

Search for **exotic Higgs boson decays** $H \rightarrow AA$

with $A \rightarrow \gamma\gamma$ in events with a **semi-merged topology**
in proton-proton collisions at 13 TeV

15/06/26

15th Course on Physics at the LHC 2026

Inês Moreira



Motivation | Combined results of Higgs boson decays by CMS

Nature - A portrait of the Higgs Boson by CMS

Article
A portrait of the Higgs boson by the CMS experiment ten years after the discovery

https://doi.org/10.1038/s41586-022-04892-4
Received: 21 March 2022
Accepted: 23 May 2022
Published online: 4 July 2022

The CMS Collaboration^{1*}

In July 2012, the ATLAS and CMS collaborations at the CERN Large Hadron Collider announced the observation of Higgs boson at around 125 GeV in proton-proton collisions. Ten years later, and with the data corresponding to the production of a 30 times larger number of Higgs bosons, we have learnt much more about the properties of the Higgs boson. The CMS experiment has observed the Higgs boson in numerous fermionic and bosonic decay channels, established its spin-parity quantum numbers, determined its mass and measured its production cross-sections in various modes. Here the CMS Collaboration reports the most up-to-date combination of results on the properties of

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)



CMS-HIG-21-018



CERN-EP-2026-009
2026/02/24

Combined measurements and interpretations of Higgs boson production and decay in proton-proton collisions at $\sqrt{s} = 13$ TeV

The CMS Collaboration¹

CMS - Combine measurements of Higgs

CMS - Exotic decay dark photons

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)



CMS-HIG-22-003



CERN-EP-2023-223
2024/04/09

Search for an exotic decay of the Higgs boson into a Z boson and a pseudoscalar particle in proton-proton collisions at $\sqrt{s} = 13$ TeV

CMS - Exotic decay Z + pseudoscalar

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)



CMS-EXO-19-007



CERN-EP-2019-159
2019/10/29

Search for dark photons in decays of Higgs bosons produced in association with Z bosons in proton-proton collisions at $\sqrt{s} = 13$ TeV

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)



CMS-EXO-24-025



CERN-EP-2025-251
2026/01/08

Search for exotic Higgs boson decays $H \rightarrow AA$ with $A \rightarrow \gamma\gamma$ in events with a semi-merged topology in proton-proton collisions at $\sqrt{s} = 13$ TeV

The CMS Collaboration¹

CMS - Exotic decay $H \rightarrow AA \rightarrow 4\gamma$

Latest CMS Higgs combination results (Run 2, 2015–2018):

- Most precise picture to date of Higgs production and decay rates

Constraints on BSM Higgs decays:

- $B(H \rightarrow \text{invisible}) < 13\%$ at 95% CL
- $B(H \rightarrow \text{undetected}) < 25\%$ at 95% CL

Sizeable room remains for exotic decays

Motivates dedicated direct searches for specific BSM final states

- \rightarrow e.g. $H \rightarrow AA \rightarrow 4\gamma$ (this analysis)

Motivation | A portrait of the Higgs boson by CMS

The Brout-Englert-Higgs mechanism



- Complex scalar field permeating the universe
- Gives mass to W, Z and fermions;
- Higgs boson = quantum excitation of the field

Mass as a free parameter

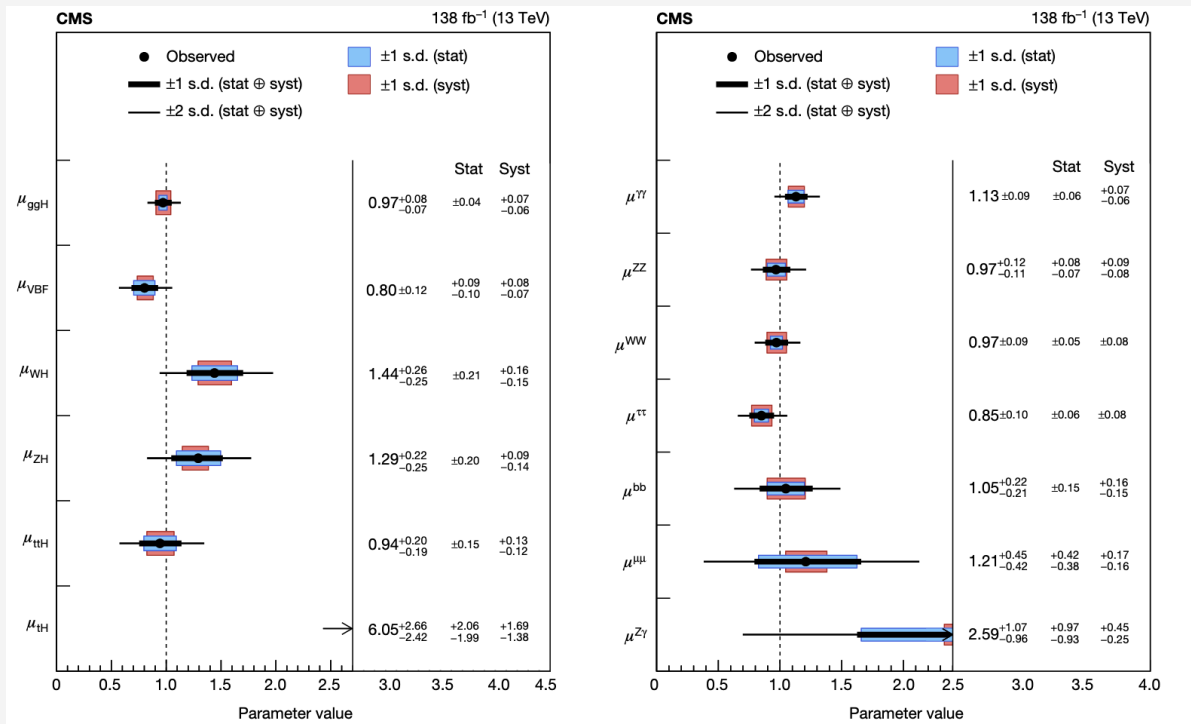
- Once m_H is fixed, other properties are predicted by the SM
- m_H puzzlingly small vs. BSM scale → motivates continued searches

Discovered in 2012 by ATLAS and CMS

- More than 10 years of precision measurements since



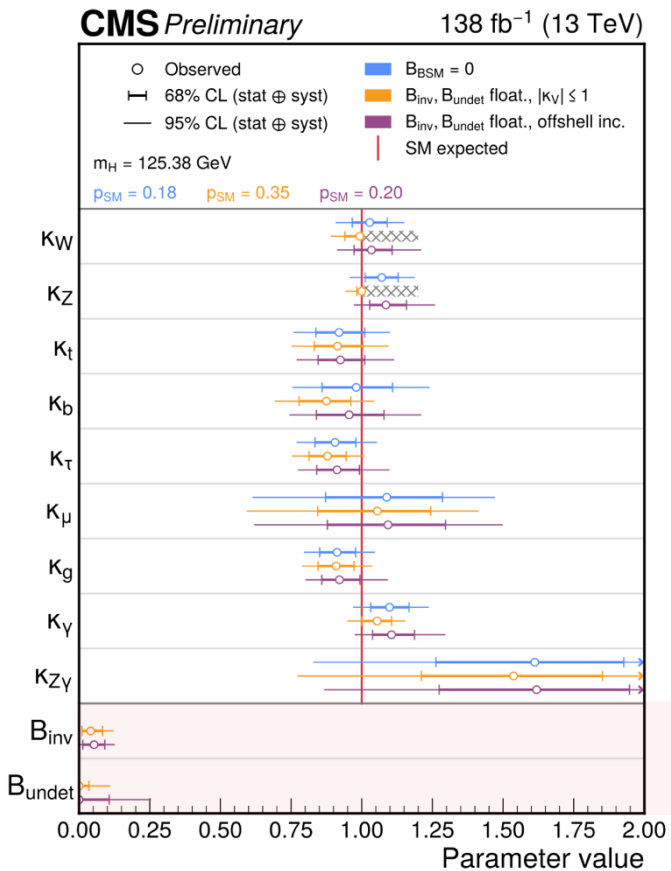
Motivation | The μ framework for signal strengths




- Fit the data with a model containing **signal-strength parameter μ**
- Scales observed yields vs. SM prediction
- $\mu = 1 \rightarrow$ perfect SM agreement
- Shape of distributions unchanged

- **Production:** $\mu_i = \sigma_i / \sigma_i^{SM}$ (left)
- **Decay:** $\mu_f = \mu_f = B_f / B_f^{SM}$ (right)
- All measured values consistent with the SM within uncertainties

Motivation | The κ framework for coupling modifiers



BSM physics affects production modes + decay channels in a correlated way if governed by similar interactions



A modified H - W or H - t coupling changes $H \rightarrow WW$ and $H \rightarrow \gamma\gamma$ rates and ggH , WH , VBF cross-sections

- SM quantities (σ_i , Γ_f , Γ_H) scaled by coupling modifiers κ_i
- In the SM: all $\kappa = 1$

B_{BSM} = 0

no BSM contributions

B_{inv}, B_{undet} float., |κ_V| ≤ 1

allows BSM; constrains W/Z couplings

B_{inv}, B_{undet} float., offshell inc.

offshell analysis regions included, W/Z constraint removed

B_{inv} < 13% (95% CL)


B_{undet} < 25% (95% CL)

→ Motivates dedicated searches like $H \rightarrow AA \rightarrow 4\gamma$

Context | Exotic $H \rightarrow AA$ decay searches


- Hypothetical new spin-0 particle A \rightarrow pair of SM particles
- Simplest Higgs sector extension; e.g. **2HDM+S**

For A masses above the heavy-fermion thresholds




- If $m_a > 2m_\tau, 2m_b$
- $A \rightarrow \tau\tau$ or $A \rightarrow bb$ are dominant
- Predicted in certain 2HDM+S extensions

Below heavy-fermion thresholds:



- Leptonic mode $A \rightarrow \mu^+\mu^-$ accessible
- Provides a clean experimental signature




- $A \rightarrow \gamma\gamma$ mode
- Loop induced, generally subdominant
- Significant in scenarios where A couples to uncoloured new charged particles

Context | Exotic $H \rightarrow AA \rightarrow 4\gamma$ topologies

Regime description

Boosted regime



- $m_A \lesssim 1 \text{ GeV}$
- A with large Lorentz boost
- Photons highly collimated
- Separation equal or less to ECAL crystal size

Resolved regime



- $m_A \gtrsim 10 \text{ GeV}$
- Photons well isolated and fully resolved
- Diphoton invariant mass reconstructed

Semi-merged regime

- One $A \rightarrow \gamma\gamma$ produces 2 well separated photon objects
- The other $A \rightarrow \gamma\gamma$ produces a pair reconstructed as a single photon object
- Due to **highly collimated** pair OR due to **γ reconstruction threshold**

Previous CMS searches by topology

[CMS - Exotic decay two merged diphotons](#)

2023

- $m_A \in [0.1, 1.2] \text{ GeV}$
- novel deep-learning mass reconstruction from ECAL energy deposits

[CMS - Exotic decay with four photons](#)

2023

- $m_A \in [15, 62] \text{ GeV}$
- Diphoton invariant mass from kinematics of both photon objects

[CMS - Exotic decay semi-merged](#)

2026

- $m_A \in [1, 15] \text{ GeV}$
- Fills the gap between the two existing CMS analyses
- Deep-learning mass regression recovers m_A from ECAL energy patterns

Topology of $H \rightarrow AA \rightarrow 4\gamma$ depends on m_A

Reconstruction strategy based on selected regime

Detector | CMS detector

Reconstructed objects:

- Electrons, muons, τ leptons, **photons**, jets, p_{T}^{miss}


$$[p_{T}^{\text{miss}} = -\sum p_{T} \text{ of all reconstructed particles}]$$

(transverse to the beam)

Subdetector structure

- All-silicon inner tracker (pixel + microstrips)
- ECAL**: ~76,000 PbWO_4 scintillating crystals
- HCAL: brass + scintillator
- 3.8 T superconducting solenoid
- Gas-ionization muon chambers

Multipurpose detector for pp collisions at the LHC



Selection | Datasets , Object Reconstruction & Event Selection

Datasets

- **Data:** pp collisions, $\sqrt{s} = 13$ TeV, Run2 (2015-2018)
- **Signal MC:** $H \rightarrow AA \rightarrow 4\gamma$ generated with MadGraph + PYTHIA + GEANT4
- Mass points: $m_A = 1, 3, 5, 10, 15$ GeV (covers the targeted mass range)
- **Backgrounds:** QCD multijet, γ +jet, prompt diphoton \rightarrow **modeled entirely from data (no MC)**

Photon reconstruction

- **Energy deposits clustered in ECAL crystals;** grouped into superclusters (accounts for shower spread + magnetic-field sweeping)
- **Threshold: $E_T > 10$ GeV** (Secondary photons below 10 GeV stay unreconstructed)

Mass reconstruction

- m_A reconstructed directly from ECAL crystal energy patterns
- **GNN improves on the CNN of the boosted analysis**
- Details on the **appendix**

Photon Selection


- Leading photon $p_T > 30$ GeV, subleading $p_T > 18$ GeV
- **Loose isolation + electron veto** to suppress jet/electron fakes
- **Discriminator** to separate prompt photons from EM-enriched jets (EM-enriched jet = hadronic jet with large EM fraction, typically from neutral-meson decays)

Mass reconstruction | GNN for signal mass regression

- Overlapping photon showers in ECAL
- Standard reconstruction, one photon-like object, no m_A handle
- Need end-to-end reconstruction

Direct m_A reconstruction from ECAL crystal patterns

- Why a GNN (vs. CNN)
- ECAL deposits are sparse:
 - GNN uses only crystals with energy
 - Permutation & rotational invariance
 - Learned naturally by GNN
 - Substantial reduction in problem size vs. CNN's 2D grid



Inputs per crystal (node features)
 $E_T = E \sin \theta$, $E_z = E \cos \theta$
Position $(x, y, z) \rightarrow$ Network learns η -dependent shower shapes
Input window
128 \times 128 ECAL crystals around the most energetic crystal ($\Delta R = 1.1$)
Captures >90% of secondary photons for $m_A < 5$ GeV; 60-80% up to 15 GeV

Architecture
Input MLP \rightarrow 3 \cdot (EdgeConv + clustering/pooling) with residual connections
 \rightarrow global max pool \rightarrow output MLP $\rightarrow m_A$




Regions | Signal Regions and Sideband Regions

Three preselected photons \rightarrow grouped into A_1 (merged) + A_2 (two resolved)

- m_{A_2} : from photon four-momenta (standard reconstruction)
- m_{A_1} : from GNN mass regression (direct measurement not possible) $m_{A_2}[\text{GeV}]$

The 2D m_A plane

- Each event fills a 2D (m_{A_1} , m_{A_2}) template
- Under the identical-mass hypothesis $m_{A_1} = m_{A_2}$, signal sits on the diagonal



Triphoton mass ($m_{\Gamma\Gamma\Gamma}$)

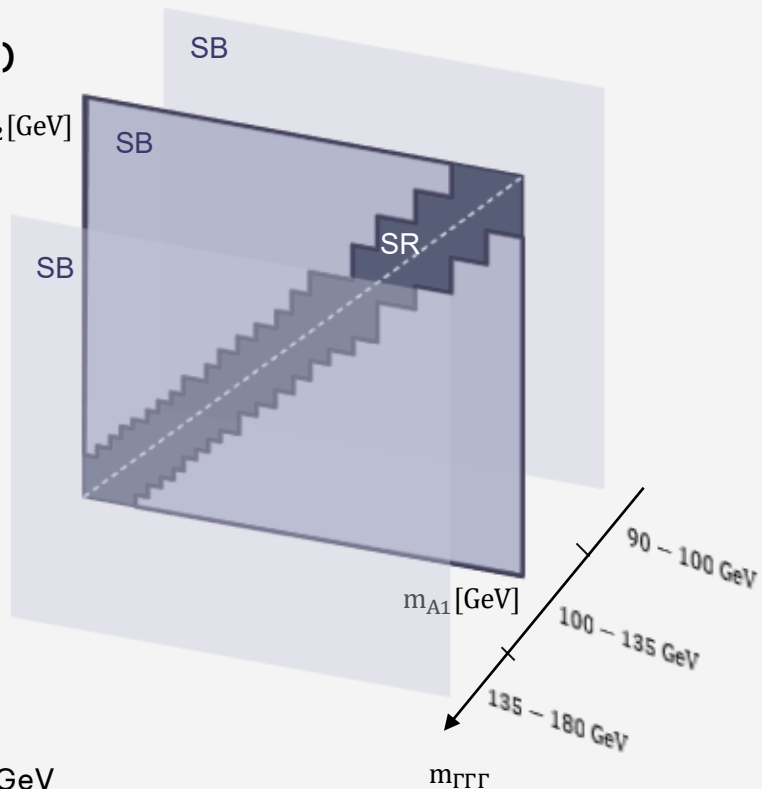
Assuming the three photon-like objects come from H decay (no invisibles): $m_{\Gamma\Gamma\Gamma} \rightarrow$ Higgs mass window

Along the 2D m_A plane:

- m_A -SR : $|\Delta m_A| < 2 \text{ GeV}$ (on-diagonal)
- m_A -SB : $|\Delta m_A| > 2 \text{ GeV}$ (off-diagonal)

Along $m_{\Gamma\Gamma\Gamma}$:

- m_H -SB_low : $90 < m_{\Gamma\Gamma\Gamma} < 100 \text{ GeV}$
- m_H -SR : $100 < m_{\Gamma\Gamma\Gamma} < 135 \text{ GeV}$
- m_H -SB_high : $135 < m_{\Gamma\Gamma\Gamma} < 180 \text{ GeV}$



Models | Signal and Background Model

Signal Model

A model needs:

- **Shape** of the 2D m_A distribution
- **Total yield** (normalization) for a given integrated luminosity

Shape

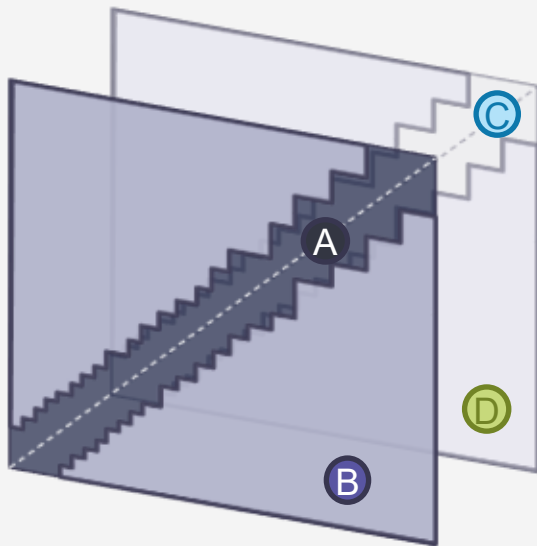
- Derived from simulated $H \rightarrow AA \rightarrow 4\gamma$ samples
- Mass points: $m_A = 1, 3, 5, 10, 15$ GeV

Yield

- $N = \sigma \times B \times \epsilon \times \int L dt$
 - For ease of calculation normalize to $\sigma \times B = 1$ pb
- allows the result of the analysis to be directly a limit on $\sigma \times B$

Final signal region is $m_{TT} \in \text{Higgs window AND } |\Delta m_A| < 2 \text{ GeV}$ **A**

Every other region **B** **C** **D** feeds the data-driven background model



Background Model

Yield

- Assumption: on/off-diagonal ratio is the same in m_{H-SR} and m_{H-SB}
- $A / B = C / D$
- $A = B \times (C / D)$

Shape

- Assumption: 2D m_A shape in A is a weighted average of the shapes in C_{low} and C_{high}
- $\text{shape}(A) = f_{low} \cdot \text{shape}(C_{low}) + f_{high} \cdot \text{shape}(C_{high})$

$$p(m_A | m_{H-SR}) = f_{SB-low} p(m_A | m_{H-SB_{low}}) + f_{SB-high} p(m_A | m_{H-SB_{high}})$$

Both background yield and shape inside A are predicted entirely from data - no MC needed

Models | Validation of Background Model

Validates if the **interpolation** used to get the **shape** on the background model describes the data

Test region: B – off-diagonal, Higgs window

$$\text{shape}(B) = f_{\text{low}} \cdot \text{shape}(D_{\text{low}}) + f_{\text{high}} \cdot \text{shape}(D_{\text{high}})$$

Same assumption used to predict A, applied to B where it can be checked safely

Uncertainties

- **Dominant overall:** Statistical uncertainty on the background model

Signal model systematics:




- **Luminosity** - 1.6%
- **Trigger efficiency** - 1–20% (largest at low m_A)
- **Photon ID** - 1–3% (data/MC scale factors)
- **Photon energy scale & resolution** - < 3% bin by bin
- **Mass regression scale & smearing** - shifts up to 6%, smearing up to 260 MeV

Background model systematics:

- **Shape mismodeling** - residual data/prediction differences in m_A -SB

Region B validation of Background Model

Observed distribution vs. expected



- **Top:** unrolled 2D spectrum (m_{A1} - m_{A2} off-diagonal flattened into 1D bins)
- **Bottom:** 1D projections onto m_{A1} and m_{A2}
- **Data agrees with prediction** across all bins; ratio panels consistent with 1 within uncertainty
- **2D shape prediction validated**

Results | Predicted background distribution after fitting to the data

Maximum-likelihood fit to the 2D m_A distribution

- Model: $\mu S + B$ (signal-plus-background)
- μ = signal strength
- S, B = signal & background models
- Systematic & statistical uncertainties : nuisance params.

$$\mathcal{L}(\mu) = P(\text{data}|\mu S + B)$$



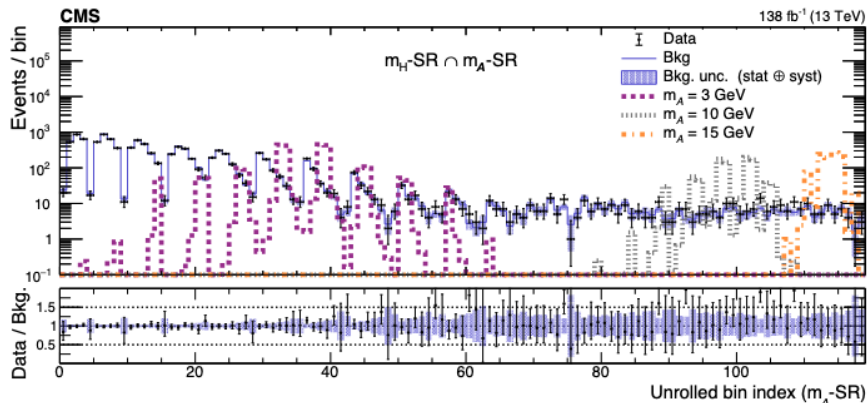
- The fit simultaneously adjusts μ and all nuisance parameters

Signal model normalized to $\sigma \times B = 1$ pb by convention,
 μ and $\sigma \times B$ (in pb) are numerically equal,
 fit directly delivers limits on $\sigma(H) \cdot B(H \rightarrow AA \rightarrow 4\gamma)$


- Blue histogram: best-fit background ($\mu S + B$ fitted to data)
- Blue band: total uncertainty (stat \oplus syst) on background prediction
- Dashed colored lines: signal templates at $m_A = 3, 10, 15$ GeV

Top: unrolled 2D spectrum in m_H -SR \cap m_A -SR (region A)

Bottom: 1D projections



Results | Observed upper limits at 95% CL on $\sigma(pp \rightarrow H) B(H \rightarrow AA \rightarrow 4\gamma)$



Observed limit (solid black)

From the fit to actual data: $\hat{\mu}$ compatible with observed dataset
→ translates to an upper bound on $\sigma \times B$ at each mass point

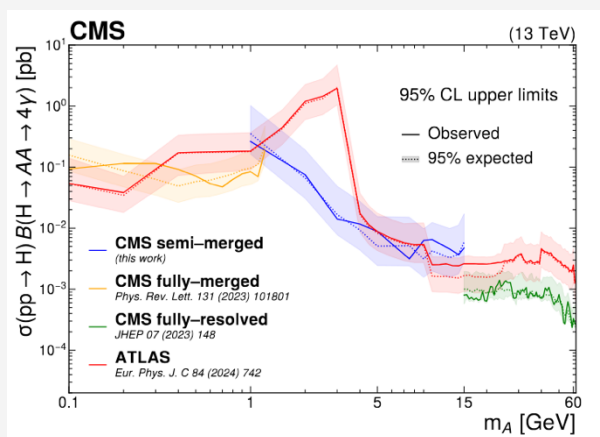
Expected limit (dashed black)

Same fit, but applied to a background-only dataset
Gives the limit we would have set if there were exactly zero signal (null hypothesis)

Green / Yellow bands

68% / 95% statistical fluctuations of the expected limit

Observed agrees with expected within the bands
no excess; limits set



Combined exclusion across m_A

- **CMS Semi-merged**: $1 < m_A < 15$ GeV
- CMS fully merged: 0.1 – 1.2 GeV
- CMS fully resolved: 15 – 62 GeV
- ATLAS inclusive: 0.1 – 62 GeV

This search enables the extraction of **stringent limits** across the entire mass spectrum, **bridging the gap** in sensitivity between the two previous CMS searches,

Improvement in this analysis compared with ATLAS search (e.g, at $m_A = 3$ GeV observed limit x100 more stringent than ATLAS result)

Appendix

Motivation | Production modes of Higgs boson



Production rate = luminosity * cross-section

Five main production modes at the LHC:

- **ggH** – gluon fusion via top quark loop · ~87% of total σ
- **VBF** – vector boson fusion · two forward high- p_T jets
- **VH** – associated with W/Z ('Higgsstrahlung') · extra leptons / $p_{T^{\text{miss}}}$
- **ttH, tH** – associated with top quarks · b-tagged jets
- **bbH** – associated with bottom quarks

Events categorized by production-mode signature → drives the analysis strategy

Production mode	Cross section (pb)
ggH	48.31 ± 2.44
VBF	3.771 ± 0.807
WH	1.359 ± 0.028
ZH	0.877 ± 0.036
ttH	0.503 ± 0.035
bbH	0.482 ± 0.097
tH	0.092 ± 0.008



Motivation | Decay channels of Higgs boson

Higgs couplings set by mass

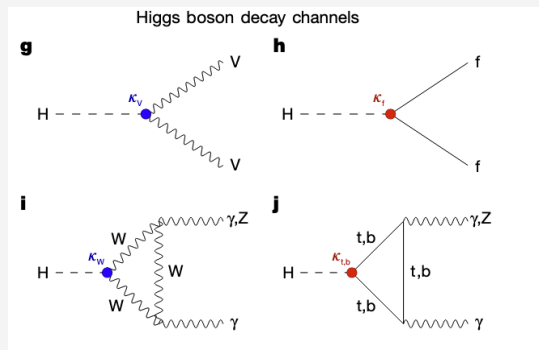
- To vector bosons: amplitude $\propto m_V^2$
- To fermions: amplitude $\propto m_f$
- Third-generation fermions dominate
- second-generation rare

Decays to SM particles test the theory


- Measured branching fractions vs. SM prediction
- Sizeable deviations \rightarrow hint of BSM physics

Main decay channels probed at CMS:

- Bosonic: $H \rightarrow \gamma\gamma \cdot ZZ \rightarrow 4\ell \cdot WW \rightarrow \ell\nu\ell\nu \cdot Z\gamma$
- Fermionic (3rd gen): $H \rightarrow bb \cdot \tau\tau$
- Fermionic (2nd gen, rare): $H \rightarrow \mu\mu$
- In $t\bar{t}H$ / tH : multilepton final states



Decay channel	Branching fraction (%)
bb	57.63 \pm 0.70
WW	22.00 \pm 0.33
gg	8.15 \pm 0.42
$\tau\tau$	6.21 \pm 0.09
cc	2.86 \pm 0.09
ZZ	2.71 \pm 0.04
$\gamma\gamma$	0.227 \pm 0.005
Z γ	0.157 \pm 0.009
ss	0.025 \pm 0.001
$\mu\mu$	0.0216 \pm 0.0004



Mass reconstruction | GNN for signal mass regression

- Overlapping photon showers in ECAL
- Standard reconstruction . one photon-like object, no mA handle
- Need end-to-end reconstruction




Direct mA reconstruction from
ECAL crystal patterns



Why a GNN (vs. CNN)

- ECAL deposits are sparse:
- GNN uses only crystals with energy
- **Permutation & rotational invariance**
- Learned naturally by GNN
- Substantial reduction in problem size vs. CNN's 2D grid



Inputs per crystal (node features)

$$E_T = E \sin \theta, E_z = E \cos \theta$$

Position $(x, y, z) \rightarrow$ Network learns η -dependent shower shapes


Input window

128 × 128 ECAL crystals around the most energetic crystal ($\Delta R \approx 1.11$)

Captures >90% of secondary photons for mA < 5 GeV; 60–80% up to 15 GeV

Architecture

Input MLP \rightarrow 3 × (EdgeConv + clustering/pooling) with residual connections
 \rightarrow global max pool \rightarrow output MLP \rightarrow mA



Mass reconstruction | GNN for signal mass regression

Training set

600k $A \rightarrow \gamma\gamma$ events + 70k single-photon events
Flat in $m_A \in (0, 18)$ GeV and $p_T \in (10, 200)$ GeV
ECAL barrel only, $|\eta| < 1.44$
Loss: mean absolute error

Two key training tricks

Phase-space flattening - counters bias toward boosted A from selection cuts

Domain continuation - single photons assigned random negative masses in $[-3, 0]$ GeV

→ background single photons regressed to $m < 0$ and removed

→ linearizes response near $m_A = 0$

Validation

True vs. predicted: tight diagonal across full mass range

Clean peaks for $m_A = 1, 3, 5, 10, 15$ GeV

Robust across 2016/17/18 data-taking conditions

Data/MC calibration via $Z \rightarrow e^+e^-$ (tag-and-probe) → mass-scale + smearing systematics

