

Search for non-resonant di-Higgs production in $bb + \tau\tau$ final states in pp collisions at CMS

Simon Daigler

Master Thesis

at the Department of Physics
Institute of Experimental Particle Physics (ETP)

Advisor: Prof. Dr. Markus Klute

Co-Advisor: Dr. Roger Wolf

start date 11/15/2024 – end date 11/15/2025

Karlsruher Institut für Technologie
Fakultät für Physik
D-76128 Karlsruhe

I declare that I have developed and written the enclosed thesis completely by myself and have not used any sources or resources without declaring them in the text. For spelling and phrasing, I used AI tools. Their assistance was limited to language editing and did not contribute to the conception, analysis, or interpretation of the work.

Karlsruhe, 11/15/2025

.....
(Simon Daigler)

Abstract

A search for nonresonant Higgs-boson pair production (HH) in the $b\bar{b}\tau^+\tau^-$ final state is presented at $\sqrt{s} = 13$ TeV using the 2018 CMS proton–proton dataset corresponding to 59.8 fb^{-1} after a common re-optimization of the reconstruction and calibration of the data by the CMS Collaboration. The resolved topology is targeted and three $\tau^+\tau^-$ final states $\tau_h\tau_h$, $\mu\tau_h$ and $e\tau_h$ are taken into account, where τ_h refers to the decay of the τ lepton into hadrons. Object reconstruction and identification are based on Particle Flow, DEEJET and DEETAU v2.5 algorithms. A multi-class classification neural network is used for event discrimination and category definition. The background from QCD multijet events is estimated from data via an ABCD method using the charges and DEETAU vsJet discriminants to define corresponding control regions. The di- τ invariant mass, a key observable for discriminating signal from background, is reconstructed using the FASTMTT algorithm.

Expected upper limits are derived from a profile-likelihood fit to the Asimov dataset. The median expected 95 % CL upper limit on the signal strength $\mu = \sigma(\text{pp} \rightarrow \text{HH})/\sigma_{\text{SM}}$ is found to be $\mu_{95}^{\text{exp}} = 10.9$, about 2% weaker than the result obtained when considering only statistical uncertainties ($\mu_{95}^{\text{exp}} = 10.7$). Projections obtained from scaling rates with integrated luminosity while keeping shapes and the size of nominal systematics yield median expected 95 % CL limits of $\mu_{95}^{\text{exp}} = 7.2$ for the full Run 2 corresponding to an integrated luminosity of 138 fb^{-1} and $\mu_{95}^{\text{exp}} = 4.7$ for a Run 2 + Run 3 scenario with 500 fb^{-1} .

Acknowledgements

First and foremost, I am grateful to Dr. Nikita Shadskiy. His ongoing support and work on the search for resonant di-Higgs production [1] provided the foundation for this thesis, and several of its concepts and ideas were adopted and further developed here. I also thank Dr. Artur Gottmann for his support with computing tasks, including the preparation of MC datasets. I am grateful to my supervisor and our group leader, Dr. Roger Wolf, for his steady guidance and for always making time to discuss questions, even during the busiest periods. His constructive feedback, patience, and encouragement were essential throughout this work. Finally, my thanks go to the entire τ working group, including Artur Monsch, Ralf Schmieder, Christian Winter, Tim Voiglaende, Moritz Molch and Jan Voss, for many helpful discussions and their continuous support.

Contents

Abstract	i
Acknowledgements	iii
1. Introduction	1
2. The Standard Model	3
2.1. Higgs Mechanism	4
2.2. Gluon–Gluon Fusion Di Higgs Boson Production	6
2.3. Higgs Boson Decays	7
3. The Experiment	9
3.1. The Large Hadron Collider (LHC)	9
3.2. The CMS detector	10
4. Event Reconstruction and Identification	13
4.1. Particle Flow	13
4.2. Jet Reconstruction	14
4.3. Electron Identification	14
4.4. Muon Identification	15
4.5. b Identification with DEEPJET	15
4.6. τ Identification with the DEEPTAU v2.5	16
4.6.1. Input Features	16
4.6.2. Architecture	17
4.6.3. Training Strategy	19
4.6.4. Performance and Comparison of DEEPTAU v2.5 vs. v2.1	19
4.6.5. Scale Factors and Validation in Data	21
5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$	23
5.1. Overview of Current Results	23
5.2. Signal Process	25
5.3. Background Processes	25
5.3.1. Top–Antitop Production	25
5.3.2. Drell-Yan $Z/\gamma^* + \text{jets}$	27
5.3.3. Single Top	27
5.3.4. W+jets	28
5.3.5. Single Higgs	29
5.3.6. Other Backgrounds	29
5.3.7. QCD Multijet Production	29

5.4.	Event Simulation	31
5.4.1.	Matrix Element (ME) event generators	31
5.4.2.	PYTHIA 8 parton shower and the CP5 tune	32
5.4.3.	Pileup, reconstruction and object calibration	32
5.5.	Event Selection	33
5.5.1.	$\tau^+\tau^-$ -Pair Selection	33
5.5.2.	$b\bar{b}$ Pair Selection	35
5.6.	$\tau^+\tau^-$ Pair Mass Reconstruction with FastMTT	35
5.7.	Control shapes	36
6.	Analysis	39
6.1.	Event Classification with Neural Networks	39
6.1.1.	Architecture	39
6.1.2.	Training	39
6.1.3.	Training Data & Classification Results	43
6.2.	Systematic Uncertainties	45
6.3.	Results	49
6.4.	Extrapolation to higher luminosities	50
7.	Conclusion	51
A.	Appendix	53
A.1.	Control Shapes	53
A.2.	Training Datasets	68
A.3.	Training Metrics	69
A.4.	Classification Confusion Matrices	71
A.5.	Analysis Shapes	75
A.6.	Upper Limit Confidence Signal Strength	80
	Bibliography	91

1. Introduction

Particle physics asks simple questions with deep impact: What are the basic building blocks of matter and how do they interact? The Standard Model (SM) of particle physics answers much of this by unifying the electromagnetic, weak, and strong interactions in a quantum field-theoretic framework. The 2012 discovery of the Higgs boson [2, 3] completed the SM particle content and established the mechanism of electroweak symmetry breaking experimentally, fixing the Higgs mass of $m_H \approx 125$ GeV. But key parts of the Higgs sector are still not measured directly. In particular, we do not yet have a precise experimental measurement of the Higgs boson self-coupling. Measuring it tests the shape of the Higgs potential and is therefore a central check of the SM.

This thesis uses data from the Compact Muon Solenoid (CMS) detector. CMS is a large, general-purpose particle detector installed at the Large Hadron Collider (LHC). The LHC is the proton-proton accelerator at CERN, the largest European laboratory for particle physics. CMS surrounds one of the LHC collision points with tracking, calorimetry and muon systems, all inside of a 3.8 T magnetic field [4].

Higgs-pair (HH) production gives direct access to the Higgs self-coupling. At the LHC, the main production mode is gluon-gluon fusion (ggF). Two leading-order diagrams contribute: a “box” diagram with a top-quark loop and a “triangle” diagram with the trilinear self-coupling. These amplitudes interfere destructively, which makes the total cross section very small (about 31 fb at $\sqrt{s} = 13$ [5]). Because the rate is tiny, any improvement in analysis sensitivity has high impact.

Among all accessible final states, $HH \rightarrow b\bar{b}\tau^+\tau^-$ provides a good balance between branching fraction and purity. One Higgs boson decays to a b-quark pair ($b\bar{b}$) and the other to a τ -lepton pair ($\tau^+\tau^-$), giving a combined branching fraction $\sim 7.3\%$. After data recording, events are reconstructed and particles are identified using a set of algorithms and neural networks (NNs).

On the following pages I study nonresonant HH production via ggF in the $b\bar{b}\tau^+\tau^-$ final state using Run 2 Ultra Legacy (UL) data from 2018 based on an ultimate re-reconstruction of the data performed by the CMS Collaboration. Three final states of the $\tau^+\tau^-$ pair are considered: full hadronic $\tau_h\tau_h$ and the two semi leptonic channels $\mu\tau_h$ and $e\tau_h$, resulting in a statement about the 95 % confidence level (CL) upper limit on the signal strengths $\sigma(\text{pp} \rightarrow \text{HH})/\sigma_{\text{theory}}$ at $\sqrt{s} = 13$ TeV. This thesis evaluates only the expected sensitivity of the analysis. All fits are performed on an Asimov dataset [6] to obtain the median expected limits in the absence of the process.

2. The Standard Model

The SM is a quantum field theory to describe the dynamics of elementary particles and three of their fundamental interactions: the strong, weak and electromagnetic interactions. Its underlying gauge symmetry is

$$SU(3)_C \times SU(2)_L \times U(1)_Y. \quad (2.1)$$

Matter is built from quarks and leptons, arranged in three generations. Quarks carry color charge and interact via all three forces. Leptons are color neutral and therefore do not participate in the strong interaction. Forces are transmitted by gauge bosons: eight gluons related to the $SU(3)_C$, the W^\pm and Z bosons related to the $SU(2)_L$, and the photon for the $U(1)_{em}$ symmetry, after electroweak symmetry breaking. The Higgs boson is the quantum excitation of the Higgs field Φ . Through the Higgs mechanism, particles that couple to Φ acquire mass. Figure 2.1 gives an overview of the elementary particles in the SM.




Figure 2.1.: Elementary particles in the SM. Quarks are shown in purple and leptons in green. Together they form the elementary matter particles. The gauge bosons (shown in red) carry the fundamental forces. The Higgs boson is the quantum excitation of the Higgs field Φ . Through the Higgs mechanism, particles that couple to Φ acquire their mass. [7]

The SM took shape in the 1960s–1970s. The quark model, developed by Murray Gell-Mann [8] and Georg Zweig [9], and the introduction of color led to Quantum Chromo-

dynamics (QCD) as an $SU(3)_C$ gauge theory. Electroweak unification resulting in the symmetry $SU(2)_L \times U(1)_Y$ was formulated by Glashow [10], Weinberg [11] and Salam [12]. The Brout–Englert–Higgs mechanism [13–15] gave masses to the weak gauge bosons and, via Yukawa couplings, to fermions. The Glashow–Iliopoulos–Maiani (GIM) mechanism [16] predicted the charm quark, discovered in 1974. Key milestones include the discovery of the W^\pm and Z 1983 [17–20], the top quark 1995 [21] and finally the Higgs boson 2012 [2, 3], consolidating the SM as a consistent theory of elementary particle interactions potentially up to the Planck scale.

Today the SM has been tested to high precision. At the same time, established phenomena such as neutrino oscillations, the baryon asymmetry of the Universe and the presumed existence of dark matter point beyond the SM and call for precise studies of the Higgs sector.

2.1. Higgs Mechanism

A central challenge in constructing a gauge theory of the weak and electromagnetic interactions is to generate particle masses without explicitly breaking the gauge invariance. Proca mass terms for the weak gauge bosons violate local symmetry and Dirac mass terms for chiral fermions are not gauge invariant under electroweak group transformations. The Higgs mechanism provides a consistent solution: a scalar field with a nonzero vacuum expectation value. This triggers the spontaneous breaking of the electroweak symmetry and thereby generating masses for the weak gauge bosons, and via Yukawa interactions for the fermions, while preserving the gauge invariance.

The electroweak sector of the SM is a gauge theory invariant under the $SU(2)_L \times U(1)_Y$ gauge symmetry with couplings g and g' . The symmetry is broken by a single complex scalar doublet field with hypercharge $Y = 1$,

$$\Phi(x) = \begin{pmatrix} \phi^+(x) \\ \phi^0(x) \end{pmatrix}. \quad (2.2)$$

The dynamic behavior of Φ is described by the Lagrangian density

$$\mathcal{L}_H = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi), \quad D_\mu = \partial_\mu - i g \frac{\tau^i}{2} W_\mu^i - i g' \frac{Y}{2} B_\mu, \quad (2.3)$$

with the scalar potential

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad \lambda > 0, \quad \mu^2 < 0. \quad (2.4)$$

Here, $\lambda > 0$ ensures that the potential is bounded from below and therefore stable. Choosing $\mu^2 < 0$ makes the symmetric point at $\Phi = 0$ unstable and induces a nonzero vacuum expectation value. The potential is minimized at $\langle \Phi^\dagger \Phi \rangle = v^2/2$ with $v = \sqrt{-\mu^2/\lambda} \simeq 246$ GeV and it is convenient to choose

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \quad (2.5)$$

which spontaneously breaks the $SU(2)_L \times U(1)_Y$ to the $U(1)_{EM}$ symmetry. The unbroken generator is the electric charge Q , ensuring a massless photon.

Before electroweak symmetry breaking, the complex doublet Φ carries four real degrees of freedom. When Φ acquires its vacuum expectation value, the symmetry of the vacuum is reduced and three of these degrees of freedom become the would-be Nambu–Goldstone modes [22] [23] [24]. Because the electroweak symmetry is local, these modes are not observed as particles. Instead they are absorbed by the weak gauge fields $W^{\mu,a}$ and provide the longitudinal polarizations of the massive vector bosons. The remaining scalar fluctuation can be identified with the physical Higgs boson H .

The gauge fields that appear naturally in the symmetry bases, W^1, W^2, W^3 for the $SU(2)_L$ and B for the $U(1)_Y$ symmetry, are not yet the particles to be observed in detectors. The charged bosons W^\pm arise from specific linear combinations of W^1 and W^2 . The neutral fields W^3 and B are mixed by the Weinberg angle to form the massive Z boson and the massless photon γ . Electromagnetism remains unbroken, so the photon has no mass, while the W and Z bosons acquire masses set by the gauge couplings and the Higgs vacuum expectation value.

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp i W_\mu^2), \quad \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & -\sin \theta_W \\ \sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix}, \quad (2.6)$$

one obtains the masses

$$m_W = \frac{1}{2}g v, \quad m_Z = \frac{1}{2}\sqrt{g^2 + g'^2} v, \quad m_\gamma = 0, \quad (2.7)$$

with $\tan \theta_W = g'/g$.

Weak interactions are parity-violating: the W^\pm couple only to left-handed fermions and right-handed antifermions, while the Z boson couples to both chiralities with different strengths. Consequently, the left-handed fermions form in $SU(2)_L$ doublets, whereas right-handed fermions are introduced as $SU(2)_L$ singlets. Writing a Dirac mass term would directly couple these fields and break gauge invariance. The SM avoids this by introducing gauge-invariant Yukawa interactions: the Higgs doublet field bridges left-handed doublets and right-handed singlets. When the Higgs field takes its vacuum value, these interactions turn into fermion mass terms. The resulting masses are proportional to the corresponding Yukawa couplings. The couplings of the physical Higgs boson to fermions are proportional to the fermion masses.

For Quarks the Yukawa interactions are matrices in flavor space. Diagonalizing the up-type and down-type mass matrices does not align the weak and mass eigenstates. This mismatch shows up in the charged weak interactions as the unitary Cabibbo–Kobayashi–Maskawa (CKM) matrix [25, 26], which tells how quark flavors mix and how likely transitions are. In the minimal SM, neutrinos have no right-handed partners and remain massless.

In the SM, the Higgs potential fixes the trilinear self-coupling to $\lambda_{HHH} = 3m_H^2/v$. This trilinear term enters the triangle amplitude in ggF di-Higgs production and interferes with the top-quark box amplitude, so measurements of $pp \rightarrow HH$ constrain $\kappa_\lambda \equiv \lambda_{HHH}/\lambda_{HHH}^{\text{SM}}$.

2.2. Gluon–Gluon Fusion Di Higgs Boson Production




Figure 2.2.: Nonresonant Higgs-boson pair production through ggF. **Left:** Triangle diagram, where two gluons couple to a top-quark loop that produces a single off-shell Higgs boson, which then splits via the trilinear self-coupling into two Higgs bosons. **Right:** Box diagram, where the top-quark loop directly emits two Higgs bosons. In the SM these amplitudes interfere destructively. [27]

The dominant production mode for both single-Higgs and Higgs-pair production at the LHC is ggF. It is a loop-induced process: gluons couple to the Higgs boson through heavy quarks, with the top quark giving the leading contribution. The strength of the interaction is controlled by the top Yukawa coupling and by the Higgs self-coupling.

At leading order, two quantum amplitudes contribute with similar size but different structure. The two diagrams are shown in Figure 2.2. A triangle diagram contains the trilinear Higgs self-coupling and a box diagram contains two top–Higgs Yukawa couplings. In the SM the two interfere destructively, which suppresses the total rate and changes the kinematic properties of the decay products such as the invariant mass of the Higgs boson pair.

For $m_H = 125$ GeV at $\sqrt{s} = 13$ TeV, the inclusive SM prediction for nonresonant $gg \rightarrow HH$ is

$$\sigma^{\text{SM}}(gg \rightarrow HH) = 31.05^{+6\%}_{-23\%} \pm 3\% (\text{PDF}+\alpha_s) \text{ fb} \quad (\text{NNLO})[28] [29], \quad (2.8)$$

For comparison, the inclusive single-Higgs boson production cross section via gluon fusion at the same energy is about

$$\sigma^{\text{SM}}(gg \rightarrow H) = 48.5^{+4.6\%}_{-6.7\%} \pm 3.2\% (\text{PDF}+\alpha_s) \text{ pb} \quad (\text{N}^3\text{LO})[30], \quad (2.9)$$

Numerically, the pair-production rate is tiny compared with single-Higgs boson production:

$$\frac{\sigma(gg \rightarrow HH)}{\sigma(gg \rightarrow H)} = \frac{31.05 \text{ fb}}{48.5 \text{ pb}} \approx 6.4 \times 10^{-4},$$

which means roughly one HH event per 1.6×10^3 single-Higgs boson events. This explains why HH events have not yet been detected.

2.3. Higgs Boson Decays

Table 2.1.: Branching fraction of the SM Higgs boson at $m_H = 125.09$ GeV [31, 32]

Decay mode	Branching fraction [%]
$H \rightarrow b\bar{b}$	58.2
$H \rightarrow WW$	21.5
$H \rightarrow gg$	8.2
$H \rightarrow \tau^+\tau^-$	6.3
$H \rightarrow c\bar{c}$	2.9
$H \rightarrow ZZ$	2.6
$H \rightarrow \gamma\gamma$	0.23
$H \rightarrow Z\gamma$	0.15
$H \rightarrow \mu^+\mu^-$	0.022

In the SM the Higgs boson couples proportionally to the mass of fermions, which leads to a large branching fraction to bottom quarks. The branching fractions are summarized in Table 2.1. Consequently, the most abundant HH final state is $b\bar{b}b\bar{b}$ with about 34%. This channel is not used here because, it suffers from very large QCD multijet backgrounds and severe combinatorial uncertainties when pairing the resulting four b jets. Similar issues affect $HH \rightarrow b\bar{b}WW$. On the other hand, modes such as $HH \rightarrow b\bar{b}\gamma\gamma$ or $HH \rightarrow 4\gamma$ provide very distinctive signatures while suffering from its vanishingly small branching fractions.

A practical compromise is the $HH \rightarrow b\bar{b}\tau^+\tau^-$ final state: with a branching fraction of approximately 7.3% it combines the large $H \rightarrow b\bar{b}$ branching fraction with a distinctive $\tau^+\tau^-$ signature for background rejection.

Bottom-quark decay ($H \rightarrow b\bar{b}$). At $m_H = 125$ GeV, $H \rightarrow b\bar{b}$ is the dominant decay. Bottom quarks cannot be measured directly. Instead they hadronize into jets. Jets that likely come from bottom quarks are experimentally identified (“b-tagged”) e.g. through the small displacement of their decay tracks. A Higgs boson candidate is built from two b-tagged jets. The strength of the b tags and the mass of the jet pair, together with characteristic properties from the remaining decay of the $\tau^+\tau^-$ system, are used to separate the signal from any remaining backgrounds. The main backgrounds are top–antitop quark ($t\bar{t}$) and Z+jets production, including cases where light-quark or gluon induced jet is mistakenly tagged as a b quark induced jet.

Also the τ pair cannot be measured directly due to the short lifetime of the τ . The branching fractions of τ pair decay modes are shown in Table 2.2. For the analysis the

Table 2.2.: $\tau^+\tau^-$ final states sorted by branching fraction [32].

Final state	Branching fraction [%]
$\tau_h\tau_h$	42.0
$e\tau_h$	23.1
$\mu\tau_h$	22.5
$e\mu$	6.2
ee	3.2
$\mu\mu$	3.0

following decay modes are used:

$$\text{HH} \rightarrow \text{b}\bar{\text{b}} \tau^+\tau^- \quad \text{with} \quad \tau^+\tau^- \in \{\tau_h\tau_h, \mu\tau_h, e\tau_h\},$$

which cover 87.6 % of all $\tau^+\tau^-$ decays while providing good background control. A tau decay always includes at least one neutrino. Therefore the visible four-momenta underestimate the true τ energy, and the event's missing transverse momentum (p_T^{miss}) carries information about the neutrinos.

Fully hadronic ($\tau_h\tau_h$). Both taus decay to narrow hadronic jets plus neutrinos. This decay is giving the highest $\tau^+\tau^-$ signal rate among the decay modes. The dominant backgrounds are $t\bar{t}$ and Drell-Yan $Z/\gamma^* + \text{jets}$. In the this analysis the τ_h candidates are classified in decay modes (DM). The most important single hadronic τ decay are:

- **1-prong (net charge ± 1):**

- $\tau^\pm \rightarrow \pi^\pm \nu_\tau$ (1-prong $0\pi^0$, DM=0)
- $\tau^\pm \rightarrow \pi^\pm \pi^0 \nu_\tau$ [$\rho^\pm \rightarrow \pi^\pm \pi^0$] (1-prong $1\pi^0$, DM=1)
- $\tau^\pm \rightarrow \pi^\pm \pi^0 \pi^0 \nu_\tau$ (1-prong $2\pi^0$, DM=2)

- **3-prong (net charge ± 1):**

- $\tau^\pm \rightarrow \pi^\pm \pi^\mp \pi^\pm \nu_\tau$ [$a_1^\pm \rightarrow 3\pi$] (3-prong, DM=10)
- $\tau^\pm \rightarrow \pi^\pm \pi^\mp \pi^\pm \pi^0 \nu_\tau$ (3-prong $1\pi^0$, DM=11)

Semi Leptonic ($\ell\tau_h$ with $\ell \in \{\mu, e\}$). One τ decays leptonically to a prompt, isolated muon or electron ($\tau^- \rightarrow \ell^- \bar{\nu}_\ell \nu_\tau$) with a clean signature in the detector. The other τ decays hadronically. The leading background is $t\bar{t}$ with additional contributions from $Z \rightarrow \tau^+\tau^-$ and $\text{jet} \rightarrow \tau_h$ fakes. Tight lepton isolation and τ_h identification criteria help reducing these backgrounds.

3. The Experiment

3.1. The Large Hadron Collider (LHC)

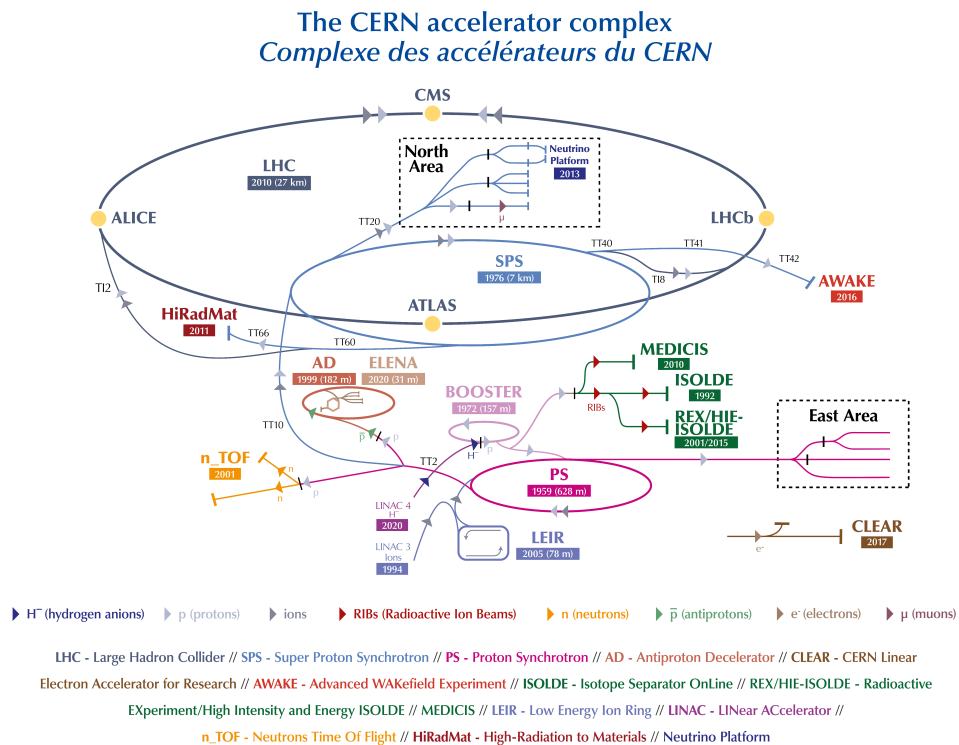


Figure 3.1.: Overview of the accelerators and detectors at CERN. Protons are produced by ionizing hydrogen gas. Before injection into the LHC, the beam is accelerated in smaller accelerators: LINAC4, the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS) [33].

The LHC is a 27 km circular proton-proton collider at CERN with two counter-rotating beams that intersect at four interaction points hosting the experiments ATLAS, CMS, LHCb, and ALICE. During Run 2 it was operated at a center-of-mass energy of $\sqrt{s} = 13$ TeV with 25 ns bunch spacing. Protons are extracted from a hydrogen source, accelerated in the linear accelerator (LINAC), and then boosted stepwise through the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS). From the SPS they are injected into the LHC at a momentum of 450 GeV and subsequently ramped to the collision energy by superconducting dipole magnets and RF cavities. Superconducting

3. The Experiment

quadrupole magnets guide and focus the beams. High instantaneous luminosity leads to multiple pp interactions per bunch crossing, which affects reconstruction.

3.2. The CMS detector




Figure 3.2.: Schematic structure of the CMS detector [4]

The Compact CMS is a general-purpose detector built to study high-momentum particles at the LHC. The subdetectors are arranged in a barrel around one of the LHC interaction points, providing nearly complete 4π solid-angle coverage. A 3.8 T superconducting solenoid magnetic field surrounds a silicon pixel and strip tracker that measure charged particles close to the beam line with high precision. Next to the tracker are two calorimeters: a crystal electromagnetic calorimeter (ECAL) for electrons and photons and a sampling hadron calorimeter (HCAL) to measure the energy of hadrons. Additional forward calorimeters are extending the reach at small angles to the beam. Beyond the solenoid coil, muons are detected in dedicated gas chambers embedded in the steel return yoke, providing robust muon identification and momentum measurements.

CMS uses a right-handed coordinate system with origin at the interaction point, the z axis along the beam, the x axis pointing to the LHC center, and the y axis upward. The azimuthal angle ϕ is measured in the x - y plane from the positive x axis, while the polar

angle θ is measured from the positive z axis. The pseudorapidity is defined as

$$\eta \equiv -\ln \left(\tan \frac{\theta}{2} \right). \quad (3.1)$$

The absolute pseudorapidity $|\eta|$ of a subdetector means it can cover particles with $-\eta_{\max} < \eta < \eta_{\max}$. Small θ (forward direction) corresponds to large $|\eta|$, while $\theta \rightarrow \pi/2$ (central) gives $\eta \rightarrow 0$. For highly relativistic particles, η approximates the true rapidity. The angular distance ΔR measures the separation between two objects in the detector and is defined as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \quad (3.2)$$

Tracker [34] The silicon tracker is used to reconstruct charged-particle trajectories and interaction vertices with high precision. It comprises an inner pixel detector and an outer strip detector, both covering $|\eta| < 2.5$. Pixel sensors (about $100 \mu\text{m} \times 150 \mu\text{m}$) give three-dimensional hits with typical resolutions of around $10 \mu\text{m}$ transversely and $20\text{--}40 \mu\text{m}$ along z . This precision enables robust identification of the primary interaction vertex and secondary vertices. Around the pixel detector there is the strip system which uses $100\text{--}200 \mu\text{m}$ pitch sensors, including stereo layers to recover 3D information. It reaches a hit resolutions of about $15\text{--}45 \mu\text{m}$.

ECAL [35] The ECAL is a homogeneous calorimeter made of $\sim 75,000$ PbWO_4 crystals that measure the energy of electrons and photons. Its barrel covers $|\eta| < 1.479$ and the endcaps extend to $|\eta| < 3$, with a preshower in front. There is a small transition gap at $1.479 < |\eta| < 1.653$. Silicon avalanche photodiodes (APDs) are used in the barrel and (vacuum phototriodes) VPTs in the endcaps to readout the relatively low light yield of the PbWO_4 crystals.

HCAL [36] The HCAL is a sampling calorimeter with brass absorbers and scintillator tiles that measures particles which were not fully absorbed by the ECAL, mainly hadrons. Its information is primarily used for jet reconstruction. It includes a barrel ($|\eta| < 1.5$), endcap ($1.5 < |\eta| < 3.0$) and forward calorimeters ($3 < |\eta| < 5$) with an additional outer HCAL outside the solenoid to catch late-developing showers. As a sampling device its resolution is poorer than the ECAL.

Solenoid The superconducting solenoid coil surrounds the tracker and calorimeters and provides an almost uniform 3.8 T magnetic field inside the coil. Charged particles follow curved paths with radius $r = p_T/(qB)$ in the plane perpendicular to the beam axis, enabling precise momentum and charge measurements. The coil is made of NbTi and operated near 4.6 K to achieve superconductivity. The magnetic flux is returned by a 12 kt steel yoke outside the coil.

Muon System [37, 38] The muon system lies outside the solenoid coil and measures muons up to $|\eta| < 2.4$ using gaseous detectors embedded in the steel return yoke. Drift tubes (DTs) in the barrel provide precise position measurements. Cathode strip chambers (CSCs) in the

3. The Experiment

endcaps add excellent spatial and timing resolution. The resistive plate chambers (RPCs) in both, the endcaps and barrel, offer fast timing for triggering. The combination provides redundant coverage and robust standalone and global muon reconstruction. In combination the muon detectors provide good coverage and allows standalone measurements of muons, and they yield higher precision when combined with inner-tracker tracks (global muons).

Trigger [39] Events are selected by a two-tiered trigger system. A hardware Level-1 (L1) trigger reduces the collision rate from 40 MHz to $\mathcal{O}(100 \text{ kHz})$ using calorimeter and muon information. A software High-Level Trigger (HLT) then performs a fast reconstruction resulting in an rate of interesting events of $\mathcal{O}(1 \text{ kHz})$ for permanent storage.

4. Event Reconstruction and Identification

Collider detectors do not observe quarks, gluons or Higgs bosons directly. They only record tracks and energy deposits from long-lived final-state particles in a high-pileup environment. Event reconstruction converts these raw signals into a consistent, calibrated set of physics objects with four-momenta. Object identification then assigns particle hypotheses based on working points with quantified efficiencies and misidentification rates. In other words: Object identification assesses how likely it is that a given candidate is, for example, an electron.

The core ingredients for CMS event reconstruction are the Particle Flow (PF) algorithm [40] and jet clustering [41] based on PF candidates. Lepton and heavy-flavor identification are eventually applied to these objects.

4.1. Particle Flow

The PF algorithm combines measurements from the silicon tracker, ECAL, HCAL and muon system to reconstruct a list of charged and neutral particles.

Electrons and Photons Electrons and photons are reconstructed together. As they pass through the detector material, electrons radiate photons via bremsstrahlung, and photons can convert into e^+e^- pairs. Electron tracks in the silicon tracker are therefore fit with a Gaussian-sum filter (GSF) to handle this energy loss and possible track kinks. In the ECAL, energy deposits are grouped into superclusters that collect the electromagnetic showers caused by such radiation. This is needed because the electromagnetic showers often spread over several nearby clusters, so a supercluster combines them to capture the full energy and a stable direction estimate. The tracking and ECAL information are then linked to form PF electron and photon candidates. A multivariate regression calibrates the supercluster energy and improves the electron momentum estimate. After building electrons and photons, their tracks and clusters are removed from the pool of not reconstructed data, so they are not reused when reconstructing hadrons.

Muons Muon reconstruction in CMS is performed with the dedicated muon detectors (DT, CSC, RPC) and is therefore built outside the PF algorithm. Three categories are used: (i) standalone muons from muon-chamber hits only with lower momentum resolution, (ii) global muons that match a tracker track together with muon hits with best momentum resolution, especially at high p_T and (iii) tracker muons built from tracker tracks that are extrapolated outward to the muon system. If the extrapolated track matches at least one muon segment or hit, the track is classified as a tracker muon. In practice, almost all muons are global or tracker. PF then applies additional quality and isolation selections.

Hadrons After removing electrons, photons and identified muons, hadrons are reconstructed from the remaining tracks and calorimeter clusters. Neutral hadrons are built from HCAL energy deposits without matching tracks. Charged hadrons are formed by linking tracker tracks to calorimeter deposits. For each PF charged-hadron candidate, the track momentum and the calibrated ECAL/HCAL energy must match. If they do not, the candidate is adjusted or split. In rare cases, a second pass of the muon reconstruction recovers missed global muons, which are then added to the PF candidate list.

4.2. Jet Reconstruction

Quarks and gluons from the hard scatter do not appear as isolated particles in the detector. As they lose energy, they hadronize into many hadrons that travel in roughly the same direction. In CMS, jets are clustered from PF candidates with sequential-recombination algorithms. The generalized distance measures are

$$d_{ij} = \min(p_{T,i}^{2n}, p_{T,j}^{2n}) \frac{\Delta R_{ij}^2}{R^2}, \quad \Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2, \quad (4.1)$$

$$d_{iB} = p_{T,i}^{2n}, \quad (4.2)$$

and pairs are successively recombined until all $d_{ij} > d_{iB}$, at which point the remaining pseudo-jets are declared as reconstructed jets. The parameter R sets the characteristic jet size. In CMS, $R = 0.4$ (AK4) is standard for resolved topologies.

Specific clustering algorithms arise from the choice of n . The default anti- k_T algorithm [41, 42] uses $n = -1$. The inverse p_T weighting pulls soft radiation in the direction of hard seeds and yields jets that are close to conical in shape. This robustness makes anti- k_T the preferred choice for most analyses.

Although calorimeter clusters are already calibrated, reconstructed jets receive dedicated energy-scale and resolution corrections that vary with pseudorapidity η and transverse momentum p_T . These corrections are derived from momentum-balance techniques using precisely measured reference objects like $Z \rightarrow \mu\mu/ee$ decays and isolated photons, so that the measured p_T of the Z or γ is calibrating the recoiling jet. To extend the calibration, especially for gluon-initiated jets, QCD dijet events are employed, transferring the scale from one jet already anchored by the Z/γ +jet method.

4.3. Electron Identification

Electron candidates from PF can be true prompt electrons, but also come from photon conversions or from hadrons that were misreconstructed as electrons. To suppress such backgrounds, quality selections are applied, most importantly an isolation requirement. The isolation is the scalar sum of transverse momenta of PF candidates in a cone around the electron (typically $\Delta R < 0.3$), excluding the electron's own deposits. The primary vertex (PV) is defined as the reconstructed pp interaction point, chosen as the vertex with the largest $\sum p_T^2$ of its associated tracks. Charged particles are counted only if they are

associated with the PV, while neutral particles receive a pileup correction using the median event energy density and an effective area.

In addition, a BDT-based electron identification [43] is used, fed by variables such as track–supercluster matching, shower-shape information and the inferred bremsstrahlung fraction. The BDT outputs a score in $[-1, 1]$, where -1 corresponds to a misidentified electron and 1 to a real prompt electron. For this thesis the working point WP90 is chosen that keeps about 90 % of prompt electrons with a misidentification rate of roughly 3 %.

4.4. Muon Identification

CMS muon identification [37] mirrors the electron case and uses PF-based isolation in a cone of $\Delta R < 0.4$, but with a different pileup correction. The isolation working point is analysis dependent and may be loosened or tightened for specific final states. Independent of the isolation, a dedicated muon ID defines working points from tracking and muon-system quality variables. This thesis employs the medium WP, requiring a tracker or global muon with a track through at least 80% of tracker layers, global-muon fit quality ($\chi^2 < 12$ and $\chi^2/\text{ndof} < 3$) and a kink-veto ($\chi_{\text{split}}^2 < 20$), yielding about 99.5% efficiency for simulated $W \rightarrow \mu$ and $Z \rightarrow \mu\mu$ decays.

4.5. b Identification with DEEPJET

The goal of b-tagging is to identify jets originating from b quarks. Compared with jets initiated by light-flavor quarks, charm quarks and gluons, b quark induced jets typically exhibit a long-lived b hadron whose decay leaves displaced tracks and often a reconstructed secondary vertex. They also exhibit a higher charged-particle multiplicity and larger track p_T with respect to the jet axis, reflecting the heavy-hadron mass and harder fragmentation. Additional signatures include occasional soft leptons from semileptonic b-hadron decays and a larger jet mass. DEEPJET [44] is a neural-network tagger that exploits low-level, per-particle information from the full jet to separate b jets from charm- and light-flavor-initiated jets.

Per-object features are first transformed with small 1×1 convolutional layers, with one branch each for charged candidates, neutral candidates and secondary vertices. The objects in each branch are then processed as sequences by recurrent layers to capture correlations among constituents. The ordering is physics-motivated (not temporal): charged tracks are sorted in descending order by impact-parameter significance, secondary vertices in descending order by flight-distance significance, and neutral PF candidates by the shortest angular distance to a secondary vertex, with a fallback to descending p_T when no secondary vertex is present. This places the most b-informative constituents first, allowing the recurrent neural network to exploit long-range dependencies. The outputs of these branches are combined with the global features in fully connected layers, and the final layer produces six class probabilities corresponding to b, $b\bar{b}$, leptonic-b (lep b), c, light-quark and gluon jets via a softmax activation. The standard DEEPJET b-tagging discriminator

used in CMS is defined as the sum of the three b-like outputs, $P(b) + P(b\bar{b}) + P(\text{lep}b)$.

The model is trained on simulated jets (anti- k_T , $R = 0.4$) from QCD multijet and fully hadronic $t\bar{t}$ samples. Jets are labeled as followed: if a jet contains at least one b hadron it is labeled as a b jet; if it contains no b but at least one c hadron it is labeled as a c jet; otherwise it is labeled as light flavor. During preprocessing, flavor and kinematic properties of the jet are decorrelated so the tagger does not simply learn differences in p_T or η differences.

Like in the μ and e identification a working point is applied to balance signal efficiency and mistag rate. In this analysis the medium working point of the DEEPTAU b-tagger is used, which means that about 1 % of all true light jets are incorrectly tagged as b jets. The b-tagging efficiency depends on p_T but is in the range of 75 – 85 %.

4.6. τ Identification with the DEEPTAU v2.5

Hadronic tau lepton (τ_h) identification in CMS has been significantly improved with the introduction of the DEEPTAU algorithm [45], a deep convolutional neural network (CNN) that discriminates genuine τ_h decays against background objects like jets, electrons, and muons. In this analysis the latest version, DEEPTAU v2.5 [46], is used. This is an enhanced iteration of the previous DEEPTAU v2.1 algorithm [45], providing better physics performance and improved agreement between data and simulation. This section describes the architecture, training and performance of DEEPTAU v2.5 and contrasts it with DEEPTAU v2.1.

4.6.1. Input Features

The visible decay products of τ_h are first reconstructed as PF candidates, which are then clustered into anti- k_T jets (Section 4.2). The Hadron-Plus-Strips (HPS) algorithm [47, 48] runs on these jets and combines their constituents into τ_h candidates. Each τ_h candidate provides a set of high-level features and is surrounded by many low-level detector signals which DEEPTAU utilizes as inputs. It uses global features as well as detailed per-particle information from the detector:

- Global event and τ_h candidate features: One global event variable, the average energy deposition density in the event ρ is included to quantify the amount of pileup in the environment. In addition, 42 high-level τ_h features are used, like kinematic quantities like p_T , η , ϕ , the τ_h charge, and number of charged/neutral hadrons in the τ decay, which are largely derived from the HPS reconstruction.
- Particle-by-particle low-level inputs: For each particle found within the τ_h candidate's signal cone ($\Delta R < 0.1$) or isolation cone ($\Delta R < 0.5$), a set of detailed features is provided. Depending on the object type, each particle contributes between 7 and 27 features per candidate, and 37 features for each fully reconstructed electron or muon. These features describe the particle's four-momentum, its spatial separation from the tau axis ($\Delta\eta$, $\Delta\phi$ relative to the τ_h direction), its identification information

(e.g. track quality, number of hits, calorimeter energy deposits) and its relationship to interaction vertices, like impact parameters and association with primary or secondary vertices. Additionally, each particle’s likelihood of originating from pileup interactions is included via the pileup per-particle identification (PUPPI) weight [49, 50].

- **Geometric arrangement:** To exploit spatial correlations as in image recognition, the particles in the τ_h signal and isolation regions are arranged into two structured grids in the (η, ϕ) plane. The inner grid covers the signal cone with fine granularity cells of size 0.02×0.02 (11×11 cells) and is intended to capture the core τ decay products, typically one or three charged pions and up to several π^0 's. The outer grid covers the isolation cone with cells of size 0.05×0.05 (21×21 cells) to capture surrounding particles and energy deposits. For each grid cell and each particle type, the feature vector of the highest- p_T particle in that cell is used as input, while cells without particles are zero-padded. In this way, the low-level inputs form two multi-channel “images” centered on the τ_h candidate, one for the core and one for the isolation region, allowing the CNN to analyze the spatial distribution of energy and tracks around the τ_h candidate.

4.6.2. Architecture

The neural network architecture of DEEPTAU v2.5 is illustrated in Figure 4.1. It is a deep CNN composed of multiple sub-networks that process different input streams. First, the high-level τ_h features are passed through a small stack of fully-connected dense layers to extract a high-level feature vector. In parallel, each cell of the inner and outer η - ϕ grids is processed by a small fully-connected network (embedding layer) to condense the per-cell information. These embedded cell features are then fed into convolutional layers: the inner grid is processed by a set of 5 convolutional layers with 3×3 kernels, while the larger outer grid passes through 10 convolutional layers (also 3×3 kernels) to capture the broader isolation pattern. Each convolutional stack produces a set of 64 output features. The outputs of the inner and outer convolutional blocks are then flattened and concatenated together with the high-level feature vector, forming a combined deep representation of the tau candidate. This combined representation is propagated through a series of five fully-connected layers which integrate all information. The network ends in an output layer that uses a softmax activation to produce four output scores (probabilities) corresponding to the hypotheses that the τ_h candidate is: (i) a real τ_h , (ii) an electron, (iii) a muon or (iv) a jet originating from a light quark- or gluon.

Like for DEEPTAU (Section 4.5) CMS defines multiple working points such as VLoose, Loose, Medium and Tight, by choosing thresholds on the DEEPTAU outputs to achieve certain efficiency or misidentification rates. The multiclass structure allows DEEPTAU to simultaneously discriminate against all background types in one go, rather than needing separate classifiers for electrons, muons, and jets as was common for earlier tau ID approaches.

4. Event Reconstruction and Identification

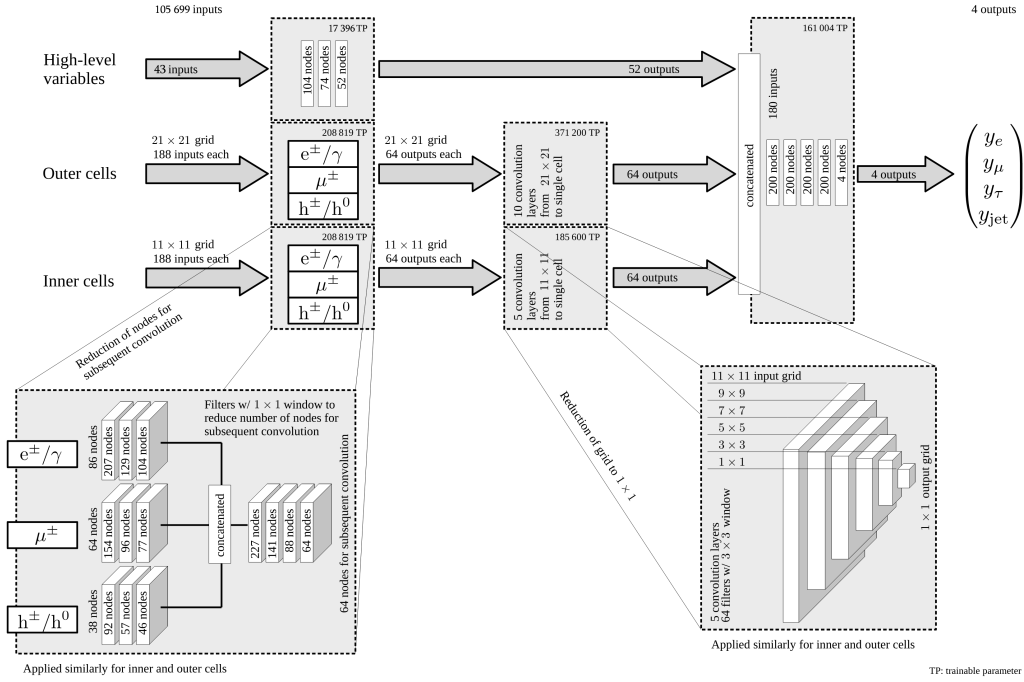


Figure 4.1.: DEEPTAU v2.5 architecture. High-level features and the inner/outer η - ϕ grids are processed in separate branches with dense layers. The two grids then pass through convolutional blocks that compress them to 1×1 feature vectors. All streams are concatenated and fed to additional dense layers, followed by a softmax function that returns class probabilities for electron, muon, τ_h , and light quark- or gluon-induced jets [51].

4.6.3. Training Strategy

The DEEPTAU v2.5 model was trained using a large, diverse sample of simulated and real data events from 2018 Run 2. This is a notable change from the previous version (v2.1), which was trained exclusively on simulation.

The simulated training samples for v2.5 span multiple processes to cover each class of τ_h candidate: genuine τ_h examples are taken from simulated $H \rightarrow \tau^+ \tau^-$ events or dedicated single- τ gun samples, with very high p_T up to ~ 3 TeV to improve performance in the boosted regime. Quark/gluon jet background examples are taken from simulated QCD and $t\bar{t}$ events, electrons from $Z/\gamma^* \rightarrow e^+ e^-$ (Drell-Yan) events and muons from $Z/\gamma^* \rightarrow \mu^+ \mu^-$ events.

The training dataset is balanced across τ_h p_T , η , decay modes and background types by applying weighting, shuffling and merging procedures so that the network does not over-fit to dominant kinematic regions. This balancing ensures that e.g. rare high- p_T τ_h examples are sufficiently represented.

The network is trained to minimize a multi-part loss function that includes terms for correctly classifying τ_h candidates and for reducing data-MC discrepancies. The primary component is the standard categorical cross-entropy loss $L_{\text{class}} = -\sum_i t_i \ln y_i$, which quantifies the error in predicting the true class i (τ_h , jet, e , or μ). In DEEPTAU v2.5, an additional adaptation term is incorporated to improve data/MC agreement. This is implemented via an adversarial network attached to the main CNN. This auxiliary sub-network takes as input the deep features from the CNN and attempts to distinguish whether a given τ_h candidate came from real data or from simulation. It effectively learns a binary classifier output y_{adv} with $y_{\text{adv}} = 0$ for data and 1 for MC, using a binary cross-entropy loss L_{adv} . Crucially, during training the gradient from this adversarial component is reversed so that the CNN is penalized if it produces features that allow the adversary to tell data vs. simulation apart. This forces the CNN to learn features that are as data-MC-invariant as possible. The total loss can be written as:

$$L_{\text{total}} = L_{\text{class}} + \lambda L_{\text{adv}} \quad \lambda < 0 \quad (4.3)$$

By tuning the hyperparameter λ , one controls the strength of data-MC adaptation. In practice, λ is set via hyperparameter optimization to maximally suppress data/MC differences without sacrificing classification performance.

The network was trained using modern deep learning libraries and techniques, and care was taken to avoid over-training. The resulting DEEPTAU v2.5 model is of similar computational complexity as v2.1, allowing it to be deployed in CMS data analyses and even considered for real-time use in the HLT for Run 3.

4.6.4. Performance and Comparison of DEEPTAU v2.5 vs. v2.1

The performance of the τ_h identification is quantified by the signal efficiency, defined as the fraction of genuine τ_h that are correctly identified, versus the misidentification probability, defined as the fraction of background objects that are mistakenly identified as

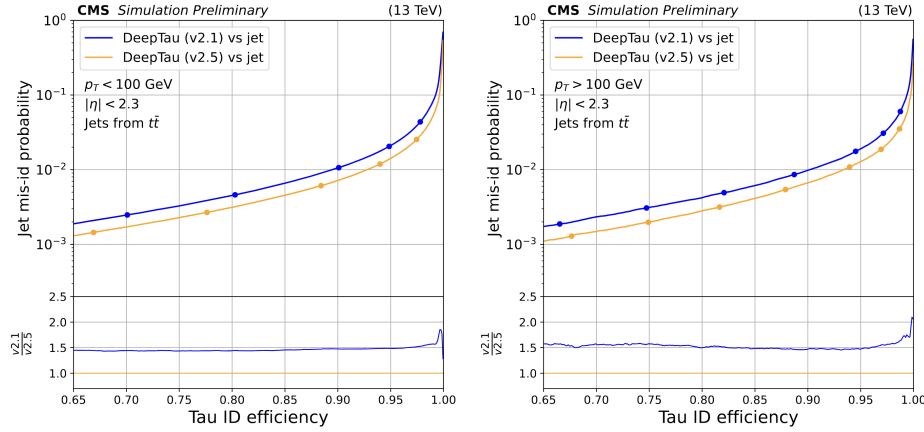


Figure 4.2.: Receiver–operating–characteristic (ROC) curves for τ_h identification evaluated on $t\bar{t}$ events. Left: $p_T^\tau < 100$ GeV; right: $p_T^\tau > 100$ GeV. The blue curve shows DEEPTAU v2.1 and the orange curve shows DEEPTAU v2.5. Compared to v2.1, v2.5 achieves a significantly better rejection of jets, reducing the jet misidentification by about 50% at fixed τ_h efficiency [46].

τ_h . Higher efficiency for a given misidentification probability or lower misidentification probability for a given efficiency, indicates a better classifier. DEEPTAU v2.5 was found to yield substantial improvements in τ_h identification performance compared to DEEPTAU v2.1 across all major background categories.

Figure 4.2 shows representative Receiver Operating Characteristic (ROC) curves for DEEPTAU v2.5 and v2.1, evaluated on simulation, plotting the jet mis-ID probability versus tau efficiency. For every working points, the v2.5 curve lies well below of the v2.1 curve, indicating higher τ_h identification efficiency at any given background rate. In particular, the probability for a $t\bar{t}$ jet to be misidentified as a τ_h is reduced by approximately 50% at the same signal efficiency with DEEPTAU v2.5, as compared to v2.1, significantly enhancing the purity of τ_h samples in analyses.

The performance of the discriminator against electron misidentification, where electrons are misidentified as narrow one-prong τ_h candidates, is shown in Figure 4.3. DEEPTAU v2.5 performs similarly to v2.1 with only minimal differences. In the low- p_T region, the electron misidentification rate versus τ_h efficiency is better in v2.5, while for electrons with high p_T ($p_T > 100$ GeV) a small degradation in v2.5 is observed. This small loss at high- p_T is negligible for practical purposes.

For muon misidentification, the performance of v2.5 is essentially unchanged relative to v2.1 for $p_T < 100$ GeV and a little bit worse for $p_T > 100$ GeV. But this is again negligible because muons are already efficiently rejected by the dedicated CMS muon identification (see Section 4.4)

In summary, the major gains of v2.5 lie in the τ_h vs. jet separation, while τ_h vs. electron and τ_h vs. muon was already very effective before. This results in significantly fewer QCD jets misidentified as τ_h in the selected data, improving the signal purity in the $b\bar{b}\tau^+\tau^-$ final state. In analyses limited by jet \rightarrow τ fakes, such as $H \rightarrow \tau^+\tau^-$ or $HH \rightarrow b\bar{b}\tau^+\tau^-$, a

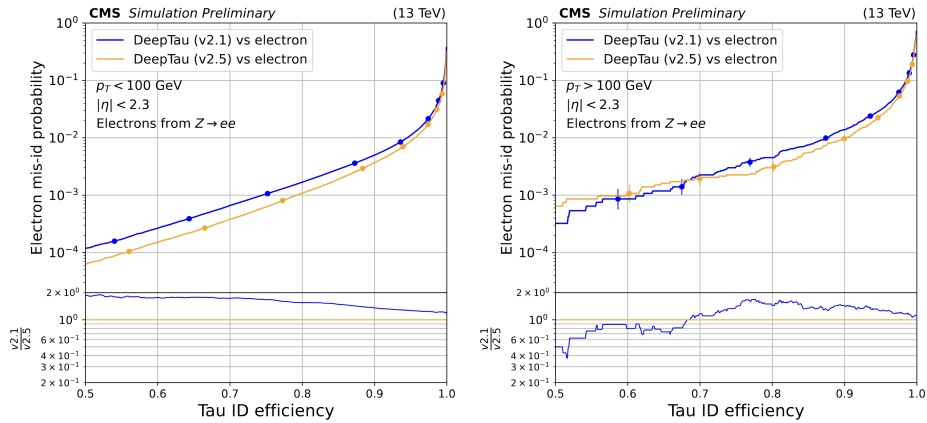


Figure 4.3.: ROC curves for τ_h identification against electrons, evaluated on Drell-Yan events. Left: $p_T^{\text{th}} < 100$ GeV; right: $p_T^{\text{th}} > 100$ GeV. Overall, v2.5 maintains good rejection with a small degradation at high p_T .

30-50% reduction in misidentification rate can boost the attainable sensitivity or reduce background systematics.

4.6.5. Scale Factors and Validation in Data

Any identification algorithm in CMS must be validated with real data and assigned scale factors (SFs) to correct the simulation if needed. These SFs are defined as the ratio of identification efficiency in data to that in simulation, $\text{SF} = \epsilon_{\text{data}}/\epsilon_{\text{MC}}$, typically measured in a clean sample of genuine τ leptons (e.g. from $Z \rightarrow \tau^+\tau^-$ decays) using a tag-and-probe technique. Thanks to the data-MC adaptation training, the data-to-simulation agreement for τ_h ID has markedly improved relative to v2.1. When tested on an independent 2018 data sample (not used in training), the observed τ_h identification efficiency with DEEPTAU v2.5 was found to agree with the MC-predicted efficiency to within about 10%, across multiple working points. This is a notable improvement over DEEPTAU v2.1, where differences were larger and required significant corrections.

CMS measures separate τ_h ID scale factors in bins of τ_h decay mode (DM) and p_T to account for any residual differences that might depend on the τ 's decay topology or kinematic properties. For v2.5, the simulation-to-data correction factors (SFs) were found to be much closer to unity (1.0) than for v2.1, indicating a smaller need for adjustments. Figure 4.4 illustrates a comparison of the measured SFs for DEEPTAU v2.5 vs. v2.1 as a function of τ_h p_T . One observes that for v2.1, especially at low p_T (~ 20 – 40 GeV), the scale factors deviate from 1 by a significant amount, reflecting that the MC overestimated the τ_h ID efficiency in that regime. In contrast, the SFs for DEEPTAU v2.5 are closer to 1.0 across the entire p_T spectrum.




Figure 4.4.: Data-to-simulation scale factors (SFs) for τ_h identification. The combined values are weighted by the τ decay branching fractions and by the HPS reconstruction efficiencies per decay mode (DM). DEEPTAU v2.5 scale factors are closer to unity than v2.1, indicating that the training improvements in v2.5 reduce the sensitivity to differences between data and MC simulations. [46]

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$

5.1. Overview of Current Results

The latest CMS publication on nonresonant Higgs boson pair production in the $HH \rightarrow b\bar{b}\tau^+\tau^-$ final state uses the full Run 2 data set recorded in proton-proton collisions between 2016 and 2018 at a center-of-mass energy of 13 TeV, corresponding to an integrated luminosity of 138 fb^{-1} . [5]. It combines the three $\tau^+\tau^-$ -decay channels mentioned in Section 2.3 ($e\tau_h, \mu\tau_h, \tau_h\tau_h$). The analysis uses the DEEPJET algorithm for b-jet identification (see Section 4.5), the DEEPTAU v2.1 algorithm (see Section 4.6) for τ_h -identification and a deep neural network (DNN) based event categorization. In addition to the production channel via ggF, described in Section 2.2, the production via vector boson fusion (VBF) is used with a cross section of $\sigma_{HH}^{\text{VBF}} = 1.726^{+0.03\%}_{-0.04\%}(\text{scale}) \pm 2.1\%(\text{PDF} + \alpha_s)$.

Using the combined 2016–2018 data, CMS sets an 95 % CL upper limit on $\sigma_{\text{ggF+VBF}}(\text{pp} \rightarrow \text{HH})$ of 3.3 (5.2) times the SM expectation, where the value in parenthesis refers to the expected median limit in the absence of signal. The dataset recorded in 2018, corresponding to an integrated luminosity of 59.8 fb^{-1} provides the single-year best sensitivity with an observed (expected) upper limit of 5.5 (8.2). The 95 % CL upper limits and their uncertainties are shown in Figure 5.1.

Interpreting the result in terms of coupling modifiers, the analysis constrains the trilinear Higgs self-coupling to $-1.7 < \kappa_\lambda < 8.7$ (95 % CL) and the contact interaction between two Higgs bosons and two vector bosons to $-0.4 < \kappa_{2V} < 2.6$ (95 % CL), under SM assumptions for all other couplings.

The CMS Run 2 publication establishes the benchmark sensitivity for $(HH \rightarrow b\bar{b}\tau^+\tau^-)$ with current data and methods. For this thesis the 2018 data, the ggF production mode and resolved topologies in the three $(\tau^+\tau^-)$ channels are used. The CMS combination quoted above (2016–2018, ggF+VBF) provides the reference against which the performance of this analysis can be compared.

Furthermore, ATLAS published a study of nonresonant Higgs–boson pair production [52, 53]. It combines the $b\bar{b}\tau^+\tau^-$, $b\bar{b}\gamma\gamma$, $b\bar{b}b\bar{b}$, multilepton and $b\bar{b}l\bar{l} + E_T^{\text{miss}}$ channels using the full Run 2 dataset at $(\sqrt{s} = 13 \text{ TeV})$ to constrain HH production via ggF and VBF. The combined fit yields a observed (expected) 95 % CL upper limit on the signal strength of 2.9 (2.4). For the $b\bar{b}\tau^+\tau^-$ channel, the observed (expected) 95 % CL upper limit is 5.9 (3.3) and is shown in Figure 5.2.

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$




Figure 5.1.: Observed and expected 95 % CL upper limits on the Higgs boson pair production rate in the $b\bar{b}\tau^+\tau^-$ final state, expressed as $\sigma(pp \rightarrow HH)/\sigma_{SM}$. Results are shown for 2016, 2017, 2018, and their combination using the full Run 2 data with $\mathcal{L}_{int} = 138 \text{ fb}^{-1}$ at $\sqrt{s} = 13 \text{ TeV}$. SM couplings are assumed ($\kappa_\lambda = \kappa_t = \kappa_V = \kappa_{2V} = 1$) [5]




Figure 5.2.: Observed and expected 95 % CL upper limits on the Higgs boson pair production rate in the $b\bar{b}\tau^+\tau^-$ final state, expressed as $\sigma(pp \rightarrow HH)/\sigma_{SM}$. Results are shown for the $e\tau_h$, $\mu\tau_h$, $\tau_h\tau_h$ final states and their combination using the full Run 2 data with an integrated luminosity of $\mathcal{L}_{int} = 140 \text{ fb}^{-1}$ at $\sqrt{s} = 13 \text{ TeV}$. SM couplings are assumed ($\kappa_\lambda = \kappa_t = \kappa_V = \kappa_{2V} = 1$) [53]

5.2. Signal Process

The signal considered in this thesis is nonresonant Higgs-boson pair production via ggF,

$$pp \rightarrow HH \rightarrow b\bar{b}\tau^+\tau^- \quad (5.1)$$

with $m_H = 125$ GeV. Of the τ -pair the final states, mentioned in Section 2.3, $e\tau_h$, $\mu\tau_h$, and $\tau_h\tau_h$ are considered. The analysis targets the resolved regime of non-boosted Higgs bosons, where the $H \rightarrow b\bar{b}$ decay is reconstructed with two separate small-radius anti- k_T jets and the τ decay products form isolated, narrow objects.

At leading order (LO) in the SM, $gg \rightarrow HH$ proceeds through two one-loop amplitudes: a triangle diagram proportional to the trilinear Higgs self-coupling λ_{HHH} and a box diagram proportional to the top Yukawa coupling y_t (see Section 2.2). The two amplitudes interfere destructively, which reduces the total production rate. Near the kinematic threshold, the Higgs bosons carry moderate transverse momenta, so their decay products are typically well separated ($\Delta R > 0.4$). So two b-tagged anti- k_T jets from $H \rightarrow b\bar{b}$ and two τ objects forming the $H \rightarrow \tau^+\tau^-$ candidate are expected.

5.3. Background Processes

The search for $HH \rightarrow b\bar{b}\tau^+\tau^-$ is affected by several SM processes that can mimic the signal topology, either through genuine τ leptons and b jets, through the misidentification of jets as τ_h , and/or through the misidentification of light- or charm-flavor jets as b jets. The relative background composition depends on the $\tau^+\tau^-$ final state.

The background is estimated from multiple simulated samples normalized to their theoretical cross sections at highest available accuracy. Background from QCD multijet events, where purely light quark- or gluon-induced jets are misidentified as τ_h or b jet, further on referred to as QCD multijet production, is determined in a data-driven way using the ABCD method [54].

5.3.1. Top–Antitop Production

Top–antitop quark ($t\bar{t}$) production is the dominant background in the semileptonic channels of this analysis. At $\sqrt{s} = 13$ TeV the inclusive cross section is about 832 pb, dominated by ggF, with a smaller contribution from $q\bar{q}$ annihilation [55]. The LO Feynman diagrams are shown in Figure 5.3.

Each top quark almost always decays as $t \rightarrow bW$, other modes are CKM-suppressed [56]. The two W bosons then decay either hadronically ($W \rightarrow q\bar{q}'$) or leptonically ($W \rightarrow \ell\nu$). For a t quark pair, this results in (i) a fully hadronic final state, (ii) semileptonic final states with one leptonic and one hadronic W decay, and (iii) fully leptonic final states.

In the $\ell\tau_h$ channels, $t\bar{t}$ production is particularly abundant because a prompt e/μ is produced in the $W \rightarrow \ell\nu$ decay and a τ_h can arise either from the $W \rightarrow \tau\nu$ decay or from

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$

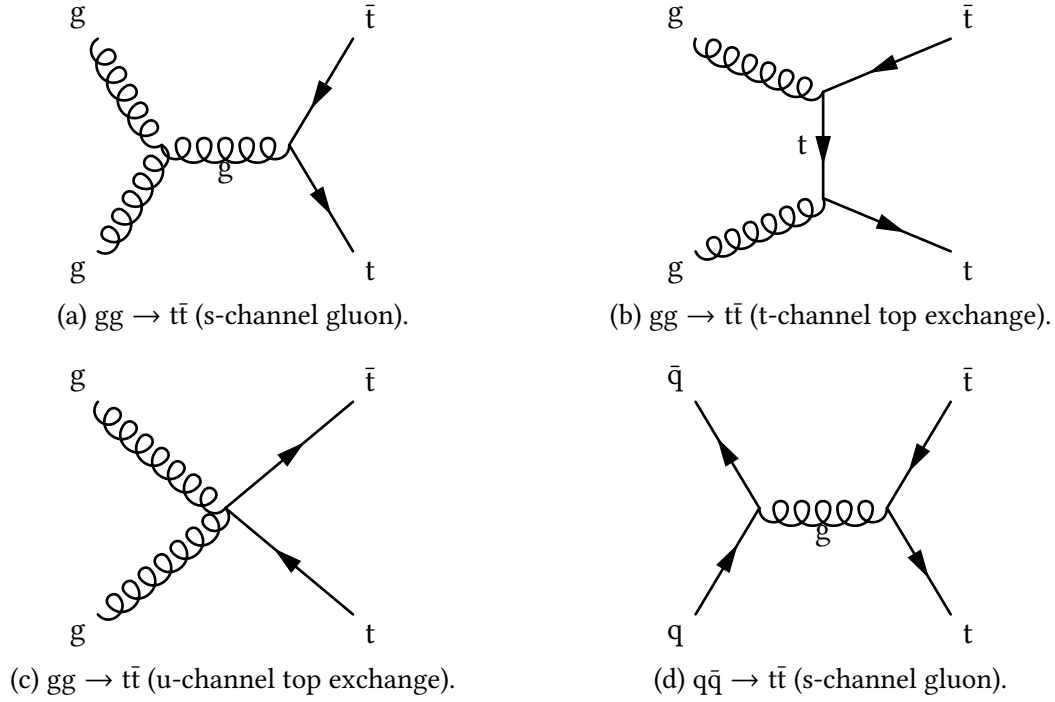


Figure 5.3.: Leading-order QCD diagrams for $t\bar{t}$ production at the LHC: ggF via s-, t-, and u-channel (top/gluon exchange) and quark–antiquark annihilation via an s-channel gluon.

a jet misidentified as τ_h . In the $\tau_h \tau_h$ channel the yield is smaller but still with Drell–Yan resembles the biggest contribution: events enter through the $W \rightarrow \tau \nu$ decay on one or both top legs or through $\ell, \text{jet} \rightarrow \tau_h$ misidentification. Lepton vetoes, opposite-sign charge, and the requirement of two identified τ_h candidates reduce this background, but its large cross section still leads to significant contamination in the baseline selection.

Despite the similar visible content (two b-jets + $\tau^+ \tau^-$ decay products + p_T^{miss}), $t\bar{t}$ production exhibits characteristic features that differ from the signal:

- *Resonance structure:* In HH events the b-jet pair and the $\tau^+ \tau^-$ system each originate from a Higgs boson, so they yield values of $m_{b\bar{b}} \simeq m_H$ and $m_{\tau\tau} \simeq m_H$. Because of the challenging neutrino reconstruction in the $H \rightarrow \tau^+ \tau^-$ decay, the $m_{\tau\tau}$ distribution can form a broad peak for signal events. In $t\bar{t}$ production, the b jets arise from different top decays and any τ leptons come from W bosons on separate legs; consequently neither $m_{b\bar{b}}$ nor $m_{\tau\tau}$ peaks at m_H .
- *Global kinematics:* variables like angular correlations ($\Delta R_{b\bar{b}}, \Delta R_{\tau\tau}$) follow different patterns in $t\bar{t}$ production due to the presence of two heavy top quarks and their decay topology.

The analysis exploits these differences with multivariate discriminants (Neural Network).

5.3.2. Drell-Yan $Z/\gamma^* + \text{jets}$




Figure 5.4.: Representative DY Z/γ^* +jets topologies with $Z/\gamma^* \rightarrow \tau^+\tau^-$: (a) Compton process $qg \rightarrow Zq$; (b) $Z+b\bar{b}$ from gluon splitting $g \rightarrow b\bar{b}$. These illustrate sources of light- and heavy-flavor jets in DY events.

At $\sqrt{s} = 13$ TeV the cross section times branching fraction for Drell-Yan (DY) production with subsequent decay into a pair of leptons is of order 6000 pb, produced dominantly via $q\bar{q} \rightarrow Z$ with additional jets from QCD radiation. The production of genuine $\tau^+\tau^-$ pairs, as well as ee and $\mu\mu$ pairs, contributes when either a lepton or an additional jet is misidentified as τ_h .

The b jets can arise from radiated gluons that split into a $b\bar{b}$ pair or from misidentification of light-flavor, c quark and gluon jets. Two representative Feynman diagrams are shown in Figure 5.4.

After requiring two b-tagged jets, the DY rate is reduced but remains relevant, especially in the $\tau_h\tau_h$ final state. Like for the background from $t\bar{t}$ production, $m_{\tau\tau}$ provides separation, since DY populates a broad peak near the Z mass (~ 90 GeV), while the signal prefers values around m_H .

5.3.3. Single Top

Single top quarks are produced in electroweak processes with a summed cross section of about 217 pb at $\sqrt{s} = 13$ TeV [57]. The dominant modes are t-channel and the associated tW production, which are shown in the Feynman diagrams in Figure 5.5.

At LO, t-channel production yields one top quark, one b jet and a characteristic forward light jet, while tW production leads to a top quark together with a W boson. After the $t \rightarrow bW$ and $W \rightarrow q\bar{q}'$ or $W \rightarrow \ell\nu$ ($\ell = e, \mu, \tau$) decays, typical configurations contain one b jet and either additional light-flavor or c quark jets or a single prompt lepton. None of these LO final states provide both a $b\bar{b}$ pair and a $\tau^+\tau^-$ pair, so the selection requiring two b-tagged jets and a $\tau^+\tau^-$ candidate suppresses single t quark more strongly than $t\bar{t}$ production.

Leakage into the signal regions occurs mainly through:

- extra heavy-flavor radiation or gluon splitting that yields a second b jet,

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$

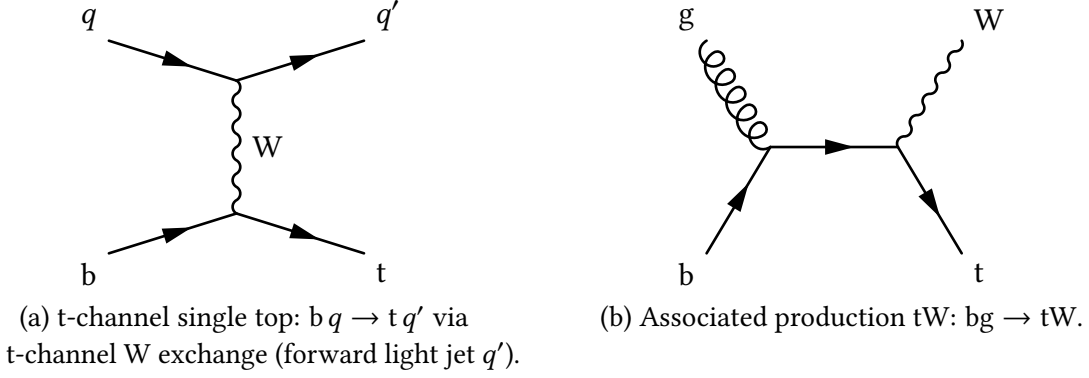


Figure 5.5.: Leading-order electroweak single-top quark production topologies: (a) t-channel production with a forward light jet, (b) associated tW production. Subsequent decays $t \rightarrow bW$ and $W \rightarrow q\bar{q}'/\ell\nu$ (not shown) yield one b jet and either additional light jets or a prompt lepton.

- misidentification of gluon, light-flavor or c quark jets as b jets,
- $W \rightarrow \tau\nu$ providing one genuine τ in the $\ell\tau_h$ final state together with jet $\rightarrow \tau_h$ misidentification.

Compared to the background from $t\bar{t}$ production and to the $HH \rightarrow b\bar{b} \tau^+ \tau^-$ signal, single t production events typically show: lower jet and b-jet multiplicities, a forward light-flavor or c quark in the t-channel and m_{bb} , $m_{\tau\tau}$, which do not peak at m_H . This reduces the background from single t quark production to a subdominant level after the final event selections.

5.3.4. W+jets

W boson production in association with jets (W+jets) is kinematically similar to DY production but occurs with a substantially larger rate (about an order of magnitude higher at $\sqrt{s} = 13$ TeV). It proceeds predominantly via quark-antiquark annihilation with additional jets from QCD radiation, followed by the leptonic $W \rightarrow \ell\nu$ ($\ell = e, \mu, \tau$) decay.

For this analysis only the leptonic W decay is relevant. Because the selection requires a $\tau^+ \tau^-$ candidate and two b-tagged jets, genuine W+jet production is strongly suppressed and the contamination of the signal region is mainly driven by:

- jet $\rightarrow \tau_h$ misidentification, together with the prompt e/μ from the $W \rightarrow \ell\nu$ decay, mimics the $\ell\tau_h$ final states, and double jet $\rightarrow \tau_h$ fakes mimic the $\tau_h\tau_h$ final state,
- light-flavor and c quark jets \rightarrow b misidentification,
- genuine heavy flavor in $W+b\bar{b}$ or $W+c\bar{c}$ that yields b-tagged jets in otherwise W+jets-like events.

Kinematically, the W+jets process is dominant at $m_T(\ell, p_T^{\text{miss}})$ values close to m_W . The distributions of m_{bb} and $m_{\tau\tau}$ are not peaking. In the $\tau_h\tau_h$ channel the fact that two jets would have to be misidentified as τ_h reduces the yield to a minor level.

5.3.5. Single Higgs

Single Higgs boson production (ggF, VBF, VH, $t\bar{t}H$, and tHq) yields final states with b jets and τ leptons that overlap with $HH \rightarrow b\bar{b} \tau^+ \tau^-$:

- ZH production with ($H \rightarrow \tau^+ \tau^-$, $Z \rightarrow b\bar{b}$) or ($H \rightarrow b\bar{b}$, $Z \rightarrow \tau^+ \tau^-$) mimics the two-resonance structure and can populate high-purity regions.
- $t\bar{t}H$ and tHq produce multiple b jets and leptons. Selections with exactly two b-tagged jets and the resolved topology reduce these contributions.

Kinematically, single Higgs boson production events often feature only one mass peak near m_H ($m_{b\bar{b}}$ or $m_{\tau\tau}$), whereas the signal prefers both quantities peaking close to m_H . Together with the low cross section, this process accounts for only a small proportion of the background.

5.3.6. Other Backgrounds

Other, minor background processes are grouped here. After requiring two b tags and a $\tau^+ \tau^-$ pair, their rates are very small:

- **Diboson production** (WW, WZ, ZZ): Events contribute when accompanied by, occasionally heavy flavor jets or via object misidentification. In particular, $ZZ \rightarrow b\bar{b} \tau^+ \tau^-$ shares the same final state as the signal while it has a very small yield and a different resonance pattern (both $m_{b\bar{b}}$ and $m_{\tau\tau}$ peak at m_Z).
- **Triboson production** (WWW, WWZ, WZZ, ZZZ): These processes have very small cross sections and negligible yields after selection and are therefore included only for completeness.
- **Electroweak V + 2 jets**: Events from electroweak vector boson production in association with two VBF-like jets are characterized by two forward/backward jets, large dijet mass and rapidity separation, and very little heavy flavor jets. Requiring two b-tagged jets suppresses the contribution of this background to a negligible level.

5.3.7. QCD Multijet Production

QCD multijet production enters the signal selection when one or more jets are misidentified as τ_h . In the $\ell\tau_h$ final state this appears as a single jet misidentified as τ_h , with the accompanying lepton possibly nonprompt or misidentified, while in the $\tau_h\tau_h$ final state both τ_h candidates may originate from $\text{jet} \rightarrow \tau_h$ misidentification. The necessary b tagged jets come, either through gluon, light-flavor or c quark to b misidentification or via genuine heavy-flavor production ($b\bar{b}$, $c\bar{c}$).

The background from multijet QCD production is obtained from data with an ABCD method based on two axes that are approximately uncorrelated: (i) the τ_h identification working point and (ii) the $\tau_h\tau_h$ or $\ell\tau_h$ charges.

- **Charge**: same sign (SS) if $q_1 q_2 > 0$

- **Charge:** opposite sign (OS) if $q_1 q_2 < 0$
- τ_h **isolation (“Iso”):** DEEPTAU VSjet Medium
- τ_h **anti-isolation (“Anti-Iso”):** fails Medium but passes VVLoose

Because jet charges are essentially random, SS and OS events occur at roughly similar rates. The explicit definitions of the regions are given in Tables 5.1 for the $\tau_h \tau_h$ and in 5.2 for the $\ell \tau_h$ final state.

Table 5.1.: ABCD regions for the $\tau_h \tau_h$ final state. Iso τ_h pass the DEEPTAU VSjet Medium working point; Anti-Iso τ_h fail the Medium but pass the DEEPTAU VSjet VVLoose working point.

Region	Charge	$\tau_{h,1}$ isolation	$\tau_{h,2}$ isolation
A (target)	OS ($q_1 q_2 < 0$)	Iso	Iso
B (shape)	OS ($q_1 q_2 < 0$)	Iso	Anti-Iso
C (norm)	SS ($q_1 q_2 > 0$)	Iso	Iso
D (norm)	SS ($q_1 q_2 > 0$)	Iso	Anti-Iso

Table 5.2.: ABCD regions for the $\ell \tau_h$ final states. Iso τ_h pass the DEEPTAU VSjet Medium working point; Anti-Iso τ_h fail the Medium but pass the DEEPTAU VSjet VVLoose working point. Lepton ID/iso follows the nominal selection.

Region	Charge	τ_h isolation
A (target)	OS ($q_1 q_2 < 0$)	Iso
B (shape)	OS ($q_1 q_2 < 0$)	Anti-Iso
C (norm)	SS ($q_1 q_2 > 0$)	Iso
D (norm)	SS ($q_1 q_2 > 0$)	Anti-Iso

First, in each region $X \in \{A, B, C, D\}$, non-QCD backgrounds (e.g. $t\bar{t}$ and DY production) are subtracted from data:

$$N_X^{\text{QCD}} = N_X^{\text{data}} - \sum_{p \neq \text{QCD}} N_X^p. \quad (5.2)$$

Assuming charge and isolation are nearly uncorrelated for events where the τ_h candidate results from a misidentified jet, the yield in the signal region A is calculated with

$$N_A^{\text{QCD}} = \frac{N_B^{\text{QCD}} N_C^{\text{QCD}}}{N_D^{\text{QCD}}}. \quad (5.3)$$

The kinematic distribution of events from QCD multijet production of an observable x is taken from region B and scaled to N_A^{QCD} :

$$\left(\frac{dN}{dx}\right)_A^{\text{QCD}} = \alpha \left(\frac{dN}{dx}\right)_B^{\text{QCD}} \quad \text{with} \quad \alpha = \frac{N_A^{\text{QCD}}}{N_B^{\text{QCD}}}. \quad (5.4)$$

The resulting distributions are used in the final fit together with the simulated backgrounds and the signal model.

5.4. Event Simulation

All processes apart from QCD multijet production are estimated from simulation. The samples have been produced based on the 2018 Ultra Legacy conditions and are processed through the full CMS chain: the hard scattering process at parton level, parton showering and hadronization, a GEANT4-based detector simulation [58] and offline reconstruction identical to data. The underlying event is modeled with the CP5 tune [59] and additional inelastic interactions simulated with the pythia event generator [60] are overlaid. Simulated events are reweighted to reproduce the pileup profile observed in data.

5.4.1. Matrix Element (ME) event generators

ME event generators compute the hard-scattering amplitudes at fixed order, LO and next-to-leading-order (NLO) in perturbation theory, in QCD and EWK, integrate them over phase space with PDFs at a chosen factorization scale and produce unweighted parton-level events. Scale and PDF choices are provided as event weights. When extra jets are taken into account already at the parton level, matching or merging schemes are used to avoid double counting between partons in the final state of the ME calculation and the subsequent parton shower.

The following ME generators are used in this analysis:

- **MadGraph5 (LO) [61]:** used for $gg \rightarrow HH \rightarrow b\bar{b} \tau^+ \tau^-$ and for W +jets production. It generates tree-level amplitudes with optional extra partons in the ME. Multiple jet multiplicities in the matrix element are consistently combined using the MLM merging prescription, whereby, after parton showering, events are retained only if each matrix-element parton can be matched to a reconstructed jet and no additional jet above the matching scale is present. This avoids double counting between ME and shower emissions and yields a smooth transition between their respective descriptions. Without merging, additional jets beyond the matrix-element multiplicity are produced purely through the parton shower model (see Section 5.4.2).
- **MadGraph5_aMC@NLO (NLO+PS) [62]:** shortened as aMC@NLO. Used for $DY Z/\gamma^* + \text{jets}$, diboson, and triboson production. This event generator employs the MC@NLO matching scheme to combine NLO matrix elements with the parton shower model, which can introduce negative event weights, and provides per-event weights for scale and PDF variations. When several jet multiplicities at NLO are combined, the FxFx NLO merging procedure [63] is used, with a merging scale that separates ME and shower emissions, preserving the inclusive NLO cross section and improving the description of events with one or more hard jets. A significant fraction of negative weights is expected.
- **POWHEG (NLO+PS) [64]:** used for $t\bar{t}$, single t quark, and single Higgs boson production. Implements NLO matching by generating the hardest emission directly from the

real ME dressed with a Sudakov form factor. This usually leads to positive-weight events and a controlled first-emission spectrum. Spin correlations in heavy-particle decays are preserved, and process-specific damping parameters regulate very hard radiation.

In this analysis, both LO and NLO simulated samples are used, while the NLO generators provide an improved description of the kinematic distributions. Table 5.3 summarizes the ME event generators and perturbative orders used for each process. In the column *ME generator*, the merging scheme used to combine ME partons with the parton shower are given in parentheses.

With POWHEG, the first (hardest) emission is generated directly from the real matrix element and regulated by a Sudakov form factor that vetoes harder shower emissions. Events therefore almost always yield a positive weight, and the first-emission tail is often a bit harder (tunable, e.g. with `h_damp`). In MadGraph5_aMC@NLO, the MC@NLO method makes two event types. S (“Born-like”) events start without an extra parton and let the shower create the first emission. H (“hard”) events already include one extra parton from the real matrix element. To avoid double counting, the shower approximation is subtracted in the event weights, which may turn some weights negative. Combined, S and H events reproduce the full NLO prediction. Typically, POWHEG gives a slightly harder first emission, while aMC@NLO follows the shower more closely. Negative weights are rare for POWHEG but may occur at a level of 10% for aMC@NLO and even higher with FxFx merging. Both simulations provide per-event renormalization and factorization scale, and PDF reweighting.

5.4.2. PYTHIA 8 parton shower and the CP5 tune

The event simulation program PYTHIA 8[60] takes the parton-level events from the ME generators and turns them into realistic final states made of stable particles. It does this in a few steps: it adds initial-state radiation (ISR) from the incoming partons, adds final-state radiation (FSR) from the outgoing partons, simulates extra semi-soft scatterings in the pp collision (multiparton interactions, MPI) and finally hadronizes the partons into hadrons using the Lund string model [65, 66]. Unstable particles are then eventually decayed. The result is a collection of stable particles that are passed to the detector simulation.

The CP5 tune [59, 67] sets the many shower and underlying-event parameters used by PYTHIA 8. It is a CMS Run 2 tune at $\sqrt{s} = 13$ TeV, built to describe minimum-bias and underlying-event data. In practice, CP5 controls the overall activity from MPI, the amount and hardness of ISR/FSR, how color reconnection takes place, and details of the beam-remnant treatment. Using one common tune for all samples keeps the modeling consistent across processes.

5.4.3. Pileup, reconstruction and object calibration

Pileup from additional pp interactions in the same or nearby bunch crossings is included by overlaying minimum-bias collisions to the hard-scattering event. The number of interactions is reweighted to match the distribution measured in data. An uncertainty is

Table 5.3.: Event generators and perturbative order in α_s used for signal and background samples in the 2018 UL simulation. All samples are showered and hadronized with PYTHIA 8 using the CP5 tune.

Process	Order	ME generator
$gg \rightarrow HH \rightarrow b\bar{b} \tau^+ \tau^-$	LO	MadGraph5
$t\bar{t}$	NLO	POWHEG
DY+jets	NLO	aMC@NLO (FxFx)
Single t	NLO	POWHEG
W+jets	LO	MadGraph5 (MLM)
Diboson	NLO	aMC@NLO (FxFx)
Triboson	NLO	aMC@NLO
Single H	NLO	POWHEG

assigned by varying the assumed inelastic cross section and the effect is propagated to all downstream observables.

Before reconstruction, the detector response is simulated with a detailed GEANT4-based model [58]. Generator-level particles are transported through the full detector geometry. GEANT4 simulates electromagnetic and hadronic showers, ionization energy loss, multiple scattering and decays in flight in each subdetector. The resulting energy deposits and hits are digitized with realistic electronics response, timing and noise. Run-dependent conditions, like alignment, calibrations and dead/noisy channels, are applied so that simulated events reproduce the data-taking conditions as closely as possible.

The events are then filtered in the triggers (L1 + HLT) and reconstruction follows the PF approach (see Chapter 4). Data–simulation differences are corrected with per-object scale factors and energy-scale/resolution calibrations, and their uncertainties are propagated to the final observables.

5.5. Event Selection

5.5.1. $\tau^+ \tau^-$ -Pair Selection

For this analysis three $\tau^+ \tau^-$ decay channels ($e\tau_h$, $\mu\tau_h$ and $\tau_h\tau_h$) in the decay modes 0, 1, 10 and 11 of τ_h are used (see Section 2.3). Electron identification uses the BDT-based no-isolation working point WP90 (Section 4.3) and muons must satisfy the cut-based Medium ID (Section 4.4). For τ_h candidates, DEEPTAU v2.5 working points are applied per channel as listed in Table 5.4. The vsJet Medium point defines the well identified τ_h leg, while looser vsEle/vsMu points suppress misidentification from electrons and muons. Additional working points are stored for control regions, e.g. the estimation of the background from QCD multijet production, with the ABCD method.

Requirements on p_T $|\eta|$, $|d_z|$, $|d_{xy}|$, lepton isolation, and the transverse mass m_T are applied to select well-measured, prompt objects and to suppress backgrounds. For a visible

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$

object $X \in \{e, \mu, \tau_h\}$, m_T is defined as

$$m_T(X, \vec{p}_T^{\text{miss}}) = \sqrt{2 p_T^X p_T^{\text{miss}} (1 - \cos \Delta\phi(\vec{p}_T^X, \vec{p}_T^{\text{miss}}))}, \quad (5.5)$$

where p_T^X and p_T^{miss} denote the magnitudes of the transverse-momentum vectors \vec{p}_T^X and \vec{p}_T^{miss} , and $\Delta\phi(\vec{p}_T^X, \vec{p}_T^{\text{miss}})$ is the azimuthal angle between them. The $|\eta|$ requirement confines electrons, muons, and τ_h candidates to detector regions with reliable tracking and calorimeter coverage, reducing edge effects and mismodelling. The longitudinal and transverse impact parameters, $|d_z|$ and $|d_{xy}|$, enforce compatibility with the primary interaction vertex and therefore suppress pileup tracks and nonprompt leptons from heavy-flavor decays. The isolation requirement further reduces jets misidentified as leptons and leptons produced inside jets by limiting additional activity in a cone around the object. For electrons and muons a relative isolation requirement $I_{\text{rel}} < 0.15$ is used, while for τ_h candidates the DEEPTAU v2.5, vsJet Medium discriminator is used as an effective isolation criterion. In the $e\tau_h$ and $\mu\tau_h$ channels, a transverse-mass requirement $m_T^{e,\mu} < 70$ GeV is applied to suppress W +jet events. Because genuine $W \rightarrow \ell\nu$ decays yield large values of m_T up to around m_W , this criterion removes much of the W +jets background while keeping most of the signal events.

The thresholds on these quantities for each object type are shown in Table 5.5

Table 5.4.: Identification working points of DeepTau v2.5 used for τ_h identification

Channel	Object	vsEle	vsMu	vsJet
$e\tau_h$	τ_h	Tight	VLoose	Medium
$\mu\tau_h$	τ_h	VVLoose	Tight	Medium
$\tau_h\tau_h$	τ_h (each)	VVLoose	VLoose	Medium

Table 5.5.: Per-channel object requirements used in the $\tau^+ \tau^-$ selection.

Channel	Object	p_T [GeV]	$ \eta $	$ d_z $ [cm]	$ d_{xy} $ [cm]	$I_{\text{rel}}^{e/\mu}$	m_T [GeV]
$e\tau_h$	e	> 33	< 2.1	< 0.2	< 0.045	< 0.15	< 70
	τ_h	> 30	< 2.3	< 0.2			
$\mu\tau_h$	μ	> 25	< 2.1	< 0.2	< 0.045	< 0.15	< 70
	τ_h	> 30	< 2.3	< 0.2			
$\tau_h\tau_h$	τ_h (each)	> 40	< 2.1	< 0.2			

To keep the three $\tau^+ \tau^-$ channels mutually exclusive, a dedicated lepton veto is applied using very loose “veto-lepton” definitions. Veto electrons and muons are required to have $p_T > 10$ GeV and a loose isolation requirement of $I_{\text{rel}}^{e/\mu} < 4.0$. These thresholds are intentionally looser than the nominal selections so that almost any additional prompt lepton in an event is caught by the veto. In the $e\tau_h$ final state, exactly one selected electron is allowed and no extra veto electrons or veto muons may be present. In the $\mu\tau_h$ final state, exactly one selected muon is allowed and no extra veto muons or veto electrons may be

present. In the $\tau_h\tau_h$ final state, no veto electrons nor veto muons are allowed. This strategy prevents overlap between channels and suppresses backgrounds with extra prompt leptons (e.g. dileptonic $t\bar{t}$ or $Z \rightarrow \ell\ell$) while minimally affecting the signal acceptance.

5.5.2. $b\bar{b}$ Pair Selection

Compared to ref. [5] mentioned in Section 5.1, which is using a dedicated neural network for assigning candidate jets to the Higgs boson decay, called HH-bTag, the $b\bar{b}$ pair selection in this analysis is done in a simpler way.

Jets from b quarks are identified by the DEEPJET algorithm when fulfilling the medium working point (see Section 4.5). The selection is limited to the central tracker region with $p_T^{\text{jet}} > 20$ GeV and $|\eta^{\text{jet}}| < 2.5$.

The $b\bar{b}$ pair is constructed in two cases: (i) If at least two jets pass the Medium working point, the leading b-tagged jet is taken as b_1 and b_2 is the first other b-tagged jet in p_T order that satisfies $\Delta R(b_1, b_2) > 0.4$ (if the next-to-leading jet fails the separation, the next one is tried, and so on). (ii) If exactly one b-tagged jet is found, this jet is b_1 , and b_2 is chosen among the remaining jets with $p_T^{\text{jet}} > 30$ GeV, $|\eta^{\text{jet}}| < 2.5$, as the one with the largest DEEPJET discriminant value, again requiring $\Delta R(b_1, b_2) > 0.4$. If no such partner exists, no b-jet pair is defined and the event is rejected.

5.6. $\tau^+\tau^-$ Pair Mass Reconstruction with FastMTT

Reconstructing the $\tau^+\tau^-$ mass is compromised by the fact that each τ decay produces at least one neutrino that is not detected. The visible decay products (e, μ and τ_h) therefore carry only part of the original τ momenta. The missing transverse momentum, \vec{p}_T^{miss} , provides partial information about the neutrinos in the transverse plane, but by itself does not fully determine the event. FASTMTT [68] is a faster variant of the SVfit algorithm [69]: it evaluates a likelihood that combines the measured visible objects, \vec{p}_T^{miss} and the known \vec{p}_T^{miss} resolution to estimate the full τ four-vectors and hence $m_{\tau\tau}$.

The fit uses three inputs: (i) the four-vectors p_1^{vis} and p_2^{vis} of the two visible τ decay products, together with the encoded decay mode for τ_h ; (ii) the measured missing transverse momentum vector \vec{p}_T^{miss} ; and (iii) the 2×2 covariance matrix \mathbf{V} that encodes the resolution of \vec{p}_T^{miss} in Cartesian coordinates. A τ -mass constraint, $m_\tau = 1.777$ GeV [32], and simple models for the invisible part of each decay relate the unknown neutrino momenta to a small set of fit parameters. In particular, two visible energy fractions z_1 and z_2 are introduced so that

$$p(\tau_i) = \frac{1}{z_i} p_i^{\text{vis}}, \quad 0 < z_i < 1 \quad (i = 1, 2). \quad (5.6)$$

The strictly collinear approximation is imposed ($\theta_{\text{GJ}} = 0$ [70]): the invisible system is taken to be collinear with the visible decay products, and no additional angular parameters are fitted.

With this, the neutrino transverse momenta \vec{p}_T^{vv} can be written as

$$\vec{p}_T^{\text{vv}} = \vec{p}_{T,1}^{\text{inv}} + \vec{p}_{T,2}^{\text{inv}} \quad \text{with} \quad p_i^{\text{inv}} = p(\tau_i) - p_i^{\text{vis}}, \quad (5.7)$$

so that \vec{p}_T^{vv} represents the sum of the unobserved neutrinos from both τ decays: a single ν_τ for τ_h and two neutrinos treated as one effective invisible system for leptonic τ decays.

The FASTMTT algorithm defines a likelihood $\mathcal{L}(z_1, z_2)$, that factorizes into a term describing the compatibility of the fitted neutrino system with the measured \vec{p}_T^{miss} and a term encoding the τ -decay kinematics. The former is modeled as a bivariate Gaussian in $\vec{p}_T^{vv} - \vec{p}_T^{\text{miss}}$ with covariance matrix \mathbf{V} , while the latter, denoted $I(z_1, z_2)$, represents the decay mode dependent phase-space weight for the two τ legs. This can be written as

$$\mathcal{L}(z_1, z_2) \propto \exp\left[-\frac{1}{2} \left(\vec{p}_T^{vv} - \vec{p}_T^{\text{miss}}\right)^\top \mathbf{V}^{-1} \left(\vec{p}_T^{vv} - \vec{p}_T^{\text{miss}}\right)\right] I(z_1, z_2). \quad (5.8)$$

The fit scans over the physically allowed domain in (z_1, z_2) , defined by $0 < z_i < 1$ and additional decay-mode-dependent constraints, and selects the point that maximizes $\mathcal{L}(z_1, z_2)$. From the corresponding best-fit parameters, the algorithm reconstructs the τ four-momenta $p^\mu(\tau_1)$ and $p^\mu(\tau_2)$, and returns these together with the reconstructed di- τ mass

$$m_{\tau\tau}^{\text{FastMTT}} = \sqrt{(p(\tau_1) + p(\tau_2))^2}. \quad (5.9)$$

Compared to simple approximations such as the collinear approximation, FASTMTT achieves a narrower and less biased $m_{\tau\tau}$ distribution by using the full \vec{p}_T^{miss} information together with constraints on the τ -decay kinematics. The resolution improves with the boost of the $\tau^+ \tau^-$ system, i.e. for larger $p_T^{\tau^+ \tau^-}$. In back-to-back configurations, where the two τ leptons are nearly opposite in the transverse plane and the system has low p_T , the collinearity assumptions are weaker and the reconstructed mass distribution becomes broader. In this regime the \vec{p}_T^{miss} covariance \mathbf{V} is essential to keep the fit stable.

5.7. Control shapes

Figures 5.6 ($\tau_h \tau_h$) and 5.7 ($\ell \tau_h$) show two key control distributions for each final state: on the left side the reconstructed di- τ mass $m_{\tau\tau}^{\text{FastMTT}}$ and on the right the transverse momentum of the leading b-jet, $p_T(b_1)$. Backgrounds are stacked in the categories DY $Z/\gamma^* \rightarrow \ell\ell$, diboson, $t\bar{t}$, single t , QCD multijet production, and minor sources. The grey band shows the statistical uncertainty of the summed background. Data are shown as points with error bars. The $HH \rightarrow b\bar{b} \tau^+ \tau^-$ signal is overlaid and scaled by a factor of $5000\times$ in the $\ell \tau_h$ final states and $1000\times$ in the $\tau_h \tau_h$ final state to make the expected shape visible. Ratio panels display the data divided by the background expectation with the corresponding statistical uncertainties.

The $p_T(b_1)$ spectrum peak at moderate p_T and then fall smoothly, as expected for resolved $b\bar{b}$ production. As defined in Section 5.5.2, only jets with $p_T > 20$ GeV enter the b-pair selection. Consequently the $p_T(b_1)$ distributions start at 20 GeV. The agreement between data and the expectation is generally good over most of the range. In the $e\tau_h$ final state a modest overestimation of the expectation by about 10–20% is observed, mainly attributable to the $t\bar{t}$ component. Differences at very high p_T are limited by statistics. In the $\ell \tau_h$ final states the dominant background is $t\bar{t}$, whereas in $\tau_h \tau_h$ other processes (QCD

multijet, $Z/\gamma^* \rightarrow \tau^+\tau^-$) are relatively more important. The overall event yield is largest in $\mu\tau_h$, followed by $e\tau_h$, with the fewest events in $\tau_h\tau_h$.

For $m_{\tau\tau}^{\text{FastMTT}}$, a broad structure from $Z/\gamma^* \rightarrow \tau^+\tau^-$ appears near the Z mass and is well reproduced, which validates both the τ reconstruction and the \vec{p}_T^{miss} modeling. In the $e\tau_h$ final state the total MC prediction is about 10–20% higher than the data, which is again due to a small overestimation of the $t\bar{t}$ background. The $\text{HH} \rightarrow b\bar{b}\tau^+\tau^-$ signal peaks around ~ 115 GeV rather than exactly at 125 GeV. This small downward shift is expected: the $\tau^+\tau^-$ mass is reconstructed with finite resolution and a likelihood regularization (see Section 5.6). In the FastMTT algorithm a strictly collinear approximation is imposed for each τ decay, so only the transverse missing momentum is constrained by the fit while the longitudinal neutrino components remain underconstrained. Detector effects (MET scale/resolution and τ_h energy scale) further smear the mass.

These figures are a small subset of the control plots used in the analysis. Other relevant distributions are provided in the Appendix A.1.

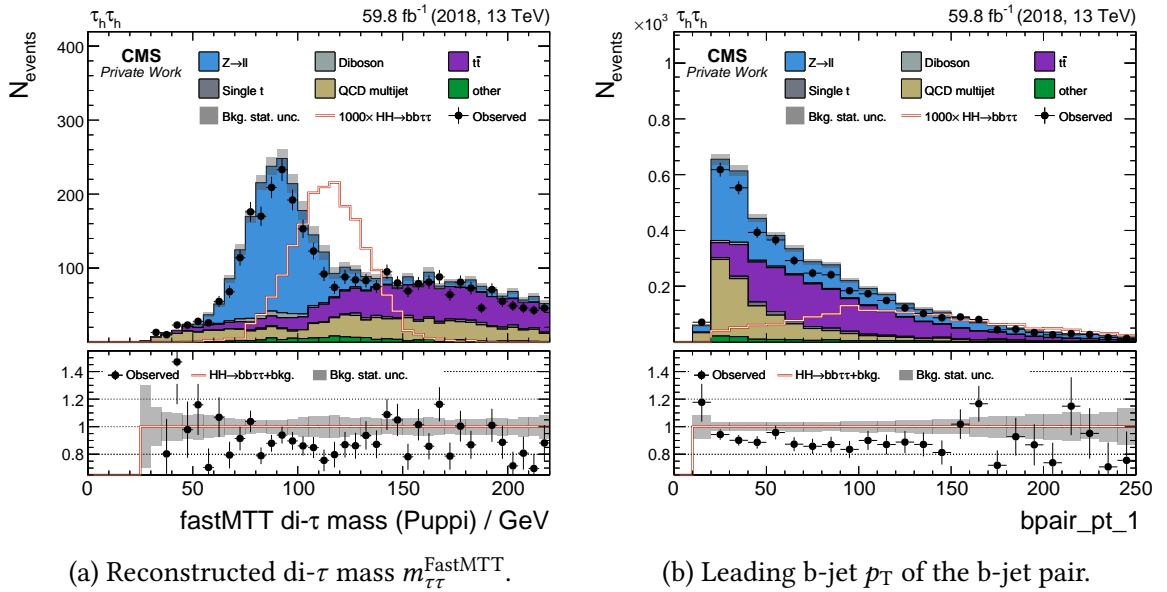


Figure 5.6.: $\tau_h\tau_h$ **final state**: Data–background comparisons. The stacked histograms represent the expected contributions from DY Z/γ^* , diboson, $t\bar{t}$, single t, and other minor processes, while the QCD multijet production contribution is taken from data control regions. The observed data are shown as points with error bars. The $\text{HH} \rightarrow b\bar{b}\tau^+\tau^-$ signal is overlaid and scaled for visibility. The lower panel in each plot shows the data-to-background ratio, with the shaded band indicating the statistical uncertainty of the total background prediction.

5. Search for $HH \rightarrow b\bar{b} \tau^+ \tau^-$

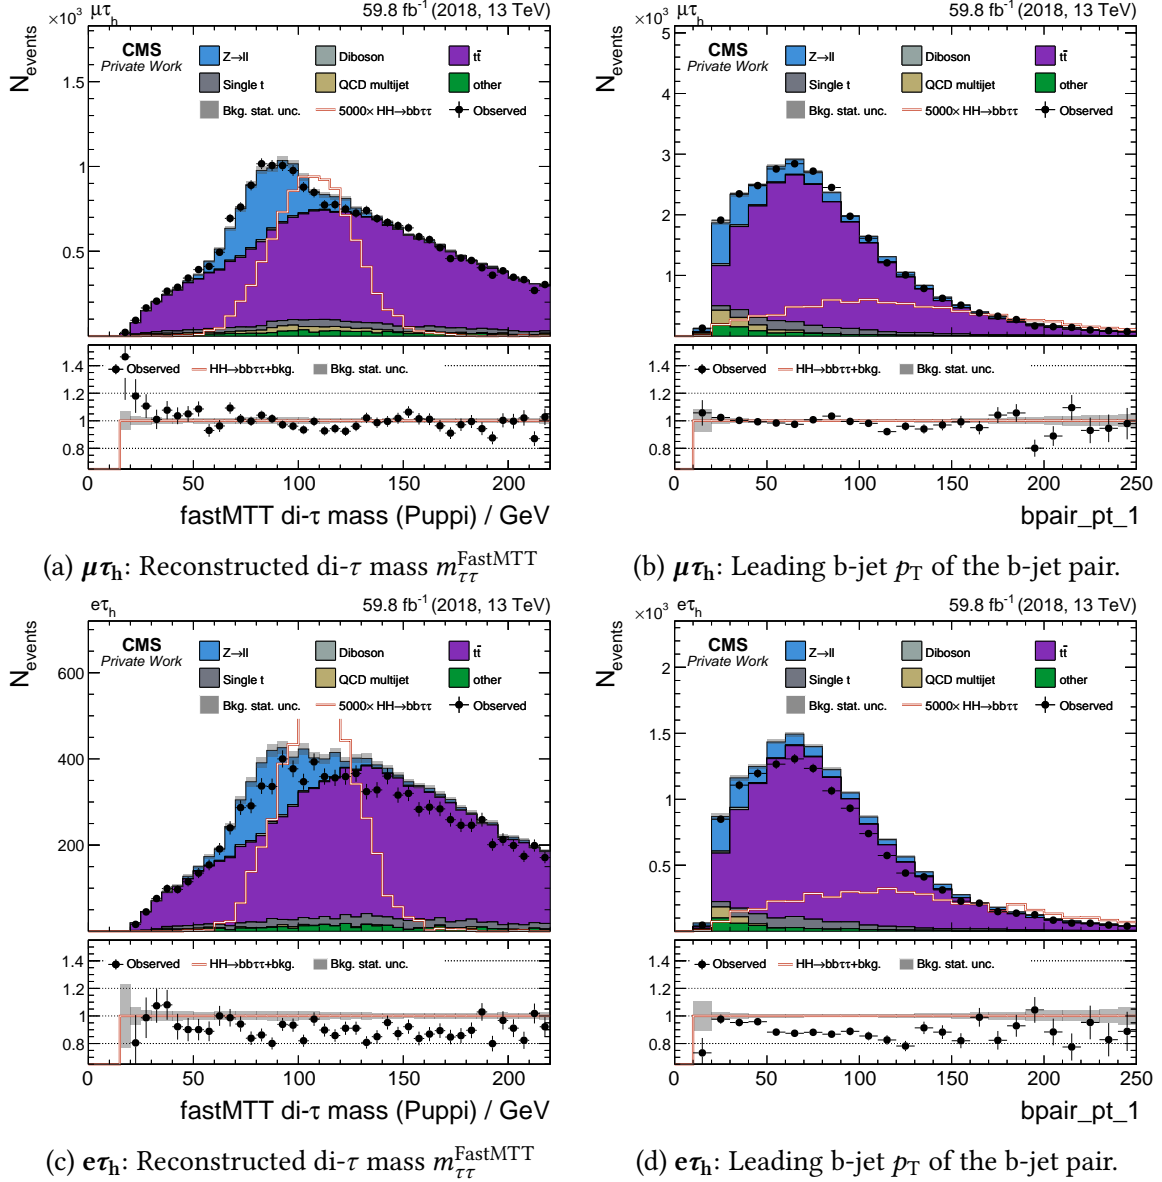


Figure 5.7.: $\ell\tau_h$ final states: Data–background comparisons. The stacked histograms represent the expected contributions from DY Z/γ^* , diboson, $t\bar{t}$, single t, and other minor processes, while the QCD multijet production contribution is taken from data control regions. The observed data are shown as points with error bars. The $HH \rightarrow b\bar{b} \tau^+ \tau^-$ signal is overlaid and scaled for visibility. The lower panel in each plot shows the data-to-background ratio, with the shaded band indicating the statistical uncertainty of the total background prediction.

6. Analysis

6.1. Event Classification with Neural Networks

6.1.1. Architecture

For event classification a compact fully connected network is chosen. The network operates on 20 high-level kinematic inputs which are shown in Table 6.1. In the control plots those observables showed the best signal-background separation. All inputs to the NN are standardized to zero mean and unit variance using the training sample and the same μ, σ are applied to validation and test sets. The architecture is illustrated in Figure 6.1. It has two hidden layers with 50 nodes each and rectified-linear (ReLU) activation functions, together with batch normalization and dropout for regularization purposes.

The six output units correspond to the mutually exclusive event categories. The softmax layer returns class probabilities

$$p(y = k | \mathbf{x}) = \frac{\exp(z_k)}{\sum_{j=1}^6 \exp(z_j)}, \quad k = 1, \dots, 6, \quad (6.1)$$

where z are the pre-activation logits produced by the final linear layer.

6.1.2. Training

The training of the NNs is performed separately for each final state ($e\tau_h, \mu\tau_h, \tau_h\tau_h$) and for two statistically independent folds (`fold_0`, `fold_1`), yielding in total six networks with identical architectures and preprocessings. For a given fold, the events are split into a training and a validation subset of equal size within the fold (50%/50%), while the opposite fold is used as test sample. The loss function is a weighted categorical cross-entropy,

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^6 w_k y_{ik} \log p(y = k | \mathbf{x}_i), \quad (6.2)$$

where N is the number of training examples, k runs over the six classes, w_k denotes the class weight for class k , y_{ik} is the target label, and $p(y = k | \mathbf{x}_i)$ is the predicted probability for class k given the input features \mathbf{x}_i . The loss is optimized with the Adam algorithm [71], using L2 regularization with a weight decay coefficient of 10^{-3} to control overfitting. Batch normalization layers stabilize the training and accelerate convergence. A dropout probability of 0.2 after the first hidden layer adds further regularization. Training makes use of all training events in the fold, so a single parameter update per epoch is computed from

Table 6.1.: Kinematic and event-level observables used in the analysis.

Variable	Description
m_{inv}^{jj}	Invariant mass of the two hardest jets; expected to peak near m_{H} .
$m_{\text{inv}}^{\text{b}\bar{\text{b}}}$	Invariant mass of the two b -tagged jets; highly correlated with m_{inv}^{jj} .
$m_{\tau\tau}^{\text{FastMTT}}$	Reconstructed $\tau^+\tau^-$ mass using FastMTT (Sec. 5.6) to account for neutrinos in the τ decays.
$m_{\tau\tau}^{\text{CollApprox}}$	Reconstructed $\tau^+\tau^-$ mass using the collinear $m_{\tau\tau}$ approximation.
m_{vis}	Invariant mass of the visible decay products of the $\tau^+\tau^-$ -system.
$m_{\text{T}}^{\text{tot}}$	Total transverse mass, which is the sum of the transverse masses of all final state particles.
$\Delta R_{\text{b}\bar{\text{b}}}, \Delta R_{\tau\tau}$	Angular separations capturing the production kinematics.
$p_{\text{T}}^{j,1}, p_{\text{T}}^{j,2}$	Transverse momenta (p_{T}) of the leading and subleading jet.
$p_{\text{T}}^{\text{bpair},1}$	p_{T} of the leading b -tagged jet; highly correlated with $p_{\text{T},1}^j$.
$p_{\text{T}}^{\text{dijet}}$	Combined p_{T} of the two hardest jets.
$p_{\text{T}}^{\text{bpair,dijet}}$	Combined p_{T} of the b -jet pair; highly correlated with $p_{\text{T}}^{\text{dijet}}$.
p_{T}^1	p_{T} of the electron, muon or the p_{T} -leading τ_{h} .
$p_{\text{T}}^{\text{FastMTT}}$	p_{T} of the $\tau^+\tau^-$ system reconstructed with FastMTT (Sec. 5.6) to account for neutrinos in the τ decays.
$p_{\text{T}}^{\text{vis}}$	p_{T} of the visible decay products of the $\tau^+\tau^-$ -system.
n_{btag}	Number of b -tagged jets.
n_{jets}	Number of reconstructed jets.
$\text{btag}^1, \text{btag}^2$	DeepJet discriminants of the selected candidate jet to form the $\text{b}\bar{\text{b}}$ pair.

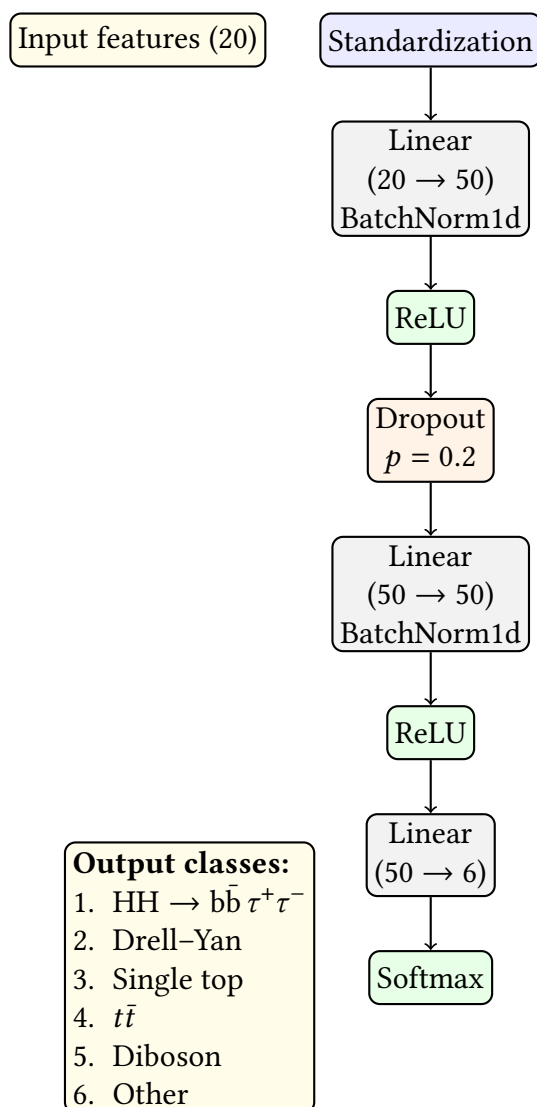


Figure 6.1.: Top-down schematic of the neural-network architecture used for event classification. The model takes 20 standardized inputs, applies two hidden blocks (Linear-BatchNorm-ReLU with a Dropout layer after the first block), and outputs six class probabilities via a Softmax layer.

the entire training set. After each parameter update the loss is also evaluated on the validation data set. Early stopping based on the validation loss is used to select the final model.

Each simulated event carries a physical weight w_i^{phys} so that the expected yield in class k is $N_k^{\text{phys}} = \sum_{i \in k} w_i^{\text{phys}}$. For example in $\tau_h \tau_h$ final state the signal sample contains ~ 5800 simulated events but corresponds to only two expected physical events for the given \mathcal{L}_{int} . Four class-weighting prescriptions were tested in the definition of \mathcal{L}_{CE} :

- Balance by MC counts: $w_k^{\text{cls}} \propto 1/N_k^{\text{MC}}$ with N_k^{MC} the number of simulated training examples in class k .
- Balance by physical yields: $w_k^{\text{cls}} \propto 1/N_k^{\text{phys}}$.
- 50/50 balance by MC counts: target priors of 50% signal and 10% per background based on the number of simulated training samples.
- 50/50 balance by physical yields: the same target priors as given in the previous item based on the physical yields instead of the number of simulated training samples.

The first scheme, balancing by MC counts, yielded the best-performing models across final states and folds. Training metrics balancing weights by MC counts are summarized in Appendix A.3. For each decay channel the plots show the categorical cross-entropy for the training and validation sets as a function of epoch and purity, efficiency and F_1 score on the validation data, with

$$\text{Purity} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Efficiency} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad F_1 = \frac{2 \text{Purity} \cdot \text{Efficiency}}{\text{Purity} + \text{Efficiency}}. \quad (6.3)$$

where, TP (true positives) denotes correctly identified signal events, FP (false positives) denotes background events incorrectly classified as signal, and FN (false negatives) denotes signal events that are incorrectly rejected.

Convergent training is reflected by a steadily decreasing validation loss and a saturating F_1 score. Figure 6.2 show this behavior for the $\tau_h \tau_h$ final state. The training metrics for the $\ell \tau_h$ final states are shown in the Appendix A.3. At very early epochs the training loss can exceed the validation loss because dropout is applied only during training, injecting noise into the forward pass. As optimization proceeds, the validation loss typically reaches a minimum and then flattens, while the training loss continues to decrease – an indication that the network is starting to rely on properties which are unique for the training sample (overfitting). Early stopping is therefore configured to find the minimum of the validation loss.

The trained networks provide per-class probabilities. The signal probability or suitable combinations thereof are used as the final discriminants in the statistical analysis.

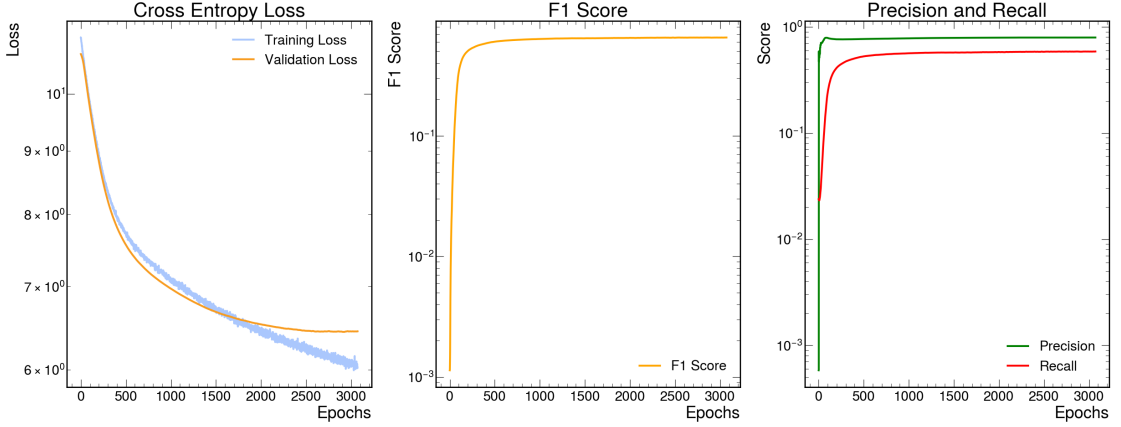


Figure 6.2.: NN training metrics for the $\tau_h \tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.

Table 6.2.: Class composition in the $\tau_h \tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 244,726$ and $\sum w^{\text{phys}} = 3606$.

Class	$N^{\text{MC}} [10^3]$	MC fraction	$\sum w^{\text{phys}}$	Physical fraction
Signal	5.8	2.4%	1.87	0.05%
DY	42.0	17.2%	1700.0	47.1%
Single t	0.6	0.3%	120.0	3.3%
$t\bar{t}$	45.3	18.5%	1587.0	44.0%
Diboson	7.6	3.1%	67.7	1.9%
other	143.4	58.6%	129.1	3.6%

6.1.3. Training Data & Classification Results

The Tables 6.2 for the $\tau_h \tau_h$ final state and in the Appendix A.2 for the $\ell \tau_h$ final states show for each class k the nominal number of simulated events N_k^{MC} and the sum of physical weights $\sum w_k^{\text{phys}}$ that represents the expected physical yield. The physically expected mixture depends on the $\tau^+ \tau^-$ final state. In the $\tau_h \tau_h$ final state, DY and $t\bar{t}$ contribute comparably to the expected with 47% and 44%, while the signal fraction is 0.05%. In the $\ell \tau_h$ final states the expected yield is strongly $t\bar{t}$ dominated ($\sim 80\%$), with DY at the 10–12% level and signal at the 0.007% level. These priors are the baseline against which the classifier has to operate. The strong imbalance makes the signal–background separation particularly challenging.

6. Analysis

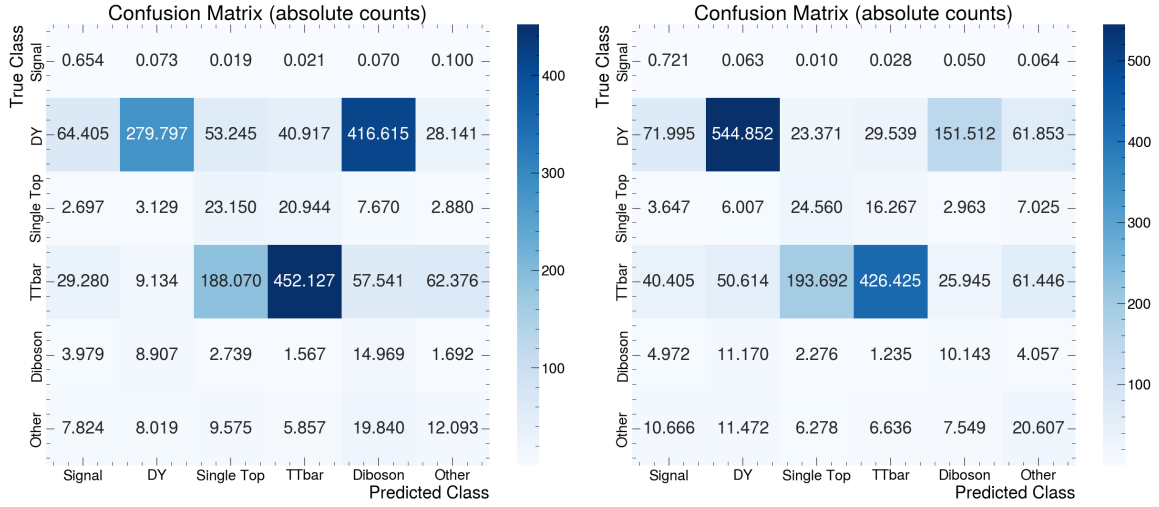


Figure 6.3.: $\tau_h\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.

The trained NNs are applied to the test dataset. For each event, the corresponding confusion-matrix entry N_{ij} , where i is the true class and j the predicted class, is updated by adding the event's physical weight w^{phys} . The matrix is optionally normalized by rows (efficiency) or columns (purity) for interpretation. Efficiency is used because it measures the efficiency for each true class and quantifies how many events of class i are correctly retained, which is crucial for signal acceptance and for understanding losses to specific background categories. Purity is used because it measures the contamination of each predicted class and indicates how clean a category can be selected.

Figure 6.3 shows the confusion matrices with physical event counts for the $\tau_h\tau_h$ final state. In Appendix A.4 the remaining confusion matrices for each final state with absolute counts and normalization by rows and columns are shown. Each view is provided for the trainings with balancing by MC events (left) and by physical events (right).

In the $\tau_h\tau_h$ final state, predominantly DY (59%) and, to a lesser extent, $t\bar{t}$ (27%) events leak into the signal class. With physical yield balancing, the signal efficiency increases, but the purity of the predicted signal class decreases. In particular, the $t\bar{t}$ contamination of the signal column grows. A striking difference between the two trainings is the large DY to diboson event misidentification seen with MC-balanced training, which is strongly reduced under physical yield balancing. This aligns with expectations, as diboson and DY events share low b -tag content, and MC-balancing tends to equalize classes at the expense of prior information. A migration between the $t\bar{t}$ and single- t classes is observed in both trainings, which is expected given their similar top-like event topologies. Compared to the training sample where the signal fraction is only 0.05%, the NN raises the signal fraction within its predicted signal category to 0.6%, an enrichment by a factor of 12, which is useful but still underscores the difficulty of signal-background separation.

In the $\ell\tau_h$ final states, the NN operates in a $t\bar{t}$ -dominated environment. As expected the dominant contamination of the signal arises from the $t\bar{t}$ class ($\sim 80\%$), with DY ($\sim 10\%$) as

the subleading and single t as a minor ($\sim 5\%$) contribution in the signal class. Moving from MC-balanced to physical yield balancing increases the signal efficiency but decreases the purity of the predicted signal column, mirroring the behavior seen in $\tau_h\tau_h$. The migrations between $t\bar{t}$ and single top, as well as between DY and diboson, are also visible and are driven by the similar event topologies of these processes. Like in $\tau_h\tau_h$ the signal within the predicted signal class is enriched by a factor of ~ 12 .

Across final states, the misclassifications of HH events follow the expected class composition: in the $\ell\tau_h$ final states HH events are mostly confused with $t\bar{t}$, while in the $\tau_h\tau_h$ final state they are mostly confused with DY. This is not a failure of the network but a consequence of overlapping kinematics after the selections. It motivates using the full multi-class output in the fit instead of a single binary score, since the non-signal probabilities help constraining the normalizations and shapes of the main backgrounds.

The analysis uses the per-event class score $p_{\text{class}} \in [0, 1]$ as a one-dimensional discriminant in six exclusive categories defined by the maximum score of the multi-class output. Therefore, with six classes there are no entries below $1/6 \approx 0.17$, since a smaller maximum would imply assignment to another class. Figure 6.4 shows the p_{class} distributions in the $\tau_h\tau_h$ final state. Appendix A.5 presents the p_{class} distributions for the $\ell\tau_h$ final states. Data are blinded in the HH category, and the background expectations are shown with statistical uncertainties only. In the HH category the background spectra fall steeply toward $p_{\text{HH}} \rightarrow 1$, leaving a sparse but relatively clean tail that concentrates sensitivity. In this tail the purity is rising up to a few %. Compared to the underlying dataset this is a signal enrichment of more than $\times 100$. In the DY category the spectrum peaks at intermediate scores with little contamination in the extreme tail. The $t\bar{t}$ category exhibits a broad spectrum in the $\ell\tau_h$ final states. The single t category traces a similar but lower shape and is naturally constrained with $t\bar{t}$.

6.2. Systematic Uncertainties

All systematic uncertainties in the nominal fit are implemented as log-normal ($\ln\mathcal{N}$) rate nuisances. An $\ln\mathcal{N}$ constraint rescales the nominal yield of the affected process by a factor k for a $+1\sigma$ shift of the given nuisance parameter. Uncertainties can be specified as symmetric or asymmetric. The yield of each process i can be calculated as

$$N_i = \sigma_i \times \mathcal{L}_{\text{int}} \times A_i \times \varepsilon_i \quad (6.4)$$

with the process cross section σ_i , integrated Luminosity \mathcal{L}_{int} , acceptance A_i and efficiency ε_i . To each of this factors an uncertainty is assigned.

Object efficiencies ε . Object efficiencies quantify the probability for a physics object (e, μ or τ_h) being triggered, reconstructed, isolated and identified. These probabilities as determined from simulation differ slightly from the values as obtained from data. Therefore per-object scale factors (SFs) are applied to align the simulation to the data. The corresponding uncertainties cover the limited precision of the data-driven measurements and kinematic mismodelling. Lepton efficiencies only affect the $\ell\tau_h$ final states. Efficiencies for the τ_h identification are factorized into against leptons and against jets components,

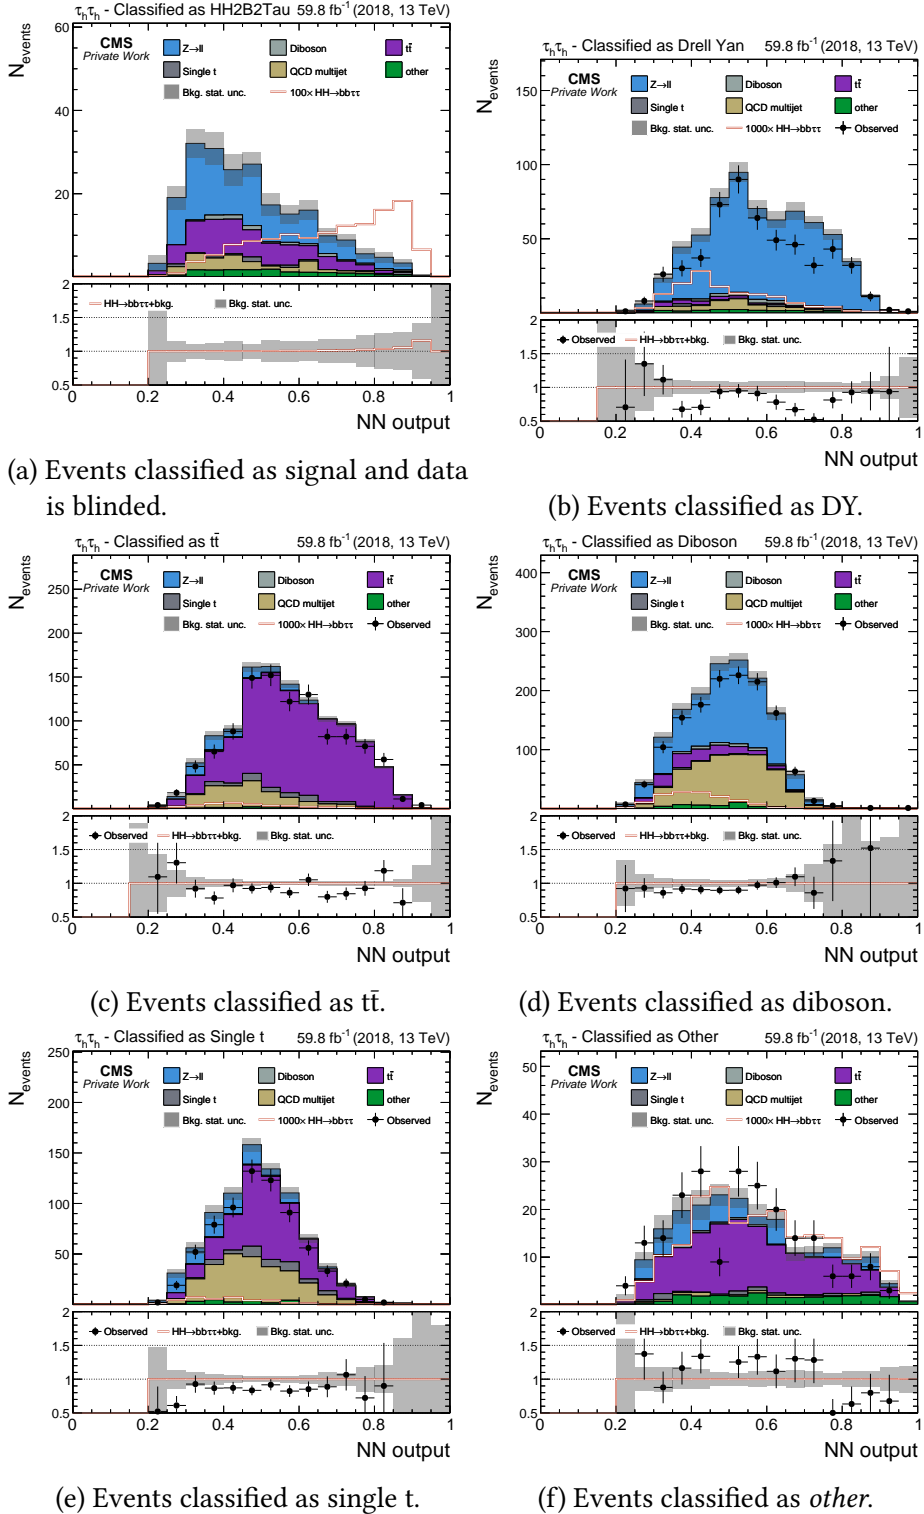


Figure 6.4.: $\tau_h \tau_h$ final state: NN score of events assigned to each output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

reflecting the structure of the DEEPTAU algorithm. Each component comes with its own uncertainty to account for different control samples and uncertainties. In the $\tau_h \tau_h$ final state, the total rate effect is applied per object, so the factor is squared.

Integrated luminosity \mathcal{L}_{int} . The simulated yields (N_i) scale with \mathcal{L}_{int} as shown in eq. 6.4. So any imprecision in \mathcal{L}_{int} propagates directly to the normalization of all MC processes that are normalized by cross sections. The luminosity is measured and calibrated e.g. via van der Meer scans [72, 73]. Residual uncertainties arise from the absolute calibration, time stability, detector modeling and operational conditions. In the fit this is implemented as a single $\ln N$ nuisance. It affects only fully simulated processes. Purely data-driven backgrounds like QCD multijet production are not scaled by \mathcal{L}_{int} and thus are left unaffected by this nuisance parameter.

Background normalizations σ_i . For simulated backgrounds, the nominal rate is set by the inclusive cross section σ_i , and theoretical and modelling uncertainties in σ_i are propagated as rate nuisance parameters. The dominant sources are missing higher-order corrections, the choice of parton distribution functions (PDFs), and the value of the strong coupling constant α_s . For some processes, a single $\ln N$ uncertainty summarizes the total cross-section uncertainty, while for others it is decomposed into orthogonal components: `QCDscale`, which varies the renormalization and factorization scales to cover missing higher orders; `pdf`, which propagates PDF uncertainties using the corresponding error sets; and `alphaS`, which varies the α_s used in the calculation or in the PDF determination.

Acceptance A . In this analysis no per-event corrections are applied to the acceptance A , because A is defined at generator level (fiducial and kinematic coverage before detector effects) and cannot be measured directly in data. Uncertainties on A are instead covered by theory and modelling variations, as detailed in the previous paragraph. Detector related effects are calibrated with data-driven object scale factors and belong to the efficiency ϵ , not to A .

All systematic uncertainties used for this analysis are listed in Table 6.3. This choice keeps the fit simple while reflecting common sources of uncertainty. The luminosity and all uncertainties related to the signal shifts only the overall normalization by a few percent and are subdominant to statistical uncertainties in the high-score tail of the HH event category.

Table 6.3.: Summary of $\ln N$ rate uncertainties used in the nominal fit. Correlations are across final states within the same era unless noted. [†] The uncertainty in the vsJet working point of the τ_h identification discriminant is typically a shape effect; here a conservative $\ln N$ proxy is used.

Nuisance	Processes	Channels	Size (up/down)
eff_e	all	et	1.02
eff_m	all	mt	1.02
eff_t_vsLep	all	et,mt	1.01
eff_t_vsLep	all	tt	1.02
eff_t_vsJet [†]	all	et,mt	1.02
eff_t_vsJet [†]	all	tt	1.04
lumi_13TeV	all	all	1.025
htt_vvXsec	diboson	all	1.05
htt_stXsec	single t	all	1.05
htt_tjXsec	t \bar{t}	all	1.06
htt_wjXsec	W	all	1.04
htt_zjXsec	DY	all	1.02
QCDNorm	QCD	et, mt	1.108
QCDNorm	QCD	tt	1.216
QCDscale_ttH, pdf_ttH, alphaS_ttH	ttH	all	(1.058/0.908), 1.030, 1.020
QCDscale_ggH, pdf_ggH, alphaS_ggH	ggH	all	1.039, 1.019, 1.026
QCDscale_qqH, pdf_qqH, alphaS_qqH	qqH	all	(1.004/0.997), 1.021, 1.005
QCDscale_VH, pdf_VH, alphaS_VH	VH(WH+ZH)	all	(1.038/0.970), 1.017, 1.009
QCDscale_HH	signal	all	(1.022/0.950)
PDF_alphas_HH	signal	all	1.030
mtop_HH	signal	all	1.026
BR_h_bb	signal	all	(1.0125/0.9873)
BR_h_tautau	signal	all	1.0165

6.3. Results

The limits are obtained with a profile-likelihood fit using the COMBINE package described in Ref [74] using the asymptotic CL_s prescription with the signal strength $\mu = \sigma(pp \rightarrow HH)/\sigma_{SM}$ as parameter of interest. The templates are first cleaned by removing processes with zero expected yield and are then rebinned to avoid bins with very low statistics and large relative fluctuations. The effect of the finite sample sizes used to build the input distributions is incorporated as Poisson-like uncertainties for each individual bin, following the Barlow–Beeston lite prescription [75]. Expected 95 % CL upper limits are derived from a profile-likelihood fit to the Asimov dataset constructed under the background-only hypothesis.




Figure 6.5.: Expected 95 % CL upper limits on μ for each final state and their combination, with the median expectation and the 68 % and 95 % uncertainty bands shown. **Left:** statistical-only fit; **Right:** full fit with all systematic uncertainties included.

Figure 6.5 shows the median expected limits with their central 68 % and 95 % confidence intervals (CIs), indicated by the dark and light green bands. Including systematics slightly degrades the median sensitivity and broadens the CL, as the fit profiles luminosity, object efficiencies, background normalizations, and theory inputs in addition to counting the statistics. Numerically, the combined median limit increases by 2 % from $\mu_{95}^{\text{exp}} \approx 10.7$ (stat-only) to ≈ 10.9 (full).

The $\tau_h \tau_h$ final state is the most sensitive, followed by $\mu \tau_h$, while $e \tau_h$ contributes the least. This ordering reflects both the classifier performance and the control power in the fit: $\tau_h \tau_h$ exhibits the cleanest high-score tail where DY and $t\bar{t}$ are concurrently constrained and $\mu \tau_h$ benefits from large yields and strong μ triggering, reconstruction and identification. The combined fit constrains shared nuisance parameters across final states and achieves the best expected limit.

6.4. Extrapolation to higher luminosities




Figure 6.6.: Expected 95 % CL upper limits on μ for each final state and their combination with luminosity scaling. **Left:** extrapolation to the full Run 2 $\mathcal{L}_{\text{int}} = 138 \text{ fb}^{-1}$. **Right:** projection to a combined Run 2 + Run 3 expected $\mathcal{L}_{\text{int}} = 500 \text{ fb}^{-1}$. Rates are scaled with \mathcal{L}_{int} while shapes and systematic models are kept unchanged.

To project the expected sensitivity to larger datasets, all rates are scaled by the ratio of integrated luminosities,

$$N'_b = N_b \times \frac{\mathcal{L}'_{\text{int}}}{\mathcal{L}_{\text{int}}},$$

while the shapes and the systematic model are kept unchanged. Auto-rebinning and bin-by-bin MC statistical uncertainties remain enabled so that finite-template statistics are correctly propagated after the rescaling. In particular, scaling the luminosity by a given factor does not correspond to having that many more simulated events, and the MC statistical uncertainties therefore stays the same in relative terms. The expected limits are evaluated on the Asimov dataset with the same COMBINE configuration as in Sec. 6.3.

Figure 6.6 shows the 95 % CL upper limits for the full Run 2 dataset and the Run 2+Run 3 projections, obtained using the same systematic uncertainty model as in the results shown in the right panel of Figure 6.5. Relative to the 2018-only result, the combined full Run 2 median improves from 10.9 to 7.2 by a factor $10.9/7.2 = 1.51$. This is essentially the naive statistical scaling of $\sqrt{138/59.8} \approx 1.52$, indicating that the fit remains dominated by statistical uncertainties. For the Run 2 + Run 3 projection to the expected $\mathcal{L}_{\text{int}} = 500 \text{ fb}^{-1}$, the naive statistical gain would be $\sqrt{500/59.8} \approx 2.89$, whereas the profiled fit reaches a factor $10.9/4.7 = 2.32$. The weaker scaling reflects the systematic uncertainties with increasing integrated luminosity. In practice, further improvements in background modelling and object calibrations are required to approach the statistical limit at very high luminosity.

7. Conclusion

This thesis presented a search for nonresonant Higgs–boson pair production in the $HH \rightarrow b\bar{b} \tau^+ \tau^-$ final state using the CMS 2018 proton–proton collision dataset at $\sqrt{s} = 13$ TeV, corresponding to an integrated luminosity of 59.8 fb^{-1} . The analysis targets the resolved $\tau^+ \tau^-$ and $b\bar{b}$ topology and combines the $\tau_h \tau_h$, $\mu \tau_h$ and $e \tau_h$ final states of the $\tau^+ \tau^-$ system. Object reconstruction relies on the Particle Flow, DEEPTAU and the DEEPTAU v2.5 algorithms. A multi-class classification neural network provides the final event discrimination and category definition. Expected limits are derived based on the Asimov dataset using a profile-likelihood fit.

Dedicated control studies validate the key ingredients of the selection. In particular, the reconstructed di- τ mass $m_{\tau\tau}^{\text{FastMTT}}$ reproduces the $Z/\gamma^* \rightarrow \tau^+ \tau^-$ peak, providing a cross-check of τ reconstruction and \vec{p}_T^{miss} , while jet kinematics agree with expectations for the resolved $b\bar{b}$ system within statistical precision.

Using the multi-class classification NN score as a one-dimensional discriminant in six exclusive event categories, the combined expected 95 % CL upper limit on the signal strength,

$$\mu \equiv \frac{\sigma(\text{pp} \rightarrow \text{HH})}{\sigma_{\text{SM}}}, \quad (7.1)$$

is $\mu_{95}^{\text{exp}} = 10.9$ for 2018-only, with a small ($\sim 2\%$) degradation relative fit including only statistical uncertainties $\mu_{95}^{\text{exp}} = 10.7$. Among all final states, $\tau_h \tau_h$ is most sensitive, followed by the final states $\mu \tau_h$ and $e \tau_h$.

To gauge the benefit of additional data, the analysis was extrapolated by scaling event yields with luminosity while retaining shapes and the nominal systematic model. The projected median limits improve to $\mu_{95}^{\text{exp}} = 7.2$ for the full Run 2 luminosity (138 fb^{-1}) and to $\mu_{95}^{\text{exp}} = 4.7$ for a Run 2+Run 3 scenario (500 fb^{-1}). The Run 2 projection follows near-statistical scaling, indicating a statistics-limited regime, while at 500 fb^{-1} profiled systematics begin to curb the naive stochastic gain.

Overall, the study establishes a complete 2018 analysis chain searching for $HH \rightarrow b\bar{b} \tau^+ \tau^-$ in the resolved regime, covering object identification, NN-based categorization, control validations and a full statistical interpretation. It quantifies the achievable expected sensitivity with current selections. The results underscore the importance of larger datasets and continued improvements in modelling and object performance for approaching the sensitivity to observe or exclude HH production with the expected SM rate at the 95

7. Conclusion

Several targeted refinements can enhance future sensitivity. First, incorporating the remaining Run 2 years together with Run 3 data and adding the VBF production mode should provide immediate statistical gains and complementary kinematics. Second, continued improvements in background modelling, especially for the $t\bar{t}$ and DY processes in the high-score tails of the discriminating distributions, and in object calibrations will help preserve statistical scaling at higher luminosities. Finally, optimizing the analysis shape binning used in the fit may help improving the signal to background discrimination while controlling statistical fluctuations.

Relative to the latest CMS di-Higgs result of $\mu_{95}^{\text{exp}} = 8.2$ for 2018-only [5], this analysis ($\mu_{95}^{\text{exp}} = 10.9$) is less sensitive but competitive. With the targeted refinements and the inclusion of VBF production this gap may be closed. The $b\bar{b}\tau^+\tau^-$ channel remains a central path toward probing the Higgs self-coupling at the LHC.

A. Appendix

A.1. Control Shapes




Figure A.1.: Visible mass of object 1. In the $\ell\tau_h$ channels object 1 is the lepton, so the distribution peaks near zero. In the $\tau_h\tau_h$ channel object 1 is a τ_h .

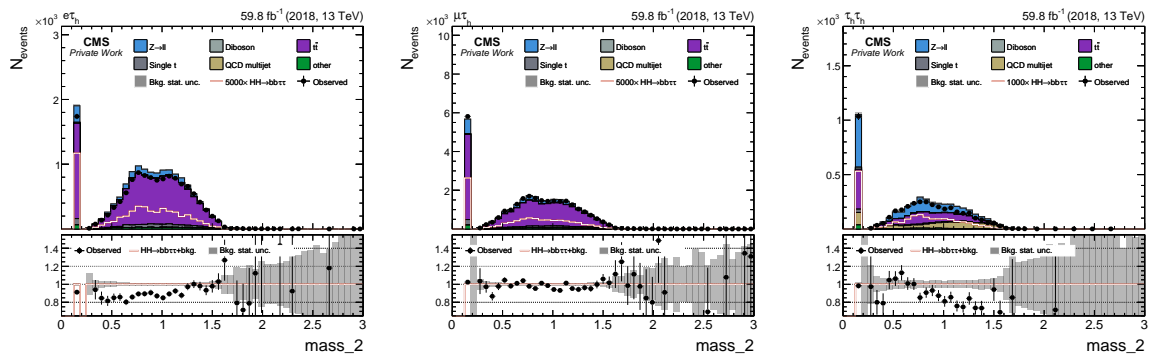


Figure A.2.: Visible mass of object 2, which is a τ_h in all final states.




Figure A.3.: Transverse momentum of object 1: lepton p_T in $\ell\tau_h$ and p_T^τ in $\tau_h\tau_h$.




Figure A.4.: Transverse momentum of object 2, which is a τ_h in all final states.

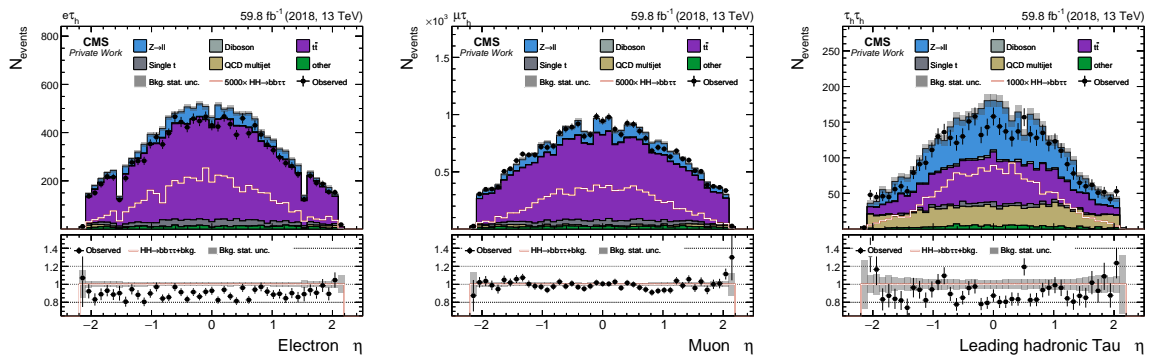


Figure A.5.: Pseudorapidity of object 1: lepton η in $\ell\tau_h$ and η^τ in $\tau_h\tau_h$.

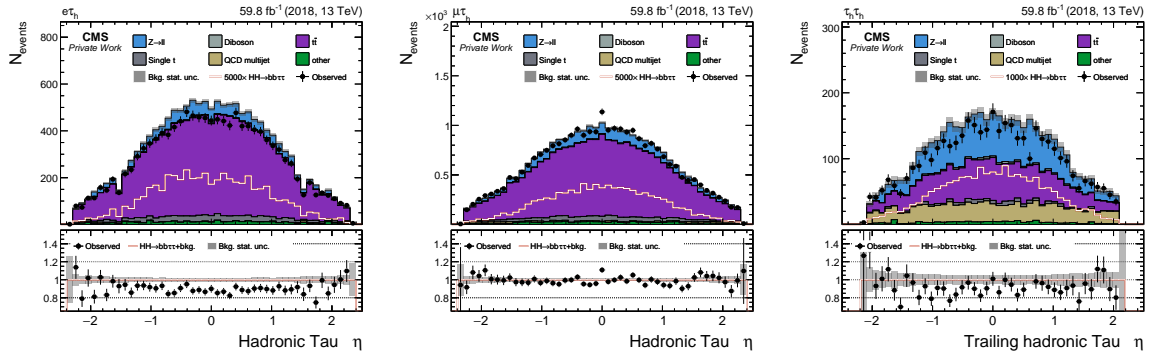


Figure A.6.: Pseudorapidity of object 2, which is a τ_h in all final states.




Figure A.7.: τ_h decay mode of object 1. In the $\ell\tau_h$ final states object 1 is the lepton, so these distributions are empty.




Figure A.8.: τ_h decay mode of object 2.




Figure A.9.: Visible di- τ mass, m_{vis} of the $\tau^+\tau^-$ system.

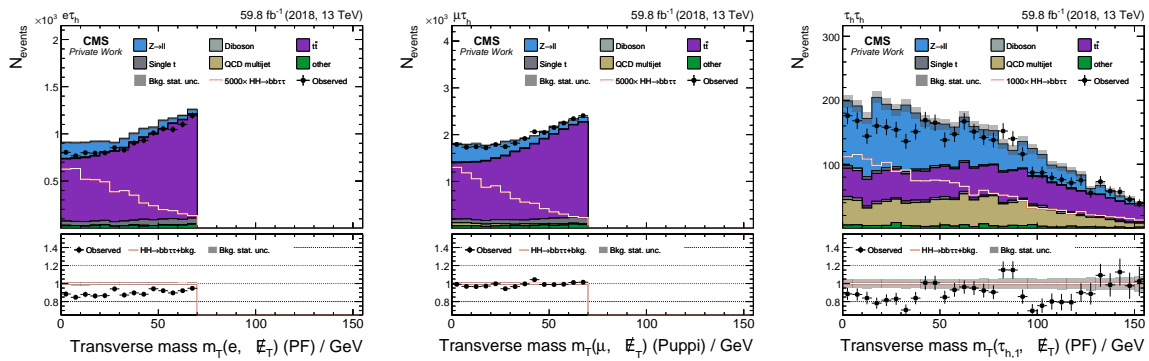


Figure A.10.: Transverse mass m_T of object 1 with \vec{p}_T^{miss} .

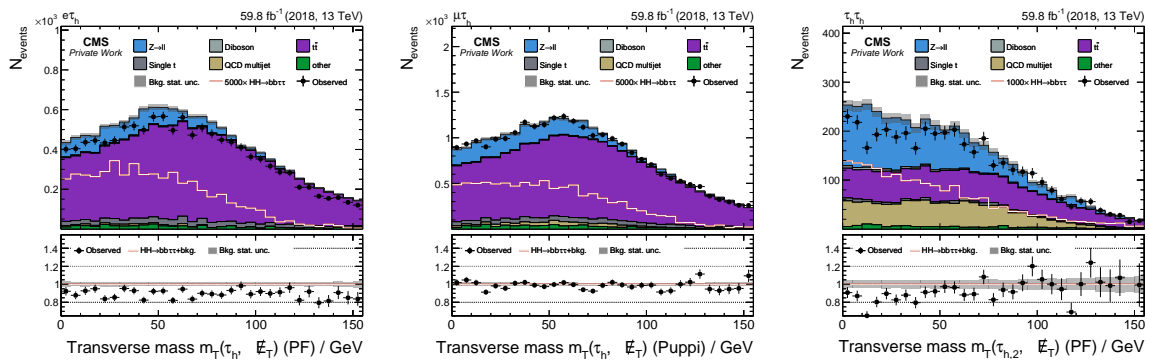


Figure A.11.: Transverse mass m_T of object 2 with \vec{p}_T^{miss} .

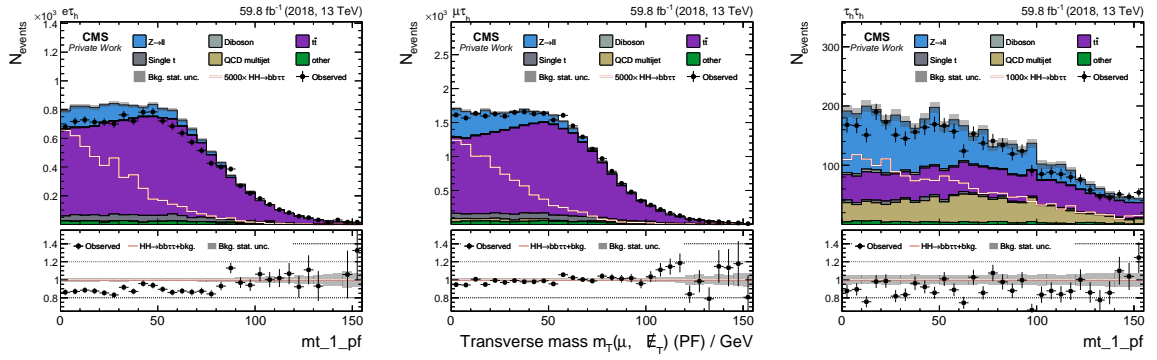


Figure A.12.: PF-based transverse mass m_T of object 1 with \vec{p}_T^{miss} .

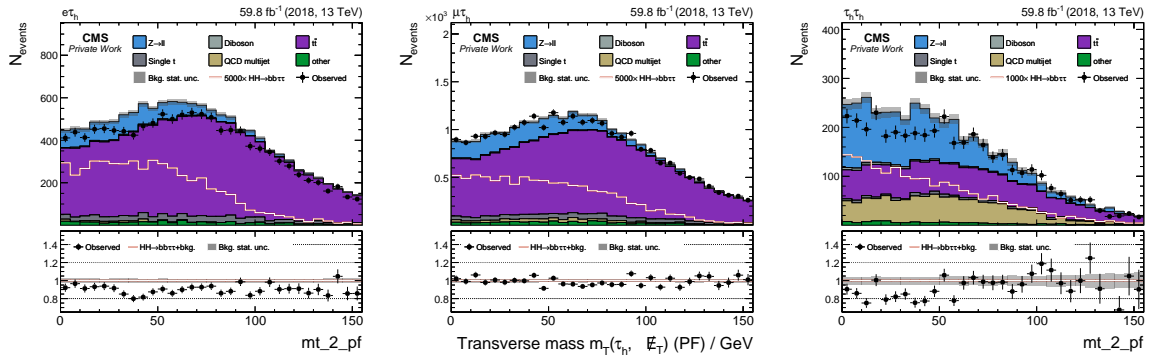


Figure A.13.: PF-based transverse mass m_T of object 2 with \vec{p}_T^{miss} .

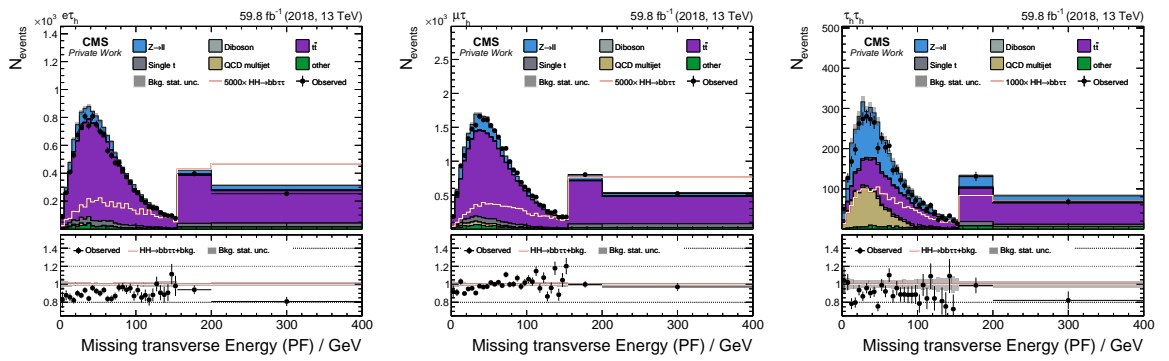


Figure A.14.: PF-based missing transverse momentum, p_T^{miss} .

A. Appendix

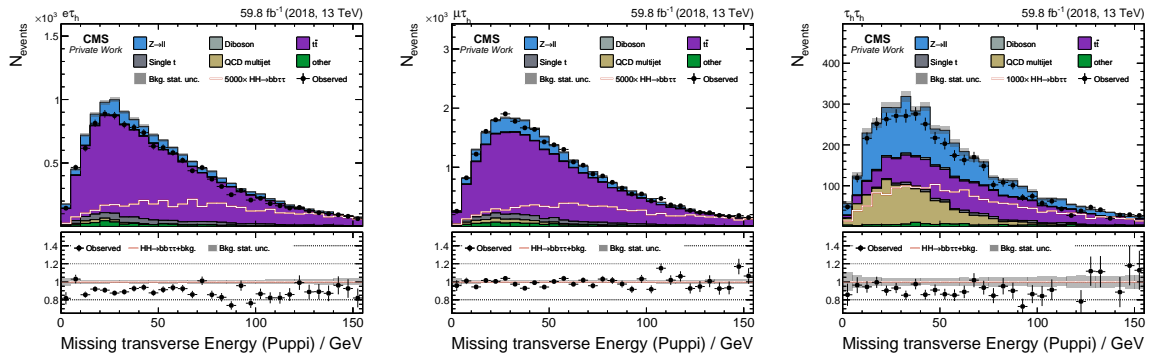


Figure A.15.: Baseline missing transverse momentum, p_T^{miss} .




Figure A.16.: ζ variable: projection of the sum of visible and missing transverse momenta onto the ζ axis, $p_\zeta^{\text{miss+vis}}$.


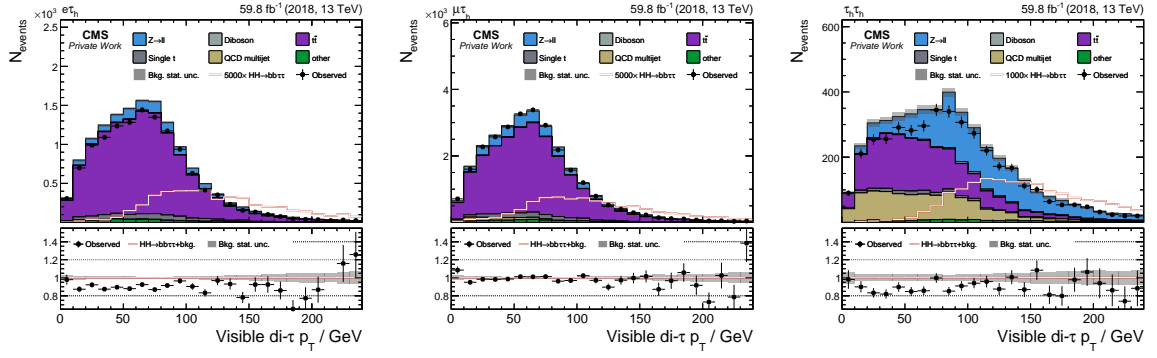

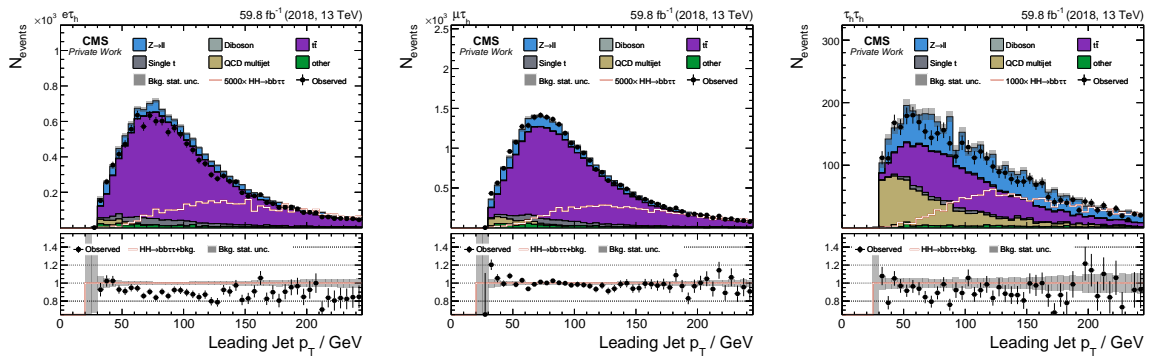


Figure A.17.: Angular separation of the two τ candidates, $\Delta R(\tau_1, \tau_2)$.


Figure A.18.: Transverse momentum of the visible $\tau^+\tau^-$ system, p_T^{vis} .

Figure A.19.: Collinear-approximation di- τ mass, $m_{\tau\tau}^{\text{coll}}$.

Figure A.20.: Transverse momentum of the leading jet, $p_T^{j,1}$.

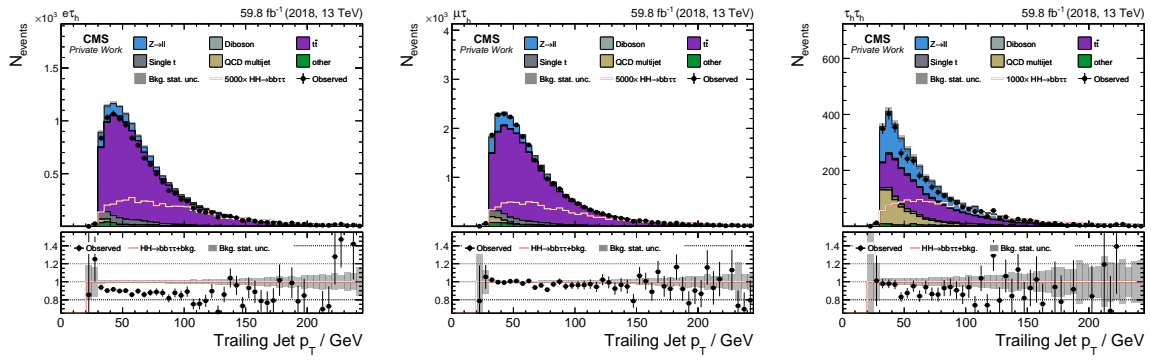


Figure A.21.: Transverse momentum of the subleading jet, $p_T^{j,2}$.

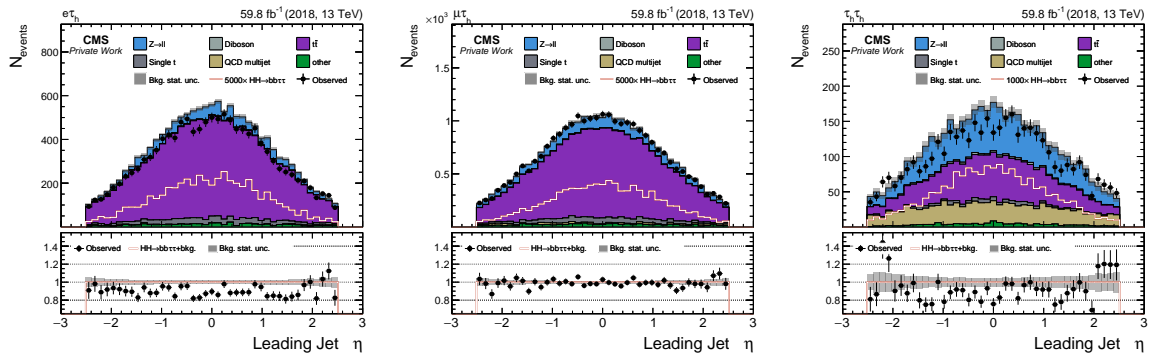


Figure A.22.: Pseudorapidity of the leading jet, $\eta(j_1)$.

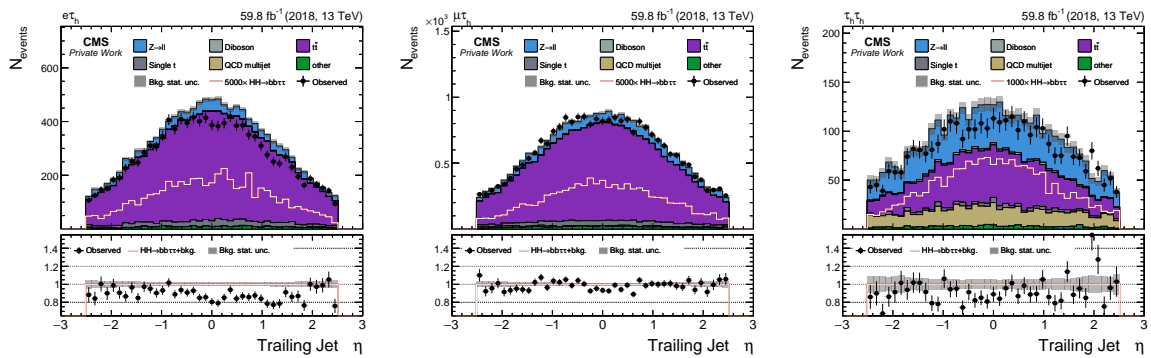


Figure A.23.: Pseudorapidity of the subleading jet, $\eta(j_2)$.

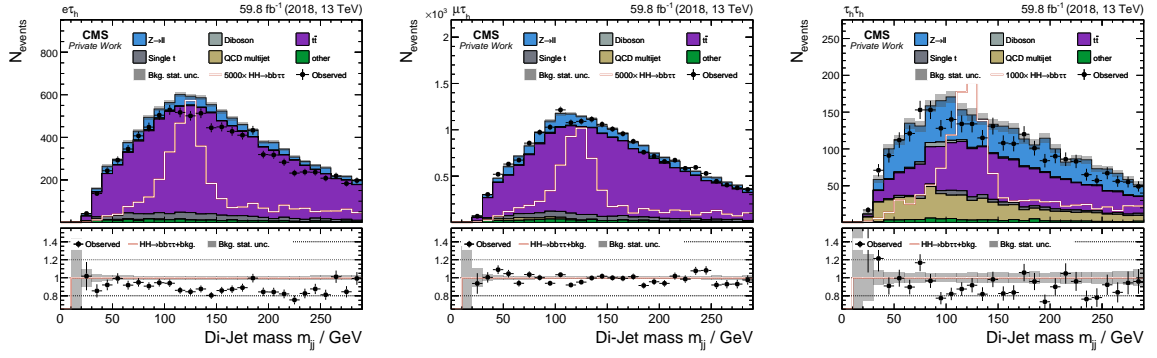


Figure A.24.: Invariant mass of the two leading jets, m_{inv}^{jj} .

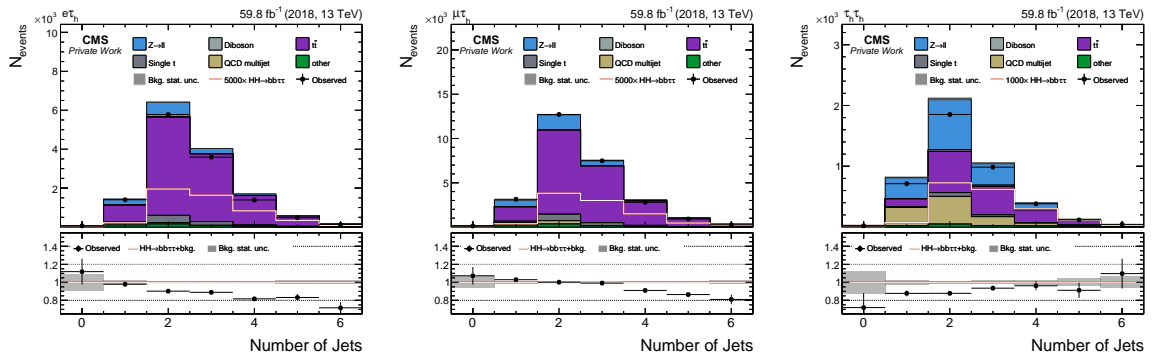


Figure A.25.: Jet multiplicity, n_{jets} .




Figure A.26.: Transverse momentum of the dijet system, p_T^{dijet} .

A. Appendix

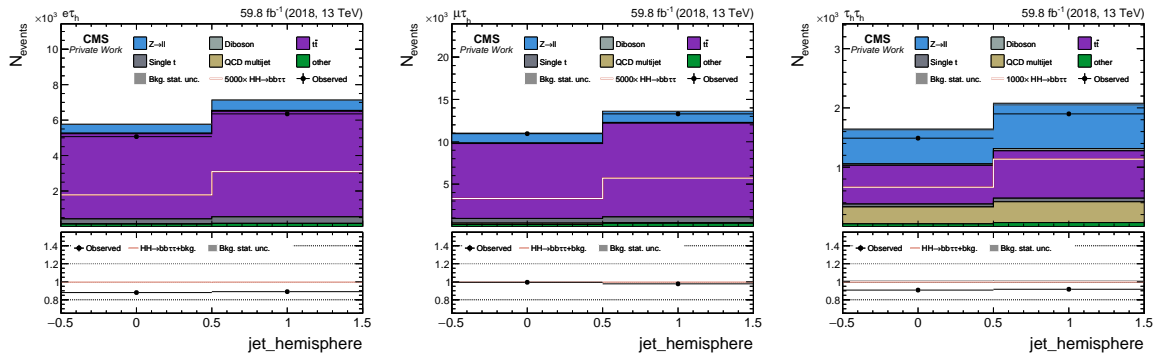


Figure A.27.: Hemisphere configuration of the two leading jets (forward/backward).




Figure A.28.: Number of b-tagged jets, n_{btag} .




Figure A.29.: Transverse momentum of the leading b jet, $p_{\text{T}}^{\text{bpair},1}$.

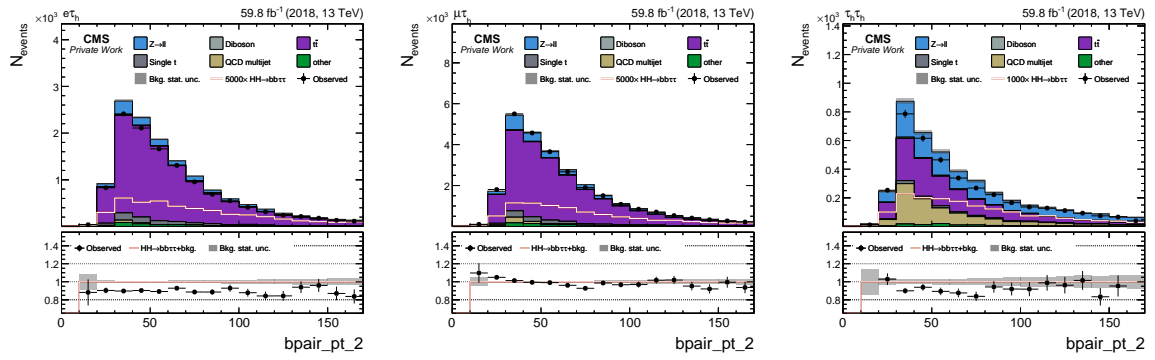


Figure A.30.: Transverse momentum of the subleading b jet, $p_T^{bpair,2}$.

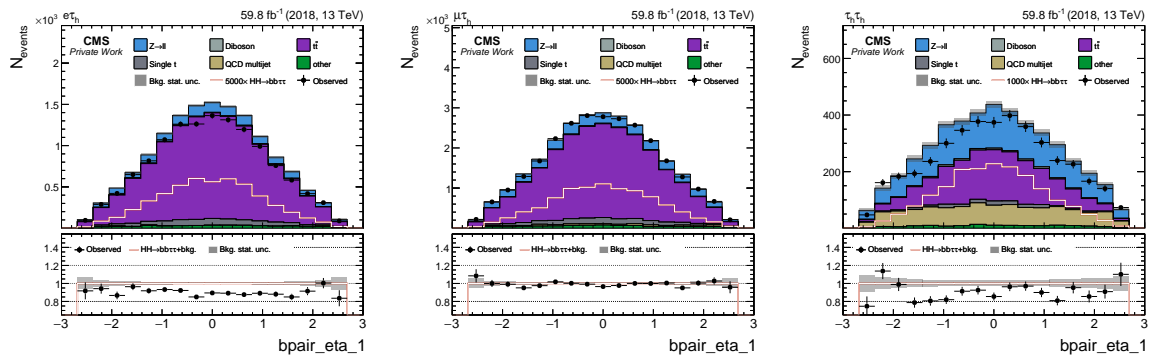


Figure A.31.: Pseudorapidity of the leading b jet, $\eta(b_1)$.

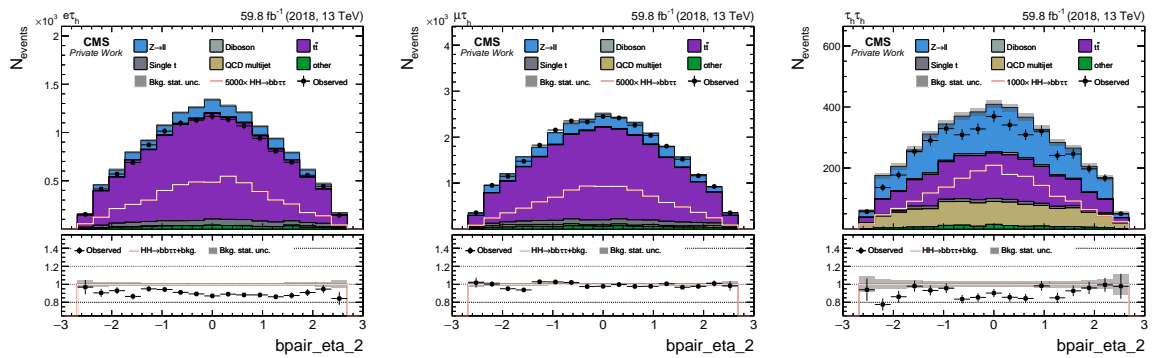


Figure A.32.: Pseudorapidity of the subleading b jet, $\eta(b_2)$.

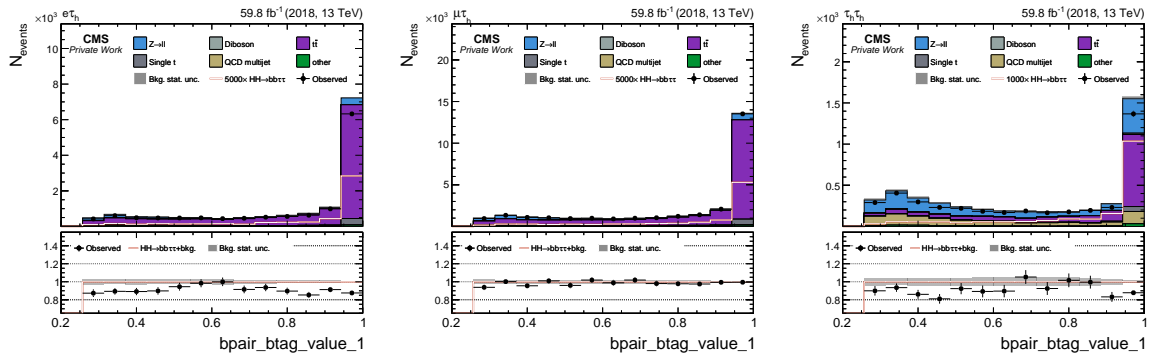


Figure A.33.: DEEPJET b-tag discriminator value of the leading b jet, $btag^1$.




Figure A.34.: DEEPJET b-tag discriminator value of the subleading b jet, $btag^2$.




Figure A.35.: Invariant mass of the b-jet pair, m_{inv}^{bb} .




Figure A.36.: Transverse momentum of the $b\bar{b}$ system, $p_T^{b\bar{b}}$.

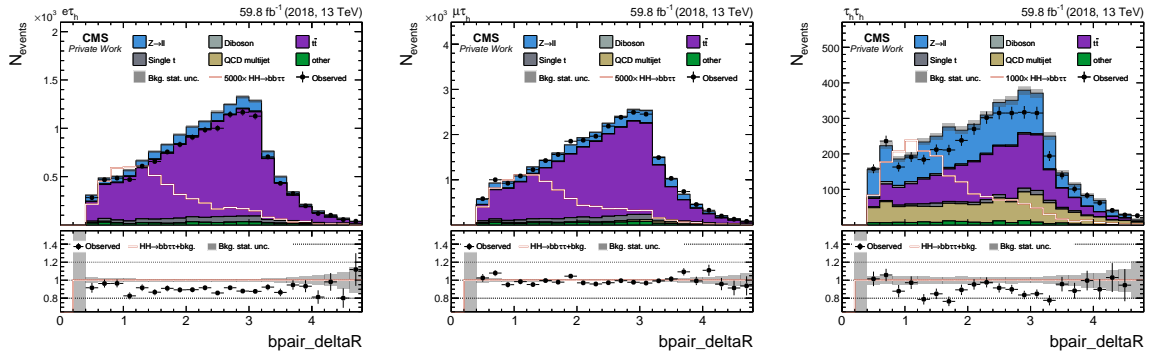


Figure A.37.: Angular separation of the b jets, $\Delta R(b_1, b_2)$.




Figure A.38.: Total transverse mass, m_T^{tot} , defined as the sum of the transverse masses of all final-state objects.

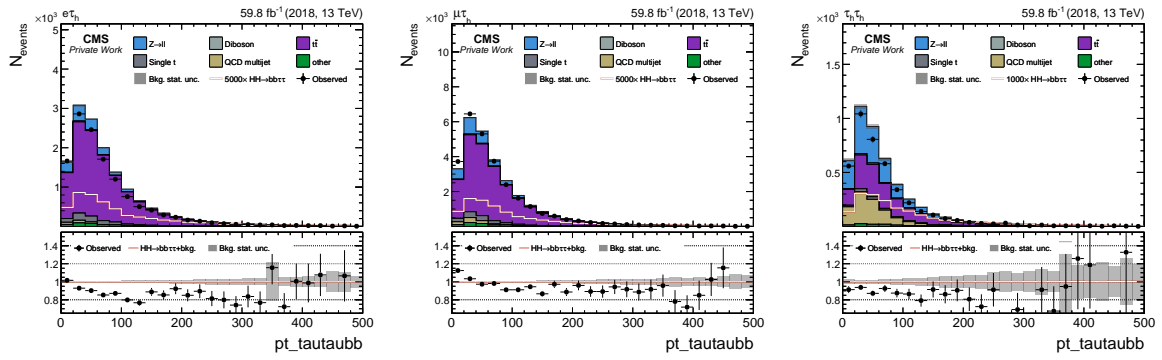


Figure A.39.: Transverse momentum of the $\tau^+\tau^- + b\bar{b}$ system, $p_{T\tau b\bar{b}}$.




Figure A.40.: Invariant mass of the $\tau^+\tau^- + b\bar{b}$ system, $m_{\tau\tau b\bar{b}}$.

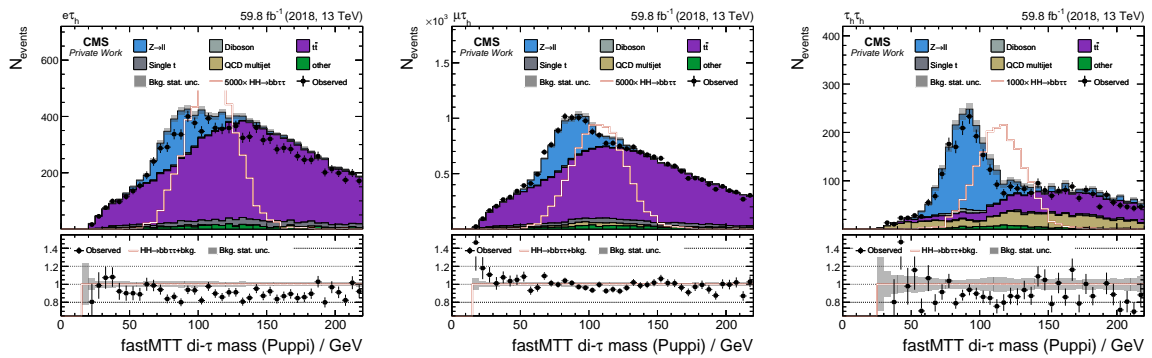


Figure A.41.: Reconstructed di- τ mass from FASTMTT, $m_{\tau\tau}^{\text{FastMTT}}$.

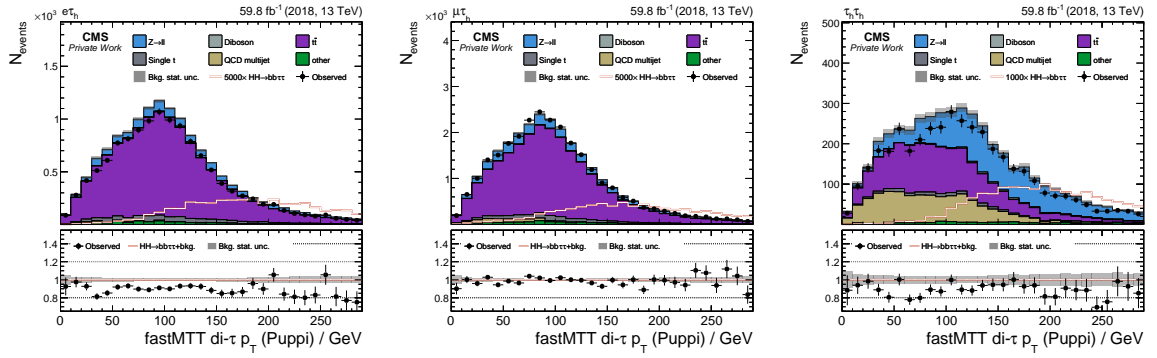


Figure A.42.: Transverse momentum of the FASTMTT di- τ system, p_T^{FastMTT} .

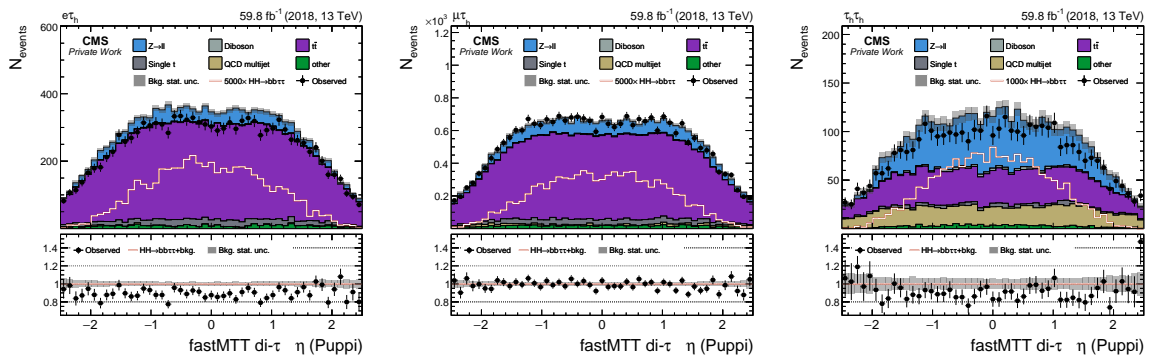


Figure A.43.: Pseudorapidity of the FASTMTT di- τ system, η^{FastMTT} .

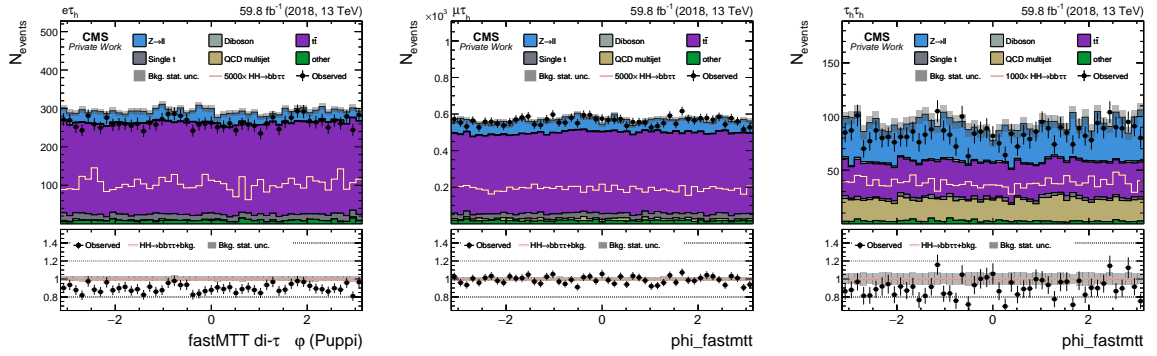


Figure A.44.: Azimuthal angle of the FASTMTT di- τ system, ϕ^{FastMTT} .

A.2. Training Datasets

Table A.1.: Class composition in the $\mu\tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 1,088,740$ and $\sum w^{\text{phys}} = 27,243$.

Class	$N^{\text{MC}} [10^3]$	MC fraction	$\sum w^{\text{phys}}$	Physical fraction
Signal	6.4	0.6%	1.87	0.0069%
DY	59.4	5.5%	3288	12.1%
Single t	7.7	0.7%	1384	5.1%
$t\bar{t}$	654.0	60.1%	21 582	79.2%
Diboson	13.6	1.2%	172	0.6%
other	347.7	31.9%	816	3.0%

Table A.2.: Class composition in the $e\tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 566,562$ and $\sum w^{\text{phys}} = 14,288$.

Class	$N^{\text{MC}} [10^3]$	MC fraction	$\sum w^{\text{phys}}$	Physical fraction
Signal	3.3	0.6%	1.02	0.0071%
DY	30.7	5.4%	1407	9.8%
Single t	3.9	0.7%	746	5.2%
$t\bar{t}$	338.2	59.7%	11 644	81.5%
Diboson	6.5	1.1%	88	0.6%
other	184.1	32.5%	401	2.8%

A.3. Training Metrics




Figure A.45.: NN training metrics for the $\mu\tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.




Figure A.46.: NN training metrics for the $\epsilon\tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.

A.4. Classification Confusion Matrices




Figure A.47.: $\tau_h\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure 6.3 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.




Figure A.48.: $\tau_h\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure 6.3 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.

A. Appendix




Figure A.49.: $\mu\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.




Figure A.50.: $\mu\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.49 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.




Figure A.51.: $\mu\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.49 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.




Figure A.52.: $\epsilon\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.

A. Appendix




Figure A.53.: $e\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.52 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.




Figure A.54.: $e\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.52 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.

A.5. Analysis Shapes

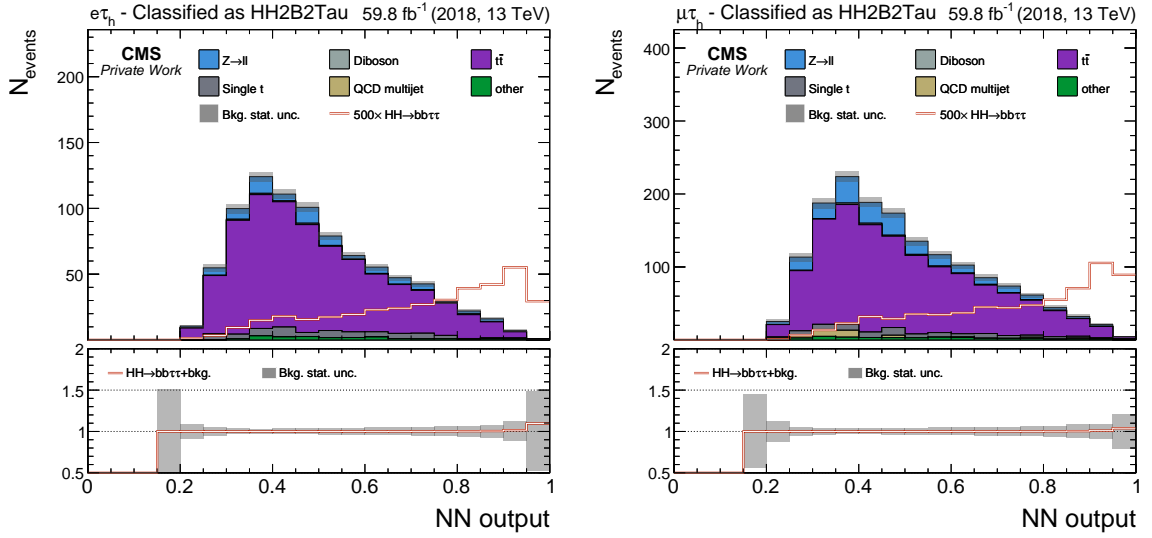


Figure A.55.: $\ell\tau_h$ final states: NN score of events assigned to the signal output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

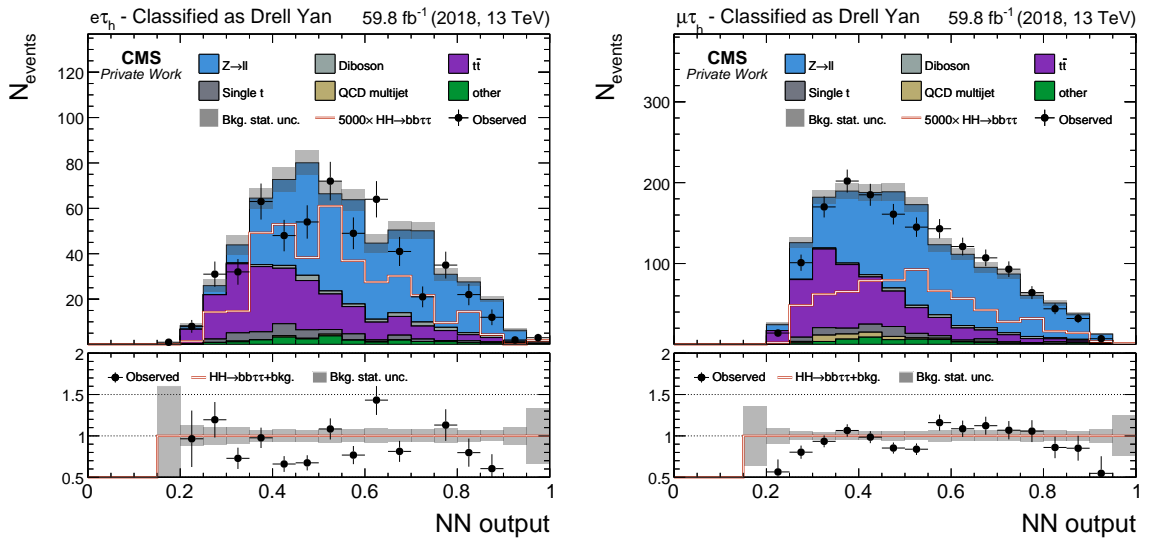


Figure A.56.: $\ell\tau_h$ final states: NN score of events assigned to the DY output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

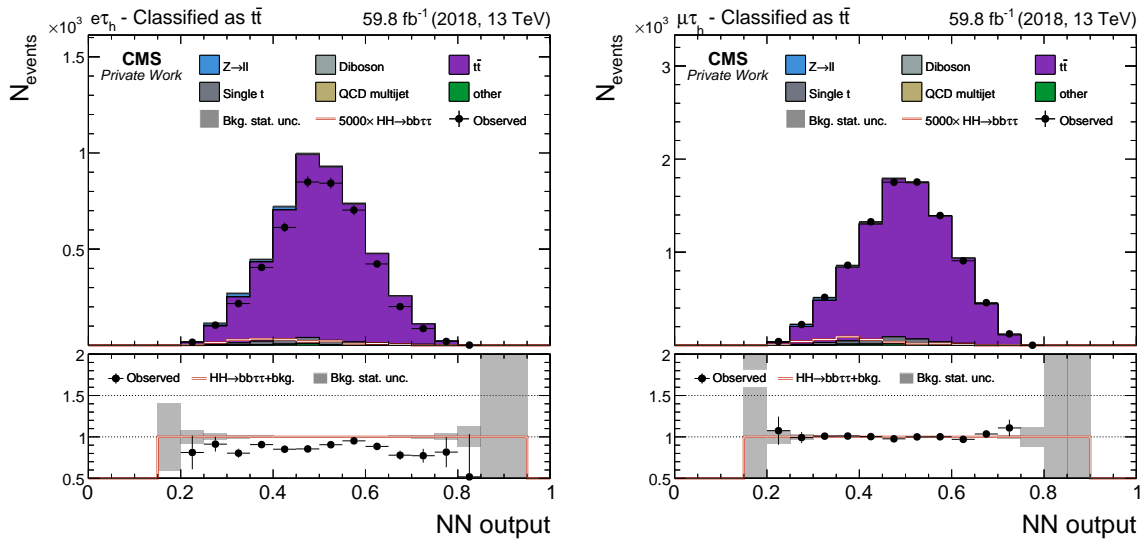


Figure A.57.: $\ell\tau_h$ final states: NN score of events assigned to the $t\bar{t}$ output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

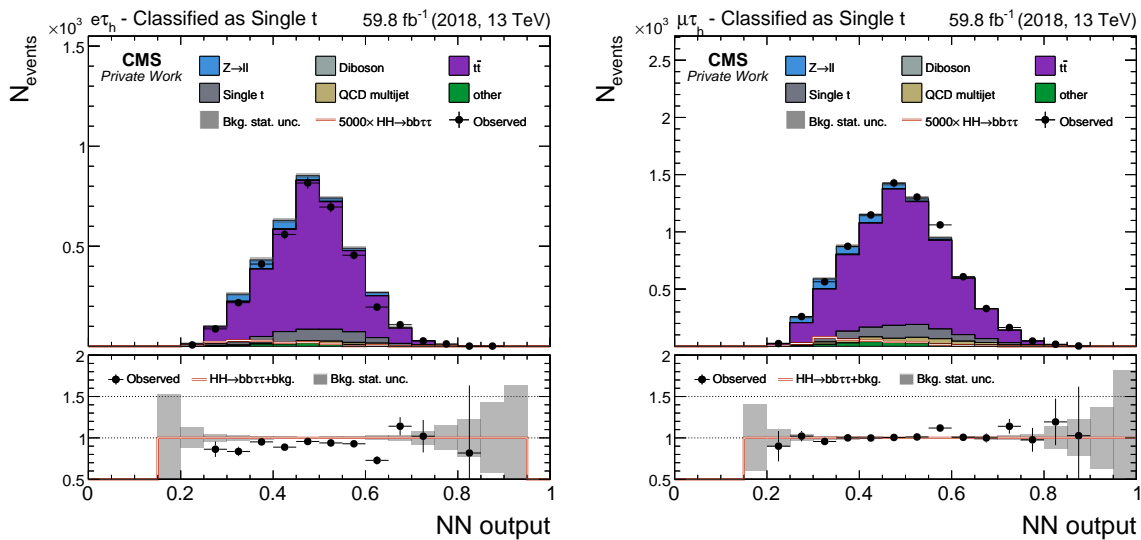


Figure A.58.: $\ell\tau_h$ final states: NN score of events assigned to the single t output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

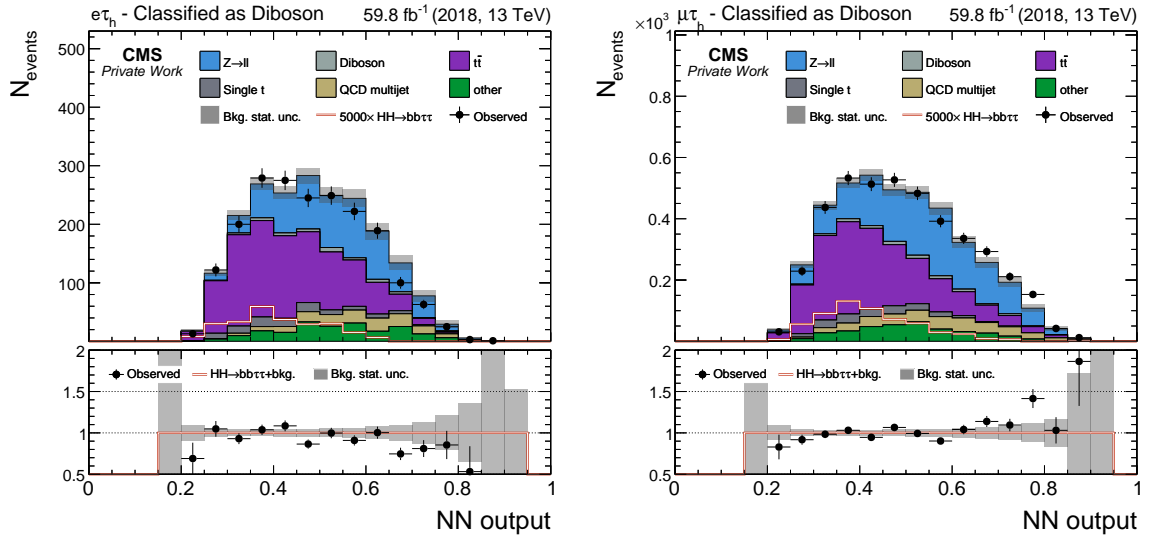


Figure A.59.: $\ell\tau_h$ final states: NN score of events assigned to the diboson output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

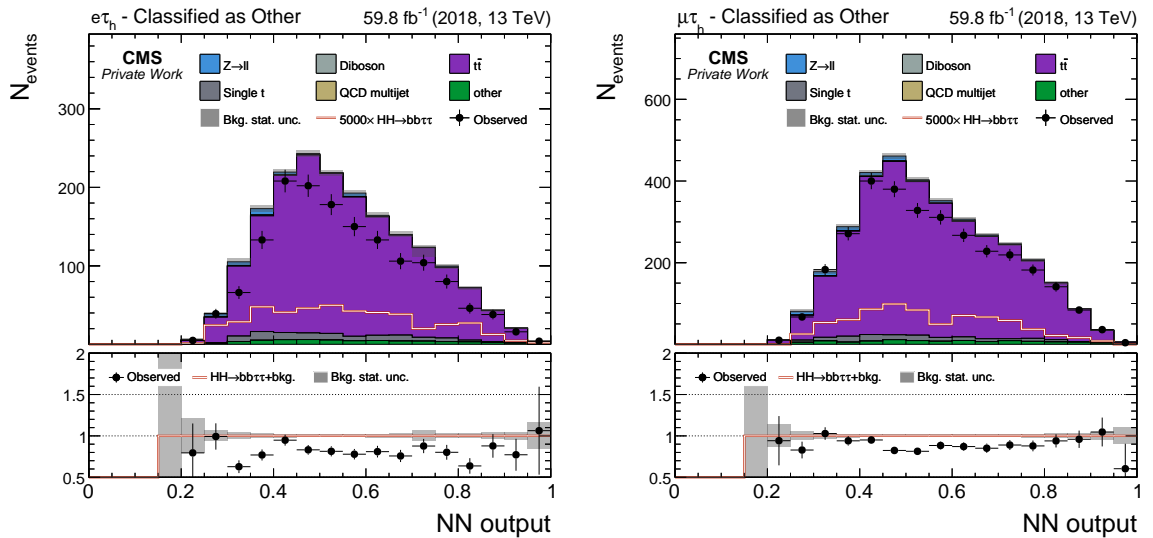


Figure A.60.: $\ell\tau_h$ final states: NN score of events assigned to the *other* output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.

A. Appendix

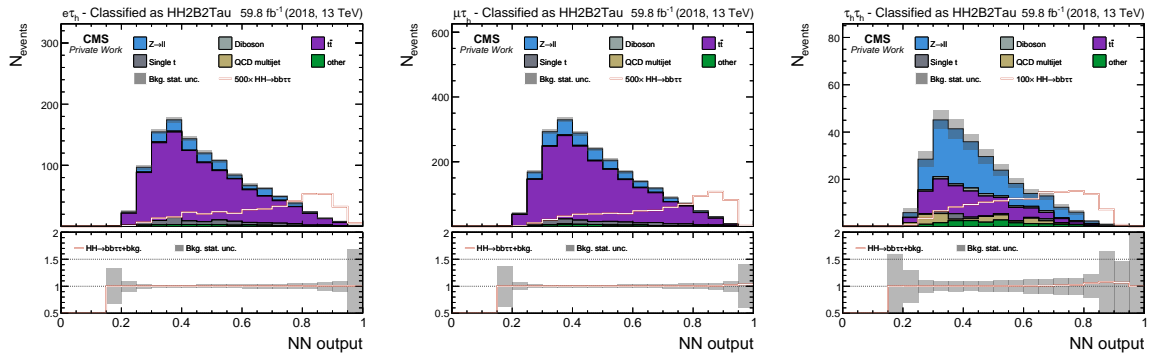


Figure A.61.: NN score of events assigned to the signal output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

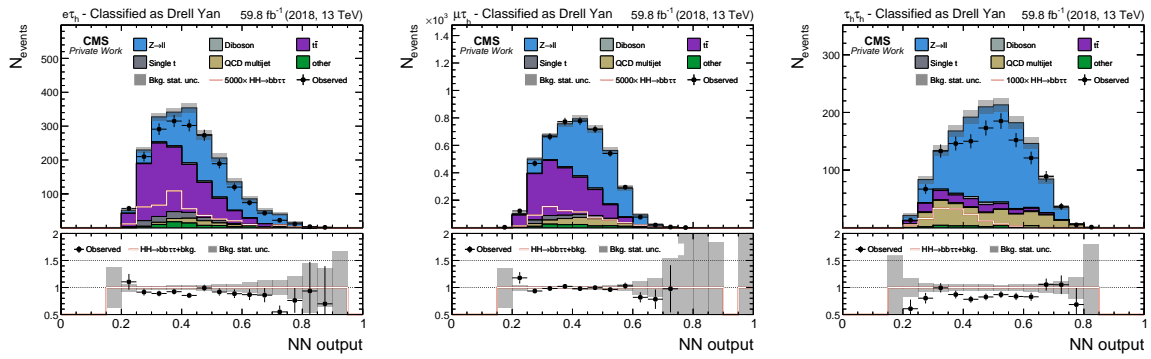


Figure A.62.: NN score of events assigned to the DY output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

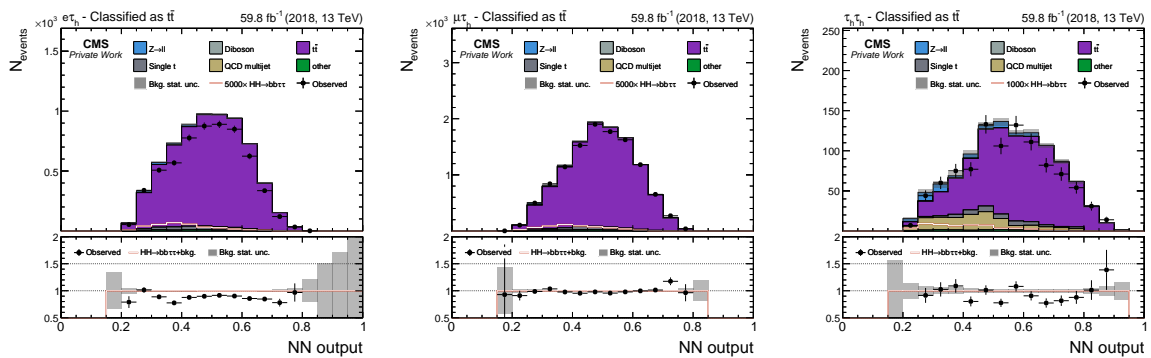


Figure A.63.: NN score of events assigned to the $t\bar{t}$ output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

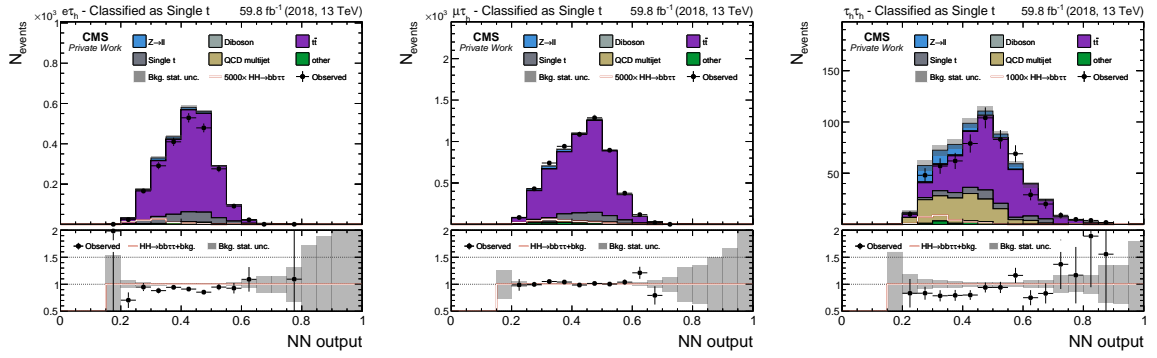


Figure A.64.: NN score of events assigned to the single t output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

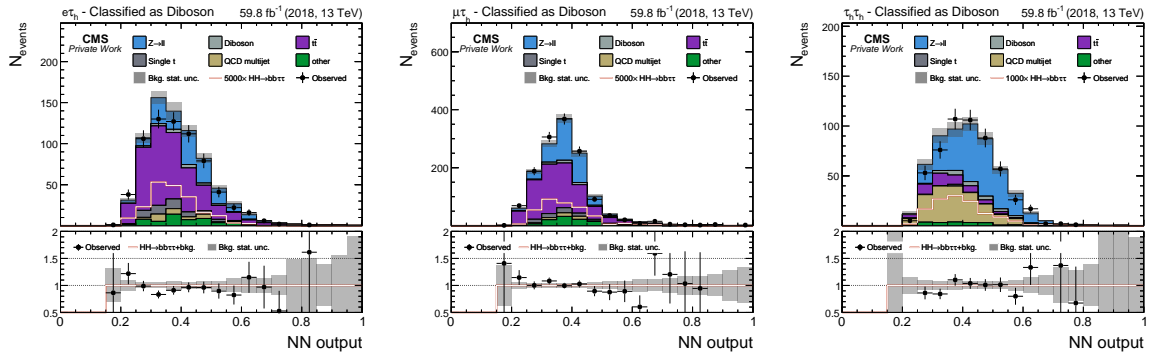


Figure A.65.: NN score of events assigned to the diboson output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

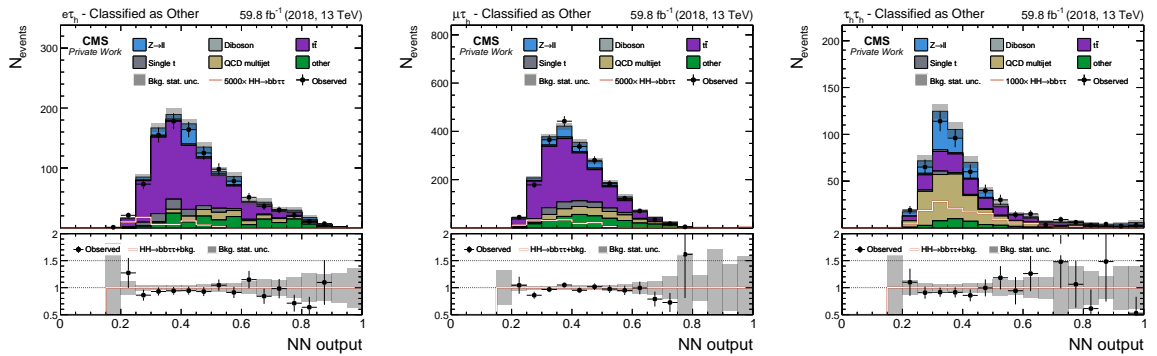


Figure A.66.: NN score of events assigned to the *other* output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.

A.6. Upper Limit Confidence Signal Strength




Figure A.67.: Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{\text{SM}}$ for each final state and their combination. **Left:** statistical-only fit; **Right:** full fit with systematic uncertainties included. NN trained with balancing weights by physical weights.




Figure A.68.: Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{\text{SM}}$ for each final state and their combination. **Left:** statistical-only fit; **Right:** full fit with systematic uncertainties included. NN trained with 50 % signal and 50 % background balancing by MC events.




Figure A.69.: Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{SM}$ for each final state and their combination. **Left:** statistical-only fit; **Right:** full fit with systematic uncertainties included. NN trained with 50 % signal and 50 % background balancing by physical weights.

List of Figures

2.1.	Elementary particles in the SM. Quarks are shown in purple and leptons in green. Together they form the elementary matter particles. The gauge bosons (shown in red) carry the fundamental forces. The Higgs boson is the quantum excitation of the Higgs field Φ . Through the Higgs mechanism, particles that couple to Φ acquire their mass. [7]	3
2.2.	Nonresonant Higgs-boson pair production through ggF. Left: Triangle diagram, where two gluons couple to a top-quark loop that produces a single off-shell Higgs boson, which then splits via the trilinear self-coupling into two Higgs bosons. Right: Box diagram, where the top-quark loop directly emits two Higgs bosons. In the SM these amplitudes interfere destructively. [27]	6
3.1.	Overview of the accelerators and detectors at CERN. Protons are produced by ionizing hydrogen gas. Before injection into the LHC, the beam is accelerated in smaller accelerators: LINAC4, the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS) [33].	9
3.2.	Schematic structure of the CMS detector [4]	10
4.1.	DEEPTAU v2.5 architecture. High-level features and the inner/outer η - ϕ grids are processed in separate branches with dense layers. The two grids then pass through convolutional blocks that compress them to 1×1 feature vectors. All streams are concatenated and fed to additional dense layers, followed by a softmax function that returns class probabilities for electron, muon, τ_h , and light quark- or gluon-induced jets [51].	18
4.2.	Receiver–operating–characteristic (ROC) curves for τ_h identification evaluated on $t\bar{t}$ events. Left: $p_T^\tau < 100$ GeV; right: $p_T^\tau > 100$ GeV. The blue curve shows DEEPTAU v2.1 and the orange curve shows DEEPTAU v2.5. Compared to v2.1, v2.5 achieves a significantly better rejection of jets, reducing the jet misidentification by about 50% at fixed τ_h efficiency [46].	20
4.3.	ROC curves for τ_h identification against electrons, evaluated on Drell-Yan events. Left: $p_T^{\tau_h} < 100$ GeV; right: $p_T^{\tau_h} > 100$ GeV. Overall, v2.5 maintains good rejection with a small degradation at high p_T	21

4.4.	Data-to-simulation scale factors (SFs) for τ_h identification. The combined values are weighted by the τ decay branching fractions and by the HPS reconstruction efficiencies per decay mode (DM). DEEPTAU v2.5 scale factors are closer to unity than v2.1, indicating that the training improvements in v2.5 reduce the sensitivity to differences between data and MC simulations. [46]	22
5.1.	Observed and expected 95 % CL upper limits on the Higgs boson pair production rate in the $b\bar{b}\tau^+\tau^-$ final state, expressed as $\sigma(\text{pp} \rightarrow \text{HH})/\sigma_{\text{SM}}$. Results are shown for 2016, 2017, 2018, and their combination using the full Run 2 data with $\mathcal{L}_{\text{int}} = 138 \text{ fb}^{-1}$ at $\sqrt{s} = 13 \text{ TeV}$. SM couplings are assumed ($\kappa_\lambda = \kappa_t = \kappa_V = \kappa_{2V} = 1$) [5]	24
5.2.	Observed and expected 95 % CL upper limits on the Higgs boson pair production rate in the $b\bar{b}\tau^+\tau^-$ final state, expressed as $\sigma(\text{pp} \rightarrow \text{HH})/\sigma_{\text{SM}}$. Results are shown for the $e\tau_h, \mu\tau_h, \tau_h\tau_h$ final states and their combination using the full Run 2 data with an integrated luminosity of $\mathcal{L}_{\text{int}} = 140 \text{ fb}^{-1}$ at $\sqrt{s} = 13 \text{ TeV}$. SM couplings are assumed ($\kappa_\lambda = \kappa_t = \kappa_V = \kappa_{2V} = 1$) [53] .	24
5.3.	Leading-order QCD diagrams for $t\bar{t}$ production at the LHC: ggF via s-, t-, and u-channel (top/gluon exchange) and quark-antiquark annihilation via an s-channel gluon.	26
5.4.	Representative DY $Z/\gamma^* + \text{jets}$ topologies with $Z/\gamma^* \rightarrow \tau^+\tau^-$: (a) Compton process $qg \rightarrow Zq$; (b) $Z + b\bar{b}$ from gluon splitting $g \rightarrow b\bar{b}$. These illustrate sources of light- and heavy-flavor jets in DY events.	27
5.5.	Leading-order electroweak single-top quark production topologies: (a) t-channel production with a forward light jet, (b) associated tW production. Subsequent decays $t \rightarrow bW$ and $W \rightarrow q\bar{q}'/\ell\nu$ (not shown) yield one b jet and either additional light jets or a prompt lepton.	28
5.6.	$\tau_h\tau_h$ final state: Data-background comparisons. The stacked histograms represent the expected contributions from DY Z/γ^* , diboson, $t\bar{t}$, single t, and other minor processes, while the QCD multijet production contribution is taken from data control regions. The observed data are shown as points with error bars. The $\text{HH} \rightarrow b\bar{b}\tau^+\tau^-$ signal is overlaid and scaled for visibility. The lower panel in each plot shows the data-to-background ratio, with the shaded band indicating the statistical uncertainty of the total background prediction.	37
5.7.	$\ell\tau_h$ final states: Data-background comparisons. The stacked histograms represent the expected contributions from DY Z/γ^* , diboson, $t\bar{t}$, single t, and other minor processes, while the QCD multijet production contribution is taken from data control regions. The observed data are shown as points with error bars. The $\text{HH} \rightarrow b\bar{b}\tau^+\tau^-$ signal is overlaid and scaled for visibility. The lower panel in each plot shows the data-to-background ratio, with the shaded band indicating the statistical uncertainty of the total background prediction.	38

6.1.	Top-down schematic of the neural-network architecture used for event classification. The model takes 20 standardized inputs, applies two hidden blocks (Linear–BatchNorm–ReLU with a Dropout layer after the first block), and outputs six class probabilities via a Softmax layer.	41
6.2.	NN training metrics for the $\tau_h\tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.	43
6.3.	$\tau_h\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	44
6.4.	$\tau_h\tau_h$ final state: NN score of events assigned to each output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	46
6.5.	Expected 95 % CL upper limits on μ for each final state and their combination, with the median expectation and the 68 % and 95 % uncertainty bands shown. Left: statistical-only fit; Right: full fit with all systematic uncertainties included.	49
6.6.	Expected 95 % CL upper limits on μ for each final state and their combination with luminosity scaling. Left: extrapolation to the full Run 2 $\mathcal{L}_{\text{int}} = 138 \text{ fb}^{-1}$. Right: projection to a combined Run 2 + Run 3 expected $\mathcal{L}_{\text{int}} = 500 \text{ fb}^{-1}$. Rates are scaled with \mathcal{L}_{int} while shapes and systematic models are kept unchanged.	50
A.1.	Visible mass of object 1. In the $\ell\tau_h$ channels object 1 is the lepton, so the distribution peaks near zero. In the $\tau_h\tau_h$ channel object 1 is a τ_h	53
A.2.	Visible mass of object 2, which is a τ_h in all final states.	53
A.3.	Transverse momentum of object 1: lepton p_T in $\ell\tau_h$ and p_T^τ in $\tau_h\tau_h$	54
A.4.	Transverse momentum of object 2, which is a τ_h in all final states.	54
A.5.	Pseudorapidity of object 1: lepton η in $\ell\tau_h$ and η^τ in $\tau_h\tau_h$	54
A.6.	Pseudorapidity of object 2, which is a τ_h in all final states.	55
A.7.	τ_h decay mode of object 1. In the $\ell\tau_h$ final states object 1 is the lepton, so these distributions are empty.	55
A.8.	τ_h decay mode of object 2.	55
A.9.	Visible di- τ mass, m_{vis} of the $\tau^+\tau^-$ system.	56
A.10.	Transverse mass m_T of object 1 with \vec{p}_T^{miss}	56
A.11.	Transverse mass m_T of object 2 with \vec{p}_T^{miss}	56
A.12.	PF-based transverse mass m_T of object 1 with \vec{p}_T^{miss}	57
A.13.	PF-based transverse mass m_T of object 2 with \vec{p}_T^{miss}	57
A.14.	PF-based missing transverse momentum, p_T^{miss}	57
A.15.	Baseline missing transverse momentum, p_T^{miss}	58
A.16.	ζ variable: projection of the sum of visible and missing transverse momenta onto the ζ axis, $p_\zeta^{\text{miss+vis}}$	58

A.17. Angular separation of the two τ candidates, $\Delta R(\tau_1, \tau_2)$	58
A.18. Transverse momentum of the visible $\tau^+\tau^-$ system, p_T^{vis}	59
A.19. Collinear-approximation di- τ mass, $m_{\tau\tau}^{\text{coll}}$	59
A.20. Transverse momentum of the leading jet, $p_T^{j,1}$	59
A.21. Transverse momentum of the subleading jet, $p_T^{j,2}$	60
A.22. Pseudorapidity of the leading jet, $\eta(j_1)$	60
A.23. Pseudorapidity of the subleading jet, $\eta(j_2)$	60
A.24. Invariant mass of the two leading jets, m_{inv}^{jj}	61
A.25. Jet multiplicity, n_{jets}	61
A.26. Transverse momentum of the dijet system, p_T^{dijet}	61
A.27. Hemisphere configuration of the two leading jets (forward/backward).	62
A.28. Number of b-tagged jets, n_{btag}	62
A.29. Transverse momentum of the leading b jet, $p_T^{\text{bpair},1}$	62
A.30. Transverse momentum of the subleading b jet, $p_T^{\text{bpair},2}$	63
A.31. Pseudorapidity of the leading b jet, $\eta(b_1)$	63
A.32. Pseudorapidity of the subleading b jet, $\eta(b_2)$	63
A.33. DEEJET b-tag discriminator value of the leading b jet, btag^1	64
A.34. DEEJET b-tag discriminator value of the subleading b jet, btag^2	64
A.35. Invariant mass of the b-jet pair, $m_{\text{inv}}^{\text{bb}}$	64
A.36. Transverse momentum of the bb system, p_T^{bpair}	65
A.37. Angular separation of the b jets, $\Delta R(b_1, b_2)$	65
A.38. Total transverse mass, m_T^{tot} , defined as the sum of the transverse masses of all final-state objects.	65
A.39. Transverse momentum of the $\tau^+\tau^- + \text{bb}$ system, $p_{\tau\tau\text{bb}}$	66
A.40. Invariant mass of the $\tau^+\tau^- + \text{bb}$ system, $m_{\tau\tau\text{bb}}$	66
A.41. Reconstructed di- τ mass from FASTMTT, $m_{\tau\tau}^{\text{FastMTT}}$	66
A.42. Transverse momentum of the FASTMTT di- τ system, p_T^{FastMTT}	67
A.43. Pseudorapidity of the FASTMTT di- τ system, η^{FastMTT}	67
A.44. Azimuthal angle of the FASTMTT di- τ system, ϕ^{FastMTT}	67
A.45. NN training metrics for the $\mu\tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.	69
A.46. NN training metrics for the $e\tau_h$ final state with balancing class weights by MC events. Left: cross-entropy loss vs. epochs for training and validation, showing steady convergence without divergence. Middle: F_1 score, rises rapidly and saturates. Right: precision and recall stabilize, with high precision and lower recall, consistent with class imbalance.	70
A.47. $\tau_h\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure 6.3 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	71

A.48. $\tau_h\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure 6.3 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	71
A.49. $\mu\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	72
A.50. $\mu\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.49 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	72
A.51. $\mu\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.49 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	73
A.52. $e\tau_h$ final state: Confusion matrices on the test sample. Numbers are physical weighted event counts. Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	73
A.53. $e\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.52 but normalized by predicted class (column-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	74
A.54. $e\tau_h$ final state: Confusion matrices on the test sample. Showing the same data as Figure A.52 but normalized by true class (row-normalized). Left: training with balancing class weights by MC counts; Right: training with balancing class weights by physical weights.	74
A.55. $\ell\tau_h$ final states: NN score of events assigned to the signal output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	75
A.56. $\ell\tau_h$ final states: NN score of events assigned to the DY output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	75
A.57. $\ell\tau_h$ final states: NN score of events assigned to the $t\bar{t}$ output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	76
A.58. $\ell\tau_h$ final states: NN score of events assigned to the single t output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	76
A.59. $\ell\tau_h$ final states: NN score of events assigned to the diboson output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	77
A.60. $\ell\tau_h$ final states: NN score of events assigned to the <i>other</i> output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by the number of simulated training samples.	77

A.61. NN score of events assigned to the signal output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	78
A.62. NN score of events assigned to the DY output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	78
A.63. NN score of events assigned to the $t\bar{t}$ output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	78
A.64. NN score of events assigned to the single t output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	79
A.65. NN score of events assigned to the diboson output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	79
A.66. NN score of events assigned to the <i>other</i> output class. The signal is overlaid and scaled for better visibility. The NN training is done with class balancing by physical weights.	79
A.67. Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{SM}$ for each final state and their combination. Left: statistical-only fit; Right: full fit with systematic uncertainties included. NN trained with balancing weights by physical weights.	80
A.68. Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{SM}$ for each final state and their combination. Left: statistical-only fit; Right: full fit with systematic uncertainties included. NN trained with 50 % signal and 50 % background balancing by MC events.	80
A.69. Expected 95 % CL upper limits on $\mu = \sigma(pp \rightarrow HH)/\sigma_{SM}$ for each final state and their combination. Left: statistical-only fit; Right: full fit with systematic uncertainties included. NN trained with 50 % signal and 50 % background balancing by physical weights.	81

List of Tables

2.1.	Branching fraction of the SM Higgs boson at $m_H = 125.09$ GeV [31, 32]	7
2.2.	$\tau^+\tau^-$ final states sorted by branching fraction [32].	8
5.1.	ABCD regions for the $\tau_h\tau_h$ final state. Iso τ_h pass the DEEPTAU VSjet Medium working point; Anti-Iso τ_h fail the Medium but pass the DEEPTAU VSjet VVLoose working point.	30
5.2.	ABCD regions for the $\ell\tau_h$ final states. Iso τ_h pass the DEEPTAU VSjet Medium working point; Anti-Iso τ_h fail the Medium but pass the DEEPTAU VSjet VVLoose working point. Lepton ID/iso follows the nominal selection.	30
5.3.	Event generators and perturbative order in α_s used for signal and background samples in the 2018 UL simulation. All samples are showered and hadronized with PYTHIA 8 using the CP5 tune.	33
5.4.	Identification working points of DeepTau v2.5 used for τ_h identification	34
5.5.	Per-channel object requirements used in the $\tau^+\tau^-$ selection.	34
6.1.	Kinematic and event-level observables used in the analysis.	40
6.2.	Class composition in the $\tau_h\tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 244,726$ and $\sum w^{\text{phys}} = 3606$	43
6.3.	Summary of $\ln N$ rate uncertainties used in the nominal fit. Correlations are across final states within the same era unless noted. [†] The uncertainty in the vsJet working point of the τ_h identification discriminant is typically a shape effect; here a conservative $\ln N$ proxy is used.	48
A.1.	Class composition in the $\mu\tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 1,088,740$ and $\sum w^{\text{phys}} = 27,243$	68
A.2.	Class composition in the $e\tau_h$ final state. Shown are the number of nominal MC events and the sum of physical weights which represents the expected physical yield in a class. Totals after selection: $N^{\text{MC}} = 566,562$ and $\sum w^{\text{phys}} = 14,288$	68

Bibliography

- [1] Nikita Shadskiy. “Search for resonant di-Higgs production in $bb + \tau\tau$ final states in pp collisions at $\sqrt{s} = 13$ TeV”. PhD thesis. Karlsruhe Institute of Technology (KIT), 2025.
- [2] S. Chatrchyan et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 30–61. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.021. URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.021>.
- [3] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020. URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.020>.
- [4] CMS Collaboration. *Detector*. CMS Experiment (CERN) — overview of the CMS detector. 2025. URL: <https://cms.cern/detector> (visited on 10/25/2025).
- [5] A. Tumasyan et al. “Search for nonresonant Higgs boson pair production in final state with two bottom quarks and two tau leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Physics Letters B* 842 (July 2023), p. 137531. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2022.137531. URL: <http://dx.doi.org/10.1016/j.physletb.2022.137531>.
- [6] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *The European Physical Journal C* 71.2 (Feb. 2011). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0. URL: <http://dx.doi.org/10.1140/epjc/s10052-011-1554-0>.
- [7] Cush. *Standard Model of Elementary Particles*. URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg (visited on 10/22/2025).
- [8] M. Gell-Mann. “A schematic model of baryons and mesons”. In: *Physics Letters* 8.3 (1964), pp. 214–215. ISSN: 0031-9163. DOI: [https://doi.org/10.1016/S0031-9163\(64\)92001-3](https://doi.org/10.1016/S0031-9163(64)92001-3). URL: <https://www.sciencedirect.com/science/article/pii/S0031916364920013>.
- [9] G. Zweig. “An SU(3) model for strong interaction symmetry and its breaking. Version 1”. In: (Jan. 1964). DOI: 10.17181/CERN-TH-401.
- [10] Sheldon L. Glashow. “Partial-symmetries of weak interactions”. In: *Nuclear Physics* 22.4 (1961), pp. 579–588. ISSN: 0029-5582. DOI: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2). URL: <https://www.sciencedirect.com/science/article/pii/0029558261904692>.

- [11] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.19.1264>.
- [12] A. Salam and J.C. Ward. “Electromagnetic and weak interactions”. In: *Physics Letters* 13.2 (1964), pp. 168–171. ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)90711-5](https://doi.org/10.1016/0031-9163(64)90711-5). URL: <https://www.sciencedirect.com/science/article/pii/0031916364907115>.
- [13] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.321>.
- [14] P.W. Higgs. “Broken symmetries, massless particles and gauge fields”. In: *Physics Letters* 12.2 (1964), pp. 132–133. ISSN: 0031-9163. DOI: [https://doi.org/10.1016/0031-9163\(64\)91136-9](https://doi.org/10.1016/0031-9163(64)91136-9). URL: <https://www.sciencedirect.com/science/article/pii/0031916364911369>.
- [15] Peter W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [16] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. In: *Phys. Rev. D* 2 (7 Oct. 1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285. URL: <https://link.aps.org/doi/10.1103/PhysRevD.2.1285>.
- [17] G. Arnison et al. “Experimental observation of isolated large transverse energy electrons with associated missing energy at $s=540$ GeV”. In: *Physics Letters B* 122.1 (1983), pp. 103–116. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(83\)91177-2](https://doi.org/10.1016/0370-2693(83)91177-2). URL: <https://www.sciencedirect.com/science/article/pii/0370269383911772>.
- [18] M. Banner et al. “Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider”. In: *Physics Letters B* 122.5 (1983), pp. 476–485. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(83\)91605-2](https://doi.org/10.1016/0370-2693(83)91605-2). URL: <https://www.sciencedirect.com/science/article/pii/0370269383916052>.
- [19] G. Arnison et al. “Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider”. In: *Physics Letters B* 126.5 (1983), pp. 398–410. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(83\)90188-0](https://doi.org/10.1016/0370-2693(83)90188-0). URL: <https://www.sciencedirect.com/science/article/pii/0370269383901880>.
- [20] P. Bagnaia et al. “Evidence for $Z \rightarrow e^+e^-$ at the CERN pp collider”. In: *Physics Letters B* 129.1 (1983), pp. 130–140. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(83\)90744-X](https://doi.org/10.1016/0370-2693(83)90744-X). URL: <https://www.sciencedirect.com/science/article/pii/037026938390744X>.
- [21] S. Abachi et al. “Observation of the Top Quark”. In: *Physical Review Letters* 74.14 (Apr. 1995), pp. 2632–2637. ISSN: 1079-7114. DOI: 10.1103/physrevlett.74.2632. URL: <http://dx.doi.org/10.1103/PhysRevLett.74.2632>.

-
- [22] Y. Nambu and G. Jona-Lasinio. “Dynamical Model of Elementary Particles Based on an Analogy with Superconductivity. I”. In: *Phys. Rev.* 122 (1 Apr. 1961), pp. 345–358. DOI: 10.1103/PhysRev.122.345. URL: <https://link.aps.org/doi/10.1103/PhysRev.122.345>.
- [23] Y. Nambu and G. Jona-Lasinio. “Dynamical Model of Elementary Particles Based on an Analogy with Superconductivity. II”. In: *Phys. Rev.* 124 (1 Oct. 1961), pp. 246–254. DOI: 10.1103/PhysRev.124.246. URL: <https://link.aps.org/doi/10.1103/PhysRev.124.246>.
- [24] J. Goldstone. “Field Theories with Superconductor Solutions”. In: *Nuovo Cim.* 19 (1961), pp. 154–164. DOI: 10.1007/BF02812722.
- [25] Nicola Cabibbo. “Unitary Symmetry and Leptonic Decays”. In: *Phys. Rev. Lett.* 10 (12 June 1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.10.531>.
- [26] Makoto Kobayashi and Toshihide Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657. ISSN: 0033-068X. DOI: 10.1143/PTP.49.652. eprint: <https://academic.oup.com/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf>. URL: <https://doi.org/10.1143/PTP.49.652>.
- [27] CMS Collaboration. *How does the Higgs boson interact with itself?* URL: <https://cms.cern/news/how-does-higgs-boson-interact-itself> (visited on 10/23/2025).
- [28] Massimiliano Grazzini et al. *Higgs boson pair production at NNLO with top quark mass effects*. JHEP 05 (2018) 059; arXiv:1803.02463. May 2018. URL: [https://doi.org/10.1007/JHEP05\(2018\)059](https://doi.org/10.1007/JHEP05(2018)059) (visited on 10/23/2025).
- [29] J. Baglio et al. *gg → HH: Combined uncertainties*. Phys. Rev. D 103 (2021) 056002; arXiv:2008.11626v3. Mar. 2021. URL: <https://doi.org/10.1103/PhysRevD.103.056002> (visited on 10/23/2025).
- [30] LHC Higgs Cross Section Working Group. *SM Higgs production cross sections at $\sqrt{s} = 13$ TeV (CERN Yellow Report page)*. CERN TWiki; living page linked to CERN Yellow Report 4. 2016. URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageAt13TeV> (visited on 10/23/2025).
- [31] LHC Higgs Cross Section Working Group. *SM Higgs Branching Ratios and Total Decay Widths (CERN Yellow Report page)*. CERN TWiki; living page linked to CERN Yellow Report 4. 2016. URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageBR> (visited on 10/23/2025).
- [32] S. Navas et al. “Review of particle physics”. In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.
- [33] Ewa Lopienska. *The CERN accelerator complex, layout in 2022*. General Photo. 2022. URL: <https://cds.cern.ch/record/2800984>.

- [34] The CMS Collaboration. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. In: *Journal of Instrumentation* 9.10 (Oct. 2014), P10009. DOI: 10.1088/1748-0221/9/10/P10009. URL: <https://doi.org/10.1088/1748-0221/9/10/P10009>.
- [35] A Benaglia. “The CMS ECAL performance with examples”. In: *Journal of Instrumentation* 9.02 (Feb. 2014), p. C02008. DOI: 10.1088/1748-0221/9/02/C02008. URL: <https://doi.org/10.1088/1748-0221/9/02/C02008>.
- [36] J Mans et al. *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*. Tech. rep. Additional contact persons: Jeffrey Spalding, Fermilab, spalding@cern.ch, Didier Contardo, Universite Claude Bernard-Lyon I, contardo@cern.ch. CERN, 2012. URL: <https://cds.cern.ch/record/1481837>.
- [37] A.M. Sirunyan et al. “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $s=13$ TeV”. In: *Journal of Instrumentation* 13.06 (June 2018), P06015. DOI: 10.1088/1748-0221/13/06/P06015. URL: <https://doi.org/10.1088/1748-0221/13/06/P06015>.
- [38] CMS Collaboration. *Detecting Muons*. CMS Experiment (CERN) — muon detection overview. 2025. URL: <https://cms.cern/detector/detecting-muons> (visited on 10/25/2025).
- [39] V. Khachatryan et al. “The CMS trigger system”. In: *Journal of Instrumentation* 12.01 (Jan. 2017), P01020. DOI: 10.1088/1748-0221/12/01/P01020. URL: <https://doi.org/10.1088/1748-0221/12/01/P01020>.
- [40] A. M. Sirunyan et al. “Particle-flow reconstruction and global event description with the CMS detector”. In: *JINST* 12.10 (2017), P10003. DOI: 10.1088/1748-0221/12/10/P10003. arXiv: 1706.04965 [physics.ins-det].
- [41] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti- k_T jet clustering algorithm”. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. URL: <https://doi.org/10.1088/1126-6708/2008/04/063>.
- [42] S. Catani et al. “Longitudinally-invariant k_T -clustering algorithms for hadron-hadron collisions”. In: *Nuclear Physics B* 406.1 (1993), pp. 187–224. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(93\)90166-M](https://doi.org/10.1016/0550-3213(93)90166-M). URL: <https://www.sciencedirect.com/science/article/pii/055032139390166M>.
- [43] “Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $s=8$ TeV”. In: *Journal of Instrumentation* 10.06 (June 2015), P06005. DOI: 10.1088/1748-0221/10/06/P06005. URL: <https://doi.org/10.1088/1748-0221/10/06/P06005>.
- [44] E. Bols et al. “Jet flavour classification using DeepJet”. In: *Journal of Instrumentation* 15.12 (Dec. 2020), P12012–P12012. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/12/p12012. URL: <http://dx.doi.org/10.1088/1748-0221/15/12/P12012>.

-
- [45] A. Tumasyan et al. “Identification of hadronic tau lepton decays using a deep neural network”. In: *Journal of Instrumentation* 17.07 (July 2022), P07023. ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/07/p07023. URL: <http://dx.doi.org/10.1088/1748-0221/17/07/P07023>.
- [46] “Performance of the CNN-based tau identification algorithm with Domain Adaptation using Adversarial Machine Learning for Run 3”. In: (2024). URL: <https://cds.cern.ch/record/2904699>.
- [47] CMS Collaboration. “Performance of τ -lepton reconstruction and identification in CMS”. In: *Journal of Instrumentation* 7.01 (Jan. 2012), P01001–P01001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/7/01/p01001. URL: <http://dx.doi.org/10.1088/1748-0221/7/01/P01001>.
- [48] A.M. Sirunyan et al. “Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $s=13$ TeV”. In: *Journal of Instrumentation* 13.10 (Oct. 2018), P10005–P10005. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/10/p10005. URL: <http://dx.doi.org/10.1088/1748-0221/13/10/P10005>.
- [49] Daniele Bertolini et al. “Pileup Per Particle Identification”. In: *JHEP* 10 (2014), p. 059. DOI: 10.1007/JHEP10(2014)059. arXiv: 1407.6013 [hep-ph].
- [50] Albert M Sirunyan et al. “Pileup mitigation at CMS in 13 TeV data”. In: *JINST* 15.09 (2020), P09018. DOI: 10.1088/1748-0221/15/09/P09018. arXiv: 2003.00503 [hep-ex].
- [51] Lucas Russell. *Identification of Hadronic Tau Lepton Decays with Domain Adaptation using Adversarial Machine Learning at CMS*. Master’s thesis, CERN-THESIS-2022-126; Imperial College London and Ecole Polytechnique Lausanne. Oct. 2022. URL: <https://repository.cern/records/1ebs0-dzp03> (visited on 10/28/2025).
- [52] G. Aad et al. “Combination of Searches for Higgs Boson Pair Production in pp Collisions at $\sqrt{s} = 13\text{TeV}$ with the ATLAS Detector”. In: *Phys. Rev. Lett.* 133 (10 Sept. 2024), p. 101801. DOI: 10.1103/PhysRevLett.133.101801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.101801>.
- [53] Georges Aad et al. “Search for the nonresonant production of Higgs boson pairs via gluon fusion and vector-boson fusion in the $bb\bar{\tau}\tau$ - final state in proton-proton collisions at $s=13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 110.3 (2024), p. 032012. DOI: 10.1103/PhysRevD.110.032012. arXiv: 2404.12660 [hep-ex].
- [54] CDF Collaboration. *Measurement of $\sigma \cdot \mathcal{B}(W \rightarrow e\nu)$ and $\sigma \cdot \mathcal{B}(Z^0 \rightarrow e^+e^-)$ in $\bar{p}p$ collisions at $\sqrt{s} = 1800$ GeV*. *Phys. Rev. D* 44 (1991) 29–52. July 1991. URL: <https://doi.org/10.1103/PhysRevD.44.29> (visited on 11/07/2025).
- [55] A. M. Sirunyan et al. “Measurement of the $t\bar{t}$ production cross section using events with one lepton and at least one jet in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 2017.9 (Sept. 2017). ISSN: 1029-8479. DOI: 10.1007/jhep09(2017)051. URL: [http://dx.doi.org/10.1007/JHEP09\(2017\)051](http://dx.doi.org/10.1007/JHEP09(2017)051).

- [56] Ulrich Husemann. “Top-quark physics: Status and prospects”. In: *Progress in Particle and Nuclear Physics* 95 (July 2017), pp. 48–97. ISSN: 0146-6410. DOI: 10.1016/j.pnnp.2017.03.002. URL: <http://dx.doi.org/10.1016/j.pnnp.2017.03.002>.
- [57] LHC Top Working Group. *NLO single-top channel cross sections (ATLAS–CMS recommended predictions)*. CERN TWiki — reference cross sections for single-top production. 2017. URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SingleTopRefXsec> (visited on 11/07/2025).
- [58] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [59] A. M. Sirunyan et al. “Extraction and validation of a new set of CMS pythia8 tunes from underlying-event measurements”. In: *The European Physical Journal C* 80.1 (Jan. 2020). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-7499-4. URL: <http://dx.doi.org/10.1140/epjc/s10052-019-7499-4>.
- [60] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. arXiv: 2203.11601 [hep-ph]. URL: <https://arxiv.org/abs/2203.11601>.
- [61] J. Alwall et al. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. In: *Journal of High Energy Physics* 2014.7 (July 2014). ISSN: 1029-8479. DOI: 10.1007/jhep07(2014)079. URL: [http://dx.doi.org/10.1007/JHEP07\(2014\)079](http://dx.doi.org/10.1007/JHEP07(2014)079).
- [62] Stefano Frixione and Bryan R Webber. “Matching NLO QCD computations and parton shower simulations”. In: *Journal of High Energy Physics* 2002.06 (June 2002), pp. 029–029. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2002/06/029. URL: <http://dx.doi.org/10.1088/1126-6708/2002/06/029>.
- [63] Rikkert Frederix and Stefano Frixione. “Merging meets matching in MC@NLO”. In: *Journal of High Energy Physics* 2012.12 (Dec. 2012). ISSN: 1029-8479. DOI: 10.1007/jhep12(2012)061. URL: [http://dx.doi.org/10.1007/JHEP12\(2012\)061](http://dx.doi.org/10.1007/JHEP12(2012)061).
- [64] Simone Alioli et al. *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*. JHEP 06 (2010) 043; arXiv:1002.2581. June 2010. URL: [https://doi.org/10.1007/JHEP06\(2010\)043](https://doi.org/10.1007/JHEP06(2010)043) (visited on 11/04/2025).
- [65] B. Andersson et al. “Parton fragmentation and string dynamics”. In: *Physics Reports* 97.2 (1983), pp. 31–145. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7). URL: <https://www.sciencedirect.com/science/article/pii/0370157383900807>.
- [66] Bo Andersson, G. Gustafson, and B. Soderberg. “A General Model for Jet Fragmentation”. In: *Z. Phys. C* 20 (1983), p. 317. DOI: 10.1007/BF01407824.

-
- [67] A. Tumasyan et al. “CMS pythia 8 colour reconnection tunes based on underlying-event data”. In: *The European Physical Journal C* 83.7 (July 2023). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-11630-8. URL: <http://dx.doi.org/10.1140/epjc/s10052-023-11630-8>.
- [68] Artur Kalinowski and Wiktor Matyszkiewicz. *Efficient tau-pair invariant mass reconstruction with simplified matrix element techniques*. 2025. arXiv: 2509.26069 [hep-ph]. URL: <https://arxiv.org/abs/2509.26069>.
- [69] Lorenzo Bianchini et al. “Reconstruction of the Higgs mass in events with Higgs bosons decaying into a pair of τ leptons using matrix element techniques”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 862 (2017), pp. 54–84. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2017.05.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900217305259>.
- [70] K. Gottfried and John David Jackson. “On the Connection between production mechanism and decay of resonances at high-energies”. In: *Nuovo Cim.* 33 (1964), pp. 309–330. DOI: 10.1007/BF02750195.
- [71] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG]. URL: <https://arxiv.org/abs/1412.6980>.
- [72] S. van der Meer. “Calibration of the Effective Beam Height in the ISR”. In: (June 1968).
- [73] Vladislav Balagura. “Van der Meer scan luminosity measurement and beam–beam correction”. In: *Eur. Phys. J. C* 81.1 (2021), p. 26. DOI: 10.1140/epjc/s10052-021-08837-y. arXiv: 2012.07752 [hep-ex].
- [74] A. Hayrapetyan et al. “The CMS Statistical Analysis and Combination Tool: Combine”. In: *Computing and Software for Big Science* 8.1 (Nov. 2024). ISSN: 2510-2044. DOI: 10.1007/s41781-024-00121-4. URL: <http://dx.doi.org/10.1007/s41781-024-00121-4>.
- [75] Roger Barlow and Christine Beeston. “Fitting using finite Monte Carlo samples”. In: *Computer Physics Communications* 77.2 (1993), pp. 219–228. ISSN: 0010-4655. DOI: [https://doi.org/10.1016/0010-4655\(93\)90005-W](https://doi.org/10.1016/0010-4655(93)90005-W). URL: <https://www.sciencedirect.com/science/article/pii/001046559390005W>.