CMS Physics Analysis Summary

Combination of searches for nonresonant Higgs boson pair production in proton-proton collisions at $\sqrt{s} = 13$ TeV

The CMS Collaboration

Abstract

This note presents a comprehensive overview and statistical combination of searches for the nonresonant production of Higgs boson pairs (HH) using data from proton-proton collisions collected by the CMS experiment at the LHC from 2016 to 2018 at a centre-of-mass energy of 13 TeV, corresponding to a total integrated luminosity of 138 fb^{-1}. Upper limits at 95% confidence level are set on the rate of the HH production. The observed (expected) upper limit on the inclusive production cross section relative to the standard model expectation is found to be 3.5 (2.5). Assuming all other Higgs boson couplings are equal to their values in the standard model, we exclude HH production at 95% confidence level for values of the Higgs boson trilinear self-coupling modifier κ_λ outside the range between -1.39 and 7.02. Similarly, for the coupling modifier κ_{2V} affecting the interaction between two vector bosons and two Higgs bosons, we exclude HH production for values outside the range between 0.62 and 1.42. This work also studies HH production in new physics scenarios, using the Higgs effective field theory parametrisation. An extrapolation of the results to the integrated luminosity expected after the high-luminosity upgrade of the LHC is also presented.

This document has been revised with respect to the version dated November, 16, 2024.

1 Introduction

In 2012, the ATLAS and CMS Collaborations at the LHC discovered a new particle with a mass of approximately 125 GeV in searching for a Higgs boson [1–3]. Extensive measurements of its properties, including spin and parity, couplings with other standard model (SM) particles and width, have revealed that this particle is compatible with the SM Higgs boson (H) [4, 5]. The Higgs boson self-interaction is an important pending test of the Brout–Englert–Higgs mechanism. The Higgs boson trilinear (λ_3) and quartic self-couplings (λ_4) determine the shape of the Higgs potential. Precise measurements of these couplings provide valuable information about electroweak symmetry breaking and the electroweak phase transition [6]. Several cosmological models predict that the shape of the Higgs field could be related to inflation and baryogenesis [7]. While the quartic coupling is challenging to measure at the LHC and requires future colliders [8–10], the trilinear coupling can be directly accessed by studying Higgs boson pair production at the LHC. The measurement of the Higgs boson trilinear self-coupling will be among the most significant results at the LHC alongside the Higgs boson discovery.

At the LHC, pairs of SM Higgs bosons (HH) are primarily produced via gluon-gluon fusion (ggF), with a cross section of $31.1^{+2.1}_{-7.2}$ fb [11, 12] at next-to-next-to-leading order (NNLO) in quantum chromodynamics (QCD) for a centre-of-mass energy of 13 TeV and a Higgs boson mass of 125 GeV, which is the assumed value for the studies in this note. At leading order (LO), two amplitudes destructively interfere, namely the Higgs boson self-interaction ('triangle diagram') and the top-quark loop ('box diagram'), represented by Feynman diagrams in Fig. 1. The triangle-diagram amplitude is proportional to the Higgs boson trilinear self-coupling, or in the framework described in Section 3, the coupling modifier κ_λ, where $\kappa_\lambda = \lambda_3/\lambda_3^{\mathrm{SM}}$.

The secondary production mechanism for HH events is vector boson fusion (VBF) with a cross section of 1.726 ± 0.036 fb [13] in the SM at 13 TeV and next-to-NNLO ($\mathrm{N}^3\mathrm{LO}$). The VBF production mode, shown in Fig. 2, gives access to the coupling modifier of the interaction between two vector bosons and two Higgs bosons ($\kappa_{2\mathrm{V}}$) and the interaction between a Higgs boson and a vector boson (κ_{V}). The characteristic signature of VBF, comprising two jets in the forward direction of the detector with large gap in rapidity between them, allows us to identify the process despite the small cross section. Higgs boson pairs can also be produced in association with a vector boson V = W or Z (VHH) with a cross section of approximately half that of the VBF mode, as shown in Fig. 3. However, VHH production is less affected by the destructive interference because of the presence of additional diagrams at next-to-LO (NLO), offering complementary constraints on κ_λ.

In new physics models, additional diagrams that include couplings not present in the SM may contribute to HH production. These can be studied in the context of the Higgs effective field theory (HEFT) [14]. The anomalous couplings studied in the present note are denoted c and shown in Fig. 4. Here c_2 corresponds to the coupling between two top quarks and two Higgs bosons, c_g corresponds to the coupling between a Higgs boson and a gluon, and c_{2g} corresponds to the coupling between two Higgs bosons and two gluons.

The ATLAS and CMS Collaborations have performed previous HH searches at $\sqrt{s} = 8$ and 13 TeV. ATLAS has combined HH searches in the $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, and $b\bar{b}b\bar{b}$ final states at 13 TeV and set an upper limit on the inclusive production cross section of 2.9 times the SM expectation at 95% confidence level (CL), with 2.4 expected under the background-only hypothesis [15]. A combined search in the $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, $b\bar{b}b\bar{b}$, $b\bar{b}WW$, and multilepton final states by CMS at 13 TeV constrained (expected to constrain) the inclusive HH production cross section to less than 3.4 (2.5) times the SM expectation at 95% CL [16]. This note introduces additional channels compared to previous results by including the $b\bar{b}WW$, $WW\gamma\gamma$, and $\tau\tau\gamma\gamma$ final states and

VHH production mode in the $b\bar{b}b\bar{b}$ final state, to perform a variety of parameter scans as well as searches for anomalous Higgs boson couplings in the HEFT approach. The results presented in this note are based on the data set of proton-proton collisions at 13 TeV collected by the CMS experiment from 2016 to 2018 (Run 2), corresponding to an integrated luminosity of 138 fb^{-1}.


Figure 1: Leading-order Feynman diagrams of nonresonant Higgs boson pair production via gluon-gluon fusion in the standard model. The modifiers of the Higgs boson coupling with the top quark and the Higgs boson trilinear self-coupling are shown as κ_t and κ_λ, respectively.


Figure 2: Leading-order Feynman diagrams of nonresonant Higgs boson pair production via vector boson fusion in the standard model. The modifiers of the Higgs boson coupling with a vector boson, the Higgs boson trilinear self-coupling, and the coupling between two Higgs bosons and two vector bosons are shown as κ_V, κ_λ, and κ_{2V}, respectively.


Figure 3: Leading-order Feynman diagrams of nonresonant Higgs boson pair production via associated production with a vector boson in the standard model. The modifiers of the Higgs boson coupling with a vector boson, the Higgs boson trilinear self-coupling, and the coupling between two Higgs bosons and two vector bosons are shown as κ_V, κ_λ, and κ_{2V} respectively.

2 The CMS detector and event reconstruction

The CMS apparatus [17, 18] is a multipurpose, nearly hermetic detector, designed to trigger on [19–21] and identify electrons, muons, photons, and (charged and neutral) hadrons [22–24]. Its central feature is a superconducting solenoid of 6 m internal diameter, providing a magnetic field of 3.8 T. Within the solenoid volume are a silicon pixel and strip tracker, a lead

Figure 4: Leading-order Feynman diagrams of nonresonant Higgs boson pair production via gluon-gluon fusion with anomalous Higgs boson couplings c_2, c_g, and c_{2g}. The Higgs boson trilinear self-coupling modifier is shown as κ_λ.

tungstate crystal electromagnetic calorimeter (ECAL), and a brass and scintillator hadron calorimeter (HCAL), each composed of a barrel and two endcap sections. Forward calorimeters extend the pseudorapidity coverage provided by the barrel and endcap detectors. Muons are reconstructed using gas-ionisation detectors embedded in the steel flux-return yoke outside the solenoid. More detailed descriptions of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in Refs. [17, 18].

Events of interest are selected using a two-tiered trigger system. The first level (L1), composed of custom hardware processors, uses information from the calorimeters and muon detectors to select events at a rate of around 100 kHz within a fixed latency of 4 μs [19]. The second level, known as the high-level trigger (HLT), consists of a farm of processors running a version of the full event reconstruction software optimised for fast processing, and reduces the event rate to a few kHz before data storage [20, 21].

A particle-flow (PF) algorithm [25] aims to reconstruct and identify each particle in an event (PF candidate), with an optimised combination of information from the various elements of the CMS detector. The primary vertex (PV) is taken to be the vertex corresponding to the hardest scattering in the event, evaluated using tracking information alone, as described in Section 9.4.1 of Ref. [26]. The energy of photons is obtained from the ECAL measurement. The energy of electrons is determined from a combination of the track momentum at the primary interaction vertex, the corresponding ECAL cluster energy, and the energy sum of all bremsstrahlung photons attached to the track. The momentum of muons is obtained from the curvature of the corresponding track. The energy of charged hadrons is determined from a combination of their momentum measured in the tracker and the matching ECAL and HCAL energy deposits, corrected for zero-suppression effects and for the response function of the calorimeters to hadronic showers. Finally, the energy of neutral hadrons is obtained from the corresponding corrected ECAL and HCAL energies. The missing transverse momentum vector $\vec{p}_T^{\text{miss}}$ is computed as the negative vector sum of the transverse momenta of all the PF candidates in an event, and its magnitude is denoted as p_T^{miss} [27]. The $\vec{p}_T^{\text{miss}}$ is modified to account for corrections to the energy scale of the reconstructed jets in the event.

Small-radius jets are reconstructed from PF candidates, using the anti-k_T clustering algorithm [28, 29] with a distance parameter $R = 0.4$ (AK4 jets). Charged particles not originating from the PV are excluded from the jet clustering. Higgs boson decays into a pair of b quarks with high transverse momentum (p_T) result in final states with large Lorentz boost, and as a result, the b jets are overlapping, forming one large merged jet ("large-radius jet") and substructure, i.e., the two overlapping jets are "subjets" of the large-radius jet. The jets are reconstructed using the anti-k_T algorithm with a distance parameter $R = 0.8$ (AK8 jets).

The searches presented in this note use a variety of object identification techniques and algo-

rithms, tailored to their specific characteristics.

3 Signal modelling

The dependence of the ggF HH cross section on κ_λ and κ_t can be written as:

$$\sigma(\kappa_\lambda, \kappa_t) = \kappa_\lambda^2 \kappa_t^2 t + \kappa_t^4 b + \kappa_\lambda \kappa_t^3 i, \tag{1}$$

where $t = |T|^2$, $b = |B|^2$, and $i = |TB^* + B^*T|$, and the three contributions correspond to the triangle and box diagrams shown in Fig. 1, and their interference, respectively [30].

The same formula holds for every differential cross section $d\sigma/dx$ for HH production. At higher order in the QCD perturbative expansion, T and B correspond to the sum of all diagrams of the same order in κ_λ and κ_t, and the polynomial relation in Eq. (1) remains valid. Since the ggF cross section can be expressed as a polynomial in κ_λ and κ_t with three independent terms, we can model the production of ggF over a wide range of κ_λ and κ_t by evaluating the cross section at three different $(\kappa_\lambda, \kappa_t)$ values and performing a linear combination. In our case, we use three simulated data samples and combine them linearly to perform scans in the κ_λ and κ_t parameters.

In order to study further modified values for the SM couplings as well as couplings not present in the SM we use an event-level reweighting method. The reweighting is based on a parameterisation of the differential cross section in the generator-level invariant mass of the HH system and the angular distance between the two Higgs bosons in the azimuthal plane, which are sufficient to characterise the hard scattering because it only has two degrees of freedom. This allows modelling any combination of coupling modifiers (κ_λ, κ_t, c_2, c_g, c_{2g}), even for values that are not used in the sample generation. The method used to derive the reweighting factors is described in Refs. [31, 32]. The procedure is similar to the recommendation by the LHC Higgs Working Group 4 [14]. In our derivation, we use finer bins for high values of the invariant mass of the HH system. The method is used for all HH decay channels except the $b\bar{b}b\bar{b}$ channel with high Lorentz boost. Because of the limited number of simulated events in the high m_{HH} bins, the reweighting factors are not precise enough. Therefore, all necessary Monte Carlo (MC) simulated samples for each benchmark are produced.

Two sets of benchmarks are produced using the reweighting method. The first set defines thirteen benchmark scenarios [33] spanning a broad range of the coupling values in the HEFT parameterisation. The second set [34] was derived with NLO precision and considers direct and indirect constraints on the allowed range for the coupling values. The values of the coupling modifiers corresponding to each benchmark are shown in Tables 1 and 2.

The c_2 and κ_t parameters are correlated in the HEFT Lagrangian, while κ_λ and κ_t remain correlated as in the SM case. Therefore, we further study the c_2 modifier and perform scans with c_2 as the parameter of interest. To perform these scans, we use a linear combination of simulated samples to model any combination of values with a procedure similar to that used for κ_λ and κ_t for the SM-like results.

The simulated events for ggF production are produced using NLO matrix elements with the model described in Ref. [30], implemented in POWHEG v2 [35–37]. The cross section is corrected to the NNLO value as a function of κ_λ according to Refs. [38–40]. Four ggF samples are produced with different values of the trilinear coupling modifier ($\kappa_\lambda = 0$, 2.45, 5.0, and 1.0). The simulated signal samples for VBF and VHH production modes are generated at leading-order

(LO) using MADGRAPH5_aMC@NLO 2.6.5. Seven VBF samples are generated with different values of the coupling modifiers $(\kappa_V, \kappa_{2V}, \kappa_\lambda) = (1,1,1), (1, 1, 0), (1, 1, 2), (1, 0, 1), (1, 2, 1), (0.5, 1, 1)$, and $(1.5, 1, 1)$, and eight VHH samples are generated with the same coupling modifier values plus $(\kappa_V, \kappa_{2V}, \kappa_\lambda) = (1, 1, 20)$. The total cross section is corrected to the corresponding NNLO cross section [41, 42]. The parton shower and hadronisation are simulated with PYTHIA 8.2 [43] using the CUETP8M1 set of tuned parameters for 2016 simulations and the CP5 tune set in 2017 and 2018 simulations. The scans in the parameters of interest κ_λ, κ_t, κ_V, and κ_{2V} are performed using a linear combination of a set of VBF and VHH samples, similar to what is done for ggF.

Table 1: Values of the effective Lagrangian couplings for the Higgs Effective field theory benchmarks proposed in Ref. [33].

	1	2	3	4	5	6	7	8	9	10	11	12	8a
κ_λ	7.5	1.0	1.0	−3.5	1.0	2.4	5.0	15.0	1.0	10.0	2.4	15.0	1.0
κ_t	1.0	1.0	1.0	1.5	1.0	1.0	1.0	1.0	1.0	1.5	1.0	1.0	1.0
c_2	−1.0	0.5	−1.5	−3.0	0.0	0.0	0.0	0.0	1.0	−1.0	0.0	1.0	0.5
c_g	0.0	−0.8	0.0	0.0	0.8	0.2	0.2	−1.0	−0.6	0.0	1.0	0.0	$\frac{0.8}{3}$
c_{2g}	0.0	0.6	−0.8	0.0	−1.0	−0.2	−0.2	1.0	0.6	0.0	−1.0	0.0	0.0

Table 2: Values of the effective Lagrangian couplings for the Higgs effective field theory benchmarks proposed in Ref. [34].

	1	2	3	4	5	6	7
κ_λ	3.94	6.84	2.21	2.79	3.95	5.68	−0.10
κ_t	0.94	0.61	1.05	0.61	1.17	0.83	0.94
c_2	$-\frac{1}{3}$	$\frac{1}{3}$	$-\frac{1}{3}$	$\frac{1}{3}$	$-\frac{1}{3}$	$\frac{1}{3}$	1.0
c_g	0.5×1.5	0.0	0.5×1.5	-0.5×1.5	$\frac{1}{6} \times 1.5$	-0.5×1.5	$\frac{1}{6} \times 1.5$
c_{2g}	$\frac{1}{3} \times -3$	$-\frac{1}{3} \times -3$	0.5×-3	$\frac{1}{6} \times -3$	-0.5×-3	$\frac{1}{3} \times -3$	$-\frac{1}{3} \times -3$

4 Analysis strategy

In this note, we describe the statistical combination of searches for nonresonant HH production in several production modes and decay channels. The list of HH analyses considered in this combination, along with a summary of results, are listed in Table 3. A brief description of each analysis is given later in this section. More details can be found in the respective publications.

Where required, the event selection of each analysis was modified for the combination to prevent double counting of events. The strategy for this overlap removal is based upon work done for the combinations in Ref. [16]. The strategy for the correlation of systematic uncertainties is described in Section 5.

4.1 HH → b$\bar{\text{b}}$b$\bar{\text{b}}$ resolved

The b$\bar{\text{b}}$b$\bar{\text{b}}$ decay channel has the largest HH branching fraction among the HH decays to SM particles. The resolved search [46] focuses on the kinematic phase space where each Higgs boson is reconstructed from two small-radius b-tagged jets. It explores both the ggF and VBF HH production modes. The online trigger selection requires the presence of at least four jets, satisfying thresholds on jet p_T and H_T, defined as the scalar p_T sum of all the jets. These thresholds vary depending on the data collection year. Consequently, data collected in 2016 are analysed separately from those collected in 2017 and 2018.

Table 3: Summary of results for the HH analyses included in this combination. The second column is the observed (expected) 95% CL upper limit on the inclusive signal strength r. The third (fourth) column is the allowed 68% CL interval for the coupling modifier κ_λ (κ_{2V}). The last column indicates whether the analysis is included in the results using the HEFT parametrisation.

Analysis	r	κ_λ	κ_{2V}	HEFT
$b\bar{b}\gamma\gamma$ [44]	7.7 (5.2)	[−3.3, 8.5]	[−1.3, 3.5]	✓
$b\bar{b}\tau\tau$ [45]	3.3 (5.2)	[−1.7, 8.7]	[−0.4, 2.6]	✓
$b\bar{b}b\bar{b}$, resolved jets [46]	3.9 (7.8)	[−2.3, 9.4]	[−0.1, 2.2]	✓
$b\bar{b}b\bar{b}$, merged jets [47]	9.9 (5.1)	[−9.9, 16.9]	[0.62, 1.41]	✓
VHH, HH → $b\bar{b}b\bar{b}$ [48]	r_{VHH} < 294 (124)	[−37.7, 37.2]	[−12.2, 8.9]	—
$b\bar{b}VV$, VV → $\ell\nu qq/2\ell2\nu$ [49]	14 (18)	[−7.2, 13.8]	[−8.7, 15.2]	✓
$b\bar{b}VV$, VV → 4q [50]	141 (69)	—	[−0.04, 2.05]	—
HH, multilepton [51]	21.3 (19.4)	[−6.9, 11.1]	—	✓
$WW\gamma\gamma$ [52]	97 (53)	[−25.8, 14.4]	—	✓
$b\bar{b}ZZ$ [53]	32.4 (39.6)	[−8.8, 13.4]	—	—
$\tau\tau\gamma\gamma$ [54]	33 (26)	[−13, 18]	—	—

Events selected offline contain at least four jets. Jets originating from the b quark decay are identified using the DEEPJET algorithm. The four jets with the highest DEEPJET score are chosen as b jet candidates. The p_T of the four b jets is corrected using a multivariate regression method based on a DNN. At least three b jets must satisfy the DEEPJET medium working point (WP), which corresponds to a b jet identification efficiency of about 75% and a misidentification rate for light-flavour quark and gluon jets of about 1%. The VBF jet candidates are the two highest p_T non-b jets located in opposite η hemispheres. Events that contain any isolated electrons or muons are rejected.

There are three possible ways to pair the four b jets in order to reconstruct two Higgs boson candidates. For each pairing a distance parameter $d = |m_{H_1} - km_{H_2}|/\sqrt{1+k^2}$ is calculated, where m_{H_1} is the mass of the p_T-leading Higgs boson candidate, m_{H_2} is the mass of the p_T-subleading Higgs boson candidate, and $k = 1.04$. If the difference between the two smallest distance parameters is larger than 30 GeV, then the pairing with the smallest distance parameter is selected. Otherwise, between the two pairings with the smallest distance, the one maximising the p_T in the four-jet centre-of-mass frame is chosen. This procedure results in a correct jet pairing of about 82–96 (91–98)% of the selected events for the different couplings studied in ggF (VBF) signal events.

Events that do not have a VBF jet pair are assigned to the ggF event category. A boosted decision tree (BDT$_{ggF/VBF}$) discriminant is trained to separate ggF and VBF HH signals in events that contain the VBF jet pair. A threshold on the BDT$_{ggF/VBF}$ score is used to classify those events in the ggF or VBF event category. The ggF and VBF categories are each further divided into subcategories to optimise the search sensitivity to anomalous coupling hypotheses. In the ggF category, events are separated into low-mass and high-mass categories using a 450 GeV threshold on the Higgs boson pair invariant mass m_{HH}. In the VBF category, events are divided into "SM-like" and anomalous-κ_{2V} categories using a threshold on the BDT$_{ggF/VBF}$ score.

The dominant background source is from SM events composed uniquely of jets produced through the strong interaction, referred to as QCD multijet events. The analysis signal region contains at least four b jet candidates satisfying the DEEPJET medium WP. The background of this '4b' region is estimated using events from an orthogonal region in which the fourth-

highest b-tagged jet fails the medium WP (or '3b' region). Dedicated control regions are used to derive a transfer function to model the 4b background from the 3b data. This function comprises a BDT reweighting and a transfer factor for reshaping and normalisation corrections, respectively.

To remove overlapping events between this analysis and the analysis of $b\bar{b}b\bar{b}$ events in the boosted topology described in Section 4.2, an event veto is applied specifically for the combination described in this note. Events containing two massive large-radius jets with $p_T > 300\,\text{GeV}$ and at least two subjets are removed from this analysis. The p_T threshold was optimised based on the expected upper limit on SM HH production.

4.2 HH → $b\bar{b}b\bar{b}$ boosted

The search in the $b\bar{b}b\bar{b}$ boosted final state [47] focuses on the phase space region where both Higgs bosons are highly Lorentz boosted so that each Higgs boson is reconstructed as a large-radius jet. It is sensitive to modified coupling scenarios that enhance the production of highly boosted Higgs boson pairs, including through the VBF production.

A combination of several trigger algorithms is used with requirements on the total H_T or a large-radius jet p_T to be above a given threshold. This search uses the soft-drop (SD) algorithm [55] to reconstruct the mass of the two Higgs boson candidates.

The graph neural network (GNN) algorithm [56–59] known as PARTICLENET [57, 60] is used to discriminate between H → $b\bar{b}$ and QCD-induced jets. To improve the jet mass estimation, a regression algorithm based on the PARTICLENET GNN architecture is introduced to predict the jet mass (m_{reg}) [61].

Events are required to have at least two large-radius jets and then grouped into mutually exclusive ggF and VBF categories. The VBF categories are designed to select VBF signal events with the characteristic signature of two additional small-radius jets in the opposite forward regions of the detector. Three VBF event categories (high, medium, and low purity) are defined based on the PARTICLENET discriminant ($D_{b\bar{b}}$) scores of the two Higgs boson candidate jets.

Events not sorted into the VBF categories are considered for the ggF categories. A BDT is trained to discriminate between the HH signal and QCD multijet or $t\bar{t}$ background processes. Three search regions (SRs) targeting the ggF production are constructed depending on the BDT output score and the $D_{b\bar{b}}$-subleading large-radius jet $D_{b\bar{b}}$ score.

The dominant SM backgrounds are $t\bar{t}$ and QCD multijet production. Control regions (CRs) enriched in QCD multijet events are selected by changing the requirement on the $D_{b\bar{b}}$ discriminant. The QCD multijet background in the SRs is estimated using the data in the CRs and fitted transfer factors [62–64]. An auxiliary sample enriched in $t\bar{t}$ events containing one leptonically decaying W boson is used to extract corrections for the $t\bar{t}$ background estimate [65].

4.3 VHH, HH → $b\bar{b}b\bar{b}$

This analysis [48] focuses on the final state with both Higgs bosons decaying into a b quark-antiquark pair. All decay modes of vector boson Z and W are considered. Experimentally, the events are divided into four categories based on the presence of a light lepton ℓ (μ, e) and jets: 2 leptons, 1 lepton, p_T^{miss} and fully hadronic channel, corresponding to the Z → $\ell\ell$, W → $\ell\nu$, Z → $\nu\nu$, and Z/W → qq processes, respectively. Because of the large overlap with the ggF and VBF $b\bar{b}b\bar{b}$ channels, the VHH fully hadronic channel is removed from this combination. Two event topologies are explored in this analysis: one involving four small-radius b jets, and

the other with two large-radius jets, each from a $H \rightarrow b\bar{b}$ decay.

In the leptonic channels, events are selected using triggers requiring isolated leptons or large p_T^{miss}. In the resolved topology, at least three AK4 jets are required to pass the medium WP of the DEEPJET b-tagging algorithm. If there are more than four b-tagged jets, the ones with the highest b-tagging scores are selected. An energy-regression method is applied to the b jets in order to improve the resolution of the dijet mass. The two Higgs boson candidates are reconstructed by pairing the four jets with the same method used in the resolved $b\bar{b}b\bar{b}$ analysis. In the boosted topology, two AK8 jets are required to have PARTICLENET $D_{b\bar{b}}$ scores greater than 0.8. Given the higher sensitivity in the boosted topology, it is given priority if an event can be assigned to both topologies.

A categorisation BDT is trained to divide events into regions sensitive to κ_λ and κ_{2V}. In each region, another BDT is trained using all the signal processes in order to optimise the signal versus background discrimination. The final signal strength is extracted by fitting the BDT score distribution in the signal and control regions.

4.4   HH → $b\bar{b}\tau\tau$

In the $b\bar{b}\tau\tau$ channel [45], the events are characterised by the decay mode of the two τ leptons. This analysis studies final states where at least one of the τ leptons decays hadronically, i.e. $\tau_\mu\tau_h$, $\tau_e\tau_h$, and $\tau_h\tau_h$. The light leptons and τ_h candidates must have opposite electric charges. The $\tau_h\tau_h$ candidates are reconstructed using the DEEPTAU algorithm [66], a neural network method discriminating τ_h candidates from jets, electrons, and muons. Events are required to have at least two b jets identified using the DEEPJET algorithm.

For each event, all b jet candidates are assigned a score by a neural network trained to recognise the $H \rightarrow b\bar{b}$ topology. The two jets with the highest score are selected as the b jets from the Higgs boson decays. To identify signal events originating from VBF, the events must have two additional small-radius jets.

The events are selected by triggers requiring a single light lepton, or a lepton accompanied by a τ_h candidate or two τ_h candidates. A trigger dedicated to the VBF topology was introduced in late 2017, requiring two τ_h and two additional jets.

After the event selection, the events are divided into categories. The first set of categories splits the events according to the number and topology of the jets. Events with two small-radius jets separated by $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.4$ are classified in the boosted category. All remaining events are classified in the resolved categories. The resolved category is further split according to whether one or both jets pass the medium WP of the DEEPJET b-tagging algorithm. Finally, an additional VBF category selects events with at least two VBF jet candidates. The contamination of this category from ggF events with an extra jet is significant; therefore, a multiclassifier DNN is trained. The DNN splits the events in the VBF category into signal-like VBF events, ggF contamination, $t\bar{t}H$, $t\bar{t}$, and Drell–Yan (DY) background events.

All selected events are subject to a selection on the two-dimensional plane of the invariant mass of the dijet ($m_{b\bar{b}}$) and ditau ($m_{\tau\tau}$) systems, where the $m_{\tau\tau}$ is reconstructed using the SVFIT algorithm [67]. This selection is optimised separately for the resolved and the boosted categories. No invariant mass requirement is applied to the VBF categories.

The eight event categories defined earlier are used for the signal extraction using a DNN developed to identify $HH \rightarrow b\bar{b}\tau\tau$ events. The DNN is trained to classify the events from each event category as signal- or background-like by assigning a single prediction per event. The

maximum likelihood fit is performed using the DNN distribution simultaneously for all categories and all channels.

The main background processes for this search, $t\bar{t}$, DY, and QCD multijet, are estimated using data-driven methods. The shapes of the $t\bar{t}$ and DY processes are taken from the simulation and then scaled using background-enriched regions. The scaling is extracted with a simultaneous fit over all the control regions and applied to the simulated events in the signal region.

The QCD multijet process is entirely determined from data in jet-enriched regions where the τ pair requirements are inverted. The event yield and shape are obtained in a region where the τ_h isolation requirement is inverted. In $\tau_h\tau_h$ events, only the isolation of the τ_h with the lowest p_T is inverted. Then, the yield is scaled by the ratio of event yield in a region where the opposite charge requirement is inverted and a region where both the charge and the isolation are inverted. The contributions from other backgrounds are subtracted. The QCD multijet background is estimated after the $t\bar{t}$ background has been scaled.

4.5 HH → b$\bar{\text{b}}\gamma\gamma$

The b$\bar{\text{b}}\gamma\gamma$ final state [44] is one of the most sensitive channels to HH production despite the relatively small branching fraction because of the good mass resolution of the diphoton system. The parameters of interest of the model are extracted from a two-dimensional fit to the invariant mass distributions of the photon pair ($m_{\gamma\gamma}$) and b jet pair ($m_{b\bar{b}}$). The HH signal can be identified as a peak in the ($m_{\gamma\gamma}$, $m_{b\bar{b}}$) distributions at the value of the Higgs boson mass. Jets originating from the b quark decay are identified using the DEEPJET [68] algorithm based on a deep neural network (DNN). The background is dominated by the γ+jets continuum, and is modelled from the data sidebands of the $m_{\gamma\gamma}$ and $m_{b\bar{b}}$ distributions. Other important background sources are the single-Higgs-boson production processes, in which the Higgs boson decays into a photon pair (H → $\gamma\gamma$). The dominant H → $\gamma\gamma$ background is the $t\bar{t}$H($\gamma\gamma$) process because of the b jets in the final state from the $t\bar{t}$ decay. The H → $\gamma\gamma$ processes are modelled using MC simulations.

Exclusive categories targeting specific HH production modes are defined to maximise the sensitivity to the signal. The VBF topology is characterised by two additional jets with high invariant mass and high $\Delta\eta$. The signal region targeting the VBF mechanism has requirements on the two additional jets. The VBF event candidates are classified into two exclusive categories depending on the value of the reduced four-body mass M_X, defined as $M_X = m_{bb\gamma\gamma} - m_{\gamma\gamma} - m_{b\bar{b}} + 250\,\text{GeV}$, where $m_{bb\gamma\gamma}$ is the four-body invariant mass of the pairs of photons and b jets. A boosted decision tree (BDT) classifier is defined to isolate the VBF signal from the ggF signal and the continuum background. For the ggF-enriched categories, a BDT classifier is used to reject the continuum background. The values of the BDT score and M_X are used to classify the events in twelve ggF-enriched categories. The number and boundaries of the categories are optimised to maximise the expected significance of the SM HH signal.

Furthermore, a DNN is trained to separate the HH signal from the $t\bar{t}$H($\gamma\gamma$) process. The selection on the DNN output score provides a $t\bar{t}$H rejection between 80 and 85% and an efficiency between 90 and 95% for the HH signal, depending on the analysis category. This channel is limited by the data sample size.

4.6 HH → b$\bar{\text{b}}$VV, VV → $\ell\nu$qq/$2\ell2\nu$

The b$\bar{\text{b}}$WW decay mode has the second largest branching fraction among all HH decay modes, approximately 24%. While b$\bar{\text{b}}$WW is the dominant channel in this final state, b$\bar{\text{b}}$ZZ and b$\bar{\text{b}}\tau\tau$

decays of the HH system that result in a pair of b quarks and one or two light leptons are also included in the signal. The analysis [49] considers two final states: single-lepton (1ℓ or SL) bb$\ell\nu$qq and dilepton (2ℓ or DL) bb2ℓ2ν. The most significant background contribution is the t$\bar{\text{t}}$ process, while other backgrounds include events with misidentified leptons, and the DY process and W+jets process for the DL and SL channels, respectively.

The data are collected using single and double lepton triggers. The event selection criteria require the presence of one isolated lepton in the 1ℓ channel or two isolated leptons of opposite charge in the 2ℓ channel, originating from the W leptonic decays. The H → b$\bar{\text{b}}$ system is characterised by two small-radius jets. While both of these jets originate from b quarks, in the event selection, only one must be classified as a b jet by the DEEPJET algorithm. If the H → b$\bar{\text{b}}$ system is produced with high momentum, the two b jets are Lorentz boosted and can be reconstructed as one large-radius jet with substructure. In the SL case, the events must have at least one more small-radius jet present, corresponding to the hadronic W decay.

To avoid overlapping events with the b$\bar{\text{b}}\tau\tau$ analysis, this analysis vetoes events containing at least one τ_h candidate as defined in Ref. [45]. The events are classified using multiclass DNNs, trained separately for the SL and DL cases. The DNNs are trained using high-level features such as invariant masses, hadronic activity, and the output of a Lorentz-boost network that performs automated feature engineering using the four-momenta of selected leptons and jets. The DNN training data set includes signal samples with SM and anomalous κ_λ values to ensure good performance on the full spectrum.

The multiclass DNNs have different output nodes corresponding to several background processes and two signal processes: ggF and VBF HH production. The DNNs provide a score for each node that can be interpreted as the probability of each event to belong to that class. The events are categorised according to the node with the highest score. The resulting background categories are: DY or W+jets for the DL or SL channel, respectively, and t$\bar{\text{t}}$ for both channels. The two signal categories are further divided into subcategories according to the b jet multiplicity and topology into resolved-1b, resolved-2b, and boosted.

The background contribution from misidentified leptons is estimated using the fake-factor method [69]. In the DL channel, the DY background is estimated using the ABCD method. Events with zero b-tagged jets are weighted by transfer factors extracted from events with dilepton invariant mass inside the Z boson mass window. All other backgrounds are estimated using simulated events. The signal extraction is performed using a simultaneous maximum likelihood fit to the distributions of the DNN outputs in the signal and background event categories.

4.7 HH → b$\bar{\text{b}}$VV, VV → 4q

The all-hadronic b$\bar{\text{b}}$VV search [50] targets highly Lorentz-boosted pairs of Higgs bosons produced via gluon-gluon fusion and vector boson fusion, such that both b$\bar{\text{b}}$ quarks and all four quarks from the V decays are each contained in single large-radius jets. The online selection is based on the combination of several conditions, on H_T and the kinematics of the large-R jets. This channel targets mainly the b$\bar{\text{b}}$WW decays, however, HH decays into b$\bar{\text{b}}$ZZ in the fully hadronic final state are also included in the signal.

The PARTICLENET algorithm [57] is used to isolate the signal H → b$\bar{\text{b}}$ jets against background QCD using a discriminant, T_{Xbb}, derived from its outputs, while to identify the H → VV → 4q jet, we introduce a new attention-based neural network, based on the PARTICLETRANSFORMER architecture [70]. Both networks have been decorrelated from the mass of the jets by enforcing a

uniform distribution in jet mass and p_T in the training samples [71], to aid with their calibration. Additionally, as the jet mass resolution is crucial to the sensitivity of the search, we optimise the mass reconstruction for all AK8 jets using a PARTICLENET-based regression algorithm trained to learn the "true", or event-generator-level, jet mass [61], the output of which we refer to as m_{reg}. The jet with the higher (lower) T_{Xbb} score is considered the $b\bar{b}$- (WW $\to$ 4q-) candidate jet.

The VBF process produces two, likely forward, jets with large invariant masses and pseudorapidity separations. The BDT is optimised simultaneously for both the SM ggF and BSM VBF $\kappa_{2V} = 0$ signals, and separate "ggF" and "VBF" signal regions are defined using the BDT probabilities for the respective processes, referred to as $\mathrm{BDT}_{\mathrm{ggF}}$ and $\mathrm{BDT}_{\mathrm{VBF}}$. Concretely, the VBF region is defined by selections on the T_{Xbb} and $\mathrm{BDT}_{\mathrm{VBF}}$ discriminants, chosen to optimise the expected exclusion limit on the VBF signal. The ggF region is defined by a veto on events passing the VBF selections plus selections on the T_{Xbb} and $\mathrm{BDT}_{\mathrm{ggF}}$ discriminants, similarly chosen to optimise the limit on the ggF signal. The uncertainty on the BDT signal efficiency is dominated by that of the PARTICLETRANSFORMER tagger, which is calibrated based on a new technique [72] using the ratio of the primary Lund jet plane [73] densities of each quark subjet.

The search is performed in the signal regions using a one-dimensional likelihood model binned in $m_{\mathrm{reg}}^{\mathrm{bb}}$. The background in these regions is dominated by QCD multijet events, and is estimated using the product of data in the fail region and polynomial transfer factors [62–64], as described below. Other minor backgrounds include top quark and vector boson backgrounds, which are estimated using MC simulation and whose systematic and statistical uncertainties are incorporated in the final statistical analysis.

4.8 HH $\to$ multileptons

The HH multilepton analysis [51] focuses on HH decays into the final states of 4W, 2W2τ and 4τ. Events are selected using a set of single, double, and triple lepton triggers, as well as lepton and τ_h triggers and double τ_h triggers. The leptons and τ_h are reconstructed with a set of special identification criteria (ID) originally developed for the CMS $t\bar{t}$H multilepton analysis [74]. These IDs are customised for the use of the fake factor method [69] used for the data-driven estimation of events with misidentified leptons or τ_h. For leptons, these two IDs also make use of a specialised multivariate analysis specifically trained to recognise genuine leptons that originate from vector boson and τ lepton decays. All categories veto events containing at least one b jet identified with the medium WP of the DEEPJET algorithm and events containing at least two b jets identified with the loose WP, corresponding to a b jet selection efficiency of about 84% and a misidentification rate for light-flavour quark and gluon jets of about 11%. This reduces contributions from top-quark related backgrounds such as $t\bar{t}$ and makes this analysis orthogonal to HH analyses requiring b jets.

A maximum likelihood fit is performed simultaneously using the BDT output distributions in a set of seven mutually exclusive event categories with final states containing multiple leptons ℓ (μ, e) or hadronically decaying τ leptons (τ_h): 2ℓ with the same electric charge (SS), 3ℓ, 4ℓ, 3ℓ+1τ_h, 2ℓ+2τ_h, 1ℓ+3τ_h and 4τ_h. All categories require the specified number of final state light leptons and τ_h. The 2ℓ SS and 3ℓ categories additionally require the presence of jets originating from hadronic W boson decays. The primary background in most categories is given by events containing misidentified ℓ or τ_h. All categories also contain a sizable contribution of genuine multiboson backgrounds from ZZ and WZ decays. Therefore, two kinematic distributions in WZ and ZZ control regions are included in the fit. The categories are constructed by inverting the Z boson veto in the 3ℓ and 4ℓ categories, respectively. The inclusion of these two control

categories helps to constrain the two main diboson backgrounds as well as the corresponding systemtaic uncertainties.

The analysis makes use of a parameterised BDT tuning the sensitivity to the SM and to one of the sets of EFT benchmarks as described in Section 3. The results for each benchmark point are extracted from the corresponding BDT output. For the benchmarks not included in the BDT training, the most similar one based on the shape of the parton-level m_{HH} spectrum is used during inference.

4.9 HH $\to \tau\tau\gamma\gamma$

The search for HH production in the $\tau\tau\gamma\gamma$ final state [54] aims to cover both the hadronic and leptonic decay modes of the τ lepton. Despite the small HH $\to \tau\tau\gamma\gamma$ branching fraction in the SM, the diphoton pair offers a clean experimental signature to trigger on with a good mass resolution, while the additional tau leptons in the event help to further isolate signal from background. The main target of the search is the nonresonant HH production via ggF.

Events are selected online using a diphoton trigger with asymmetric photon p_{T} thresholds of 30 and 18 GeV. All events are required to have at least one diphoton candidate, and in the case where more than one candidate exists, the candidate with the highest scalar p_{T} sum of the photons is chosen.

Events are required to have at least one τ lepton candidate. Both hadronic and leptonic decay modes of the τ lepton are considered. A $\tau\tau$ candidate can be identified from any of the following pairs of reconstructed objects: $\tau_{\mathrm{h}}\tau_{\mathrm{h}}$, $\tau_{\mathrm{h}}\mu$, τ_{h}e, μe, $\mu\mu$, ee, τ_{h}+isolated track. The DEEPTAU discriminant is used to select τ_{h} candidates. Events consistent with a Z $\to \ell\ell$ or Z $\to \ell\ell\gamma$ decay are rejected.

A BDT classifier is used to further isolate events with signal-like characteristics from the background. Hadronic jets in the event are only used in the BDT to reject backgrounds like $t\bar{t}$. To reduce the probability of artificial peaking structure, the $m_{\gamma\gamma}$ is not included as an input feature to the BDT, and the p_{T} of the photon candidates are divided by $m_{\gamma\gamma}$ to reduce correlations between the photon p_{T} and $m_{\gamma\gamma}$. Sequential boundaries are placed on the BDT output to define event categories of different signal purity, where the boundary positions are chosen to maximise the signal sensitivity.

The data are extracted through a maximum likelihood fit of the $m_{\gamma\gamma}$ distribution in each category. The smoothly falling background (continuum) is modelled directly from the data. The signal and the single Higgs boson production, which is a resonant background, are modelled from simulation.

4.10 HH $\to$ b$\bar{\mathrm{b}}$ZZ

This analysis [53] focuses on ggF HH production where one Higgs bosons decays to a Z boson pair, which subsequently decays to 4ℓ, and the other to b$\bar{\mathrm{b}}$, hadronising into jets. The 4ℓ in the final state form a clean signature, while the high branching fraction of the b$\bar{\mathrm{b}}$ decay channel partially compensates for the small branching fraction of the 4ℓ channel. The final state consists of at least two pairs of oppositely charged isolated electrons or muons and at least two jets. The 4ℓ signal region is defined by requiring the four-lepton invariant mass $m(4\ell)$ to be within the Higgs boson mass window. Events outside this range are included in the 4ℓ control region.

The reducible background (Z+X) originates from processes with non-prompt leptons, mainly from decays of heavy-flavour hadrons, misidentified jets, and electrons from photon conversions. The Z+X background is estimated in a data-driven way, using a control region, and the

probabilities of a light lepton to be misidentified. The dominant systematic uncertainty in the reducible background estimation arises mainly from the limited number of events in both the control region and the region where the misidentification rates are computed. Uncertainties from the difference in the sample composition used to calculate the misidentification rate are also considered.

A BDT is used to further discriminate signal from backgrounds. The BDT is trained using simulated events in the 4ℓ signal region; training is done using the ggF HH process as signal and all other processes as background. A maximum likelihood fit to the BDT output distribution is performed to set constraints on the parameters of interest.

4.11   HH $\rightarrow$ WW$\gamma\gamma$

This analysis [52] studies three WW decay modes: semileptonic (1ℓ), fully hadronic, and fully leptonic (2ℓ) decays. We use the diphoton invariant mass distribution for the signal extraction across all channels.

A BDT classifier is trained for photon-jet separation, resulting in a photon identification score. The jets from the hadronisation of bottom quarks are identified using the DEEPJET algorithm. The number of light leptons is used to maintain orthogonality in all three analyses. We assign an event to the fully hadronic channel if it has no leptons, to the 1ℓ channel if it has exactly one lepton, and to the 2ℓ channel if it has two leptons. The 1ℓ channel uses a multiclassifier DNN to differentiate between signal, single Higgs, and background events. A parameterised DNN is used to extract the BSM results.

The $b\bar{b}\gamma\gamma$ signal contaminates the WW$\gamma\gamma$ phase space in the fully hadronic channel. Therefore, two binary classifiers are used for the fully hadronic channel, one acting as a $b\bar{b}\gamma\gamma$ discriminator and the other distinguishing between signal and background events. A specific challenge in the fully hadronic channel is differentiating between the ZZ$\gamma\gamma$ and the WW$\gamma\gamma$ signal samples. There are residual $b\bar{b}\gamma\gamma$ events that pass the WW$\gamma\gamma$ event selection but are not favoured by the $b\bar{b}\gamma\gamma$ analysis because of low $b\bar{b}\gamma\gamma$ discriminator scores. Therefore, the combination of WW$\gamma\gamma$, ZZ$\gamma\gamma$, and residual $b\bar{b}\gamma\gamma$ processes is considered signal in this analysis. Finally, because of the limited size of the samples available for training and validating a multivariate event classifier, a selection based on traditional variables is adopted for the 2ℓ channel.

Additionally, even after applying the $b\bar{b}\gamma\gamma$ DNN discriminator, there are residual $b\bar{b}\gamma\gamma$ events that can pass the WW$\gamma\gamma$ event selection and are not favoured by the $b\bar{b}\gamma\gamma$ channel analysis because of their DNN scores. Therefore, the final signal consist of a combination of WW$\gamma\gamma$, ZZ$\gamma\gamma$, and these residual $b\bar{b}\gamma\gamma$ events.

5   Systematic uncertainties

A number of systematic uncertainties are considered, affecting HH signal and background processes' yield (i.e. normalisation uncertainties), shape (i.e. shape uncertainties) or both. The systematic uncertainties are introduced as nuisance parameters in the maximum likelihood fit used to extract the results. Common uncertainties among different analyses are fully correlated, while others specific to each analysis are treated as uncorrelated.

5.1   Theory uncertainties

Theoretical uncertainties on the QCD scale, strong coupling constant α_S, and parton density functions (PDFs) affecting the cross section of all the simulated processes are included. Other

uncertainties concern electroweak corrections for the $t\bar{t}Z$, $t\bar{t}W$, $q\bar{q}ZZ$, and ggZZ processes. Theoretical uncertainties on the branching fractions of Higgs boson decays are applied on the HH signal and single Higgs background, which are 1.2%, 1.5%, 2.1%, and 1.6% for the H $\rightarrow$ b$\bar{\text{b}}$, H $\rightarrow$ WW and H $\rightarrow$ ZZ, H $\rightarrow$ $\gamma\gamma$ and H $\rightarrow$ $\tau\tau$ branching fractions, respectively. The uncertainty in the rate of Higgs boson associated production with heavy-flavour jets is assumed to be 50% [75].

Theoretical uncertainties on the nonresonant HH cross section via gluon-gluon fusion are applied as a function of κ_λ and include a combination of uncertainties in the renormalisation and factorisation scales, as well as the uncertainty in the top quark mass. The PDF uncertainty is 3% [76]. Uncertainties for VBF production include $+0.03\%/-0.04\%$ (scale) and 2.1% (PDF+α_S) [41, 77]. An additional scale uncertainty is applied to some analyses related to PYTHIA and the colour-correlated recoil effect. This effect was found to be negligible in other analyses.

The uncertainty in the top quark mass value assumed in the simulations of the $t\bar{t}$ background is derived by varying the mass value by $-2.7\%/+2.8\%$ [78]. Some analyses suffer from a significant $t\bar{t}$ background contribution (b$\bar{\text{b}}$WW, multilepton, and b$\bar{\text{b}}\tau\tau$) and a shape uncertainty that corresponds to the NNLO correction on the top quark p_T is applied to the $t\bar{t}$ simulated samples. The b$\bar{\text{b}}$WW and multilepton analyses include additional uncertainties on the parton shower initial- and final-state radiation, as well as factorisation scale uncertainties for $t\bar{t}$ and single top quark production.

5.2 Experimental uncertainties

The integrated luminosities for the 2016, 2017, and 2018 data-taking years have 1.2–2.5% uncertainty per year [79–81], while the overall uncertainty for the 2016–2018 period is 1.6%. The number of pp interactions is calculated from the instantaneous luminosity and an estimated inelastic pp collision cross section with 5% uncertainty. An uncertainty in the shape of the distribution of the mean number of pp interactions per bunch crossing (pielup) is also applied to all simulated samples. The trigger selection efficiency is corrected in each analysis to account for the differences between the data and simulation.

When identification criteria are applied to muons, electrons, and τ_h, the identification efficiencies in the simulated samples are corrected to match those in data. The uncertainties related to these corrections affect the shapes of the kinematic distributions. For analyses that use the same selection criteria, these uncertainties are treated as correlated. For analyses that use photons, the shape uncertainties related to the photon identification efficiency are also considered.

Similarly, the uncertainties in the efficiency of the jet selection criteria in simulation, are considered as uncertainties on the shape of the kinematic distributions. Nuisance parameters that affect the shape of the jet flavour discriminant are also included. Different types of contamination are treated with separate parameters. Uncertainties in the jet energy scale and resolution are used as shape uncertainties for all simulated samples. For analyses where E_T^miss is relevant, uncertainties related to the unclustered energy reconstruction efficiency are taken into account as E_T^miss distribution shape uncertainties in simulated samples.

During data collection in 2016 and 2017, a gradual change in the timing of the ECAL L1 trigger primitives in the region with $|\eta| > 2.0$ caused a specific trigger inefficiency. For events containing an electron or a jet with p_T larger than $\approx$50 GeV ($\approx$100 GeV), in the region with $2.5 < |\eta| < 3.0$ the efficiency loss is $\approx$10–20%, depending on p_T, η, and time. Correction factors are computed from data and applied to the acceptance evaluated by simulation. In addition, a

normalisation uncertainty is included in the statistical fit.

A correction factor is applied to simulated data for the 2018 era to account for two HCAL modules being switched off, and an associated shape uncertainty is introduced. Analyses that perform data-driven background estimation methods include uncertainties related to each method. These are treated as uncorrelated among different analyses.

Finally, uncertainties from the limited number of simulated events are taken into account using the Barlow–Beeston approach [82]. These uncertainties are not considered in the case of the HL-LHC projections, as it is assumed that we will have sufficient simulated data at the time.

6 Results

The data are interpreted in several ways. Upper limits on HH production cross sections and constraints on coupling modifiers that contribute to SM HH production are derived based on the asymptotic formulae for the profile likelihood ratio test statistic $-2\Delta \log(L)$ [83, 84] and the CL_s [85, 86] criterion. In addition, using the HEFT parametrisation, we provide a number of BSM results. In all cases, expected limits are derived under the background-only hypothesis. These results have been determined using the CMS statistical analysis tool COMBINE [87], which is based on the ROOFIT [88] and ROOSTATS [89] frameworks.

First, we set upper limits at 95% CL on the inclusive HH production cross section. The limit is observed (expected) to be 3.5 (2.5) times the SM prediction. The equivalent upper limit for the VBF alone is observed (expected) to be 79 (91) times the SM prediction. In Fig. 5, we show the limits on the inclusive HH cross section divided by the SM prediction, i.e. the inclusive signal strength r, for each contributing channel and their combination. We also show the limits on the VBF cross section, divided by the SM prediction, i.e. the VBF signal strength $r_{\text{VBF HH}}$, in Fig. 6.


Figure 5: The 95% CL upper limits on the inclusive signal strength $r = \sigma_{\text{HH}}/\sigma_{\text{HH}}^{\text{SM}}$ for each channel and their combination. The inner (green) band and the outer (yellow) band indicate the 68 and 95% CL intervals, respectively, under the background-only hypothesis. The $b\bar{b}b\bar{b}$ and $b\bar{b}WW$ contributions have been combined in order to simplify the presentation of results.
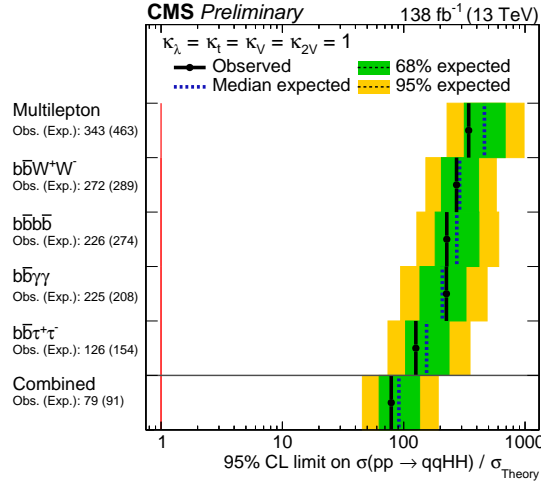
Figure 6: The 95% CL upper limits on the VBF signal strength $r_{\text{VBF HH}} = \sigma_{\text{VBF HH}}/\sigma_{\text{VBF HH}}^{\text{SM}}$ for each channel and their combination. The inner (green) band and the outer (yellow) band indicate the 68 and 95% CL intervals, respectively, under the background-only hypothesis. Not all of the eight channels are used to extract the combined VBF limit. The contributing channels are indicated in the figure. The $b\bar{b}b\bar{b}$ and $b\bar{b}WW$ contributions have been combined in order to simplify the presentation of results.

In Fig. 7, the upper limits at 95% CL on the HH cross section are shown as a function of κ_λ (left) and κ_{2V} (right), respectively. We exclude HH production at 95% CL for values of κ_λ outside the range between -1.39 and 7.02. The equivalent expected range is between -1.02 and 7.19. Similarly, we exclude HH production for the κ_{2V} coupling modifier outside the range between 0.62 and 1.42, with the expected range between 0.69 and 1.35.
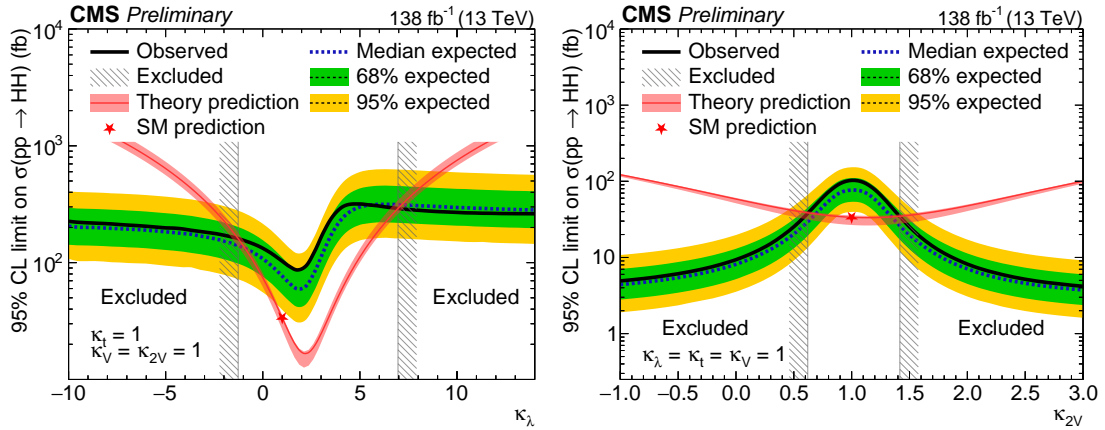

Figure 7: The 95% CL upper limits on the inclusive HH cross section as a function of κ_λ (left) and κ_{2V} (right), respectively. All other couplings are set to the values predicted by the SM. The theoretical uncertainties in the HH ggF and VBF signal cross sections are not considered because here we directly constrain the measured cross section. The inner (green) band and the outer (yellow) band indicate the 68 and 95% CL intervals, respectively, under the background-only hypothesis. The star shows the limit at the SM value for κ_λ and κ_{2V}, respectively.

The profile likelihood ratio test statistic $-2\Delta\log(L)$ as a function of κ_λ and κ_{2V} is shown in Fig. 8. Besides the ones scanned, all other couplings are set to the values expected by the SM.

All eight channels are included in the profile likelihood ratio fits. The best fit value of κ_λ is found to be 1.51 and κ_λ is constrained at $1\,\sigma$ (68.3% CL) to be within -0.08 and 4.20, while

the expected constraint is from -0.88 to 6.32. At $2\,\sigma$ (95.4% CL), κ_λ is constrained to be within -1.40 to 6.43, with an expected constraint of -2.29 to 7.95. For κ_{2V}, the best fit value is at 1.02, the $1\,\sigma$ (68.3% CL) interval is from 0.81 to 1.23 (0.77 to 1.26 expected), and the $2\,\sigma$ (95.4% CL) interval is from 0.63 to 1.40 (0.62 to 1.41 expected). The value of $\kappa_{2V} = 0$ is excluded with a significance of more than $6\,\sigma$.
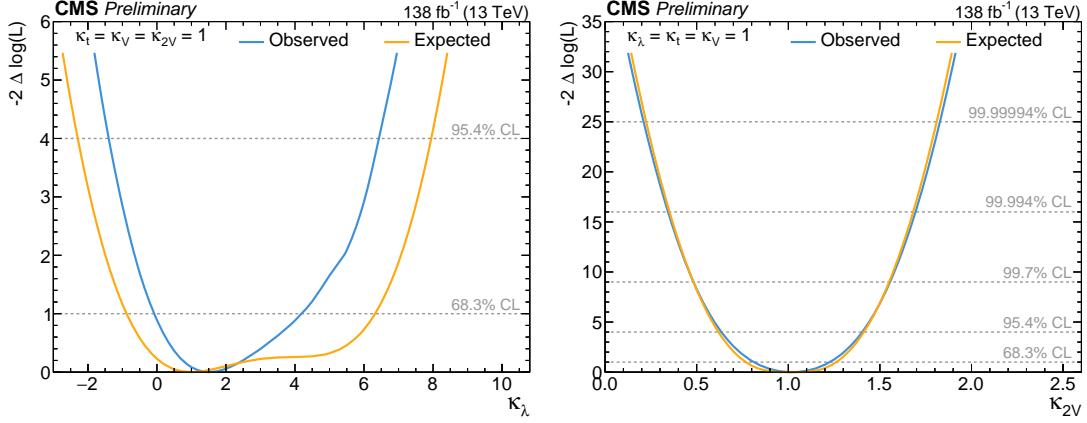

Figure 8: The profile likelihood ratio test statistic $-2\Delta \log(L)$ as a function of coupling modifiers κ_λ (left) and κ_{2V} (right) for the combination of all channels.

Figure 9 shows the two-dimensional contours of the $-2\Delta \log(L)$ in the $(\kappa_\lambda, \kappa_{2V})$, (κ_V, κ_{2V}), and $(\kappa_\lambda, \kappa_t)$ planes for the combination of all contributing channels. All the other parameters, besides the ones scanned, are set to the values expected by the SM. For the scan of (κ_V, κ_{2V}), there is a slight degeneracy in κ_V, while for the $(\kappa_\lambda, \kappa_t)$ scan, there is a degeneracy in κ_t. Only the global minima are pictured. The degeneracy is expected given that besides the interference terms, the HH production cross section has a quadratic dependence in κ_V and κ_t. No significant deviation from the SM is observed.
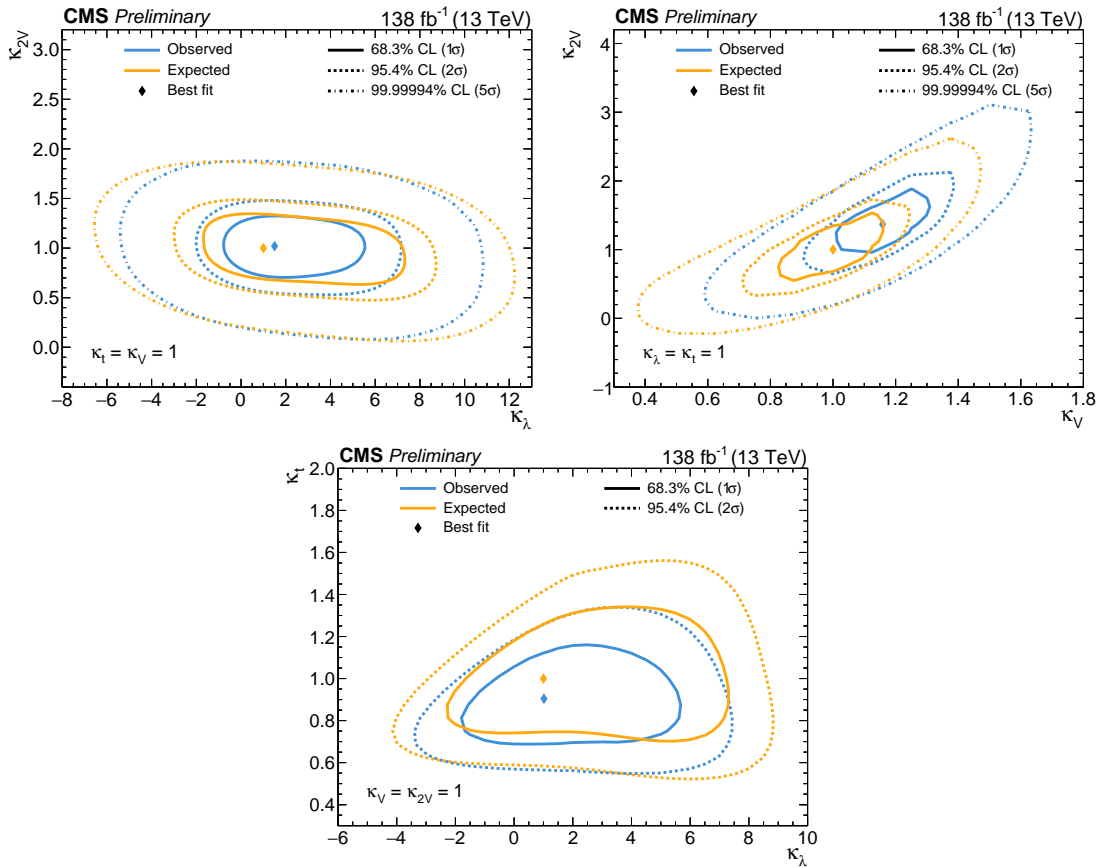
Figure 9: Profile likelihood ratio test statistic $-2\Delta\log(L)$ scans as a function of pairs of coupling modifiers $(\kappa_\lambda, \kappa_{2V})$ (top left), (κ_V, κ_{2V}) (top right), and $(\kappa_\lambda, \kappa_t)$ (bottom) for the combination of all channels when all the other parameters are fixed to their SM value.

Beyond varying κ_λ, κ_{2V}, κ_V, and κ_t to values different from the SM expectation, we interpret the data in terms of anomalous couplings that are not predicted in the SM, namely c_2, c_g, and c_{2g}. All BSM interpretations studied in this note only alter the ggF production, while the VBF production is assumed to be as predicted in the SM. The channels contributing to the BSM interpretations are $b\bar{b}\gamma\gamma$, $b\bar{b}b\bar{b}$ boosted and resolved, $b\bar{b}\tau\tau$, $b\bar{b}WW$, multilepton, $WW\gamma\gamma$, and $\tau\tau\gamma\gamma$.

First, we interpret the results in the context of two sets of benchmarks, combinations of the coupling modifiers (κ_λ, κ_t, c_2, c_g, c_{2g}) as described in Section 3. The $\tau\tau\gamma\gamma$ channel only contributes to the results for the benchmark of Ref. [33], while the rest of the channels contribute to both sets. The upper limits on the HH cross section at 95% CL are shown in Fig. 10. No significant deviations from expectations are observed, but there is an overall excess in all benchmarks, between 1 and 2 σ.
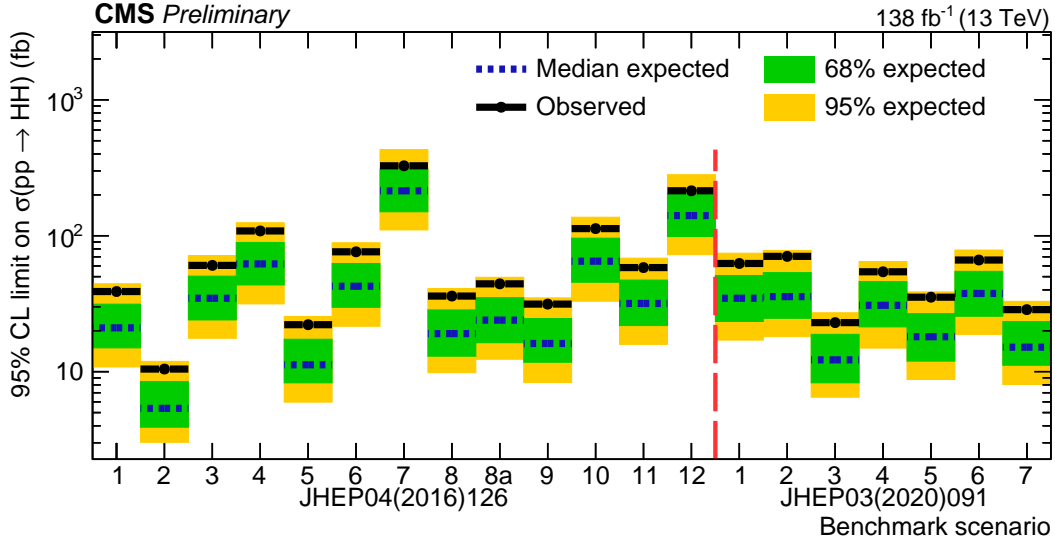

Figure 10: Upper limits on the HH production cross section at 95% CL for the two sets of HEFT benchmarks. The theoretical uncertainties in the HH ggF signal cross section are not considered because we directly constrain the measured cross section.

In the HEFT Lagrangian, the term containing κ_t is correlated with the coupling modifier c_2, which corresponds to the BSM coupling between two top quarks and two Higgs bosons. Figure 11 shows the upper limits on the HH cross section as a function of c_2 on the left and the $-2\Delta\log(L)$ scan for c_2 on the right. Taking into account the theoretical uncertainties in the HH ggF and VBF signal cross sections, we exclude HH production at 95% CL for c_2 values outside the range from -0.28 to 0.59. The corresponding expected range is between -0.17 and 0.47. For c_2, the best fit to the data is found to be 0.40, the 1 σ (68.3% CL) interval is from 0.23 to 0.51 (-0.15 to 0.41 expected), and the 2 σ (95.4% CL) interval is from -0.29 to 0.63 (-0.27 to 0.56 expected).
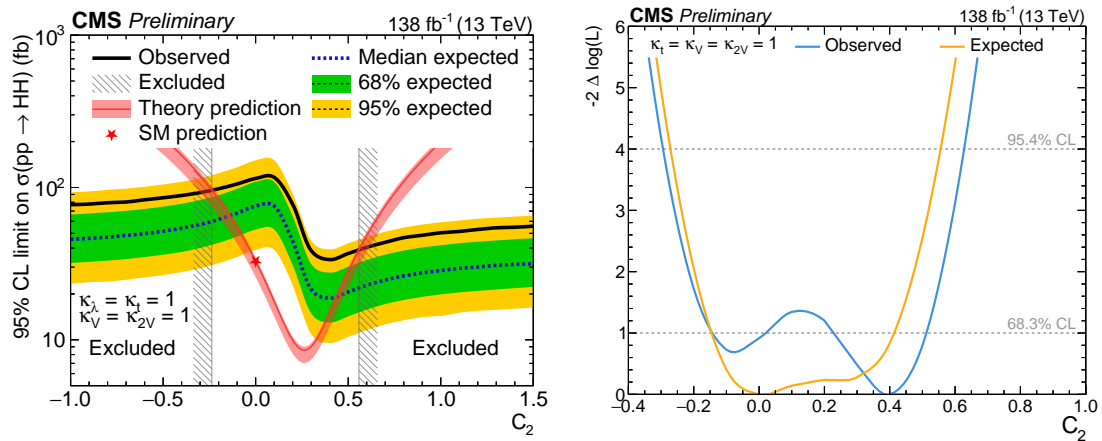
Figure 11: On the left, upper limits on the HH cross section as a function of the c_2 coupling modifier. The theoretical uncertainties in the HH ggF signal cross section are not considered because we directly constrain the measured cross section. On the right, the profile likelihood ratio test statistic $-2\Delta \log(L)$ as a function of the c_2 coupling modifier.

7 Projections to HL-LHC

To further probe the SM, with particular interest in the Higgs boson self-coupling, the High-Luminosity LHC (HL-LHC) [90] operations are currently scheduled to begin in 2030. In this new phase, the accelerator will deliver the highest ever instantaneous luminosity and energy reached at a collider. In these running conditions, an integrated luminosity up to $3000\,\mathrm{fb}^{-1}$ is expected to be collected over the anticipated ten years of data-taking. This unparalleled dataset will open a unique window on the weak-scale nature of the universe, and the study of the Higgs boson self-coupling represents one of the most important targets of the HL-LHC. It is therefore of interest to extrapolate the current results to predict the sensitivity that can be achieved at the HL-LHC.

The results are projected to different integrated luminosity values, i.e. 300, 1000, 2000, and $3000\,\mathrm{fb}^{-1}$ to track the evolution during the data taking. In the projection studies, only the $\mathrm{b\bar{b}}\gamma\gamma$, $\mathrm{b\bar{b}}\tau\tau$, $\mathrm{b\bar{b}b\bar{b}}$, multilepton, and $\mathrm{b\bar{b}}WW$ channels are included as they are the ones with the greatest sensitivity. The procedure for the statistical combination is the same as described in Sections 5 and 6.

The extrapolation of the Run 2 results to HL-LHC follows the same recipe used in previous projections of Higgs boson searches and measurements [9]. For the projections to a certain integrated luminosity L, the Run 2 signal and background yields, as well as the selected data events, are scaled up by a factor k_L equal to the increase of integrated luminosity with respect to Run 2. The scaling of the data events is necessary for those analyses that rely on the data to properly model the backgrounds in the fit. The efficiency of the physics object reconstruction and identification is assumed to be the same as in Run 2. The same assumption is made regarding the experimental energy or momentum resolution of the physics objects. This is based on the premise that the upgraded CMS detector will ensure performance comparable to Run 2 despite the larger pileup and radiation damage to the detector components.

The exact level of systematic uncertainties in CMS during HL-LHC is unknown. Therefore, we derive the projections in three different scenarios of systematic uncertainties:

- "S1": The systematic uncertainties are assumed to be at the same level of Run 2.
- "S2": The systematic uncertainties with a statistical origin, e.g. statistical uncertainties in data/MC scale factors, are reduced by a factor $\sqrt{k_L}$, until "floor" values are reached. The "floor" values prevent uncertainties from becoming unreasonably small and are based on studies in Ref. [9]. The theoretical uncertainties in the signal and background cross sections are halved to account for the expected progress in the theory calculations throughout the next years. The uncertainties originating from the limited size of the MC samples are also removed under the assumption of very large MC data sets.
- "stat. only": No systematic uncertainties are considered in the fit.

It should be noted that the systematic uncertainties related to specific issues encountered in Run 2 have been removed from all three scenarios, including S1. The nominal scenario for the HL-LHC conditions is S2.

The treatment of the uncertainties common to multiple analysis channels is summarised in Table 4. The analysis-specific uncertainties are treated case by case.

The projected upper limits on the HH signal strength for the statistical combination of the considered HH channels at different integrated luminosities are shown in the left panel of Fig. 12. Assuming SM values for HH production, we will become sensitive to excesses in HH inclus-

Table 4: Treatment of most important common systematic uncertainties in the S2 scenario.

Uncertainty	Scaling with respect to Run 2
Theory uncertainties	$1/2$
Stat. uncertainties in MC simulation	Removed
b-tag efficiency stat. component	$1/\sqrt{k_L}$
b-tag efficiency (non stat. component)	Unchanged
AK4 jet scale absolute	$\max(0.3, 1/\sqrt{k_L})$
AK4 jet scale flavor	$\max(0.5, 1/\sqrt{k_L})$
AK4 jet scale relative	$\max(0.2, 1/\sqrt{k_L})$
AK4 jet scale method	$1/\sqrt{k_L}$
AK4 jet energy resolution	$\max(0.5, 1/\sqrt{k_L})$
E_T^{miss}	$\max(0.5, 1/\sqrt{k_L})$
Luminosity	0.6
τ_h ID	Unchanged
τ_h Energy scale	Unchanged
Pileup	Unchanged
Run 2 issues	Removed

sive cross section for integrated luminosities larger than $1000\,\mathrm{fb}^{-1}$, when the expected upper limit drops below one. The projections at $3000\,\mathrm{fb}^{-1}$ under the different systematic uncerainty scenarios are shown in the right panel of Fig. 12.

The projected κ_λ likelihood scans are shown in Fig. 13. In the S2 scenario, the expected $1\,\sigma$ uncertainty on κ_λ is $+80\%/-60\%$ and $+60\%/-50\%$ for an integrated luminosity of 2000 and $3000\,\mathrm{fb}^{-1}$, respectively.

The expected significance for the HH signal strength assuming SM values for all parameters is summarised in Table 5. CMS will observe evidence for HH production by the end of HL-LHC, with significance $3.2\,\sigma$ for $2000\,\mathrm{fb}^{-1}$ at the nominal systematic uncertainty scenario and assuming that the signal will be SM-like. Figure 14 on the left shows the significance that can be achieved for the SM signal as a function of the integrated luminosity. On the right, the significance is shown as a function of κ_λ. The sensitivity to the HH signal varies with κ_λ due to the effects of the interference between the box and triangle diagrams, discussed in Section 1. The two diagrams have different kinematic properties, therefore, the interference varies not only the HH cross section but the HH signal acceptance rate as well. For $\kappa_\lambda < 1$ and $\kappa_\lambda > 5$, we would be able to observe evidence sooner. For $1 < \kappa_\lambda < 5$ the HH signal is suppressed, and the significance will be lower than expected, with minimum significance at $\kappa_\lambda = 3.4$. The projections presented here do not take into account potential improvements in triggering [91, 92], object reconstruction and selection such as b tagging, or analysis techniques. Historically, advancements in software and analysis methods have allowed us to exceed expectations, therefore the projections shown in this note are most likely conservative.

Figure 12: Expected upper limits on the HH signal strength from the combination of all the considered channels at different integrated luminosities (left), and under different assumptions on the systematic uncertainties for an integrated luminosity of $3000\,\mathrm{fb}^{-1}$ (right).

Figure 13: Expected κ_λ likelihood scan from the combination of all the considered channels projected at different integrated luminosities (left), and under different assumptions on the systematic uncertainties for an integrated luminosity of $3000\,\mathrm{fb}^{-1}$ (right).
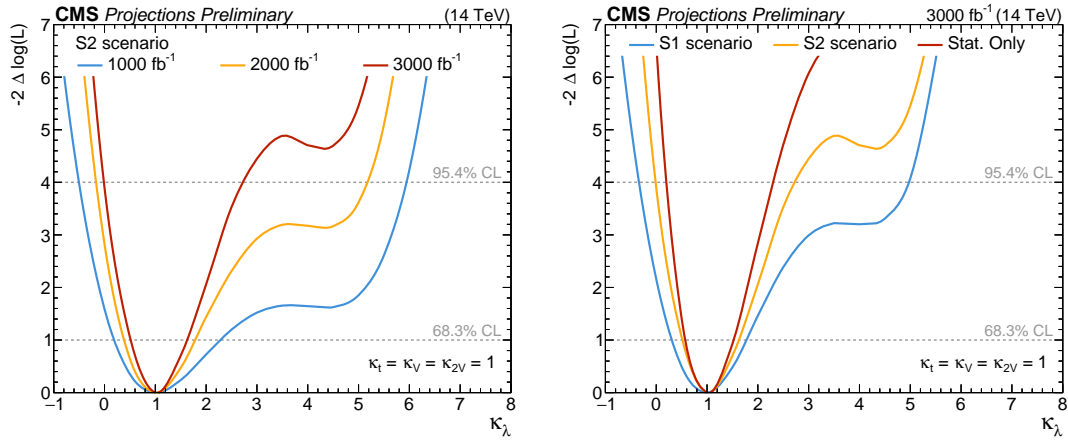
Table 5: Expected significance for the HH signal projected to 2000 or $3000\,\mathrm{fb}^{-1}$ under different assumptions of systematic uncertainties.

|	Significance (σ) at $2000\,\mathrm{fb}^{-1}$		Significance (σ) at $3000\,\mathrm{fb}^{-1}$	
	S2	Stat. only	S2	Stat. only
$b\bar{b}b\bar{b}$ resolved jets	1.0	1.3	1.4	1.6
$b\bar{b}b\bar{b}$ merged jets	1.7	1.7	2.0	2.1
$b\bar{b}\tau\tau$	1.7	1.9	2.1	2.3
$b\bar{b}WW$	0.6	0.8	0.7	0.9
$b\bar{b}\gamma\gamma$	1.8	1.9	2.2	2.3
Combination	3.2	3.6	3.8	4.3

Figure 14: Expected signal significance as a function of integrated luminosity (left). The green curve shows the significance using the same scheme for systematic uncertainties as the nominal Run 2 result for consistency with Section 6 for 138 and $300\,\text{fb}^{-1}$. The subsequent points use the S2 scenario for systematics. The magenta curve shows the significance for statistical uncertainties only. Expected signal significance as a function of κ_λ under different assumptions on the systematic uncertainties for an integrated luminosity of $3000\,\text{fb}^{-1}$ (right).

8 Summary

A combined search for nonresonant Higgs boson pair (HH) production is performed using the proton-proton collision data set produced by the LHC at $\sqrt{s} = 13\,\text{TeV}$, and collected by the CMS detector from 2016 to 2018 (Run 2), which corresponds to an integrated luminosity of $138\,\text{fb}^{-1}$. Searches for HH production via gluon-gluon (ggF) and vector boson fusion (VBF) production, are carried out in the $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, $b\bar{b}b\bar{b}$, $b\bar{b}WW$, and multilepton channels. Additionally, the gluon-gluon fusion Higgs boson pair production is also searched for in the $b\bar{b}ZZ$ with ZZ decaying to four leptons, $WW\gamma\gamma$, and $\tau\tau\gamma\gamma$ final states, which have clean signatures but relatively small branching fractions. The associated production mechanism with a vector boson is searched for in the $b\bar{b}b\bar{b}$ final state with the largest branching fraction. The analyses of these channels are combined to probe the Higgs boson trilinear self-coupling, the quartic coupling between two vector bosons and two Higgs bosons (VVHH) and to search for beyond the standard model physics scenarios in the Higgs effective field theory approach.

The observed and expected upper limits at 95% confidence level (CL) on the cross section of gluon-gluon fusion Higgs boson pair production are found to be 3.5 and 2.5 times the standard model expectations. For the vector boson fusion production, the observed and expected upper limit at 95% CL are 79 and 91 times the standard model expectations. One-dimensional scans of coupling modifiers are performed. When all other parameters are assumed to be as expected from the SM, we (expect to) exclude HH production at 95% CL when the Higgs boson trilinear self-coupling modifier κ_λ is outside the range from -1.39 to 7.02 (-1.02 to 7.19). Equivalently, HH production is excluded when the VVHH coupling modifier κ_{2V} is outside the range from 0.62 to 1.42 (0.69 to 1.35 expected).

Two-dimensional measurements are also performed, including simultaneous measurements of κ_λ and κ_{2V}, κ_λ and the modifier of the Higgs boson coupling to the top quark (κ_t), and κ_{2V} and the modifier of the Higgs boson coupling to vector bosons (κ_V). The results are in agreement with the standard model predictions.

Under an effective field theory framework, the cross section of the nonresonant ggF HH pair production is parameterised as a function of anomalous couplings of the Higgs boson, involving the contact interactions between two Higgs and two top quarks, between two gluons and two Higgs bosons, and between two gluons and a Higgs boson. We perform searches for benchmark signals under different anomalous coupling scenarios and set upper limits on their cross sections at 95% CL. We (expect to) exclude HH production at 95% CL when the coupling of the contact interaction between two Higgs and two top quarks is outside the range between -0.28 to 0.59 (-0.17 to 0.47).

These results present the most stringent limits and constraints obtained from the searches for nonresonant Higgs boson pair production using the LHC Run 2 data set collected by the CMS detector. Extrapolating our current results to the luminosity of HL-LHC it can be expected to see first evidence for Higgs boson pair production with $\approx 2000\,\text{fb}^{-1}$ of data.

References

[1] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC", *Phys. Lett. B* **716** (2012) 30, `doi:10.1016/j.physletb.2012.08.021`, `arXiv:1207.7235`.

[2] CMS Collaboration, "Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV", *JHEP* **06** (2013) 081, `doi:10.1007/JHEP06(2013)081`, `arXiv:1303.4571`.

[3] ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC", *Phys. Lett. B* **716** (2012) 1, `doi:10.1016/j.physletb.2012.08.020`, `arXiv:1207.7214`.

[4] CMS Collaboration, "Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV", *Eur. Phys. J. C* **79** (2019) 421, `doi:10.1140/epjc/s10052-019-6909-y`, `arXiv:1809.10733`.

[5] ATLAS Collaboration, "Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment", *Phys. Rev. D* **101** (2020) 012002, `doi:10.1103/PhysRevD.101.012002`, `arXiv:1909.02845`.

[6] A. Papaefstathiou and G. White, "The electro-weak phase transition at colliders: confronting theoretical uncertainties and complementary channels", *JHEP* **05** (2021) 099, `doi:10.1007/JHEP05(2021)099`, `arXiv:2010.00597`.

[7] A. Mazumdar and G. White, "Review of cosmic phase transitions: their significance and experimental signatures", *Rept. Prog. Phys.* **82** (2019) 076901, `doi:10.1088/1361-6633/ab1f55`, `arXiv:1811.01948`.

[8] V. Brigljevic et al., "HHH Whitepaper", 2024. `arXiv:2407.03015`.

[9] A. Dainese et al., "Report from Working Group 2: Higgs physics at the HL-LHC and HE-LHC", *CERN Yellow Rep. Monogr.* **7** (2019) 221, `doi:10.23731/CYRM-2019-007.221`, `arXiv:1902.00134`.

[10] T. Plehn and M. Rauch, "The quartic higgs coupling at hadron colliders", *Phys. Rev. D* **72** (2005) 053008, `doi:10.1103/PhysRevD.72.053008`, `arXiv:hep-ph/0507321`.

[11] S. Borowka et al., "Higgs boson pair production in gluon fusion at next-to-leading order with full top-quark mass dependence", *Phys. Rev. Lett.* **117** (2016) 012001, `doi:10.1103/PhysRevLett.117.079901`, `arXiv:1604.06447`. [Erratum: `doi:10.1103/PhysRevLett.117.079901`].

[12] J. Baglio et al., "Gluon fusion into Higgs pairs at NLO QCD and the top mass scheme", *Eur. Phys. J. C* **79** (2019) 459, `doi:10.1140/epjc/s10052-019-6973-3`, `arXiv:1811.05692`.

[13] F. A. Dreyer and A. Karlberg, "Vector-boson fusion Higgs pair production at n^{3}lo", *Phys. Rev. D* **98** (2018) 114016, `doi:10.1103/PhysRevD.98.114016`, `arXiv:1811.07906`.

[14] L. Alasfar et al., "Effective field theory descriptions of Higgs boson pair production", *SciPost Phys. Comm. Rep.* (2024) 2, `doi:10.21468/SciPostPhysCommRep.2`, `arXiv:2304.01968`.

[15] ATLAS Collaboration, "Combination of searches for Higgs boson pair production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", *Phys. Rev. Lett.* **133** (2024) 101801, `doi:10.1103/PhysRevLett.133.101801`, `arXiv:2406.09971`.

[16] CMS Collaboration, "A portrait of the Higgs boson by the CMS experiment ten years after the discovery", *Nature* **607** (2022) 60, `doi:10.1038/s41586-022-04892-x`, `arXiv:2207.00043`. [Erratum: `doi:10.1038/s41586-023-06164-8`].

[17] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, `doi:10.1088/1748-0221/3/08/S08004`.

[18] CMS Collaboration, "Development of the CMS detector for the CERN LHC Run 3", *JINST* **19** (2024) P05064, `doi:10.1088/1748-0221/19/05/P05064`.

[19] CMS Collaboration, "Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **15** (2020) P10017, `doi:10.1088/1748-0221/15/10/P10017`, `arXiv:2006.10165`.

[20] CMS Collaboration, "The CMS trigger system", *JINST* **12** (2017) P01020, `doi:10.1088/1748-0221/12/01/P01020`, `arXiv:1609.02366`.

[21] CMS Collaboration, "Performance of the CMS high-level trigger during LHC Run 2", 2024. `arXiv:2410.17038`. Submitted to *JINST*.

[22] CMS Collaboration, "Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC", *JINST* **16** (2021) P05014, `doi:10.1088/1748-0221/16/05/P05014`, `arXiv:2012.06888`.

[23] CMS Collaboration, "Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018) P06015, `doi:10.1088/1748-0221/13/06/P06015`, `arXiv:1804.04528`.

[24] CMS Collaboration, "Description and performance of track and primary-vertex reconstruction with the CMS tracker", *JINST* **9** (2014) P10009, `doi:10.1088/1748-0221/9/10/P10009`, `arXiv:1405.6569`.

[25] CMS Collaboration, "Particle-flow reconstruction and global event description with the CMS detector", *JINST* **12** (2017) P10003, `doi:10.1088/1748-0221/12/10/P10003`, `arXiv:1706.04965`.

[26] 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.

[27] CMS Collaboration, "Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector", *JINST* **14** (2019) P07004, `doi:10.1088/1748-0221/14/07/P07004`, `arXiv:1903.06078`.

[28] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_T jet clustering algorithm", *JHEP* **04** (2008) 063, `doi:10.1088/1126-6708/2008/04/063`, `arXiv:0802.1189`.

[29] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual", *Eur. Phys. J. C* **72** (2012) 1896, `doi:10.1140/epjc/s10052-012-1896-2`, `arXiv:1111.6097`.

[30] G. Heinrich et al., "Probing the trilinear Higgs boson coupling in di-Higgs production at NLO QCD including parton shower effects", *JHEP* **06** (2019) 066, `doi:10.1007/JHEP06(2019)066`, `arXiv:1903.08137`.

[31] D. de Florian et al., "Anomalous couplings in Higgs-boson pair production at approximate NNLO QCD", *JHEP* **09** (2021) 161, `doi:10.1007/JHEP09(2021)161`, `arXiv:2106.14050`.

[32] A. Carvalho et al., "On the reinterpretation of non-resonant searches for Higgs boson pairs", *JHEP* **02** (2021) 049, `doi:10.1007/JHEP02(2021)049`, `arXiv:1710.08261`.

[33] A. Carvalho et al., "Higgs pair production: Choosing benchmarks with cluster analysis", *JHEP* **04** (2016) 126, `doi:10.1007/JHEP04(2016)126`, `arXiv:1507.02245`.

[34] M. Capozi and G. Heinrich, "Exploring anomalous couplings in Higgs boson pair production through shape analysis", *JHEP* **03** (2020) 091, `doi:10.1007/JHEP03(2020)091`, `arXiv:1908.08923`.

[35] P. Nason, "A new method for combining NLO QCD with shower Monte Carlo algorithms", *JHEP* **11** (2004) 040, `doi:10.1088/1126-6708/2004/11/040`, `arXiv:hep-ph/0409146`.

[36] S. Frixione, P. Nason, and C. Oleari, "Matching NLO QCD computations with parton shower simulations: the POWHEG method", *JHEP* **11** (2007) 070, `doi:10.1088/1126-6708/2007/11/070`, `arXiv:0709.2092`.

[37] S. Alioli, P. Nason, C. Oleari, and E. Re, "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX", *JHEP* **06** (2010) 043, `doi:10.1007/JHEP06(2010)043`, `arXiv:1002.2581`.

[38] S. Amoroso et al., "Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report", in *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches*. 2020. `arXiv:2003.01700`.

[39] G. Heinrich, J. Lang, and L. Scyboz, "SMEFT predictions for $gg \to hh$ at full NLO QCD and truncation uncertainties", *JHEP* **08** (2022) 079, `doi:10.1007/JHEP08(2022)079`, `arXiv:2204.13045`. [Erratum: `doi:10.1007/JHEP10(2023)086`].

[40] E. Bagnaschi, G. Degrassi, and R. Gröber, "Higgs boson pair production at NLO in the POWHEG approach and the top quark mass uncertainties", *Eur. Phys. J. C* **83** (2023) 1054, `doi:10.1140/epjc/s10052-023-12238-8`, `arXiv:2309.10525`.

[41] F. A. Dreyer and A. Karlberg, "Fully differential Vector-Boson Fusion Higgs Pair Production at Next-to-Next-to-Leading Order", *Phys. Rev. D* **99** (2019) 074028, `doi:10.1103/PhysRevD.99.074028`, `arXiv:1811.07918`.

[42] J. Baglio et al., "The measurement of the Higgs self-coupling at the LHC: theoretical status", *JHEP* **04** (2013) 151, `doi:10.1007/JHEP04(2013)151`, `arXiv:1212.5581`.

[43] T. Sjostrand, S. Mrenna, and P. Z. Skands, "A brief introduction to PYTHIA 8.1", *Comput. Phys. Commun.* **178** (2008) 852, `doi:10.1016/j.cpc.2008.01.036`, `arXiv:0710.3820`.

[44] CMS Collaboration, "Search for nonresonant Higgs boson pair production in final states with two bottom quarks and two photons in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", *JHEP* **03** (2021) 257, `doi:10.1007/JHEP03(2021)257`, `arXiv:2011.12373`.

[45] CMS Collaboration, "Search for Higgs boson pair production in the $bb\tau\tau$ final state in proton-proton collisions at $\sqrt{s} = 8$ TeV", *Phys. Rev. D* **96** (2017) 072004, `doi:10.1103/PhysRevD.96.072004`, `arXiv:1707.00350`.

[46] CMS Collaboration, "Search for Higgs boson pair production in the four b quark final state in proton-proton collisions at $\sqrt{s} = 13$ TeV", *Phys. Rev. Lett.* **129** (2022) 081802, `doi:10.1103/PhysRevLett.129.081802`, `arXiv:2202.09617`.

[47] CMS Collaboration, "Search for nonresonant pair production of highly energetic Higgs bosons decaying to bottom quarks", *Phys. Rev. Lett.* **131** (2023) 041803, `doi:10.1103/PhysRevLett.131.041803`, `arXiv:2205.06667`.

[48] CMS Collaboration, "Search for Higgs boson pair production with one associated vector boson in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **10** (2024) 061, `doi:10.1007/JHEP10(2024)061`, `arXiv:2404.08462`.

[49] CMS Collaboration, "Search for Higgs boson pair production in the $b\bar{b}W^{+}W^{-}$ decay mode in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **07** (2024) 293, `doi:10.1007/JHEP07(2024)293`, `arXiv:2403.09430`.

[50] CMS Collaboration, "Search for highly energetic double Higgs boson production in the two bottom quark and two vector boson all-hadronic final state", CMS Physics Analysis Summary CMS-PAS-HIG-23-012, 2024.

[51] CMS Collaboration, "Search for Higgs boson pairs decaying to $WW^{*}WW^{*}$, $WW^{*}\tau\tau$, and $\tau\tau\tau\tau$ in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **07** (2023) 095, `doi:10.1007/JHEP07(2023)095`, `arXiv:2206.10268`.

[52] CMS Collaboration, "Search for nonresonant Higgs boson pair production in the $WW\gamma\gamma$ channel in pp collisions at $\sqrt{s} = 13$ TeV", CMS Physics Analysis Summary CMS-PAS-HIG-21-014, 2022.

[53] CMS Collaboration, "Search for nonresonant Higgs boson pair production in the four leptons plus twob jets final state in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JHEP* **06** (2023) 130, `doi:10.1007/JHEP06(2023)130`, `arXiv:2206.10657`.

[54] CMS Collaboration, "Search for the nonresonant and resonant production of a Higgs boson in association with an additional scalar boson in the $\gamma\gamma\tau\tau$ final state", CMS Physics Analysis Summary CMS-PAS-HIG-22-012, 2024.

[55] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, "Soft drop", *JHEP* **05** (2014) 146, `doi:10.1007/JHEP05(2014)146`, `arXiv:1402.2657`.

[56] I. Henrion et al., "Neural message passing for jet physics", in *Deep Learning for Physical Sciences Workshop at 31st Conf. on Neural Information Processing Systems*. Long Beach, CA, 2017.

[57] H. Qu and L. Gouskos, "Jet tagging via particle clouds", *Phys. Rev. D* **101** (2020) 056019, `doi:10.1103/PhysRevD.101.056019`, `arXiv:1902.08570`.

[58] E. A. Moreno et al., "JEDI-net: a jet identification algorithm based on interaction networks", *Eur. Phys. J. C* **80** (2020) 58, `doi:10.1140/epjc/s10052-020-7608-4`, `arXiv:1908.05318`.

[59] E. A. Moreno et al., "Interaction networks for the identification of boosted H → b$\overline{\text{b}}$ decays", *Phys. Rev. D* **102** (2020) 012010, `doi:10.1103/PhysRevD.102.012010`, `arXiv:1909.12285`.

[60] CMS Collaboration, "Performance of heavy-flavour jet identification in boosted topologies in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-BTV-22-001, 2023.

[61] CMS Collaboration, "Mass regression of highly-boosted jets using graph neural networks", CMS Detector Performance Note CMS-DP-2021-017, 2021.

[62] CMS Collaboration, "Search for a massive resonance decaying to a pair of Higgs bosons in the four b quark final state in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", *Phys. Lett. B* **781** (2018) 244, `doi:10.1016/j.physletb.2018.03.084`, `arXiv:1710.04960`.

[63] CMS Collaboration, "Search for production of Higgs boson pairs in the four b quark final state using large-area jets in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", *JHEP* **01** (2019) 040, `doi:10.1007/JHEP01(2019)040`, `arXiv:1808.01473`.

[64] CMS Collaboration, "Inclusive search for highly boosted Higgs bosons decaying to bottom quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", *JHEP* **12** (2020) 85, `doi:10.1007/JHEP12(2020)085`, `arXiv:2006.13251`.

[65] CMS Collaboration, "Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques", *JINST* **15** (2020) P06005, `doi:10.1088/1748-0221/15/06/P06005`, `arXiv:2004.08262`.

[66] CMS Collaboration, "Identification of hadronic tau lepton decays using a deep neural network", *JINST* **17** (2022) P07023, `doi:10.1088/1748-0221/17/07/P07023`, `arXiv:2201.08458`.

[67] L. Bianchini et al., "Reconstruction of the higgs mass in events with higgs bosons decaying into a pair of leptons using matrix element techniques", *Nucl. Instr. Meth. Phys. Res.* **862** (August, 2017) 54, `doi:10.1016/j.nima.2017.05.001`, `arXiv:1603.05910v3`.

[68] E. Bols et al., "Jet flavour classification using DeepJet", *JINST* **15** (2020) P12012, `doi:10.1088/1748-0221/15/12/P12012`, `arXiv:2008.10519`.

[69] CMS Collaboration, "Evidence for associated production of a Higgs boson with a top quark pair in final states with electrons, muons, and hadronically decaying τ leptons at $\sqrt{s} = 13\,\text{TeV}$", *JHEP* **08** (2018) 066, `doi:10.1007/JHEP08(2018)066`, `arXiv:1803.05485`.

[70] H. Qu, C. Li, and S. Qian, "Particle transformer for jet tagging", in *Proceedings of the 39th International Conference on Machine Learning*, K. Chaudhuri et al., eds., volume 162, p. 18281. 2022. `arXiv:2202.03772`.

[71] CMS Collaboration, "Performance of heavy-flavour jet identification in boosted topologies in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-BTV-22-001, 2023.

[72] CMS Collaboration, "Lund plane reweighting for jet substructure correction", CMS Detector Performance Note CMS-DP-2023-046, 2023.

[73] F. A. Dreyer, G. P. Salam, and G. Soyez, "The Lund jet plane", *JHEP* **12** (2018) 064, `doi:10.1007/JHEP12(2018)064`, `arXiv:1807.04758`.

[74] CMS Collaboration, "Measurement of the Higgs boson production rate in association with top quarks in final states with electrons, muons, and hadronically decaying tau leptons at $\sqrt{s} = 13\,\text{TeV}$", *Eur. Phys. J. C* **81** (2021) 378, `doi:10.1140/epjc/s10052-021-09014-x`, `arXiv:2011.03652`.

[75] S. Manzoni et al., "Taming a leading theoretical uncertainty in HH measurements via accurate simulations for b$\bar{\text{b}}$H production", *JHEP* **09** (2023) 179, `doi:10.1007/JHEP09(2023)179`, `arXiv:2307.09992`.

[76] J. Baglio et al., "$gg \to HH$: Combined uncertainties", *Phys. Rev. D* **103** (2021) 056002, `doi:10.1103/PhysRevD.103.056002`, `arXiv:2008.11626`.

[77] L.-S. Ling et al., "NNLO QCD corrections to Higgs pair production via vector boson fusion at hadron colliders", *Phys. Rev. D* **89** (2014) 073001, `doi:10.1103/PhysRevD.89.073001`, `arXiv:1401.7754`.

[78] D. de Florian, I. Fabre, and J. Mazzitelli, "Higgs boson pair production at NNLO in QCD including dimension 6 operators", *JHEP* **10** (2017) 215, `doi:10.1007/JHEP10(2017)215`, `arXiv:1704.05700`.

[79] CMS Collaboration, "Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13\,\text{TeV}$ in 2015 and 2016 at CMS", *Eur. Phys. J. C* **81** (2021) 800, `doi:10.1140/epjc/s10052-021-09538-2`, `arXiv:2104.01927`.

[80] CMS Collaboration, "CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-LUM-17-004, 2018.

[81] CMS Collaboration, "CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13\,\text{TeV}$", CMS Physics Analysis Summary CMS-PAS-LUM-18-002, 2019.

[82] R. Barlow and C. Beeston, "Fitting using finite Monte Carlo samples", *Comput. Phys. Commun.* **77** (1993) 219, `doi:10.1016/0010-4655(93)90005-W`.

[83] 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.

[84] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics", *Eur. Phys. J. C* **71** (2011) 1554, `doi:10.1140/epjc/s10052-011-1554-0`, `arXiv:1007.1727`. [Erratum: `doi:10.1140/epjc/s10052-011-1554-0`].

[85] A. L. Read, "Presentation of search results: the CL$_\text{s}$ technique", *J. Phys. G* **28** (2002) 2693, `doi:10.1088/0954-3899/28/10/313`.

[86] T. Junk, "Confidence level computation for combining searches with small statistics", *Nucl. Instrum. Meth. A* **434** (1999) 435, `doi:10.1016/S0168-9002(99)00498-2`, `arXiv:hep-ex/9902006`.

[87] CMS Collaboration, "The CMS statistical analysis and combination tool: COMBINE", *Comput. Softw. Big Sci.* **8** (2024) 19, `doi:10.1007/s41781-024-00121-4`, `arXiv:2404.06614`.

[88] W. Verkerke and D. P. Kirkby, "The RooFit toolkit for data modeling", *eConf* (2003)
`arXiv:physics/0306116`.

[89] L. Moneta et al., "The RooStats Project", *PoS* **ACAT2010** (2010) 057,
`doi:10.22323/1.093.0057,arXiv:1009.1003`.

[90] I. Zurbano Fernandez et al., "High-Luminosity Large Hadron Collider (HL-LHC):
Technical design report", *CERN Yellow Rep. Monogr.* **10** (2020)
`doi:10.23731/CYRM-2020-0010`.

[91] CMS Collaboration, "The Phase-2 upgrade of the CMS level-1 trigger", CMS Technical
Design Report CERN-LHCC-2020-004, CMS-TDR-021, 2020.

[92] CMS Collaboration, "The Phase-2 upgrade of the CMS data acquisition and high level
trigger", CMS Technical Design Report CERN-LHCC-2021-007, CMS-TDR-022, 2021.