannot possibly be rejected for certain ν , which leads to overestimates of the CI in θ . This is combated by profiling, which by construction rejects θ if $p_\theta \leq \alpha$, where the p -value is computed assuming the MLE value for a given θ , $\hat{\nu}(\theta)$. The resulting CI will give the correct coverage of θ and $\hat{\nu}(\theta)$. The *Hybrid Frequentist-Bayesian* method assumes ν adheres to a prior distribution $\Pi(\nu)$ with likelihood

$$\mathcal{L}_{\text{Hybrid}}(x|\theta) = \int L(x|\theta, \nu) \Pi(\nu) d\nu \quad (\text{B.29})$$

models the data. The p -values are then computed, and the model being tested is effectively a weighted average of possible models.

To guarantee that NPs are not over- or under-constrained by the fit, the pulls for each θ_i is measured, given by

$$\text{pull} = \frac{\hat{\theta}_i - \theta_{i,0}}{\sigma_{i,0}}, \quad (\text{B.30})$$

where $\theta_{i,0}$ and $\hat{\theta}_i$ are the pre-fit and MLE values of θ_i , respectively, and $\sigma_{i,0}$ is the pre-fit uncertainty. Another measure of the effect of NPs on μ is to compute their impacts. The impact of NP θ_i , $\Delta\mu^\pm$, on μ is obtained for each systematic uncertainty by performing the fit when the MLE $\hat{\theta}_i$ is fixed to $\pm\sigma_{\hat{\theta}_i}$ of its value, and measuring the magnitude of the shift in each direction from the MLE $\hat{\mu}$. Other NPs are allowed to vary to account for correlations amongst themselves. A detailed outline of this method is introduced in Ref. [167].

The test statistic as defined in Eq. B.20 can be scanned over a range of values for the effective estimator μ , formulated as

$$q_\mu = -2 \ln \frac{\mathcal{L}(\mathbf{x}|\mu, \hat{\boldsymbol{\theta}})}{\mathcal{L}(\mathbf{x}|\hat{\mu}, \hat{\boldsymbol{\theta}})}, \quad (\text{B.31})$$

where the NPs $\boldsymbol{\theta}$ are profiled for the best-fit $\hat{\mu}$ and a range over μ about $\hat{\mu}$, and $\mathcal{L}(\mathbf{x}|\hat{\mu}, \hat{\boldsymbol{\theta}})$ is the maximum likelihood. Likelihood scans are a good measure of how the PLRs compare between the $t\bar{t}H$ and VH categories according to the defined models, and how certain individual and sets of systematic uncertainties affect the width of the PLR. Along with the exclusion limits, the profile likelihood scans are performed both in light of the data, known as *observed* results, and fitting to data when the signal strength $\hat{\mu}$ is fixed to zero, which corresponds to the *expected* results in the presence of purely SM background processes.

Appendix C

Supporting material for the $t\bar{t}H$ and VH analysis

The following is a set of supporting material used to produce and verify the results and conclusions presented in Sec. 7. The pre-fit ℓ_{lost} , $Z \rightarrow \text{inv}$, QCD multijet, and signal event yields for each subcategory and hadronic recoil interval summed across 2016–2018 are tabulated along with observation in Table C.1. The uncertainties are inclusive of statistical and systematic contributions. The corresponding post-fit event yields predicted in a CR only, B -only fit, a CR+SR, B -only fit, and a CR+SR, $S + B$ fit are presented in Tables 7.1, C.2, and 7.2, respectively.

Toy distributions are generated to recreate a model that accurately represents the model developed in Ch. 6, but under slightly different constraints for comparison. For each constrained NP in the fit, there exists a corresponding Gaussian probability distribution that multiplies the likelihood function. This Gaussian distribution has a nuisance term and a constraint term, where the former affects the functional form of the likelihood function, while the latter is independent of it. The constraint term can be considered as an observed measurement of the NP, following the Model and Likelihood pages of Ref. [96]. Therefore, randomising the value of the constraint term while preserving the best-fit values for the nuisance term (and hence the NP values) will not affect the maximum-likelihood result but will change the underlying parametrisation. The distribution of χ^2_{sat} for these toy experiments is compared to the test statistic χ^2_{sat} measured for each year and across all years for the $t\bar{t}H$, VH, and combined categories. The toy distributions shown in Fig. C.1 correspond to the results in Table 7.4, and show the effect of statistical fluctuations as a result of disagreements between data and simulation in certain regions that share the same RP, and the worsening effect this can have on the resulting p -value.

Table C.1: Total pre-fit SR yields for each subcategory and recoil interval across 2016–2018, inclusive of statistical and systematic uncertainties, and the observed data for reference.

Subcategory	Hadronic recoil	ℓ_{lost}	$Z \rightarrow \text{inv}$	QCD	Total SM background	Data	Signal
ttH 1t1b	[200, 300)	254.4 ± 18.2	25.4 ± 3.3	23.1 ± 15.7	302.9 ± 24.2	288.0 ± 17.0	11.3 ± 0.8
	[300, 400)	243.1 ± 18.0	23.8 ± 1.9	5.2 ± 3.4	272.2 ± 18.4	257.0 ± 16.0	14.0 ± 1.0
	[400, 500)	104.1 ± 8.7	18.4 ± 1.6	0.9 ± 0.5	123.4 ± 8.9	145.0 ± 12.0	10.9 ± 0.7
	[500, ∞)	41.0 ± 4.3	19.0 ± 3.5	0.3 ± 0.0	60.3 ± 5.5	66.0 ± 8.1	10.0 ± 0.6
	[500, ∞)	339.3 ± 25.9	13.0 ± 1.4	10.9 ± 7.4	363.2 ± 26.9	298.0 ± 17.3	15.4 ± 1.2
ttH 1t2b	[200, 300)	298.1 ± 24.1	13.1 ± 1.0	2.5 ± 1.5	313.6 ± 24.2	299.0 ± 17.3	17.6 ± 1.4
	[300, 400)	105.5 ± 9.9	10.4 ± 0.9	0.4 ± 0.3	116.4 ± 10.0	136.0 ± 11.7	12.9 ± 1.0
	[400, 500)	36.8 ± 4.0	10.3 ± 1.5	0.2 ± 0.0	47.2 ± 4.3	53.0 ± 7.3	11.3 ± 0.9
	[500, ∞)	1839.3 ± 143.7	144.4 ± 25.6	18.8 ± 12.2	2002.5 ± 146.4	1819.0 ± 42.6	60.8 ± 4.1
	[500, ∞)	432.2 ± 35.8	62.4 ± 7.1	4.3 ± 2.8	498.9 ± 36.6	486.0 ± 22.0	31.3 ± 2.0
ttH 1W1b	[200, 300)	56.7 ± 5.3	24.3 ± 2.4	0.8 ± 0.4	81.8 ± 5.8	111.0 ± 10.5	10.3 ± 0.6
	[300, 400)	12.9 ± 1.7	10.1 ± 1.8	0.3 ± 0.0	23.3 ± 2.5	37.0 ± 6.1	4.6 ± 0.3
	[400, 500)	1471.9 ± 124.2	36.9 ± 4.8	5.6 ± 3.6	1514.4 ± 124.4	1276.0 ± 35.7	42.6 ± 3.6
	[500, ∞)	309.8 ± 27.8	17.9 ± 1.9	1.3 ± 0.9	329.0 ± 27.9	298.0 ± 17.3	21.2 ± 1.8
	[500, ∞)	37.2 ± 4.3	5.3 ± 0.4	0.2 ± 0.0	42.7 ± 4.3	47.0 ± 6.9	6.5 ± 0.5
ttH 1W2b	[200, 300)	7.8 ± 1.2	2.7 ± 0.5	0.1 ± 0.0	10.5 ± 1.3	17.0 ± 4.1	2.5 ± 0.0
	[300, 400)	13.8 ± 1.9	2.5 ± 0.3	0.3 ± 0.0	16.6 ± 1.9	14.0 ± 3.7	4.9 ± 0.4
	[400, 500)	7.8 ± 1.4	3.8 ± 0.3	0.1 ± 0.0	11.6 ± 1.4	15.0 ± 3.9	5.2 ± 0.4
	[500, ∞)	14.5 ± 2.1	2.4 ± 0.5	0.3 ± 0.0	17.2 ± 2.1	15.0 ± 3.9	4.2 ± 0.4
	[500, ∞)	6.1 ± 1.0	2.4 ± 0.0	0.1 ± 0.0	8.5 ± 1.0	6.0 ± 2.4	5.5 ± 0.4
ttH 2Boosted1b	[200, 300)	5009.9 ± 330.2	1562.2 ± 349.6	99.1 ± 62.6	6671.2 ± 484.9	7207.0 ± 84.9	158.3 ± 6.9
	[300, 400)	1038.5 ± 69.6	651.1 ± 82.6	22.5 ± 12.7	1712.1 ± 108.7	1907.0 ± 43.7	82.1 ± 3.6
	[400, 500)	174.4 ± 14.9	209.0 ± 18.2	4.0 ± 2.5	387.4 ± 23.7	427.0 ± 20.7	28.9 ± 1.6
	[500, ∞)	53.4 ± 7.6	123.1 ± 29.5	1.5 ± 0.8	178.1 ± 30.5	221.0 ± 14.9	16.2 ± 0.9
	[500, ∞)	1236.4 ± 91.8	315.8 ± 67.6	11.8 ± 7.9	1563.9 ± 114.3	1602.0 ± 40.0	70.0 ± 3.9
ttH 5j1b	[200, 300)	192.6 ± 15.4	133.1 ± 16.2	2.7 ± 1.6	328.4 ± 22.4	367.0 ± 19.2	32.6 ± 1.7
	[300, 400)	33.0 ± 3.4	42.0 ± 3.8	0.5 ± 0.3	75.5 ± 5.1	91.0 ± 9.5	10.3 ± 0.5
	[400, 500)	12.3 ± 1.9	21.2 ± 4.4	0.2 ± 0.0	33.6 ± 4.8	36.0 ± 6.0	5.4 ± 0.3
	[500, ∞)	3526.4 ± 234.8	725.8 ± 149.9	85.9 ± 53.7	4338.2 ± 283.6	4632.0 ± 68.1	134.4 ± 7.1
	[500, ∞)	796.8 ± 57.2	342.4 ± 40.7	19.5 ± 12.0	1158.7 ± 71.2	1371.0 ± 37.0	72.9 ± 3.5
ttH 6j1b	[200, 300)	147.2 ± 12.9	126.8 ± 11.0	3.4 ± 2.2	277.3 ± 17.1	312.0 ± 17.7	26.2 ± 1.3
	[300, 400)	55.1 ± 6.8	83.2 ± 20.2	1.3 ± 0.7	139.6 ± 21.3	197.0 ± 14.0	17.7 ± 0.8
	[400, 500)	1155.2 ± 87.9	207.7 ± 38.2	15.6 ± 10.2	1378.5 ± 96.4	1404.0 ± 37.5	87.4 ± 5.8
	[500, ∞)	213.2 ± 18.6	90.3 ± 9.7	3.6 ± 2.1	307.1 ± 21.1	341.0 ± 18.5	42.7 ± 2.8
	[500, ∞)	42.6 ± 4.4	31.7 ± 2.8	0.6 ± 0.3	74.8 ± 5.2	91.0 ± 9.5	15.2 ± 0.9
ttH 6j2b	[200, 300)	15.9 ± 2.1	18.8 ± 3.8	0.2 ± 0.0	34.9 ± 4.4	41.0 ± 6.4	8.4 ± 0.4
	[300, 400)	15753.5 ± 1788.9	27040.3 ± 5944.7	105.8 ± 74.7	42899.7 ± 6208.4	47559.0 ± 218.1	1874.6 ± 43.1
	[400, 500)	2166.6 ± 169.3	5137.6 ± 777.2	16.8 ± 10.1	7321.0 ± 795.5	8106.0 ± 90.0	446.5 ± 11.1
	[500, ∞)	214.2 ± 18.5	627.1 ± 44.8	2.8 ± 1.8	844.1 ± 48.5	938.0 ± 30.6	66.0 ± 1.8
	[500, ∞)	17.2 ± 2.7	68.0 ± 8.2	0.9 ± 0.5	86.1 ± 8.6	98.0 ± 9.9	6.0 ± 0.0
VH 2j1b	[200, 300)	2751.5 ± 155.5	2376.0 ± 468.0	26.2 ± 17.0	5153.7 ± 493.4	5883.0 ± 76.7	213.6 ± 5.7
	[300, 400)	321.1 ± 19.7	474.3 ± 58.4	3.6 ± 2.3	798.9 ± 61.7	949.0 ± 30.8	54.4 ± 1.5
	[400, 500)	31.2 ± 3.0	60.3 ± 4.5	0.6 ± 0.4	92.1 ± 5.4	120.0 ± 11.0	7.7 ± 0.3
	[500, ∞)	223.0 ± 14.9	462.1 ± 59.9	2.0 ± 1.6	687.1 ± 61.8	617.0 ± 24.8	112.1 ± 3.9
	[500, ∞)	34.9 ± 2.5	118.0 ± 9.0	0.2 ± 0.0	153.2 ± 9.4	128.0 ± 11.3	36.2 ± 1.3

Table C.2: Total post-fit SR yields for each subcategory and recoil interval across 2016–2018, obtained from a B -only fit using CR+SR predictions inclusive of statistical and systematic uncertainties, and the observed data for reference.

Subcategory	Hadronic recoil	ℓ_{test}	$Z \rightarrow \text{inv}$	QCD	Total SM background	Data
$t\bar{t}H$ 1t1b	[200, 300)	247.0 ± 8.0	32.0 ± 3.4	16.5 ± 7.6	295.5 ± 11.6	288.0 ± 17.0
	[300, 400)	228.0 ± 7.9	36.3 ± 4.1	3.8 ± 2.0	268.1 ± 9.1	257.0 ± 16.0
	[400, 500)	99.0 ± 4.7	35.8 ± 4.7	0.7 ± 0.3	135.5 ± 6.7	145.0 ± 12.0
	[500, ∞)	36.6 ± 2.6	33.3 ± 4.4	0.3 ± 0.0	70.1 ± 5.1	66.0 ± 8.1
	[500, ∞)	303.7 ± 9.6	16.6 ± 1.7	7.8 ± 4.0	328.1 ± 10.5	298.0 ± 17.3
$t\bar{t}H$ 1t2b	[200, 300)	265.1 ± 8.9	20.1 ± 2.3	1.9 ± 0.9	287.1 ± 9.3	299.0 ± 17.3
	[300, 400)	96.2 ± 4.6	19.9 ± 2.5	0.3 ± 0.0	116.5 ± 5.2	136.0 ± 11.7
	[400, 500)	35.3 ± 2.6	17.4 ± 2.3	0.1 ± 0.0	52.8 ± 3.5	53.0 ± 7.3
	[500, ∞)	1664.3 ± 36.3	177.7 ± 18.1	13.7 ± 7.5	1855.7 ± 41.2	1819.0 ± 42.6
	[500, ∞)	387.4 ± 12.0	94.4 ± 10.8	3.2 ± 1.6	485.0 ± 16.2	486.0 ± 22.0
$t\bar{t}H$ 1W1b	[200, 300)	55.6 ± 3.5	47.4 ± 6.2	0.6 ± 0.3	103.7 ± 7.1	111.0 ± 10.5
	[300, 400)	10.1 ± 1.2	19.2 ± 2.7	0.2 ± 0.0	29.5 ± 3.0	37.0 ± 6.1
	[400, 500)	1259.6 ± 29.0	47.9 ± 4.7	4.0 ± 1.9	1311.6 ± 29.4	1276.0 ± 35.7
	[500, ∞)	262.7 ± 9.4	27.7 ± 3.2	0.9 ± 0.4	291.3 ± 9.9	298.0 ± 17.3
	[500, ∞)	37.1 ± 2.8	10.3 ± 1.3	0.2 ± 0.0	47.6 ± 3.1	47.0 ± 6.9
$t\bar{t}H$ 2Boosted1b	[200, 300)	6.4 ± 0.9	4.7 ± 0.6	0.1 ± 0.0	11.3 ± 1.1	17.0 ± 4.1
	[300, 400)	16.7 ± 2.6	3.5 ± 0.3	0.2 ± 0.0	20.4 ± 2.6	14.0 ± 3.7
	[400, 500)	6.5 ± 1.4	6.3 ± 0.8	0.1 ± 0.0	12.9 ± 1.6	15.0 ± 3.9
	[500, ∞)	14.5 ± 2.3	3.3 ± 0.8	0.2 ± 0.0	18.0 ± 2.4	15.0 ± 3.9
	[500, ∞)	4.8 ± 1.0	3.8 ± 0.4	0.1 ± 0.0	8.6 ± 1.1	6.0 ± 2.4
$t\bar{t}H$ 2Boosted2b	[200, 300)	5260.3 ± 77.7	1757.6 ± 65.7	104.8 ± 76.9	7122.6 ± 127.6	7207.0 ± 84.9
	[300, 400)	1104.4 ± 23.6	834.7 ± 33.1	21.8 ± 14.7	1960.9 ± 43.2	1907.0 ± 43.7
	[400, 500)	177.3 ± 7.7	257.6 ± 14.0	3.9 ± 2.4	438.8 ± 16.2	427.0 ± 20.7
	[500, ∞)	52.9 ± 3.3	171.8 ± 10.9	1.5 ± 0.9	226.2 ± 11.5	221.0 ± 14.9
	[500, ∞)	1267.5 ± 30.4	356.3 ± 13.1	11.6 ± 7.1	1635.4 ± 33.9	1602.0 ± 40.0
$t\bar{t}H$ 5j1b	[200, 300)	187.1 ± 7.8	173.5 ± 7.2	2.7 ± 1.8	363.3 ± 10.7	367.0 ± 19.2
	[300, 400)	33.8 ± 3.2	52.0 ± 3.1	0.5 ± 0.3	86.3 ± 4.5	91.0 ± 9.5
	[400, 500)	7.9 ± 1.2	28.7 ± 2.1	0.2 ± 0.0	36.8 ± 2.5	36.0 ± 6.0
	[500, ∞)	3762.9 ± 56.9	824.2 ± 31.2	85.5 ± 58.1	4672.6 ± 87.1	4632.0 ± 68.1
	[500, ∞)	870.7 ± 20.6	441.7 ± 18.0	20.1 ± 13.3	1332.5 ± 30.4	1371.0 ± 37.0
$t\bar{t}H$ 5j2b	[200, 300)	171.2 ± 6.8	156.3 ± 8.8	3.4 ± 2.2	330.9 ± 11.4	312.0 ± 17.7
	[300, 400)	61.2 ± 3.8	116.6 ± 7.4	1.3 ± 1.0	179.1 ± 8.4	197.0 ± 14.0
	[400, 500)	1178.6 ± 26.1	238.9 ± 9.2	15.5 ± 11.4	1433.1 ± 29.9	1404.0 ± 37.5
	[500, ∞)	231.2 ± 9.3	117.1 ± 5.0	3.6 ± 2.4	351.9 ± 10.8	341.0 ± 18.5
	[500, ∞)	40.2 ± 3.6	39.1 ± 2.4	0.6 ± 0.4	79.9 ± 4.3	91.0 ± 9.5
$t\bar{t}H$ 6j1b	[200, 300)	12.5 ± 1.5	25.3 ± 1.8	0.2 ± 0.0	38.1 ± 2.4	41.0 ± 6.4
	[300, 400)	17856.0 ± 287.2	29539.4 ± 354.7	103.7 ± 63.2	47499.1 ± 460.7	47559.0 ± 218.1
	[400, 500)	2545.8 ± 57.6	5513.3 ± 89.4	16.6 ± 9.6	8075.7 ± 106.8	8106.0 ± 90.0
	[500, ∞)	278.6 ± 14.2	663.1 ± 22.6	2.7 ± 1.5	944.5 ± 26.7	938.0 ± 30.6
	[500, ∞)	19.3 ± 2.9	75.5 ± 5.9	0.9 ± 0.5	95.7 ± 6.6	98.0 ± 9.9
$t\bar{t}H$ 6j2b	[200, 300)	3060.6 ± 67.1	2718.4 ± 82.0	29.5 ± 33.4	5808.6 ± 111.1	5883.0 ± 76.7
	[300, 400)	349.9 ± 14.8	608.6 ± 26.0	3.9 ± 3.8	962.3 ± 30.1	949.0 ± 30.8
	[400, 500)	37.7 ± 4.3	72.9 ± 6.3	0.7 ± 0.7	111.3 ± 7.7	120.0 ± 11.0
	[500, ∞)	202.4 ± 12.3	415.7 ± 23.8	2.0 ± 1.3	620.1 ± 26.8	617.0 ± 24.8
	[500, ∞)	29.9 ± 3.3	100.9 ± 9.2	0.2 ± 0.0	131.1 ± 9.8	128.0 ± 11.3
VH 2j0b	[200, 300)	17856.0 ± 287.2	29539.4 ± 354.7	103.7 ± 63.2	47499.1 ± 460.7	47559.0 ± 218.1
	[300, 400)	2545.8 ± 57.6	5513.3 ± 89.4	16.6 ± 9.6	8075.7 ± 106.8	8106.0 ± 90.0
	[400, 500)	278.6 ± 14.2	663.1 ± 22.6	2.7 ± 1.5	944.5 ± 26.7	938.0 ± 30.6
	[500, ∞)	19.3 ± 2.9	75.5 ± 5.9	0.9 ± 0.5	95.7 ± 6.6	98.0 ± 9.9
	[500, ∞)	3060.6 ± 67.1	2718.4 ± 82.0	29.5 ± 33.4	5808.6 ± 111.1	5883.0 ± 76.7
VH 2j1b	[200, 300)	349.9 ± 14.8	608.6 ± 26.0	3.9 ± 3.8	962.3 ± 30.1	949.0 ± 30.8
	[300, 400)	37.7 ± 4.3	72.9 ± 6.3	0.7 ± 0.7	111.3 ± 7.7	120.0 ± 11.0
	[400, 500)	202.4 ± 12.3	415.7 ± 23.8	2.0 ± 1.3	620.1 ± 26.8	617.0 ± 24.8
	[500, ∞)	29.9 ± 3.3	100.9 ± 9.2	0.2 ± 0.0	131.1 ± 9.8	128.0 ± 11.3
	[500, ∞)	17856.0 ± 287.2	29539.4 ± 354.7	103.7 ± 63.2	47499.1 ± 460.7	47559.0 ± 218.1
VH 2j2b	[200, 300)	2545.8 ± 57.6	5513.3 ± 89.4	16.6 ± 9.6	8075.7 ± 106.8	8106.0 ± 90.0
	[300, 400)	278.6 ± 14.2	663.1 ± 22.6	2.7 ± 1.5	944.5 ± 26.7	938.0 ± 30.6
	[400, 500)	19.3 ± 2.9	75.5 ± 5.9	0.9 ± 0.5	95.7 ± 6.6	98.0 ± 9.9
	[500, ∞)	3060.6 ± 67.1	2718.4 ± 82.0	29.5 ± 33.4	5808.6 ± 111.1	5883.0 ± 76.7
	[500, ∞)	349.9 ± 14.8	608.6 ± 26.0	3.9 ± 3.8	962.3 ± 30.1	949.0 ± 30.8

A series of likelihood scans were performed with groups of systematic uncertainties frozen during the fit, to obtain post-fit profile likelihood parabolas for varying $\hat{\mu}$ values corresponding to $\mathcal{B}(\text{H} \rightarrow \text{inv})$. Observed likelihood fits from a CR+SR $S + B$ fit are presented in Fig. C.2, which are used as input for Table 7.5. The statistical uncertainty in the data is classed as *other* uncertainties in the figure, which is conventionally consistent with other $\text{H} \rightarrow \text{inv}$ analyses that contain NPs not classified into a particular group of systematic uncertainties. All NPs are grouped into systematic uncertainty groups for the $t\bar{t}\text{H}$ /resolved VH analysis. The statistical uncertainty in simulated samples, labelled MC, is classed as a systematic uncertainty given it scales with $\mathcal{L}_1$.

The largest pulls due to statistical and systematic uncertainties are also presented in Figs. C.3 and C.4, respectively, where those that feature have pulls of at least ± 0.6 in the statistical case, and at least ± 0.2 in the systematic case. The statistical uncertainties are those associated with simulated events, where the RPs are unconstrained and have no effect on the fit. The largest pulls on statistical uncertainties are in the VH 2j2b subcategory for 2017 and 2018 data sets, associated with the non-QCD background in the photon CR. These are split by process where in this particular subcategory and hadronic recoil interval, the total yield $n_{\text{tot}}^{\text{eff}} \leq n^{\text{threshold}}$. This is also true for $n_{\text{tot}}^{\text{eff}}$ calculated using the non-QCD background alone, and as such the NPs are Poisson-constrained, and sensitive to low numbers of simulated events in this case if the pre-fit statistical uncertainty is underestimated. The largest pulls on systematic uncertainties include observed in the jet flavour uncertainty within the JES, which are observed to have the largest impact on the $\mathcal{B}(\text{H} \rightarrow \text{inv})$ measurement, and the photon-to-Z $\rightarrow \text{inv}$ mapping uncertainty in the VH 2j2b subcategory for 2017 and 2018, which appears fairly constrained outside of one standard deviation of the pre-fit value. The behaviour is not unique to this particular photon-to-Z $\rightarrow \text{inv}$ uncertainty, and is expected given how conservative the assigned 40% value is, and therefore post-fit this uncertainty should be largely constrained about some point that is still consistent with the pre-fit input, although the discrepancy between the pre- and post-fit value can be due to the disagreement between the ratio of data-to-simulation in the dielectron and dimuon CRs.




Figure C.1: The toy distributions generated to calculate the GF p -values presented in Table 7.4, ordered by (left-to-right) $t\bar{t}H$, VH, combined $t\bar{t}H$ and VH, and (top-to-bottom) 2016, 2017, 2018, and all years.




Figure C.2: The profile likelihood scans as obtained to measure the impacts of each systematic group on the best fit $\mathcal{B}(H \rightarrow \text{inv})$ value, $\hat{\mu}$, in the observed case.




Figure C.3: A selection of the largest pulls on statistical uncertainties during the fit to both the $t\bar{t}H$ and VH categories.




Figure C.4: A selection of the largest pulls on systematic uncertainties during the fit to both the $t\bar{t}H$ and VH categories.

The pulls on the systematic uncertainties for the $t\bar{t}H$ and VH category individually, obtained from a CR+SR $S + B$ fit, are presented in Fig. C.5, where the correlations between the categories are easily identified. The pull value defined in Eq. 3.11 for AutoDQM is measured in this case, which is the difference in the pulls divided by the quadrature sum of the errors. Pulls lying along a given axis correspond to systematic uncertainties applicable to only the $t\bar{t}H$ or VH category. NPs with pull values greater than 0.45 are labelled, namely the uncertainty in b jet tagging and in the lepton veto for 2017, which are just under one standard deviation from their prefit value and are pulled in opposite directions for each category.



Figure C.5: The pull values as defined in Eq. 3.11 using the pulls on each NP in either the $t\bar{t}H$ and VH category following a CR+SR $S + B$ fit, with the correlation between the systematic uncertainties indicated by their presence in each quadrant. NPs with pull values greater than 0.45 are labelled, and those lying on a particular axis are applicable to only one category.

Appendix D

Supporting material for the combination

maximum-likelihood fits are performed for each Higgs boson production mode considered in the combination. In some cases, due to the data-taking era from which the likelihood functions were first defined, not all other production modes are considered in the fit, namely in analyses targeting the ggH channel, where the $t\bar{t}H$ topology is not accounted for at all as the analysis pre-dates any of the $t\bar{t}H$ searches. The relative contamination of the VBF, VH, $t\bar{t}H$, and ggH signal samples in each of the VBF, VH, $t\bar{t}H$, and ggH categories as defined by each analysis targeting these topologies is presented in Table D.1.

Table D.1: The relative signal contributions within each tagged channel by production mode. Note that not all production modes are considered for some analyses, for example the $t\bar{t}H$ topology is not considered in analyses for which ggH is the target channel.

Process	VH-tagged	ggH-tagged	VBF-tagged	$t\bar{t}H$ -tagged
VH signal	0.563	0.111	0.101	0.095
ggH signal	0.398	0.678	0.703	0.164
VBF signal	0.035	0.211	0.194	0.020
$t\bar{t}H$ signal	0.004	—	0.001	0.721

The 95% CL upper limits across all Run 1 and Run 2 data sets analysed by the CMS experiment are broken down finely into the various hadronic and leptonic final states listed in the second column of Table 8.1, and presented in Fig. D.1. The VBF

final state provides the greatest sensitivity to $\mathcal{B}(H \rightarrow \text{inv})$, as well as is the largest driver of the data excess of two standard deviations in the combined exclusion limit.




Figure D.1: Left: Observed and expected limits at 95% CL for the combination, broken down by final state according to the listing in Table 8.1.

Bibliography

- [1] Anja Beck, Thomas Blake, and Michal Kreps. “Angular distribution of $\Lambda_b^0 \rightarrow pK^-\ell^+\ell^-$ decays comprising Λ resonances with spin $\leq 5/2$ ”. In: *JHEP* 02 (2023), p. 189. DOI: 10.1007/JHEP02(2023)189. arXiv: 2210.09988 [hep-ph].
- [2] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (1964), p. 321. DOI: 10.1103/PhysRevLett.13.321.
- [3] Peter W. Higgs. “Broken symmetries, massless particles and gauge fields”. In: *Phys. Lett.* 12 (1964), p. 132. DOI: 10.1016/0031-9163(64)91136-9.
- [4] Peter W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (1964), p. 508. DOI: 10.1103/PhysRevLett.13.508.
- [5] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (1964), p. 585. DOI: 10.1103/PhysRevLett.13.585.
- [6] Peter W. Higgs. “Spontaneous Symmetry Breakdown without Massless Bosons”. In: *Phys. Rev.* 145 (1966), p. 1156. DOI: 10.1103/PhysRev.145.1156.
- [7] T. W. B. Kibble. “Symmetry breaking in non-Abelian gauge theories”. In: *Phys. Rev.* 155 (1967). Ed. by J. C. Taylor, p. 1554. DOI: 10.1103/PhysRev.155.1554.
- [8] Armen Tumasyan et al. “Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 105 (2022), p. 092007. DOI: 10.1103/PhysRevD.105.092007. arXiv: 2201.11585 [hep-ex].
- [9] Georges Aad et al. “Search for invisible Higgs-boson decays in events with vector-boson fusion signatures using 139 fb⁻¹ of proton-proton data recorded by the ATLAS experiment”. In: *JHEP* 08 (2022), p. 104. DOI: 10.1007/JHEP08(2022)104. arXiv: 2202.07953 [hep-ex].

- [10] “A search for decays of the Higgs boson to invisible particles in events with a top-antitop quark pair or a vector boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: (Mar. 2023). arXiv: 2303.01214 [hep-ex].
- [11] Armen Tumasyan et al. “Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 11 (2021), p. 153. DOI: 10.1007/JHEP11(2021)153. arXiv: 2107.13021 [hep-ex].
- [12] Albert M Sirunyan et al. “Search for direct top squark pair production in events with one lepton, jets, and missing transverse momentum at 13 TeV with the CMS experiment”. In: *JHEP* 05 (2020), p. 032. DOI: 10.1007/JHEP05(2020)032. arXiv: 1912.08887 [hep-ex].
- [13] Albert M Sirunyan et al. “Search for top squark pair production using dilepton final states in pp collision data collected at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 81.1 (2021), p. 3. DOI: 10.1140/epjc/s10052-020-08701-5. arXiv: 2008.05936 [hep-ex].
- [14] Armen Tumasyan et al. “Combined searches for the production of supersymmetric top quark partners in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 81.11 (2021), p. 970. DOI: 10.1140/epjc/s10052-021-09721-5. arXiv: 2107.10892 [hep-ex].
- [15] Albert M Sirunyan et al. “Search for dark matter produced in association with a leptonically decaying Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 81.1 (2021). [Erratum: *Eur.Phys.J.C* 81, 333 (2021)], p. 13. DOI: 10.1140/epjc/s10052-020-08739-5. arXiv: 2008.04735 [hep-ex].
- [16] Vardan Khachatryan et al. “Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7, 8$, and 13 TeV”. In: *JHEP* 02 (2017), p. 135. DOI: 10.1007/JHEP02(2017)135. arXiv: 1610.09218 [hep-ex].
- [17] S. L. Glashow. “Partial Symmetries of Weak Interactions”. In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.
- [18] Abdus Salam. “Weak and Electromagnetic Interactions”. In: *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034.
- [19] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.
- [20] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

- [21] Roel Aaij et al. “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays”. In: *Phys. Rev. Lett.* 115 (2015), p. 072001. DOI: 10.1103/PhysRevLett.115.072001. arXiv: 1507.03414 [hep-ex].
- [22] Roel Aaij et al. “Observation of the resonant character of the $Z(4430)^-$ state”. In: *Phys. Rev. Lett.* 112.22 (2014), p. 222002. DOI: 10.1103/PhysRevLett.112.222002. arXiv: 1404.1903 [hep-ex].
- [23] Eleonora Di Valentino, Stefano Gariazzo, and Olga Mena. “Most constraining cosmological neutrino mass bounds”. In: *Phys. Rev. D* 104.8 (2021), p. 083504. DOI: 10.1103/PhysRevD.104.083504. arXiv: 2106.15267 [astro-ph.CO].
- [24] C. S. Wu et al. “Experimental Test of Parity Conservation in β Decay”. In: *Phys. Rev.* 105 (1957), pp. 1413–1414. DOI: 10.1103/PhysRev.105.1413.
- [25] J. J. Thomson. “Cathode rays”. In: *Phil. Mag. Ser. 5* 44 (1897), pp. 293–316. DOI: 10.1080/14786449708621070.
- [26] E. Rutherford. “Collision of α particles with light atoms. IV. An anomalous effect in nitrogen”. In: *Phil. Mag. Ser. 6* 37 (1919), pp. 581–587. DOI: 10.1080/14786431003659230.
- [27] S. H. Neddermeyer and C. D. Anderson. “Note on the Nature of Cosmic Ray Particles”. In: *Phys. Rev.* 51 (1937), pp. 884–886. DOI: 10.1103/PhysRev.51.884.
- [28] Frederick Reines and Clyde L. Cowan. “The neutrino”. In: *Nature* 178 (1956), pp. 446–449. DOI: 10.1038/178446a0.
- [29] Elliott D. Bloom et al. “High-Energy Inelastic $e p$ Scattering at 6-Degrees and 10-Degrees”. In: *Phys. Rev. Lett.* 23 (1969), pp. 930–934. DOI: 10.1103/PhysRevLett.23.930.
- [30] Martin Breidenbach et al. “Observed behavior of highly inelastic electron-proton scattering”. In: *Phys. Rev. Lett.* 23 (1969), pp. 935–939. DOI: 10.1103/PhysRevLett.23.935.
- [31] J. J. Aubert et al. “Experimental Observation of a Heavy Particle J ”. In: *Phys. Rev. Lett.* 33 (1974), pp. 1404–1406. DOI: 10.1103/PhysRevLett.33.1404.
- [32] J. E. Augustin et al. “Discovery of a Narrow Resonance in e^+e^- Annihilation”. In: *Phys. Rev. Lett.* 33 (1974), pp. 1406–1408. DOI: 10.1103/PhysRevLett.33.1406.

- [33] Martin L. Perl et al. “Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation”. In: *Phys. Rev. Lett.* 35 (1975), pp. 1489–1492. DOI: 10.1103/PhysRevLett.35.1489.
- [34] S. W. Herb et al. “Observation of a Dimuon Resonance at 9.5-GeV in 400-GeV Proton-Nucleus Collisions”. In: *Phys. Rev. Lett.* 39 (1977), pp. 252–255. DOI: 10.1103/PhysRevLett.39.252.
- [35] F. Abe et al. “Observation of top quark production in $\bar{p}p$ collisions”. In: *Phys. Rev. Lett.* 74 (1995), pp. 2626–2631. DOI: 10.1103/PhysRevLett.74.2626. arXiv: hep-ex/9503002.
- [36] S. Abachi et al. “Observation of the top quark”. In: *Phys. Rev. Lett.* 74 (1995), pp. 2632–2637. DOI: 10.1103/PhysRevLett.74.2632. arXiv: hep-ex/9503003.
- [37] D. P. Barber et al. “Discovery of Three Jet Events and a Test of Quantum Chromodynamics at PETRA Energies”. In: *Phys. Rev. Lett.* 43 (1979), p. 830. DOI: 10.1103/PhysRevLett.43.830.
- [38] J. J. Aubert et al. “The ratio of the nucleon structure functions F_2^n for iron and deuterium”. In: *Phys. Lett. B* 123 (1983), pp. 275–278. DOI: 10.1016/0370-2693(83)90437-9.
- [39] G. Arnison et al. “Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/ c^2 at the CERN SPS Collider”. In: *Phys. Lett. B* 126 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.
- [40] Georges Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [41] Serguei Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].
- [42] R. Matsumoto et al. “Search for proton decay via $p \rightarrow \mu + K_0$ in 0.37 megaton-years exposure of Super-Kamiokande”. In: *Phys. Rev. D* 106.7 (2022), p. 072003. DOI: 10.1103/PhysRevD.106.072003. arXiv: 2208.13188 [hep-ex].
- [43] Emmy Noether. “Invariant Variation Problems”. In: *Gott. Nachr.* 1918 (1918), pp. 235–257. DOI: 10.1080/00411457108231446. arXiv: physics/0503066.
- [44] J. D. Bjorken. “Asymptotic Sum Rules at Infinite Momentum”. In: *Phys. Rev.* 179 (1969), pp. 1547–1553. DOI: 10.1103/PhysRev.179.1547.

- [45] Yuri L. Dokshitzer. “Calculation of the Structure Functions for Deep Inelastic Scattering and $e^+ e^-$ Annihilation by Perturbation Theory in Quantum Chromodynamics.” In: *Sov. Phys. JETP* 46 (1977), pp. 641–653.
- [46] Guido Altarelli and G. Parisi. “Asymptotic Freedom in Parton Language”. In: *Nucl. Phys. B* 126 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4.
- [47] V. N. Gribov and L. N. Lipatov. “Deep inelastic $e p$ scattering in perturbation theory”. In: *Sov. J. Nucl. Phys.* 15 (1972), pp. 438–450.
- [48] G. Curci, W. Furmanski, and R. Petronzio. “Evolution of Parton Densities Beyond Leading Order: The Nonsinglet Case”. In: *Nucl. Phys. B* 175 (1980), pp. 27–92. DOI: 10.1016/0550-3213(80)90003-6.
- [49] S. Moch, J. A. M. Vermaseren, and A. Vogt. “The QCD splitting functions at three loops: Methods and results”. In: *Nucl. Phys. B Proc. Suppl.* 135 (2004). Ed. by He-Sheng Chen et al., pp. 137–146. DOI: 10.1016/j.nuclphysbps.2004.09.057. arXiv: hep-ph/0408075.
- [50] “LHC Design Report Vol.1: The LHC Main Ring”. In: (June 2004). Ed. by Oliver S. Bruning et al. DOI: 10.5170/CERN-2004-003-V-1.
- [51] S. Chatrchyan et al. “The CMS experiment at the CERN LHC”. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.
- [52] Roel Aaij et al. “Measurement of the inelastic pp cross-section at a centre-of-mass energy of 13 TeV”. In: *JHEP* 06 (2018), p. 100. DOI: 10.1007/JHEP06(2018)100. arXiv: 1803.10974 [hep-ex].
- [53] Albert M Sirunyan et al. “Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS”. In: *Eur. Phys. J. C* 81.9 (2021), p. 800. DOI: 10.1140/epjc/s10052-021-09538-2. arXiv: 2104.01927 [hep-ex].
- [54] Christoph Schwick. *Public CMS Luminosity Information*. July 2022. URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>.
- [55] CMS Collaboration. *Technical proposal for the Phase-II upgrade of the Compact Muon Solenoid*. CMS Technical Proposal CERN-LHCC-2015-010, CMS-TDR-15-02. 2015. URL: <http://cds.cern.ch/record/2020886>.
- [56] Grégory Soyez. “Pileup mitigation at the LHC: A theorists view”. In: *Phys. Rept.* 803 (2019), pp. 1–158. DOI: 10.1016/j.physrep.2019.01.007. arXiv: 1801.09721 [hep-ph].

- [57] M. Aaboud et al. “Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2”. In: *Eur. Phys. J. C* 77.10 (2017), p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex].
- [58] Serguei Chatrchyan et al. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. In: *JINST* 9.10 (2014), P10009. DOI: 10.1088/1748-0221/9/10/P10009. arXiv: 1405.6569 [physics.ins-det].
- [59] Albert M Sirunyan et al. “Pileup mitigation at CMS in 13 TeV data”. In: *JINST* 15.09 (2020), P09018. DOI: 10.1088/1748-0221/15/09/P09018. arXiv: 2003.00503 [hep-ex].
- [60] Daniele Bertolini et al. “Pileup Per Particle Identification”. In: *JHEP* 10 (2014), p. 059. DOI: 10.1007/JHEP10(2014)059. arXiv: 1407.6013 [hep-ph].
- [61] CMS Collaboration. *The CMS tracker: addendum to the Technical Design Report*. Tech. rep. Geneva, 2000. URL: <https://cds.cern.ch/record/490194>.
- [62] Paolo Azzurri. “The CMS silicon strip tracker”. In: *J. Phys. Conf. Ser.* 41 (2006). Ed. by R. C. Johnson, pp. 127–134. DOI: 10.1088/1742-6596/41/1/011. arXiv: physics/0512097.
- [63] W. Adam et al. “The CMS Phase-1 Pixel Detector Upgrade”. In: *JINST* 16 (2021), P02027. DOI: 10.1088/1748-0221/16/02/P02027. arXiv: 2012.14304 [physics.ins-det].
- [64] Armen Tumasyan et al. “Strategies and performance of the CMS silicon tracker alignment during LHC Run 2”. In: *Nucl. Instrum. Meth. A* 1037 (2022), p. 166795. DOI: 10.1016/j.nima.2022.166795. arXiv: 2111.08757 [physics.ins-det].
- [65] CMS Collaboration. *Track impact parameter resolution for the full pseudo rapidity coverage in the 2017 dataset with the CMS Phase-1 Pixel detector*. CMS Detector Performance Summary CMS-DP-2020-049. 2020. URL: <https://cds.cern.ch/record/2743740>.
- [66] Cristina Biino. “The CMS Electromagnetic Calorimeter: overview, lessons learned during Run 1 and future projections”. In: *J. Phys. Conf. Ser.* 587.1 (2015), p. 012001. DOI: 10.1088/1742-6596/587/1/012001.
- [67] P. Adzic et al. “Energy resolution of the barrel of the CMS electromagnetic calorimeter”. In: *JINST* 2 (2007), P04004. DOI: 10.1088/1748-0221/2/04/P04004.

- [68] Albert M Sirunyan et al. “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”. In: *JINST* 16.05 (2021), P05014. DOI: 10.1088/1748-0221/16/05/P05014. arXiv: 2012.06888 [hep-ex].
- [69] CMS Collaboration. “CMS ECAL Response to Laser Light”. In: (2019). URL: <https://cds.cern.ch/record/2668200>.
- [70] CMS Collaboration. *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*. Tech. rep. Sept. 2012. DOI: 10.2172/1151651.
- [71] S Chatrchyan et al. “Performance of the CMS Hadron Calorimeter with Cosmic Ray Muons and LHC Beam Data”. In: *JINST* 5 (2010), T03012. DOI: 10.1088/1748-0221/5/03/T03012. arXiv: 0911.4991 [physics.ins-det].
- [72] S. Abdullin et al. “Design, performance, and calibration of CMS hadron-barrel calorimeter wedges”. In: *Eur. Phys. J. C* 55.1 (2008), pp. 159–171. DOI: 10.1140/epjc/s10052-008-0573-y.
- [73] G. Bayatian et al. “Design, performance and calibration of the CMS forward calorimeter wedges”. In: *Eur. Phys. J. C* 53 (2008), pp. 139–166. DOI: 10.1140/epjc/s10052-007-0459-4.
- [74] A. M. Sirunyan et al. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JINST* 13.06 (2018), P06015. DOI: 10.1088/1748-0221/13/06/P06015. arXiv: 1804.04528 [physics.ins-det].
- [75] CMS Collaboration. *The Phase-2 Upgrade of the CMS Muon Detectors*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: <https://cds.cern.ch/record/2283189>.
- [76] Hyejin Kwon. “Performance of the CMS Level-1 Trigger during Run 2”. In: *PoS ICHEP2020* (2021), p. 806. DOI: 10.22323/1.390.0806.
- [77] L. Cadamuro. “The CMS Level-1 trigger system for LHC Run II”. In: *JINST* 12.03 (2017). Ed. by Pietro Govoni et al., p. C03021. DOI: 10.1088/1748-0221/12/03/C03021.
- [78] N. Dev et al. “The CMS Level-1 electron and photon trigger: for Run II of LHC”. In: *JINST* 12.02 (2017), p. C02014. DOI: 10.1088/1748-0221/12/02/C02014.
- [79] B. Kreis et al. “Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger”. In: *JINST* 11.01 (2016), p. C01051. DOI: 10.1088/1748-0221/11/01/C01051. arXiv: 1511.05855 [physics.ins-det].

- [80] Hale Sert. “CMS Run 2 High Level Trigger Performance”. In: *PoS EPS-HEP2019* (2020), p. 165. DOI: 10.22323/1.364.0165.
- [81] Vardan Khachatryan et al. “The CMS trigger system”. In: *JINST* 12.01 (2017), P01020. DOI: 10.1088/1748-0221/12/01/P01020. arXiv: 1609.02366 [physics.ins-det].
- [82] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- k_t jet clustering algorithm”. In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].
- [83] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “FastJet User Manual”. In: *Eur. Phys. J. C* 72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].
- [84] A. M. Sirunyan et al. “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”. In: *JINST* 13.05 (2018), P05011. DOI: 10.1088/1748-0221/13/05/P05011. arXiv: 1712.07158 [physics.ins-det].
- [85] A. M. Sirunyan et al. “Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003. arXiv: 1706.04965 [physics.ins-det].
- [86] Chad Freer. *AutoDQM: A Statistical Tool*. Sept. 2022. URL: <https://autodqm-official.readthedocs.io/en/latest/>.
- [87] Samuel May. *AutoDQM ML Introduction*. Aug. 2022. URL: <https://autodqm-ml.readthedocs.io/en/latest/>.
- [88] The CMS Collaboration. 2023. URL: <https://github.com/cms-nanoAOD/nanoAOD-tools>.
- [89] Marco Peruzzi, Giovanni Petrucciani, and Andrea Rizzi. “The NanoAOD event data format in CMS”. In: *J. Phys. Conf. Ser.* 1525.1 (2020), p. 012038. DOI: 10.1088/1742-6596/1525/1/012038.
- [90] Ben Krikler. *Faster Analysis Software Taskforce (F.A.S.T.)* 2019. URL: <https://fast-hep.web.cern.ch/fast-hep/>.
- [91] Ben Krikler. *Files · MASTER · Fast-Hep / public / fast-curator · GITLAB*. 2019. URL: <https://gitlab.cern.ch/fast-hep/public/fast-curator/-/tree/master>.
- [92] Ben Krikler. 2019. URL: <https://fast-carpenter.readthedocs.io/en/latest/>.

- [93] Ben Krikler. *Files · MASTER · Fast-Hep / public / fast-flow · GITLAB*. 2019. URL: <https://gitlab.cern.ch/fast-hep/public/fast-flow/-/tree/master>.
- [94] Ben Krikler. *Fast-plotter*. 2019. URL: <https://fast-plotter.readthedocs.io/en/latest/>.
- [95] Olivier Davignon. *README.rst · master · Fast-Hep / public / fast-datacard · GITLAB*. 2020. URL: <https://gitlab.cern.ch/fast-hep/public/fast-datacard/blob/master/README.rst>.
- [96] Nick Wardle. *Higgs Analysis Combine Limit Introduction*. July 2022. URL: <https://cms-analysis.github.io/HiggsAnalysis-CombinedLimit/>.
- [97] 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”. In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].
- [98] Carlo Oleari. “The POWHEG-BOX”. In: *Nucl. Phys. B Proc. Suppl.* 205-206 (2010). Ed. by Johannes Blümlein, Sven-Olaf Moch, and Tord Riemann, pp. 36–41. DOI: 10.1016/j.nuclphysbps.2010.08.016. arXiv: 1007.3893 [hep-ph].
- [99] Rikkert Frederix and Stefano Frixione. “Merging meets matching in MC@NLO”. In: *JHEP* 12 (2012), p. 061. DOI: 10.1007/JHEP12(2012)061. arXiv: 1209.6215 [hep-ph].
- [100] Torbjörn Sjöstrand et al. “An introduction to PYTHIA 8.2”. In: *Comput. Phys. Commun.* 191 (2015), pp. 159–177. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].
- [101] Albert M Sirunyan et al. “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”. In: *Eur. Phys. J. C* 80.1 (2020), p. 4. DOI: 10.1140/epjc/s10052-019-7499-4. arXiv: 1903.12179 [hep-ex].
- [102] Richard D. Ball et al. “Parton distributions from high-precision collider data”. In: *Eur. Phys. J. C* 77.10 (2017), p. 663. DOI: 10.1140/epjc/s10052-017-5199-5. arXiv: 1706.00428 [hep-ph].
- [103] M. Asai. “Geant4-a simulation toolkit”. In: *Trans. Amer. Nucl. Soc.* 95 (2006), p. 757.

- [104] H. Hartanto et al. “Higgs boson production in association with top quarks in the POWHEG BOX”. In: *Phys. Rev. D* 91 (2015), p. 094003. DOI: 10.1103/PhysRevD.91.094003. arXiv: 1501.04498 [hep-ph].
- [105] P. Nason and C. Oleari. “NLO Higgs boson production via vector-boson fusion matched with shower in POWHEG”. In: *JHEP* 02 (2010), p. 037. DOI: 10.1007/JHEP02(2010)037. arXiv: 0911.5299 [hep-ph].
- [106] G. Luisoni et al. “ $HW^\pm/HZ + 0$ and 1 jet at NLO with the POWHEG BOX interfaced to GoSam and their merging within MiNLO”. In: *JHEP* 10 (2013), p. 083. DOI: 10.1007/JHEP10(2013)083. arXiv: 1306.2542 [hep-ph].
- [107] E. Bagnaschi et al. “Higgs production via gluon fusion in the POWHEG approach in the SM and in the MSSM”. In: *JHEP* 02 (2012), p. 088. DOI: 10.1007/JHEP02(2012)088. arXiv: 1111.2854 [hep-ph].
- [108] D. de Florian et al. “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector”. In: 2/2017 (Oct. 2016). DOI: 10.23731/CYRM-2017-002. arXiv: 1610.07922 [hep-ph].
- [109] Michelangelo L. Mangano et al. “Matching matrix elements and shower evolution for top-quark production in hadronic collisions”. In: *JHEP* 01 (2007), p. 013. DOI: 10.1088/1126-6708/2007/01/013. arXiv: hep-ph/0611129.
- [110] J. M. Lindert et al. “Precise predictions for $V + \text{jets}$ dark matter backgrounds”. In: *Eur. Phys. J. C* 77 (2017), p. 829. DOI: 10.1140/epjc/s10052-017-5389-1. arXiv: 1705.04664 [hep-ph].
- [111] Silvia Zanolini et al. “Next-to-next-to-leading order event generation for VH production with $H \rightarrow b\bar{b}$ decay”. In: *JHEP* 07 (2022), p. 008. DOI: 10.1007/JHEP07(2022)008. arXiv: 2112.04168 [hep-ph].
- [112] J. M. Campbell et al. “Top-Pair Production and Decay at NLO Matched with Parton Showers”. In: *JHEP* 04 (2015), p. 114. DOI: 10.1007/JHEP04(2015)114. arXiv: 1412.1828 [hep-ph].
- [113] Simone Alioli et al. “NLO single-top production matched with shower in POWHEG: s - and t -channel contributions”. In: *JHEP* (2009). [Erratum: 10.1007/JHEP02(2010)011]. DOI: 10.1088/1126-6708/2009/09/111. arXiv: 0907.4076 [hep-ph].
- [114] Michal Czakon et al. “Top-pair production at the LHC through NNLO QCD and NLO EW”. In: *JHEP* 10 (2017), p. 186. DOI: 10.1007/JHEP10(2017)186. arXiv: 1705.04105 [hep-ph].

- [115] Pierre Artoisenet et al. “Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations”. In: *JHEP* 03 (2013), p. 015. DOI: 10.1007/JHEP03(2013)015. arXiv: 1212.3460 [hep-ph].
- [116] E. Re. “Single-top Wt-channel production matched with parton showers using the POWHEG method”. In: *Eur. Phys. J. C* 71 (2011), p. 1547. DOI: 10.1140/epjc/s10052-011-1547-z. arXiv: 1009.2450 [hep-ph].
- [117] T. Melia et al. “ W^+W^- , WZ and ZZ production in the POWHEG BOX”. In: *JHEP* 11 (2011), p. 078. DOI: 10.1007/JHEP11(2011)078. arXiv: 1107.5051 [hep-ph].
- [118] A. M. Sirunyan et al. “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV”. In: *JINST* 13.10 (2018), P10005. DOI: 10.1088/1748-0221/13/10/P10005. arXiv: 1809.02816 [hep-ex].
- [119] V. Khachatryan et al. “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”. In: *JINST* 12 (2017), P02014. DOI: 10.1088/1748-0221/12/02/P02014. arXiv: 1607.03663 [hep-ex].
- [120] CMS Collaboration. *Study of Pileup Removal Algorithms for Jets*. CMS Physics Analysis Summary CMS-PAS-JME-14-001. 2014. URL: <https://cds.cern.ch/record/1751454>.
- [121] CMS Collaboration. *Pileup Jet Identification*. Tech. rep. Geneva: CERN, 2013. URL: <https://cds.cern.ch/record/1581583>.
- [122] Matteo Cacciari and Gavin P. Salam. “Pileup subtraction using jet areas”. In: *Phys. Lett. B* 659 (2008), pp. 119–126. DOI: 10.1016/j.physletb.2007.09.077. arXiv: 0707.1378 [hep-ph].
- [123] Andrew J. Larkoski et al. “Soft Drop”. In: *JHEP* 05 (2014), p. 146. DOI: 10.1007/JHEP05(2014)146. arXiv: 1402.2657 [hep-ph].
- [124] Mrinal Dasgupta et al. “Towards an understanding of jet substructure”. In: *JHEP* 09 (2013), p. 029. DOI: 10.1007/JHEP09(2013)029. arXiv: 1307.0007 [hep-ph].
- [125] Jonathan M. Butterworth et al. “Jet substructure as a new Higgs search channel at the LHC”. In: *Phys. Rev. Lett.* 100 (2008), p. 242001. DOI: 10.1103/PhysRevLett.100.242001. arXiv: 0802.2470 [hep-ph].

- [126] Jesse Thaler and Ken Van Tilburg. “Identifying Boosted Objects with N-subjettiness”. In: *JHEP* 03 (2011), p. 015. DOI: 10.1007/JHEP03(2011)015. arXiv: 1011.2268 [hep-ph].
- [127] A. M. Sirunyan et al. “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”. In: *JINST* 15 (2020), P06005. DOI: 10.1088/1748-0221/15/06/P06005. arXiv: 2004.08262 [hep-ex].
- [128] CMS Collaboration. “CMS ECAL with 2017 data”. In: (2018). URL: <https://cds.cern.ch/record/2319285>.
- [129] Albert M Sirunyan et al. “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JINST* 15.10 (2020), P10017. DOI: 10.1088/1748-0221/15/10/P10017. arXiv: 2006.10165 [hep-ex].
- [130] Glen Cowan. “Two developments in discovery tests: use of weighted Monte Carlo events and an improved measure of experimental sensitivity”. Progress on Statistical Issues in Searches, SLAC. 2012. URL: https://www-conf.slac.stanford.edu/statisticalissues2012/talks/glen_cowan_slac_4jun12.pdf.
- [131] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].
- [132] Olivier Davignon. *F.A.S.T. Caliper package*. 2020. URL: <https://gitlab.cern.ch/fast-hep/public/fast-caliper/-/tree/master>.
- [133] Vardan Khachatryan et al. “Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV”. In: *JINST* 10.06 (2015), P06005. DOI: 10.1088/1748-0221/10/06/P06005. arXiv: 1502.02701 [physics.ins-det].
- [134] T. Sakuma, H. Flaecher, and D. Smith. “Alternative angular variables for suppression of QCD multijet events in new physics searches with missing transverse momentum at the LHC”. Mar. 2018. DOI: 10.48550/arXiv.1803.07942. arXiv: 1803.07942 [hep-ph].
- [135] LHC Higgs Cross Section Working Group. *Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector*. CERN Report CERN-2017-002-M. 2016. DOI: 10.23731/CYRM-2017-002. arXiv: 1610.07922 [hep-ph].

- [136] Albert M Sirunyan et al. “Measurement of the Higgs boson production rate in association with top quarks in final states with electrons, muons, and hadronically decaying tau leptons at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 81.4 (2021), p. 378. DOI: 10.1140/epjc/s10052-021-09014-x. arXiv: 2011.03652 [hep-ex].
- [137] CMS Collaboration. *CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2621960>.
- [138] CMS Collaboration. *CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2676164>.
- [139] J. S. Conway. “Incorporating Nuisance Parameters in Likelihoods for Multi-source Spectra”. In: *PHYSTAT 2011*. 2011, pp. 115–120. DOI: 10.5170/CERN-2011-006.115. arXiv: 1103.0354 [physics.data-an].
- [140] Alexander L. Read. “Presentation of search results: The CL_s technique”. In: *J. Phys. G* 28 (2002), p. 2693. DOI: 10.1088/0954-3899/28/10/313.
- [141] Thomas Junk. “Confidence level computation for combining searches with small statistics”. In: *Nucl. Instrum. Meth. A* 434 (1999), p. 435. DOI: 10.1016/S0168-9002(99)00498-2. arXiv: hep-ex/9902006 [hep-ex].
- [142] G. Cowan. *Statistical data analysis*. 1998. ISBN: 978-0-19-850156-5.
- [143] Vardan Khachatryan et al. “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”. In: *Eur. Phys. J. C* 75 (2015), p. 212. DOI: 10.1140/epjc/s10052-015-3351-7. arXiv: 1412.8662 [hep-ex].
- [144] ATLAS and CMS Collaborations, and 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. URL: <https://cds.cern.ch/record/1379837>.
- [145] CMS Collaboration. *First constraints on invisible Higgs boson decays using $t\bar{t}H$ production at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2668677>.

- [146] G. Aad et al. “Constraints on spin-0 dark matter mediators and invisible Higgs decays using ATLAS 13 TeV pp collision data with two top quarks and missing transverse momentum in the final state”. In: (Nov. 2022). arXiv: 2211.05426 [hep-ex].
- [147] Serguei Chatrchyan et al. “Search for invisible decays of Higgs bosons in the vector boson fusion and associated ZH production modes”. In: *Eur. Phys. J. C* 74 (2014), p. 2980. DOI: 10.1140/epjc/s10052-014-2980-6. arXiv: 1404.1344 [hep-ex].
- [148] Vardan Khachatryan et al. “Search for dark matter in proton-proton collisions at 8 TeV with missing transverse momentum and vector boson tagged jets”. In: *JHEP* 12 (2016). [Erratum: *JHEP* 08, 035 (2017)], p. 083. DOI: 10.1007/JHEP12(2016)083. arXiv: 1607.05764 [hep-ex].
- [149] A. M. Sirunyan et al. “Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 97.9 (2018), p. 092005. DOI: 10.1103/PhysRevD.97.092005. arXiv: 1712.02345 [hep-ex].
- [150] Albert M Sirunyan et al. “Search for invisible decays of a Higgs boson produced through vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Lett. B* 793 (2019), pp. 520–551. DOI: 10.1016/j.physletb.2019.04.025. arXiv: 1809.05937 [hep-ex].
- [151] Albert M Sirunyan et al. “Performance of the DeepTau algorithm for the discrimination of taus against jets, electron, and muons”. In: CMS DP -2019/033 (2019). DOI: 10.23731/CYRM-2017-002.
- [152] ATLAS Collaboration. “Combination of searches for invisible decays of the Higgs boson using 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment”. In: (Jan. 2023). arXiv: 2301.10731 [hep-ex].
- [153] Abdelhak Djouadi et al. “Implications of LHC searches for Higgs–portal dark matter”. In: *Phys. Lett. B* 709 (2012), pp. 65–69. DOI: 10.1016/j.physletb.2012.01.062. arXiv: 1112.3299 [hep-ph].
- [154] Seungwon Baek et al. “Higgs Portal Vector Dark Matter : Revisited”. In: *JHEP* 05 (2013), p. 036. DOI: 10.1007/JHEP05(2013)036. arXiv: 1212.2131 [hep-ph].

- [155] Anthony DiFranzo, Patrick J. Fox, and Tim M. P. Tait. “Vector Dark Matter through a Radiative Higgs Portal”. In: *JHEP* 04 (2016), p. 135. DOI: 10.1007/JHEP04(2016)135. arXiv: 1512.06853 [hep-ph].
- [156] Mohamed Zaazoua et al. “Higgs Portal Vector Dark Matter Interpretation: Review of Effective Field Theory Approach and Ultraviolet Complete Models”. In: *LHEP* 2022 (2022), p. 270. DOI: 10.31526/lhep.2022.270. arXiv: 2107.01252 [hep-ph].
- [157] Martin Hoferichter et al. “Improved limits for Higgs-portal dark matter from LHC searches”. In: *Phys. Rev. Lett.* 119.18 (2017), p. 181803. DOI: 10.1103/PhysRevLett.119.181803. arXiv: 1708.02245 [hep-ph].
- [158] E. Aprile et al. “Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T”. In: *Phys. Rev. Lett.* 123 (2019), p. 241803. DOI: 10.1103/PhysRevLett.123.241803. arXiv: 1907.12771 [hep-ex].
- [159] P. Agnes et al. “Search for low-mass dark matter WIMPs with 12 ton-day exposure of DarkSide-50”. In: (July 2022). arXiv: 2207.11966 [hep-ex].
- [160] M. Wenbo et al. “A First Search for Solar ^8B Neutrino in the PandaX-4T Experiment using Neutrino-Nucleus Coherent Scattering”. 2022. arXiv: 2207.04883 [hep-ex].
- [161] J. Aalbers et al. “First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment”. In: (July 2022). arXiv: 2207.03764 [hep-ex].
- [162] LHC Higgs Cross Section Working Group. *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties*. Technical report. 2013. DOI: 10.5170/CERN-2013-004. arXiv: 1307.1347 [hep-ph].
- [163] CMS Collaboration. “A portrait of the Higgs boson by the CMS experiment ten years after the discovery”. In: *Nature* 607.7917 (2022), pp. 60–68. DOI: 10.1038/s41586-022-04892-x. arXiv: 2207.00043 [hep-ex].
- [164] L. Borgonovi et al. *Higgs measurements at FCC-hh*. Tech. rep. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2642471>.
- [165] Steve Baker and Robert D. Cousins. “Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms”. In: *Nucl. Instrum. Meth.* 221 (1984), pp. 437–442. DOI: 10.1016/0167-5087(84)90016-4.
- [166] R. D. Cousins. “Generalization of Chisquare Goodness-of-Fit Test for Binned Data Using Saturated Models, with Application to Histograms”. Note. 2013.

- [167] Georges Aad et al. “Search for the $b\bar{b}$ decay of the Standard Model Higgs boson in associated $(W/Z)H$ production with the ATLAS detector”. In: *JHEP* 01 (2015), p. 069. DOI: 10.1007/JHEP01(2015)069. arXiv: 1409.6212 [hep-ex].