Search for Exotic Decays of the 125 GeV Higgs Boson in the bb$\mu\mu$ Final State with the ATLAS Detector

A Dissertation presented

by

Christopher Robyn Hayes

to

The Graduate School

in Partial Fulfillment of the

Requirements

for the Degree of

Doctor of Philosophy

in

Physics

Stony Brook University

August 2019

Stony Brook University

The Graduate School

Christopher Robyn Hayes

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation

John Hobbs - Dissertation Advisor
Professor, Physics and Astronomy

Rouven Essig - Chairperson of Defense
Associate Professor, Physics and Astronomy

Abhay Deshpande - Internal Committee Member
Professor, Physics and Astronomy

George Redlinger - External Committee Member
Physicist, Department of Physics, Brookhaven National Lab

This dissertation is accepted by the Graduate School

Eric Wertheimer
Dean of the Graduate School

Abstract of the Dissertation

Search For Exotic Decays of the 125 GeV Higgs Boson in the bb$\mu\mu$ Final State With the ATLAS Detector

by

Christopher Robyn Hayes

Doctor of Philosophy

in

Physics

Stony Brook University

2019

Exotic decays of the 125 GeV Higgs boson provide a unique window for the discovery of new physics, as the Higgs may couple to hidden-sector states that do not interact under Standard Model (SM) gauge transformations. Models predicting exotic Higgs decays to additional light bosons appear in many extensions to the SM and can explain several unknowns in high energy particle physics, such as the nature of dark matter and the existence of supersymmetry. In this Dissertation, a search for the exotic decay of the 125 GeV Higgs boson is performed using proton-proton collisions from the Large Hadron Collider at $\sqrt{s} = 13$ TeV with 36.1 fb^{-1} of data collected by the ATLAS detector. In this search, we only consider the exotic decays to the final state $H \to bb\mu\mu$. The experimental signature consists of two oppositely signed muons, two b-tagged jets, and a negligible amount of missing transverse energy. Further, the two b-jet and dimuon systems are required to have nearly equivalent masses and the four-object mass of these particles must be within 15 GeV of 125 GeV. These constraints are placed on the system via a kinematic likelihood fit that adjusts the energy of the b-jets and assigns an event scores based on the compatibility of the hypothesis $m_{bb} \approx m_{\mu\mu}$. The main background contributions come from the Standard Model Drell-Yan + jets production and $t\bar{t}$ decays with two

muons; minor background contributions from W+jets, single-top, and $t\bar{t}+X$ processes are also considered. Monte Carlo simulation samples are used to estimate each background contributions, except for the Drell-Yan background, which is estimated using a data-driven template method. Experimental and theoretical uncertainties from a variety of sources are assessed for the signal and background estimations. As no significant excesses are found in the signal region selections, upper limits on the production cross-section times branching ratio $\frac{\sigma_H}{\sigma_{\mathrm{SM}}} \times \mathcal{B}(H \to aa \to bb\mu\mu)$ which range from 1.2×10^{-4} to 8.4×10^{-4} in the a-boson mass range, $20 \leq m_a \leq 60$ GeV. Additional model-independent limits, where no signal model is considered, are set on the visible production cross-section times the branching ratio of a new physics particle X to the $bb\mu\mu$ final state, $\sigma_{\mathrm{vis}}(X) \times \mathcal{B}(X \to bb\mu\mu)$, ranging from 0.1 fb to 0.73 fb in the dimuon mass range $18 \leq m_a \leq 62$ GeV. This analysis provides stringent limits on $H \to aa$ decays with comparable or greater sensitivity to other search channels performed on LHC data. Finally, we discuss the limiting factors of the 36.1 fb^{-1} analysis and present new experimental avenues needed to establish a greater chance for discovery potential using the full LHC Run-2 dataset.

Dedication Page

I dedicate this work to my family, both the Wolfe and Hayes-Marek clans. Without their help and encouragement, I would never have made it this far in my academic career.

Contents

List of Figures/Tables/Illustrations

List of Figures

List of Tables

List of Abbreviations

2HDM + S: Two Higgs Doublet Models with and additional Singlet
AFP - ATLAS Forward Proton, a subdetector in ATLAS ALFA - Absolute Luminosity for ATLAS, a subdetector in ATLAS
ALICE - A Large Ion Collider Experiment, one of the main experiments at the Large Hadron Collider
ATLAS - A Toroidal LHC ApparatuS, one of the main experiments at the Large Hadron Collider
BDT - Boosted Decision Tree
BR - Branching Ratio
BSM - Beyond the Standard Model
BW - relativistic Breit-Wigner distribution
CERN - European Center for Nuclear Research) CKM matrix - Cabibbo-Kobayashi-Maskawa matrix
CMS - Compact Muon Solenoid, one of the main experiments at the Large Hadron Collider
CP - Charge-Parity
CSC - Cathode Strip Chambers, a muon subdetector technology in ATLAS
CTP - Central Trigger Processor
DB - Direct Balance
DIS - Deep Inelastic Scattering
DQ - Data Quality
DY - Drell-Yan Process
DYCR - Drell-Yan Control Region
EM - Electromagnetic
EMB - Electromagnetic Barrel Calorimeter, a subdetector in ATLAS
EMEC - Electromagnetic Endcap Calorimeter, a subdetector in ATLAS
EW - Electroweak, with respect to the unified interactions of the electromagnetic and weak forces
EWK - Same as EW.
EWSB - Electroweak symmetry breaking
FCAL - Forward Calorimeter, a subdetector in ATLAS
FCNC - Flavor Changing Neutral Currents

FSR - Final State Radiation
GRL - Good Run List
GSF - Gaussian Sum Filter
HEC - Hadronic Endcap Calorimeter, a subdetector in ATLAS
HLT - High Level Trigger
IBL - Insertable B-Layer, a silicon detector in ATLAS
ID - Inner detector
IP - Impact Parameter
IP - Interaction Point
IR - Infrared
ISR - Initial State Radiation
JER - Jet Energy Resolution
JES - Jet Energy Scale
JVF - Jet Vertex Fraction
JVT - Jet Vertex Tagger
KL - Kinematic Likelihood
KLFit - Kinematic Likelihood Fit
L1 - Level-1 Trigger (hardware trigger of ATLAS)
L1A - Level-1 Trigger Accept
L1Topo - The L1 topological trigger in ATLAS
LAr - Liquid Argon
LH - Likelihood
LHC - Large Hadron Collider
LHCb - Large Hadron Collider beauty experiment, one of the main experiments at the Large Hadron Collider
LINAC - Linear Accelerator
LO - Leading Order
LUCID - LUminosity Cherenkov Integrating Detector
MC - Monte Carlo, with respect to Monte Carlo simulation methods
MDT - Monitored Drift Tubes, a muon subdetector technology in ATLAS
MPF - Missing Transverse Energy Projection Fraction
MS - Muon spectrometer
MSSM - Minimal Supersymmetric Models
MV2 - Multivariate Tagger 2, a b-tagging algorithm for ATLAS.
NLO - Next-to-Leading order
NNLO - Next-to-Next-to-Leading Order
N3LO - Next-to-Next-to-Next-to-leading order
NMSSM - Next-to-Minimal Supersymmetric Models

PDF - Parton Distribution Function
PDF - Probability Density Function
PFO - Particle Flow Objects
PS - Proton Synchrotron
PS - Parton Shower
PSB - Proton Synchrotron Booster
PV - Primary Vertex
QED - Quantum Electrodynamics
QCD - Quantum Chromodynamics
RF - radio-frequency
RMS - Root Mean Square or Variance
ROI - Region Of Interest
RPC - Resistive Plate Chambers, a muon subdetector technology in ATLAS
SCT - Semiconductor Tracker, a subdetector in ATLAS
SF - Scale Factor
SM - Standard Model of Particle Physics
SPS - Super Proton Synchrotron
SR - Signal Region
SU(n) - Special Unitary group of degree n
SV - Secondary Vertex
TCR - Top Control Region
TGC - Thin-Gap Chambers, a muon subdetector technology in ATLAS
TR - Template Region
TRT - Transition Radiation Tracker, a subdetector in ATLAS
TTC - Trigger, Timing, and Control
TTVA - Track-to-Vertex Association
U(n) - Unitary group of degree n
UV - Ultraviolet
VBF - Vector Boson Fusion
VEV - Vacuum Expectation Value
VH - Associated production of the Higgs boson (H) with a vector boson (V)
VR - Validation Region
ZDC - Zero Degree Calorimeter

Acknowledgements

As my graduate career comes to a close, I must recognize all of the individuals who aided me along the way.

My advisor, John Hobbs, has been a constant beacon of support in my academic career. He has fostered my innate curiosity for the physical world and has always challenged me to succeed in any new idea I have had. He has also helped me navigate how to succeed while being in a large collaboration. I have enjoyed having such a warm and supporting advisor who truly believes in the goals of high-energy particle physics: to discover and learn more about our unknown universe.

I must thank Lily Morvaj for her constant support during my Ph.D. career. She has been not only a fantastic research partner, but also a fantastic friend. She has kept me from straying too far off course in research, simply saying many times "šuti i radi". I thank her for her patience, her kindness, and her concerns for my general well-being throughout my research career.

I must also thank Ben Carlson for inspiring me to pursue new and interesting research paths during my graduate research. He's also been a wonderful friend and I have enjoyed hiking the Jura and Alps with he and his wife Jamie (and later PJ as well!). I wish their family the best and know that our paths will cross again!

My friends at Stony Brook helped me navigate through courses (and life!) at the beginning of my graduate career. They have given me a place to eat and rest when I would visit the United States from CERN. In particular, I must thank Joe, Stacy, Lupe, Nicki, Kevin, and Megan for being fantastic friends. I have missed you guys dearly while I've been across the ocean.

I must also thank the rest of the Stony Brook research group for their kindness throughout my PhD, in particular Giacinto, Michael, Dmitri, and Thomas. Also without Tracy's fantastic support, I think my life in Franco-Switzerland would have been a lot harder!

Last, but certainly not least, I must thank my dearest companion, my wife Megan Wolfe. Without her, I would not have gone down the path to pursue my dream of working in high-energy physics. She's been there to help in the lowest times, anticipating the celebration at the best of times. For her patience, love, and help I will be forever grateful.

1 Theoretical Background and Introduction

1.1 The Standard Model of Particle Physics

The Standard Model (SM) [1–4] of particle physics is an effective gauge field theory that describes the quantum mechanical nature of three fundamental forces of nature, namely the electromagnetic, strong, and weak forces. It further describes the interactions of the force carriers with the fundamental constituents of matter, quarks and leptons. The experimental evidence and verification of the SM has been tested to an exceptional precision. A general summary of the particle content of the Standard Model is given in Section 1.1.1. An overview of quantum electrodynamics (QED) is in Section 1.1.2. The theory of the strong nuclear interactions and color charge is explained in Section 1.1.3. The electroweak (EW) theory that describes the high-energy interactions that unify the electric and weak force together is described in Section 1.1.4. Finally, the Higgs mechanism that leads to electroweak symmetry breaking and mass generation is detailed in Section 1.1.5, with further details about the couplings of the Higgs boson given in Section 1.1.6.

1.1.1 Summary of the Standard Model

Before going into the details of the physical interactions described by the Standard Model, a general overview of the particle content of the Standard Model is discussed. The fundamental particles of matter consist of two classes of fermions, the quarks and the leptons, and four gauge bosons that act as mediators of the strong force and the electroweak force.

The leptons are spin-$\frac{1}{2}$ particles, and only participate in the electroweak interactions. There are six known leptons, which include the electron (e), muon (μ), tau(τ), electron neutrino(ν_e), muon neutrino (ν_μ), and tau neutrino (ν_τ). The mass and electrical charge of the leptons can be found in Table 1. Each of the charged leptons pairs with its corresponding neutrino to form a lepton generation that will undergo the electroweak force interactions (Section 1.1.4).

Quarks are also spin-$\frac{1}{2}$, electrically charged particles. The quarks take part in electroweak interactions and additionally participate in the "strong" interactions, which bind together nuclear particles such as the proton and neutron. There are

[1]Experimental evidence is abundant for non-zero neutrino masses, especially from flavor oscillations observed by several experiments [5]. The mass of the flavor state is constrained not only by the direct experimental limits given, but also from double-β decay if the neutrino is a Majorana particle) [6] and Cosmic Microwave Background constraints [7]. In this thesis, the difference between $m_\nu = 0$ and a small non-zero value is negligible.

Lepton Species	Mass [MeV]	Electric Charge
electron (e)	0.511	-1e
muon (μ)	105.6	-1e
tau (τ)	1776	-1e
electron neutrino (ν_e)	$< 2 \times 10^{-6}$ [1]	0
muon neutrino (ν_μ)	< 0.2 [1]	0
tau neutrino (ν_τ)	< 18.2 [1]	0

Table 1: The charge (in units of elementary charge) and mass of the six known leptons. The values for the charged lepton masses are taken from global fits performed in Ref. [5].

six known flavors of quarks, the up(u), down(d), strange(s), charm(c), bottom(b), and top(t), whose masses and electric charges are listed in Table 2. Each of the quark flavors falls into a general classification of up-type and down-type quarks, based on their fractional electromagnetic charge. The quarks are paired into three generations, with each up-type quark (Q_e = +2/3) paired with a down-type quark (Q_e = -1/3) partner. The generations follow a doublet structure, with the first generation composed of the lightest quarks (u, d), the second generation being the intermediate mass quarks (c, s), and the third generation being the most massive quark states (t, b).

The last pieces of the SM are the force carriers, consisting of the gauge bosons of the electroweak and strong force. These particles mediate the interactions of the forces. The photon (γ) mediates electromagnetic interactions, the $W^{\pm}$ and Z mediate the weak interactions (responsible for radioactive decay), and the gluons (g) are the mediators of the strong interaction. Finally, the Higgs boson provides the mechanism by which the W and Z bosons obtain their masses and explains the masses of the fermions via Yukawa interactions. The basic properties of these bosons are given in Table 3.

1.1.2 Quantum Electrodynamics

Quantum electrodynamics (QED) is the quantum theory used to describes the electromagnetic force, which describes charged particle interactions mediated by the

Quark Species	Mass [GeV]	Electric Charge
up (u)	$2.2^{+0.5}_{-0.4} \times 10^{-3}$	+ 2/3 e
down (d)	$4.7^{+0.5}_{-0.3} \times 10^{-3}$	- 1/3 e
strange (s)	$93^{+9}_{-3} \times 10^{-3}$	- 1/3 e
charm (c)	$1.28\,^{+0.03}_{-0.04}$	+ 2/3 e
bottom (b)	$4.18\,^{+0.04}_{-0.03}$	- 1/3 e
top (t)	173.0 ± 0.4	+ 2/3 e

Table 2: The charge (in units of fractional elementary charge) and mass of the six known flavors of quarks. The values for the mass are taken from global fits performed in Ref. [5].

Boson	Mass [GeV]	Electric Charge	Spin
photon (γ)	0	0	1
gluons (g)	0	0	1
$W^{\pm}$	80.379±0.012	±e	1
Z^0	91.1876±0.0021	0	1
H	125.18 ±0.16	0	0

Table 3: The spin, electric charge, and mass of the Standard Model gauge bosons and Higgs boson. In additional to the other properties, the gluons are also charged under the strong interaction SU(3) color charge. The values for the W, Z, and H bosons are taken from global fits performed in Ref. [5].

photon, γ. The QED Lagrangian takes the form [2]

$$\mathcal{L}_{\text{QED}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \bar{\Psi}(i\not{D} - m)\Psi \tag{1}$$

where $F_{\mu\nu}$ is the electromagnetic tensor, Ψ denotes the spinor field for fermions (spin-$\frac{1}{2}$ particles), m denotes the mass of the field Ψ, and $\not{D}$ is the gauge covariant derivative given in Feynman slash notation which contracts the γ matrices with a four-vector. Explicitly, these terms can be expanded

$$\begin{aligned} F_{\mu\nu} &= \partial_\mu A_\nu - \partial^\nu A^\mu \\ \not{D} &\equiv \gamma^\mu D_\mu \\ D_\mu &= \partial_\mu + iQ_e A_\mu \end{aligned} \tag{2}$$

where A_μ denotes the electromagnetic four-vector potential and Q_e denotes the charge of the fermion field.

The QED Lagrangian is gauge invariant under the transformations

$$A_\mu(x) \to A_\mu(x) - \partial_\mu \Gamma(x) \tag{3}$$

$$\Psi(x) \to e^{\pm iQ_e\Gamma(x)}\Psi(x) \tag{4}$$

where $\Gamma(x)$ is an arbitrary function of the spacetime coordinates. Here, the gauge invariance has been given as a local gauge invariance and this gauge invariance insures that the electromagnetic fields conserve electric charge. In the fermion field transformation given in Eq. 4, the fermions explicitly obey gauge invariance under a unitary transformation, specifically a U(1) transformation. Due to this feature, QED is referred to as a U(1) gauge theory.

The electromagnetic field Lagrangian in Eq. 1 does not have a mass term for the A_μ potentials, implying that the Lagrangian assumes that the photon is massless. The corresponding term for a massive photon is given by

$$\mathcal{L}_{\text{photon}} = \frac{1}{2}m_\gamma^2 A_\mu A^\mu \tag{5}$$

where m_γ is the associated mass of the photon. This additional term in the Lagrangian is problematic, as the U(1) transform on this term

$$A_\mu A^\mu \to (A_\mu - \partial_\mu\Gamma)(A^\mu - \partial^\mu\Gamma) \neq A_\mu A^\mu \tag{6}$$

breaks the local gauge invariance. The addition of this term is unneeded as the photon is known to be massless to a high level of precision [8].

[2]For all mathematical formulae, Einstein notation will be used to indicate that sums should be applied over the indices. The Greek and Latin alphabet will be used in the standard way for space-time components and spatial components respectively. For higher dimensionalities, Greek indices will be used. Further, all formula are given in *natural units*, i.e. $c = \hbar = 1$.

1.1.3 Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the quantum theory that describes the nuclear strong force. This force describes the interactions of the strong force mediators, known as gluons, and the fermionic quarks[3]. The quark fields are organized based on their color charge into fields of the form

$$\Psi = \begin{pmatrix} \psi_{\text{red}} \\ \psi_{\text{green}} \\ \psi_{\text{blue}} \end{pmatrix} \tag{7}$$

where each ψ_{color} corresponds to a particular color charge state.

The QCD Lagrangian, having both gauge and fermionic components, looks deceptively similar to the QED Lagrangian given in Eq. 1,

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} F^a_{\mu\nu} F^{\mu\nu}_a + \bar{\Psi}^i (i \not{D} - m\delta_{ij}) \Psi_j \tag{8}$$

where $F^a_{\mu\nu}$ is the strong force field strength tensor, Ψ_i denotes the quark fields which take the form given in Eq. 7, δ_{ij} is the Kronecker delta function, and D_μ is the covariant derivative of the strong force. The major changes, however, come about due to the large number of interactions in the QCD Lagrangian. The Latin indices on the field tensor has 8 components, where each index corresponds to a specific gluon field. The covariant derivative and field strength tensor, therefore, become more complicated than the QED partners due to the various interactions between all of the gluon fields. The field strength and covariant derivatives are given by

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu - g_{\text{s}} f^{abc} A^b_\mu A^c_\nu \tag{9}$$

$$D_\mu = \partial_\mu + i g_{\text{s}} A_\mu \tag{10}$$

where the A^a_μ terms describe the gluon field potentials, and f^{abc} denotes the structure constants of the SU(3) symmetry of the strong interactions.

The gluon fields can also be described in terms of the color charge, with states denoted by

$$r = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \qquad g = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \qquad b = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

[3]It should be noted that the quarks also participate in the electromagnetic interactions, with the interaction strength scaled to their fractional charge.

where the states are called red, green, and blue respectively.

The gluon octet states in terms of these vectors take the form

$$\begin{aligned} A^a &= \Big\{ \frac{1}{\sqrt{2}}(r\bar{g}+g\bar{r}), \frac{-i}{\sqrt{2}}(r\bar{g}-g\bar{r}), \frac{1}{\sqrt{2}}(r\bar{r}+g\bar{g}), && (11)\\ & \frac{1}{\sqrt{2}}(r\bar{b}+b\bar{r}), \frac{-i}{\sqrt{2}}(r\bar{b}-b\bar{r}), \frac{1}{\sqrt{2}}(g\bar{b}+b\bar{g}), && (12)\\ & \frac{-i}{\sqrt{2}}(b\bar{g}-g\bar{b}), \frac{1}{\sqrt{6}}(r\bar{r}+g\bar{g}-2b\bar{b}) \Big\} && (13) \end{aligned}$$

These states are naturally related to the generators of the SU(3) symmetry group, T^a. The exact field relations are called the Gell-Mann matrices λ^a and in the standard basis, $T^a = \frac{1}{2}\lambda^a$, where the λ matrices are given by

$$\begin{aligned} \lambda_1 &= \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \\ \lambda_4 &= \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \lambda_5 = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \lambda_6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \\ \lambda_7 &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \lambda_8 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix} \end{aligned} \quad (14)$$

Based on all of the field relations, the QCD Lagrangian is symmetric under the following local transformations

$$\begin{aligned} U(x) &= e^{i\alpha^a(x)T^a} \\ \Psi_i &\to U\Psi_i = \Psi_i + i\alpha^a(x)T^a_{ij}\Psi_j \\ A^a_\mu &\to UA^a_\mu U^\dagger - \frac{1}{g_s}(\partial_\mu U)U^\dagger = A^a_\mu - \frac{1}{g_s}(\partial_\mu \alpha^a(x)) - f^{abc}\alpha^b(x)A^c_\mu \end{aligned} \quad (15)$$

The symmetry of the QCD Lagrangian under the unitary transforms U leads to color conservation, meaning that all interactions in QCD must preserve the color charge. As gluons carry both a color and anti-color charge, there are several gluonic interactions that lead to color preservation, including the three gluon and four gluon self-interactions.

The strong force, in addition to the mathematical structures given here, has a peculiar property that affects the observation of quarks and gluons in high-energy collisions. This property, called confinement, implies that color charge carrying objects

cannot be observed directly. Therefore, all of the quarks and gluons are confined into color neutral matter. This property leads to the hadronization of quarks[4], forming baryons or mesons, which creates the jet objects (Section 3.6) detected in particle collisions.

Finally, in preparation for the electroweak theory in Section 1.1.4, the quarks can be organized into left and right-handed components based on their flavor eigenstates as

$$\Psi_i^L = \begin{pmatrix} u^L \\ d'^L \end{pmatrix}, \begin{pmatrix} c^L \\ s'^L \end{pmatrix}, \begin{pmatrix} t^L \\ b'^L \end{pmatrix} \quad \Psi_i^R = \left(u^R, d^R, s^R, c^R, b^R, t^R \right) \tag{16}$$

where the d',s',b' fields are the doublet partners to the up-type quarks in the electroweak doublets. These fields are not the same fields as d, s, b, instead being related via a unitary transformation $Q' \rightarrow V_{ij}Q$ where V_{ij} is the Cabibbo-Kobayashi-Maskawa (CKM) matrix given by

$$V_{ij} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \tag{17}$$

This relation will be implicitly used in Section 1.1.4, where these fields are needed in mass basis rather than in flavor basis.

1.1.4 Electroweak Force and Symmetry Breaking

At high energies, the electromagnetic force described in Section 1.1.2 unifies with the weak interactions that are responsible for radioactive decays. The unification of these two forces is called the electroweak (EW) force. This force is mediated by several gauge bosons and at low energies must describe both QED and the weak interactions as independent processes. The symmetry breaking that leads to these low energy interactions, the Higgs mechanism, is described in Section 1.1.5.

Every fermion in the Standard Model interacts via the electroweak force. From low energy electroweak measurements, only the left-handed leptons are charged under the electroweak interaction. Therefore, any description of the electroweak interactions must address the issue of chirality.

As the fermions in the electroweak theory are spin-$\frac{1}{2}$ particles, a natural representation arises using the SU(2) gauge group. In this representation, the left-handed leptons form three SU(2) doublets given by

$$L_i = \begin{pmatrix} \nu_e^L \\ e^L \end{pmatrix}, \begin{pmatrix} \nu_\mu^L \\ \mu^L \end{pmatrix}, \begin{pmatrix} \nu_\tau^L \\ \tau^L \end{pmatrix} \tag{18}$$

[4]The top quark is an exception, as it can decay to a b-quark + W boson due to its short lifetime

where i = 1,2,3 for the number of generations. Due to the parity violation experimentally present in the weak interactions [9], the right-handed leptons are singlets under the SU(2) symmetry in the Standard Model. These singlets are given by

$$e_i^R = \{e^R, \mu^R, \tau^R\} \tag{19}$$

Of particular note is the lack of right-handed neutrinos. In the SM, right-handed neutrinos are not addressed as a right-handed neutrino has not been observed in nature [10]. Any right-handed neutrino would also be a singlet under the EW gauge group and would be represented by

$$\nu_i^R = \{\nu_e^R, \nu_\mu^R, \nu_\tau^R\} \tag{20}$$

The quark sector also participates in the electroweak interaction follows the chiral structure given in Equation 16. For convenience, these fields can be rewritten as

$$Q_i = \begin{pmatrix} u^L \\ d'^L \end{pmatrix}, \begin{pmatrix} c^L \\ s'^L \end{pmatrix}, \begin{pmatrix} t^L \\ b'^L \end{pmatrix} \tag{21}$$

and the right-handed quarks singlets are denoted as

$$u_i^R = \{u^R, c^R, t^R\} \tag{22}$$
$$d_i^R = \{d^R, s^R, b^R\} \tag{23}$$

In addition to the weak SU(2) interactions, the electromagnetic U(1) interactions are also accounted for in the electroweak gauge sector. Separating out the fermion portion and neglecting the Higgs boson terms (Section 1.1.5), the EW gauge Lagrangian is given by

$$\begin{aligned} \mathcal{L}_{\text{EW Gauge}} &= -\frac{1}{4}(W_{\mu\nu}^a)^2 - \frac{1}{4}B_{\mu\nu}^2 \\ W_{\mu\nu}^a &= \partial_\mu W_\nu^a - \partial_\nu W_\mu^a - g\varepsilon^{abc} W_\mu^b W_\nu^c \\ B_{\mu\nu} &= \partial_\mu B_\nu - \partial_\nu B_\mu \end{aligned} \tag{24}$$

where W_μ^a (a=1,2,3) describes the SU(2) fields for the weak interactions, B_μ describes the $\mathrm{U(1)_Y}$ field for the electromagnetic interactions, g denotes the weak interaction coupling, and ε^{abc} denotes the structure constants of SU(2). The covariant derivative of the combined SU(2)×U(1) gauge groups is defined as

$$D_\mu = \partial_\mu + igW_\mu^a T^a + i\frac{g'}{2} Y B_\mu \tag{25}$$

where $T^a = \frac{\sigma^a}{2}$ denotes the generators of SU(2), Y denotes the U(1) charge and is referred to as the weak hypercharge, while g' denotes the U(1) coupling strength. The weak hypercharge is related to the electric charge Q

$$Y = 2(Q - T^3) \tag{26}$$

where T^3 is the third component of the weak isospin. It should be noted that these gauge fields are massless in the current formula. These massless states, referred to as Goldstone bosons, have not been "eaten" to form the massive W and Z bosons via the Higgs Mechanism.

For the fermionic sector of the SM, the kinetic terms of the electroweak Lagrangian are given by

$$\begin{aligned}\mathcal{L}_{\text{EW fermion}} = \bar{L}_i(i\not{\partial} - g\not{W}^a T^a - \frac{g'}{2}Y_L\not{B})L_i + \bar{e}_i^{\mathrm{R}}(i\not{\partial} + g'Y_e\not{B})e_i^{\mathrm{R}} \\ + \bar{Q}_i(i\not{\partial} - g\not{W}^a T^a - \frac{g'}{2}Y_Q\not{B})Q_i + \bar{u}_i^{\mathrm{R}}(i\not{\partial} + g'Y_u\not{B})u_i^{\mathrm{R}} \\ + \bar{d}_i^{\mathrm{R}}(i\not{\partial} + g'Y_d\not{B})d_i^{\mathrm{R}}\end{aligned} \tag{27}$$

where the covariant derivative of the right-handed fields are explicitly missing the weak isospin terms. The mass terms for these fields will also be addressed in the next section, as mass terms of the form $\bar{\Psi}^L\Psi^R$ would explicitly break the SU(2) symmetry.

1.1.5 The Higgs Mechanism

The Brout-Englert-Higgs Mechanism [11–13], referred to for simplicity as the Higgs mechanism, is responsible for electroweak symmetry breaking. The Higgs scalar field added to the Standard Model consists of an SU(2) doublet

$$\phi = \begin{pmatrix}\phi^+ \\ \phi^0\end{pmatrix} \tag{28}$$

where ϕ^+, ϕ^0 are complex fields. The kinetic and potential terms of this doublet can be expressed in a Lagrangian as

$$\begin{aligned}\mathcal{L}_{\text{Higgs}} &= (D^\mu\phi)^\dagger(D_\mu\phi) - V(\phi) \\ V(\phi) &= -\mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2\end{aligned} \tag{29}$$

where the λ term describes the self-interaction of the complex scalar field and the μ term gives the mass term. If $\mu^2 < 0$, the energy minimum will only occur for $\phi = 0$

and the SU(2) symmetry of the Lagrangian is preserved. However, for $\mu^2 > 0$, then the potential now has additional minima for any non-zero value of ϕ, namely

$$(\phi^\dagger \phi)_{\min} = \frac{\mu^2}{2\lambda} \equiv v^2/2 \tag{30}$$

where v is referred to as the Higgs vacuum expectation value (VEV). As this solution relies on the combination $\phi^\dagger\phi$, there are an infinite number of ϕ states which satisfy it. A reparametrization of the complex scalar field as two real scalar fields can be done for mathematical convenience (the Kibble reparametrization)

$$\phi(x) = h(x) e^{i\theta_a(x)} \tag{31}$$

where h and θ represent real fields. By expanding around the minimum of the potential, ϕ can be written as

$$\phi = e^{i\theta_a \sigma^a / 2v} \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}}[v + h(x)] \end{pmatrix} \tag{32}$$

Here, the scalar field is expressed as a distance h(x) away from the minima at $v/\sqrt{2}$ and the phase term given by the θ_a components. By using the unitary gauge transformation $\phi \to e^{-i\theta_a \sigma^a 2v}\phi$, the scalar field simplifies to

$$\phi = \begin{pmatrix} 0 \\ \frac{1}{\sqrt{2}}[v + h(x)] \end{pmatrix} \tag{33}$$

Returning to the EW gauge Lagrangian in Equation 24, the additional terms from the new scalar field are given by

$$\mathcal{L}_{\text{EW Gauge}} = -\frac{1}{4}(W^a_{\mu\nu})^2 - \frac{1}{4}B^2_{\mu\nu} + (D^\mu\phi)^\dagger(D_\mu\phi) + \mu^2\phi^\dagger\phi - \lambda(\phi^\dagger\phi)^2 \tag{34}$$

where the covariant derivative is given in Equation 25. To illustrate the effects of adding in the Higgs field, the Higgs kinetic term can be expanded using the field relations from Equation 33

$$|D_\mu\phi|^2 = g^2\frac{v^2}{8}\left[(W^1_\mu)^2 + (W^2_\mu)^2 + (\frac{g'}{g}B_\mu - W^3_\mu)^2\right] \tag{35}$$

where these terms now have masses proportional to the Higgs VEV and the weak coupling constants. To aid in diagonalization of the mass terms, it is convenient to define new fields via the rotation around an angle θ_W (the Weinberg angle)

$$Z_\mu \equiv \cos\,\theta_{\rm W} \mathrm{W}^3_\mu - \sin\,\theta_{\rm W} \mathrm{B}_\mu \tag{36}$$

$$A_\mu \equiv \sin\,\theta_{\rm W} \mathrm{W}^3_\mu + \cos\,\theta_{\rm W} \mathrm{B}_\mu \tag{37}$$

where $\tan\theta_{\rm W} = \frac{\rm g'}{\rm g}$. The W fields can be further simplified by using a linear combination of the isospins $T^{\pm} = (T^1 \pm iT^2)$. Using the isospin and new field substitutions, the mass eigenstates are then given by

$$\begin{aligned} W_\mu^{\pm} &= \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \\ Z_\mu &= \frac{1}{\sqrt{g'^2 + g^2}}(gW_\mu^3 - g'B_\mu) \\ A_\mu &= \frac{1}{\sqrt{g'^2 + g^2}}(g'W_\mu^3 + gB_\mu) \end{aligned} \tag{38}$$

with the masses of the eigenstates defined by

$$\begin{aligned} m_W &= \frac{gv}{2} \\ m_Z &= \sqrt{g'^2 + g^2}\ \frac{v}{2} = \frac{gv}{2\cos\theta_{\rm W}} \\ m_A &= 0 \end{aligned} \tag{39}$$

therefore, the masses of the W and Z boson are not independent and rely on the weak coupling and the Higgs VEV. The final couplings that come out of this diagonalization include the electromagnetic coupling e which is given by

$$e = g\sin\theta_{\rm W} = {\rm g'}\cos\theta_{\rm W} \tag{40}$$

1.1.6 Higgs Couplings

The Higgs sector also gives rise to the mass of the quarks and leptons via Yukawa interactions. The Lagrangian describing these interactions is given by

$$\mathcal{L}_{\rm Yukawa} = -Y_e\bar{L}_i\phi e_j^R - Y_d^{ij}\bar{Q}_i\phi d_j^R + Y_u^{ij}\bar{Q}_i\tilde{\phi}u_j^R + h.c. \tag{41}$$

where Y^{ij} matrices encode the couplings between the Higgs and the fermion fields, and $\phi, \tilde{\phi}$ are the Higgs field and its conjugate. The conjugate Higgs field is given by

$$\tilde{\phi} \equiv i\sigma^2\phi^\dagger = \begin{pmatrix} \phi^{0\dagger} \\ \phi^- \end{pmatrix} \tag{42}$$

By transforming these Yukawa terms to the mass basis via the CKM matrix in Eq. 16, the Y_f^{ij} are diagonalized and the new diagonal terms are invariant under the SU(2) symmetries. This invariances comes as the left and right-handed mixing

terms, for example $L_e\phi e_e$ produce a singlet under the SU(2) symmetry. The Yukawa terms, after inserting the minimum Higgs field from Eq. 33, take the form

$$\mathcal{L}_{\text{Yukawa}} = \frac{g_i v}{\sqrt{2}}(\bar{f_i^L} f_i^R + f_i^L \bar{f_i^R}) + \frac{g_i}{\sqrt{2}}(\bar{f_i^L} f_i^R + f_i^L \bar{f_i^R})h \tag{43}$$

where g_i are the coupling constants between the Higgs and the fermion and all fermions have been combined into the fields f. The first term has the form of a mass term, where

$$m_i = \frac{g_i v}{\sqrt{2}} \tag{44}$$

and the second term is the coupling between the Higgs field and the fermionic fields. In general, the couplings of the scalar field are proportional to the masses of the particles, both fermions and the gauge bosons.

The g_i are dependent on the Higgs mass, as the Higgs decay width

$$\Gamma(H) = \Gamma(H \to VV) + \Gamma(H \to ff) + \Gamma(H \to HH) \tag{45}$$

can differ significantly as the Higgs increases in mass $m_H > 2m_W, 2m_Z, 2m_t$. With the discovery of a Higgs particle at 125 GeV [14, 15], these fermionic and gauge boson widths are complete from a theoretical point of view, as the width is fixed to $\Gamma \sim 4$ MeV.

1.2 Beyond the Standard Model

Despite the extensive experimental verification of the SM, there remain several limitations to the model. These include the lack of a quantum description of gravity and its unification with general relativity; the lack of a description for dark matter and dark energy that has been verified by astrophysical measurements to be most of the energy density in the universe [17]. Other shortcomings affect the Standard Model predictions themselves; these include the non-zero masses of the neutrinos which are predicted to be massless in the SM, the potential breaking of lepton universality in rare B-meson decays, and the Higgs mass hierarchy problem.

One potential theoretical solution that aims to alleviate several of these problems, and provide gateways into a particle description of dark matter, comes from the Two-Higgs Doublet Models (2HDMs). A general description of 2HDMs is presented in Section 1.2.1.

Non-standard model decays, referred to here as exotic decays, of the Higgs boson can provide sensitivity to new physics beyond the Standard Model (BSM). The Higgs

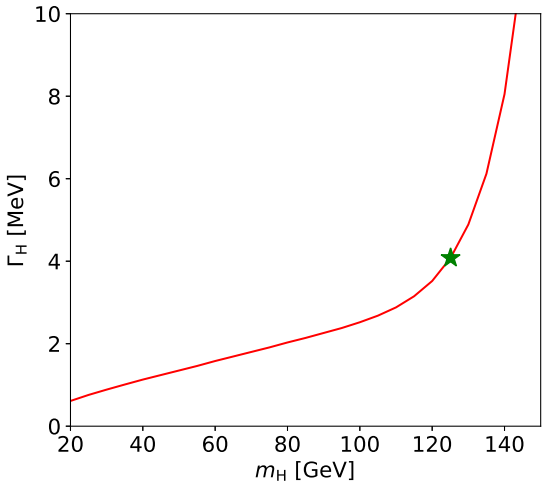

Figure 1: The Higgs decay width Γ_H as a function of the mass of the Higgs, m_H for a SM-like scalar Higgs. Adapted from Ref. [16].

boson at 125 GeV can provide an interesting portal to new physics for a variety of reasons. First, as this Higgs boson can couple via Yukawa couplings, the new physics content does not have to be charged under any SM gauge group.

Further, the experimental measurements of the Higgs sector leave a large parameter space open for the possibility of the Higgs to decay to BSM final states. The combination of the visible channels constrain BR(H→BSM)$< 22 - 47\%$ (depending on the model assumptions) using 80 fb^{-1} of data with the ATLAS experiment [18], while the CMS experiment reports similar limits for 35.9 fb^{-1} of data [19].

In the 2HDM models, the Higgs boson at 125 GeV is only one scalar among a family of fundamental scalars. In Section 1.2.1, the 2HDM with an additional singlet is discussed that gives rise to new light pseudoscalars that come from decays of the 125 GeV Higgs.

One important constraint to any model of new physics, comes from the electroweak measurements of the ρ-parameter. The ρ-parameter at tree-level in the SM is defined to be [20]

$$\rho_0 \equiv \frac{m_W^2}{m_Z^2 \cos^2\theta_W} = 1 \tag{46}$$

In the SM, this tree-level equality is guaranteed by the custodial SU(2) symmetry [21];

Parameter	Best Fit Value
ρ_0	1.00039±0.00019
S	0.02 ± 0.10
T	0.07 ± 0.12
U	0.00 ± 0.09

Table 4: The best fit values from experimental constraints on the ρ-parameter and the Peskin–Takeuchi parameters S,T, and U. Values are from global fits performed in Ref. [5].

therefore, any interaction which violates this SU(2) symmetry can modify the ρ-parameter in higher-order corrections. The Yukawa interactions in the SM give such corrections, notably in the large top Yukawa coupling. As modulations of the vacuum polarization of the gauge bosons can also affect the ρ-parameter, it is convenient to characterize any deviations by using the Peskin–Takeuchi parameters [22] S, T, U

$$\begin{aligned} \alpha_e T &= \frac{\Pi_{WW}^{\rm BSM}(0)}{M_W^2} - \frac{\Pi_{ZZ}^{\rm BSM}(0)}{M_Z^2} = \rho - 1 \\ \alpha_e S &= 4\,\sin^2\theta_{\rm W}\,\cos^2\theta_{\rm W}\,\frac{\Pi_{\rm ZZ}^{\rm BSM}({\rm M_Z^2}) - \Pi_{\rm ZZ}^{\rm BSM}(0)}{\rm M_Z^2} \\ \alpha_e (S+U) &= \frac{\Pi_{WW}^{\rm BSM}(M_W^2) - \Pi_{WW}^{\rm BSM}(0)}{M_Z^2} \end{aligned} \tag{47}$$

where $\Pi_{VV}^{\rm BSM}(q^2)$ denotes the change in the vacuum polarization for a vector boson at the given scale q^2. Electroweak precision measurements tightly constrain ρ_0, S, T, and U and the best-fit values are given in Table 4. Therefore, any new BSM interactions must be able to fit within these experimental constraints.

1.2.1 Extending the Higgs Sector: Two-Higgs-Doublet Model

In order to remain consistent with the experimental constraints on the ρ parameter, one possible extension to the Higgs sector adds a new Higgs SU(2) doublet to the SM. This extension can be done in a variety of ways, leading to a large class of models known as Two-Higgs-Doublet models (2HDMs). 2HDMs appear in several larger theoretical frameworks, such as in supersymmetric models [23, 24] and Peccei-Quinn axion models [25].

The most general potential in a 2HDM model is [26, 27]

$$\begin{aligned} V = & m_1^2|H_1|^2 + m_2^2|H_2|^2 + \frac{\lambda_1}{2}|H_1|^2 + \frac{\lambda_2}{2}|H_2|^2 + \lambda_3|H_1|^2|H_2|^2 \\ & + \lambda_4|H_1^\dagger H_2|^2 + \frac{\lambda_5}{2}\Big((H_1H_2)^2 + c.c.\Big) + m_{12}^2(H_1H_2 + c.c.) \\ & + (\lambda_6|H_1|^2(H_1H_2) + c.c.) + (\lambda_7|H_2|^2(H_1H_2) + c.c.) \end{aligned} \tag{48}$$

where m_i are the masses of each scalar doublet, H_i are the scalar doublet fields, and the λ_i terms represent the various (self-)couplings. The Higgs fields are assigned hypercharge values of ± 1 and each doublet can acquire a real VEV, v_1, v_2. The complex fields can be expanded around these minimum values

$$H_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} v_1 + H_{1,R}^0 + iH_{1,I}^0 \\ H_{1,R}^- + iH_{1,I}^- \end{pmatrix}, \quad H_2 = \frac{1}{\sqrt{2}}\begin{pmatrix} H_{2,R}^+ + iH_{2,I}^+ \\ v_2 + H_{2,R}^0 + iH_{2,I}^0 \end{pmatrix} \tag{49}$$

where there are now eight degrees of freedom, from the two complex (charged) fields and the four real fields. The diagonalization proceeds in a similar way to the single doublet SM case, with a rotation defined by the angle β which is defined by

$$\tan\,\beta = \frac{v_1}{v_2} \tag{50}$$

where $v_{1,2}$ define the VEVs of each Higgs doublet.

To preserve EWSB, one of the charged fields and one combination of the neutral fields $H_{1,2,I}^0$ are "eaten" to form the massive $W^\pm, Z^0$ gauge bosons. The other five degrees of freedom come from a charged scalar field, two neutral scalar fields, and a pseudoscalar field. These fields takes on mass eigenstates, $H^\pm, A, h, H$ [5], which are defined by

$$H^\pm = -H_{1,R}^\pm \sin\,\beta + H_{2,R}^\pm \cos\,\beta \tag{51}$$

$$A = H_{1,I}^0 \sin\,\beta - H_{2,I}^0 \cos\,\beta \tag{52}$$

$$\begin{pmatrix} h \\ H \end{pmatrix} = \begin{pmatrix} -\sin\,\alpha & \cos\,\alpha \\ \cos\,\alpha & \sin\,\alpha \end{pmatrix}\begin{pmatrix} H_{1,R}^0 \\ H_{1,R}^0 \end{pmatrix} \tag{53}$$

where α is the mixing angle between the CP-even weak eigenstates and the mass eigenstates and can take any value in the range $-\pi/2 \leq \alpha \leq \pi/2$. The SM Higgs boson then is in general given by the linear combination of h, H

$$H^{\rm SM} = h\,\sin(\alpha - \beta) - H\,\cos(\alpha - \beta) \tag{54}$$

[5] These Higgs boson states labels differ from the notation used by the ATLAS collaboration, where H is typically identified as the SM Higgs boson.

For the purposes of the discussions for exotic decays in this thesis, the charged Higgs is assumed to be heavy and does not impact the phenomenology to a large degree. The SM Higgs is also assumed to exist in the decoupling limit

$$\alpha \to \beta - \frac{\pi}{2} \tag{55}$$

such that it is uniquely identified as h in the 2HDM and the $m_H >> m_h$ such that H does participate in the interactions at the electroweak scale [28]. This limit is taken to avoid experimental constraints on the mass of the heavier Higgs boson, which approach the TeV scale in sensitivity.

These particle interactions can lead to large flavor changing neutral currents (FCNCs) at leading order, as the Yukawa interactions can take on the form [29]

$$\mathcal{L}_{\text{Yuk.2HDM}} = \lambda^1_{ij}\bar{\psi}_i\psi_j H_1 + \lambda^2_{ij}\bar{\psi}_i\psi_j H_2 \tag{56}$$

and the Yukawa terms cannot be simultaneously diagonalized. To avoid these FCNC interactions, additional Z_2 symmetries can be imposed by requiring the transform $H_1 \to H_1$ and $H_2 \to -H_2$. This Z_2 symmetry imposes that the fermions with the same quantum numbers cannot couple to more than one H field.

With all of these changes to avoid experimental limits, there are four possible avenues for the fermion couplings

- Type I 2HDM, where all of the fermions couple to H_2.
- Type II 2HDM, where the down-type right handed quarks d_R and the right-handed leptons e_R couple to H_1 and the right handed up-type quarks couple to H_2.
- Type III (or Type X) 2HDM, where the leptons couple to H_1 and the quarks couple to H_2.
- Type IV (or Type Y) 2HDM, which modifies the Type II by having the right-handed leptons couple to e_R couple to H_2.

These models have quite different properties, as different enhancements are possible in some models depending on the values of tan β. For example, in the Type III, the couplings to leptons can be enhanced for values of tan $\beta > 1$ and suppressed for tan $\beta < 1$.

To get further away from experimental constraints and provide a wide variety of exotic Higgs decays, an additional complex scalar singlet field of the form

$$S = \frac{1}{\sqrt{2}}(S_R + iS_I) \tag{57}$$

which only couples to the $H_{1,2}$ fields in the potential, is added to the 2HDM. This new field gives rise to two mass eigenstates: the CP-even scalar s which mixes with the h, H eigenstates and a CP-odd scalar a which mixes with the A eigenstate. Due to only coupling through the $H_{1,2}$ fields, there are no unique Yukawa terms for a or s bosons and the couplings to SM fermions is only through the small mixing with the SM Higgs. The exotic decays of the 125 GeV Higgs then take the form

$$\begin{aligned} h &\to aa \to f\bar{f}f'\bar{f}' \\ h &\to ss \to f\bar{f}f'\bar{f}' \\ h &\to Za \to f\bar{f}f'\bar{f}' \end{aligned} \tag{58}$$

where f, f' can be any SM fermion which satisfies $m_{a,s} > 2m_{f,f'}$. With these constraints and the choice of couplings from the SM Higgs, a variety of decays of the a, s states are possible.

As an aside, the additional singlet in these models can also be found in larger theoretical frameworks. For example, the Next to Minimal supersymmetric models (NMSSM) can include an additional pseudoscalar a which comes from its singlino (singlet) state [24].

1.3 Searching for Exotic Decays of the 125 GeV Higgs

To search for these new pseudoscalars in decays of the Higgs boson, consideration needs to given to both the production mode of the 125 GeV Higgs boson and the decays of the new a-bosons, as the produced events have to be experimentally detectable. Several of the largest production modes of the Higgs boson at 13 TeV are discussed in Section 1.3.1 and the decays of the a boson are briefly discussed in Section 1.3.2, with a primary focus on the experimental strengths of decays to $bb\mu\mu$.

1.3.1 Higgs Production and Cross Section

At the Large Hadron Collider (LHC), the Higgs boson can be produced in several different production mechanisms, as the proton supplies both quark and gluon initial states. The largest production mechanism involves a gluon-gluon interaction in the initial state, known as gluon fusion. This mechanisms involves a loop diagram, as the massless gluon does not couple directly to the Higgs boson. The cross section of this process, therefore, is dominated by the heaviest quark loops, from the top and bottom quarks, but the top quark loops dominates as $m_t >> m_b$. An interesting feature of the gluon-gluon fusion production mode is that the momentum of the Higgs is quite "soft", with the Higgs produced predominantly at (or near) rest [30].

Therefore, the particles from Higgs decays also have relatively low momentum which makes experimental detection of hadronic decays difficult.

The next largest contribution arises from the quark-anti-quark initiated diagram leading to vector boson fusion (VBF). In this process, a vector boson is radiated from a quark line in each proton and then collide to create the Higgs boson via the λ_{HVV} coupling vertex. This production mode is characterized by having a large separation between the two quarks, which leads to a large invariant mass between the quark pairs m_{qq}. This invariant mass provides a useful experimental handle for any search targeting this production mode.

Rounding out the $\mathcal{O}$(pb) cross sections is the quark-initiated diagram leading to the associated production of a Higgs with a vector boson, referred to either as the Higgs-Strahlung process or the VH production mode. In this production mode at the LHC, a quark and and an anti-quark in the initial state collide to form a virtual vector boson which then radiates a Higgs boson. The relative suppression of the VH production modes comes about due to the presence of an anti-quark in the initial state, which in the parton distribution function (Section 1.4) for a proton collider is smaller than the corresponding quark production. The leading order Feynman diagrams for these three processes are given in Figure 2.

Although outside of the scope of this thesis, the production of the Higgs boson in association with a pair of heavy flavor (top or bottom) quarks are also produced at the LHC. The production cross sections, however, are much smaller for a 125 GeV Higgs, as can be seen in the comparison of cross sections in Figure 3. However, as these production modes have direct coupling with the top and bottom quarks, they can provide sensitive tests for additional Higgs boson partners that could preferentially couple to the top or bottom quark [31, 32].

1.3.2 Exotic Decays to $bb\mu\mu$

The largest branching fraction of these exotic Higgs decays for $m_a > 3$ GeV is typically to all-hadronic final states of heavy flavor quarks or tau leptons. An example of the branching ratios for the a is given illustrated in Figure 4. For these decay modes, the ability to select these events in experimental detector triggers becomes quite difficult.

For a general search targeting all of the production modes of the 125 GeV Higgs boson, decays which contain large leptonic (e, μ) branching ratios become necessary. Of these two, the large mass of the muon implies a decay branching fraction that is much larger than the branching fraction to the electron. This results in more events in the muon final state for exotic Higgs searches.

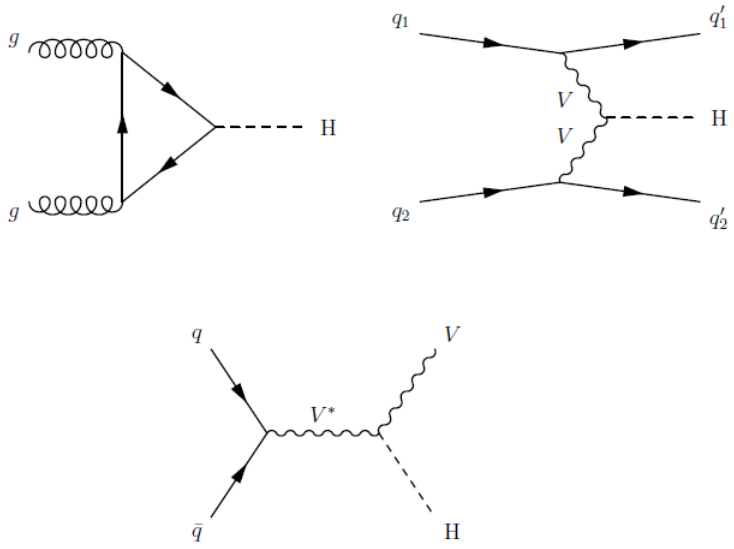

Figure 2: The leading order Feynman diagrams for the gluon-fusion, vector-boson fusion, and associated production modes of the Higgs boson at the LHC.

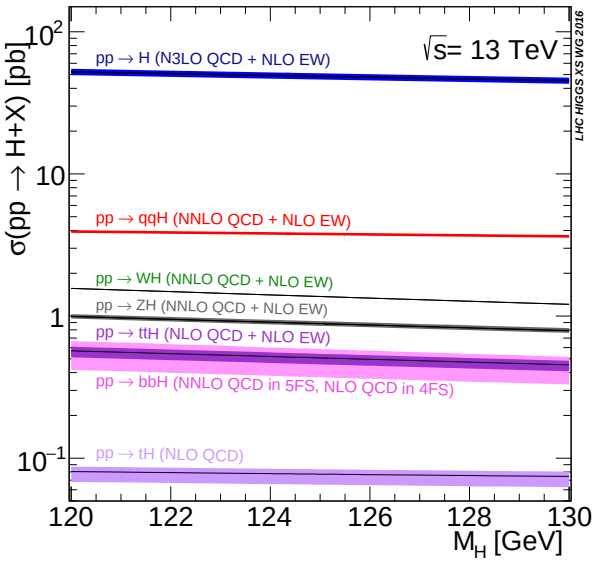

Figure 3: The largest production cross sections for a Higgs boson with a mass between 120 and 130 GeV at 13 TeV center of mass energy at the LHC. Adapted from Ref. [16].

In addition to the triggering capabilities, 4-body decays of a 125 GeV particle lead to an overall low energy being given to the decay products. In these regimes, the jet energy scale (discussed in Section 4.7.3) for the LHC experiments is not well understood and comes with a large uncertainty due to experimental effects. This scale uncertainty provides an experimental challenge, due to the difficulty in distinguishing signal jets from the much larger QCD-initiated background. Therefore, the use of low momentum leptons can also aid in determining the correct scale of the hadronic final state products.

To maximize the decay rate which enhances experimental detectability, the bb$\mu\mu$ decay mode is considered here. This decay mode is significant in most 2HDM+S models, but the largest BR(aa$\rightarrow$ bb$\mu\mu$) is found in the Type-III 2HDM+S where the decays of the a to leptons are favored. An illustration of such a 2HDM+S is given in Figure 4. The total $H \rightarrow aa \rightarrow bb\mu\mu$ in this model is given by 1.6×10^{-3}, assuming that Br$(H \rightarrow aa) = 100\%$.

The versatility of this production mode to provide a trigger using the muon side, the branching fraction of the b-quark decay, and the final discriminant of the $m_{\mu\mu}$ resolution are all attractive features for the search presented in this thesis.

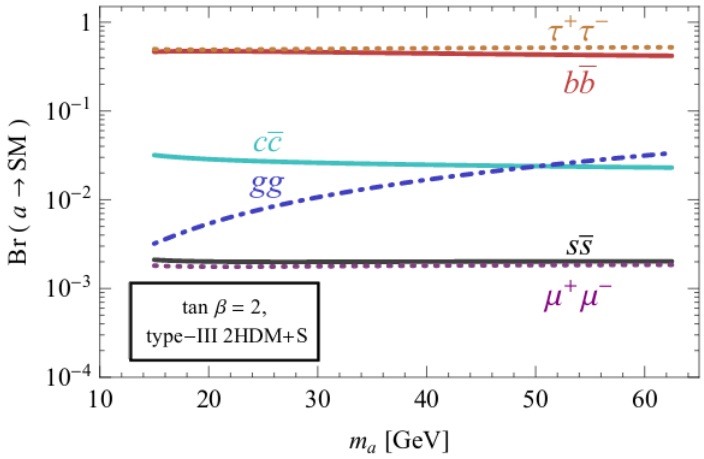

Figure 4: The branching ratios of a CP-odd a-boson in the range $14 \leq m_a \leq 62$ GeV, using predictions from a Type-III 2HDM+S model with $\tan\beta = 2$. Figure adapted from Ref. [33].

1.4 Event Generation and Monte Carlo Simulation Methods

Despite the elegant nature of the SM quantum field theory, a full description of a proton-proton collision at the LHC cannot be modeled using only these predictions. In order to fully describe the interactions and model any statistical variations, Monte Carlo (MC) methods are used to model the perturbative and non-perturbative effects arising from proton collisions. The many different processes that need to be modeled in order to describe a full proton-proton collision can be found in Figure 5.

In high energy collisions at the LHC, the colliding constituents are the quarks and gluons (the partons) inside the proton. The energy of the proton is therefore divided among these many constituents, each carrying a portion of proton's initial energy. The highest energy interactions between these partons form the *hard scatter* portion of a collision. The remaining partons can also interact and form additional, typically lower energy, collisions referred to as the *underlying event*.

The probability for a parton to carry a particular fraction of the proton momentum in an LHC collision is encoded in parton distribution functions (PDFs). PDFs are typically parametrized as a function of momentum transfer and are based on fits to experimental data, particularly using deep inelastic scattering (DIS) data and jet measurements performed at a variety of colliders. In the ATLAS simulation, a variety of PDF choices are used which incorporate different data sources and different

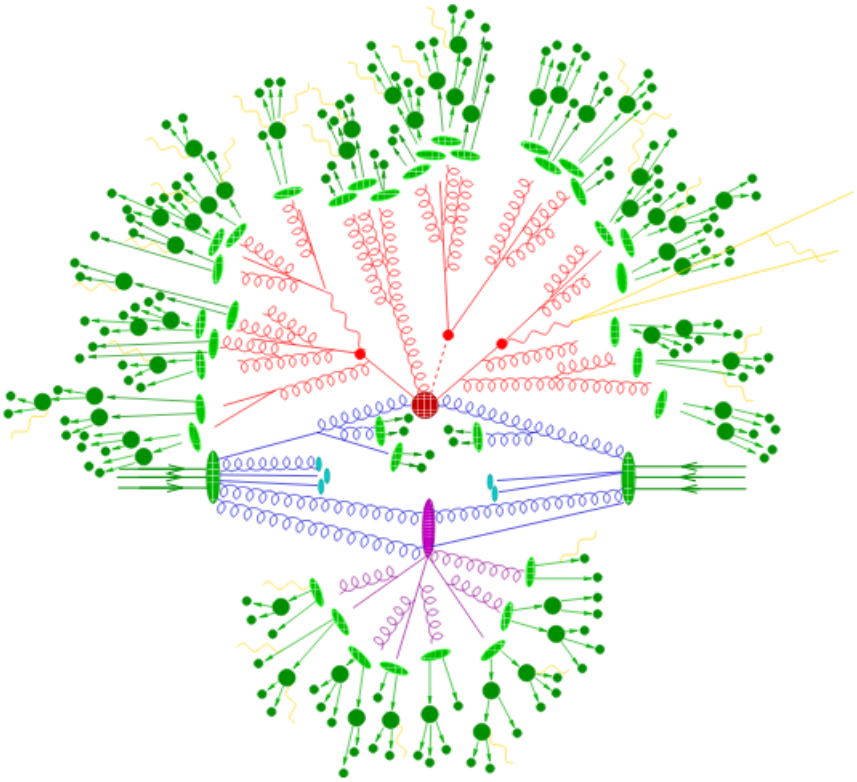

Figure 5: A representation of a typical proton-proton collision at the LHC. The hard scatter portion of the event (red circle) undergoes final state radiation and released many gluons (red lines). The gluons eventually shower into secondary particles which hadronize (light and dark green). The proton remnant, or underlying event interactions also interact and also yield hadronizing products (purple). Figure adapted from [34].

fitting techniques to obtain predictions in the leading order (LO), next-to-leading order (NLO), or next-to-next-to-leading (NNLO) order in the strong coupling constant α_s. A few examples of PDF sets relevant for this thesis include

- **The CTEQ PDF sets**: The CTEQ-TEA group provides PDF sets using a global-fit to QCD observables at high orders in perturbation theory. The CTEQ PDF sets used by ATLAS include fits up to NNLO in α_s of experimental inputs from HERA DIS measurements and Tevatron single-jet and vector boson measurements (CT10 [35]) or additionally including LHC Run-I data for the single-jet and W/Z processes (CT14 PDF set [36]).

- **The NNPDF sets**: The NNPDF collaboration determines PDF sets by converting experimental data into many MC ensembles based on the experimental uncertainties and uses neural network techniques to determine the best fit interpolation functions for the PDF prediction. The NNPDF set includes the largest number of inputs from LHC data, including both $t\bar{t}$ cross-section and W+charm data [37], and the PDFs are provided up to NNLO.

- **MSTW/MMHT PDF sets** The MSTW, and more recent MMHT14 PDF set [38], provides PDF predictions up to NNLO using a factorization method based on "leading-twist fixed-order collinear" theory in the $\bar{\text{M}}$S scheme. The most recent fits include LHC top-data and the use of Chebyshev polynomials to fit the input PDFs.

The hard scatter process in an LHC event is modeled by calculating the matrix elements, i.e. the cross-section σ for a given process. To maintain computational efficiency, these calculations are typically performed at fixed-order in perturbation theory. Most current MC event generation programs implement these calculations at LO or NLO.

In addition to the partonic constituents of the proton, radiation of additional gluons or photons can occur in proton collisions. In quantum field theory, these emissions are classified as coming from the initial state radiation (ISR) or from the final state radiation (FSR). These additional emissions are modeled by parton shower (PS) generators using splitting functions, such as the DGLAP equations, with specific phase-space cuts to avoid non-perturbative effects.

The underlying event describes the partonic interactions of the proton collisions that are not involved in the hard scatter process. These interactions are sensitive to initial-state and final-state radiation processes from the partons or additional hard scatter processes involving large momentum transfers. These contributions are often

characterized via direct measurements in low background data by looking at charged particle multiplicities and azimuthal variables [39].

In addition to the theoretical descriptions above, the effect of additional soft interactions from pileup and effects from background radiation in the experimental cavern are also necessary for an accurate description of events. Pileup event modeling often involves the generation of additional proton-proton collisions that are then overlaid on the hard-scatter process that is being generated. For the cavern and machine effects, special data-taking runs are used to collect minimum bias events. Minimum bias events are events where no hard sub-process from proton collisions is seen in the detector. These events are then inserted into the MC simulation to provide an accurate description of these low energy processes during hard scatter collisions.

Due to the varying levels of precision needed by analyses in high-energy particle physics, the various MC programs can be categorized in one of the following groups: matrix element calculation generators, parton shower programs, dedicated purpose programs, and multipurpose programs. A few examples of each type of Monte Carlo program can be found in Table 5.

Matrix element generators include any generator that provides matrix element calculations for the hard-scatter process and typically returns the four-vectors of the decaying particles from the hard-scatter. Specialized generators can sometimes be used to perform these calculations at higher orders, but the typical generators implement either LO or NLO modeling (in QCD or EWK corrections) in order to maximize computational efficiency. The predicted cross-sections from these generators are typically inconsistent with data or of limited precision that necessitate the MC events being weighted to a higher-order calculation to accurately describe the data. Most matrix element calculation programs are interfaced to another program to simulate parton showering and provide a full description of the simulated event.

Parton shower programs are used to simulate the soft gluon radiation coming from partons involved in the hard-scatter of the event and also simulate the hadronization of quarks. Special purpose generators are used to simulate specific physics processes. These include programs such as EvtGen [40], which is used for the decay of B-hadrons, or TAUOLA [41] which is used to simulate the polarization of τ lepton decay products. Multipurpose generators include any generator that can accomplish multiple functions, typically including the matrix element calculation and the parton showering under one program.

MC Generator Type	Example Programs
Matrix Element Generators	MadGraph (LO or aMC@NLO) [42] POWHEG-BOX [43, 44]
Multipurpose generators	Pythia [45], Sherpa [34, 46]
Special purpose generators	TAUOLA [41], EvtGen [40]

Table 5: Types and examples of Monte Carlo Generators used by the ATLAS collaboration. The POWHEG-BOX, Pythia, Madgraph, and Sherpa generators are used for the analysis described in this work.

2 Experimental Apparatus

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is a high-energy particle collider located at the European Center for Nuclear Research (CERN). The LHC comprises several thousand superconducting magnets, RF accelerating cavities, large scale cryogenics, and other support structures in an approximately 27 km circumference ring. Inside of this ring, two particle beams travel in opposite directions in two separate beam pipes which are kept in ultrahigh vacuum conditions. These beams cross at four interaction points (IPs), which house the ATLAS (A Toroidal LHC Apparatus), CMS (Compact Muon Solenoid), ALICE (A Large Ion Collider Experiment), and LHCb experimental detectors. ATLAS, described in greater detail in Section 2.3, and CMS are large detectors designed to precisely measure Standard Model physics decays and search for new physics phenomena. The ALICE and LHCb experiment are more specialized experiments, with ALICE's primary focus being the study of heavy ion collisions and LHCb being dedicated to physics involving the hadrons containing b-quarks. The positions of these experiments on the LHC ring can be seen in Figure 6.

After a two year shutdown, which was used for detector upgrades and preparation of the LHC magnets for higher energy running, the LHC began "Run-2" operation in 2015 at $\sqrt{s} = 13$ TeV. The LHC ran at this center of mass energy for the remainder of Run-2 which lasted until the end of 2018[6].

The LHC is only the final destination for these proton beams, which come from a larger accelerator complex that is detailed in Sec. 2.1.1.

2.1.1 Proton Accelerator Chain

In preparation for proton-proton collisions in the LHC, electrons are stripped from hydrogen atoms in a duoplasmatron located at one end of a linear accelerator known as LINAC2. The resulting protons are then accelerated to an energy of 50 MeV by radio-frequency (RF) cavities in LINAC2. Afterwards, the beam of protons is injected into the Proton Synchrotron Booster (PSB) which accelerates the protons to an energy of 1.4 GeV. The PSB consists of four independent rings and accomplishes this acceleration via multiple turns in these rings [47]. In the PSB, proton "bunches" begin to take shape as the protons are made to synchronize with the RF cavity acceleration frequency.

[6]Several special runs, including heavy ion collisions, also occurred at different energies, but are not relevant for this thesis.

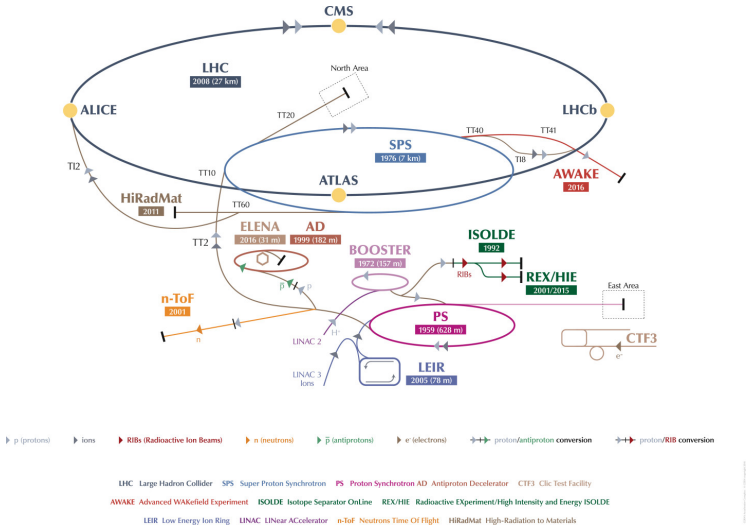

Figure 6: The accelerator complex at CERN. Protons destined for the LHC are first accelerated in LINAC2, then transferred to the PSB, PS, and SPS before finally reaching the LHC. In the LHC, these proton beams are accelerated to a maximum of 6.5 TeV per beam. Figure adapted from Ref. [49].

After the bunching of the protons and acceleration to 1.4 GeV, the proton beam is injected into the Proton Synchrotron (PS), where 8 bunches from two PSB injections are accelerated to 26 GeV. Concurrently with the acceleration, the 8 batches of protons are merged ("compressed") into four and then split into smaller bunches of protons based on the RF cavity harmonics in the PS. This splitting eventually produces 48 bunches with a spacing between the bunches of 25 ns. This bunch spacing defines final bunch spacing in the LHC and the production process is referred to as the BCMS scheme [48]. Several batches of 48 bunches are then injected into the Super Proton Synchrotron (SPS), where the beam is accelerated to 450 GeV for injection into the LHC. The proton bunches are then injected into the LHC via two transfer lines that inject protons in a clockwise and counter-clockwise direction.

Upon injection into the LHC, the proton beams undergo a ramp period that takes each proton beam from 450 GeV to the target 6.5 TeV. Then, the beams are brought closer until collisions occur in the four interaction points around the LHC experiments. Finally, the emittance and other beam parameters (Section 2.2) are adjusted to a targeted number of proton collisions desired by the experiments at each interaction point.

2.2 Luminosity and LHC Bunch Structure

Luminosity in high energy particle collisions refers to two distinct types of luminosity: instantaneous luminosity and integrated luminosity. For LHC collider experiments, the instantaneous luminosity is defined in terms of the inelastic cross section as elastic collisions leave the protons intact. The instantaneous luminosity is then given by the ratio

$$\mathcal{L}_{inst} = \frac{R_{\text{inelastic}}}{\sigma_{\text{inelastic}}} \tag{59}$$

where $R_{\text{inelastic}}$ is the rate of inelastic collisions and $\sigma_{\text{inelastic}}$ is the inelastic cross section of the particle collisions. In storage rings colliders like the LHC, this rate is

$$R_{\text{inelastic}} = n_{\text{b}} f_{\text{rev}} \mu \tag{60}$$

where f_{rev} is the revolution frequency, n_{b} is the number of colliding bunches, and μ is the average number of simultaneous inelastic interactions per bunch crossing. In terms of experimental detection, the events selected for luminosity determination must pass selection requirements; the detected luminosity is then given by

$$\mathcal{L}_{inst} = \frac{\epsilon \mu f_{\text{rev}} n_b}{\epsilon \sigma} = \frac{\mu_{\text{vis}} f_{\text{rev}} n_b}{\sigma_{\text{vis}}} \tag{61}$$

where ϵ denotes the detection efficiency for a given collision, μ_{vis} is the average number of interactions per bunch crossing passing a detection threshold, and σ_{vis} is visible inelastic cross section as determined by events passing detection thresholds. Here, ϵ refers to the efficiency of a particular detector and technique and is assumed to be the the same for the determination of μ_{vis} and σ_{vis}.

Alternatively, the instantaneous luminosity can also be determined from accelerator machine and beam parameters. For symmetric, Gaussian beams the luminosity can be expressed as [50]

$$\mathcal{L}_{inst} = \frac{n_{\text{b}} f_{\text{rev}} N_1 N_2}{4\pi \Sigma_x \Sigma_y} \tag{62}$$

where $\Sigma_{x,y}$ are the convolved beam size (e.g. $\Sigma_x = \sqrt{\sigma_{x1}^2 + \sigma_{x2}^2}$) in the horizontal and vertical directions and $N_{1,2}$ are the number of protons in each bunch. The convolved beams widths are measured using beam separation techniques known as van der Meer scans [51] during special LHC data-taking periods.

The integrated luminosity is related to $\mathcal{L}_{inst}$ via

$$\mathcal{L}_{int} = \int \mathcal{L}_{inst}(t)\ dt \tag{63}$$

Year	LHC Delivered Luminosity [fb^{-1}]
2015	4.2
2016	38.5
2017	50.2
2018	63.3
Run-2	156.2

Table 6: The total luminosity delivered by the LHC to the ATLAS experiment during Run-2. Values are derived from the ATLAS online luminosity determination and taken from Ref. [52].

The total integrated luminosity per data taking period delivered by the LHC to the ATLAS experiment is given in Table 6.

In the LHC, the protons beams are composed of proton bunches that are injected into the LHC as larger structures called "bunch trains". In each bunch train, the proton bunches are separated by a fixed gap in time called the bunch spacing. An illustration of the proton beam structure can be found in Figure 7. During Run-2, the LHC operated at its design bunch spacing of 25 ns between proton bunches for the majority of ATLAS 13 TeV data-taking periods. The LHC was also operated with a variety of filling schemes during the data-taking periods in Run-2, with the LHC having a maximum of 2556 bunches in the ring during high-luminosity fills. The number of bunches per injection during high luminosity running varied from 48 to 144 bunches.

With all of these parameters maximized to deliver higher luminosity in the LHC, unprecedented levels of pileup were also present in the collisions. Pileup refers to the presence of additional proton-proton collisions in the same 25 ns event window. These additional interactions typically arise from additional soft particle QCD interactions (Section 1.4) involving the partons. These interactions can be split into two categories depending on the source: "in-time" and "out-of-time" pileup. The in-time pileup contribution refers to the simultaneous interactions coming from one bunch-crossing (i.e. one 25 ns window). Out-of-time pileup refers to the leakage of energy signals in the ATLAS calorimeters (Section 2.3.3) from adjacent bunch crossings into a given bunch crossing [54].

The amount of pileup is typically expressed as a function of the mean number of interactions per bunch crossing $\langle\mu\rangle$, which is equivalent to μ_{vis} in Equation 61. The mean number of interactions per bunch crossing for the full Run-2 dataset is given in Figure 8.

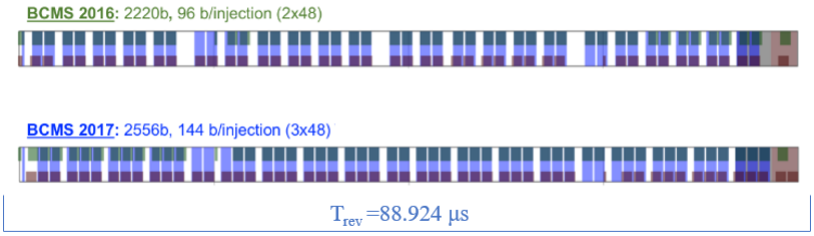

Figure 7: An illustration of the bunch structure of 25 ns proton beams in the LHC during 2016 and 2017. The gaps between filled bunches denote separations of the trains due to injection settings (small gaps) or account for the beam abort gap which protects the LHC in case of machine problems/failures (large gaps). The revolution time of the LHC is 88.924 μs and the number of bunches and bunches per injection of the LHC are added for scale. Adapted from Ref. [53].

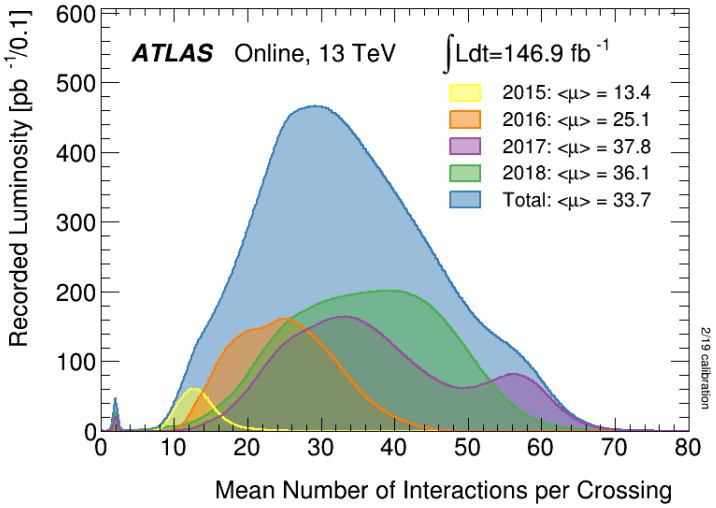

Figure 8: The amount of luminosity collected for as a function of mean number of interactions per bunch crossing $\langle\mu\rangle$ for the ATLAS Run-2 dataset [52]. A large variation in $\langle\mu\rangle$ values can be seen during the data-taking years, despite the similar averages in the 2016-2018 data-sets.

2.3 ATLAS Detector

The ATLAS detector [55] is a large general purpose high-energy particle experiment located at the first interaction point (IP) of the LHC. ATLAS is constructed to be forward-backward symmetric with respect to the LHC interaction point and covers nearly the entire 4π solid angle. The ATLAS detector components can be seen in Figure 9. The ATLAS detector is comprised of several subdetector components which measure different properties of the particles involved in proton collisions at the LHC. The closest subdetectors to the beamline are the inner tracking detectors (Section 2.3.2) which measure the trajectories of charged particles. Moving outward, the next detectors are the electromagnetic and hadronic calorimeters (Section 2.3.3) that measure the energy deposited by charged and neutral particles as they are stopped by (or if sufficiently energetic, pass through) the detection material. The outermost subdetector is the muon spectrometer (Section 2.3.4), which is designed to measure the trajectory of muons. Finally, several specialized detectors placed far away from the IP in the forward detection region (Section 2.3.6) are dedicated to luminosity determination or forward-diffractive physics. In addition to these descriptions, a brief overview of the ATLAS coordinate system and associated kinematic description of particles is given in Section 2.3.1.

2.3.1 Coordinate System and Kinematic Description

ATLAS describes physics events within the detector using a right-handed coordinate system, where the z-axis points along the beam direction, the x-axis points from the IP to the center of the LHC ring, and the y-axis points vertically (towards the surface) from the IP. Polar coordinates, however, are often more convenient to use for the event description. The angular coordinates, with azimuthal angle ϕ and polar angle θ are defined as:

$$\begin{aligned} \phi &= \arctan\left(\frac{x}{y}\right) \\ \theta &= \arctan\left(\frac{\sqrt{x^2+y^2}}{z}\right) \end{aligned} \tag{64}$$

The transverse plane (xy) is particularly important, as the proton-proton collisions are boosted along the z-direction. The rapidity, y, is defined as

$$y = \frac{1}{2}\ln\left(\frac{E+p_z}{E-p_z}\right) \tag{65}$$

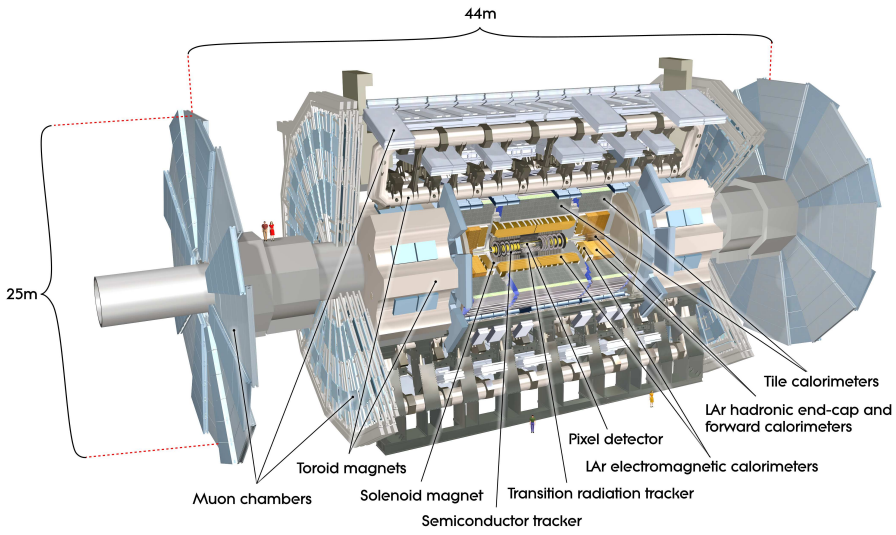

Figure 9: A figure of the ATLAS detector, illustrating all of the subdetector components of the detector centered around the interaction point. The detector is 44 m in length and 25 m in height. Figure adapted from Ref. [55].

in order to be invariant under boosts along the z-direction. The transverse energy and momentum are conserved and given by

$$\begin{aligned} p_{\mathrm{T}} &= \sqrt{p_x^2 + p_y^2} \\ E_{\mathrm{T}} &= \sqrt{E_x^2 + E_y^2} \end{aligned} \tag{66}$$

Finally, the distance measure between particles, denoted as ΔR_y is defined as

$$\Delta R_y = \sqrt{\Delta y^2 + \Delta\phi^2} \tag{67}$$

where $\Delta\phi$ must be calculated appropriately for the [0,2π] boundary condition. The rapidity is not regularly used in the detector description, as it is simpler to express geometrical units in experimental pseudorapidity, defined by

$$\eta = -\ln(\tan(\theta/2)) \tag{68}$$

which has a corresponding distance measure given by $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$. η and y are equivalent for massless particles or for highly energetic particles.

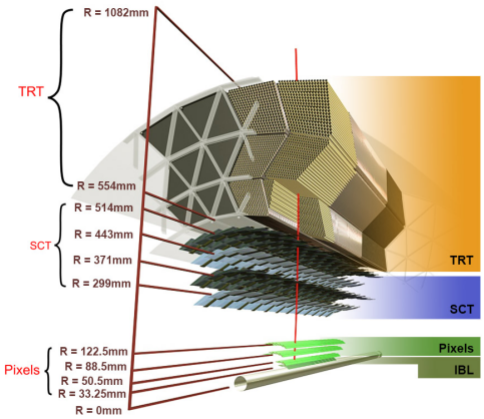

Figure 10: A drawing illustrating the sensors and structural elements of the four inner tracking detectors in the ATLAS experiment along with the LHC beam pipe. The Insertable B-layer (IBL), is the innermost silicon tracking layer, followed by three silicon-pixel layers, the silicon microstrip detector (SCT), and finally the straws of the transition radiation tracker (TRT). Figure taken from Ref. [57].

2.3.2 Inner Tracking Detectors

The inner detector of the ATLAS experiment is composed of four concentric detectors (see Figure 10. Starting from the IP and moving radially outwards, a particle cross the Insertable B-Layer (IBL) [56], the Pixel detector, the Semiconductor Tracker (SCT), and finally the Transition Radiation Tracker (TRT). The entire inner detector is immersed in a 2T field provided by a central solenoid magnet.

The IBL, Pixel, and SCT detectors all follow similar detection properties, namely the excitation of p-n junctions in silicon. As charged particles pass through a portion of the silicon and deposit energy, electron-hole pairs are created via ionization. As the detectors contain a bias voltage, the electrons then drift and a current is measured. The current depends on the energy of the incident particle, as a higher momentum particle deposits more energy and produces more electron-hole pairs. A further description of each tracking detector is given below.

IBL The IBL [56] is the first and newest layer of the inner detector and was installed in the ATLAS detector during the first long shutdown period of the LHC. The purpose of the IBL is to provide an additional radiation hard silicon layer to the pixel detector, which allows for better tracking resolution of a track transverse and

longitudinal impact parameter ($d_0, z_0 \sin\theta$). The impact parameters are particularly important as the impact parameters are crucial inputs for primary vertex reconstruction (Section 3.2) and aid in the identification of b-jets (Section 3.7). The IBL is located at an average radial distance of 33mm from the beam pipe and covers 330mm in the z-direction. The IBL consists of two types of pixel sensors, having planar sensors in the central region and 3D sensors in the high-η region. The IBL extends the active area of the tracker[7] up to $|\eta| < 3$ compared to $|\eta| < 2.5$ covered by the Pixel and SCT detectors. The planar IBL pixels have an intrinsic resolution of $8 \times 40 \mu\text{m}^2$, providing highly granular and precise measurements of charged particle momenta.

Pixel The Pixel detector is comprised of three concentric barrel layers of semiconductor staves at radial distances of 50.5, 88.5, and 122.5 mm extending to $|z| \approx$ 400mm and two end-caps sections consisting of three semiconductor disks held perpendicular to the beam axis at $|z|$ distances of 495, 580, and 650 mm. The Pixel detector provides high granularity measurements of charged particle positions in the region $|\eta| < 2.5$, with an intrinsic resolution of 10x115μm^2 [55] in (R-ϕ,z) planes. As the pixel and IBL layers are the closest to the IP, these detectors provide precise measurements of the charged particles coming from the proton-proton interaction vertex, which is crucial for identifying the hard-scatter vertex in LHC events. Further, these layers also can capture additional track decays from relatively long lived decays of B-hadrons (10^{-12}s) which provides additional measurements for b-jet identification.

SCT The SCT detector consists of four concentric cylindrical barrel layers of silicon micro-strip sensors at radial distances 284 - 498mm and extending to $|z| \approx$ 800mm and two end-cap sections with nine semiconductor disk layers each that cover the range $|z| \approx$ 854 - 2720mm. Each barrel SCT module consists of four strip sensors glued back-to-back (two on the top and two on the bottom of the layer), with the paired sensors rotated by $\pm$20 mrad from the center of the module to provide measurements of the charged particle trajectory in the z-direction. The SCT covers the full ID acceptance range of $|\eta| < 2.5$ and has an intrinsic accuracy of 17μm in the R-ϕ plane and 580μm in the z/R-direction.

TRT The TRT detector consists of carbon fiber reinforced Kapton drift tubes of 4mm diameter, known as straws, with 31μm diameter gold plated tungsten wires acting as the anode for each tube. These tubes are interleaved with polypropylene or

[7] In the absence of vertex spread from the IP, as the IBL is situated close to the LHC beam pipe. For vertices 112mm away from the IP in $|z|$ (maximal 2σ value), the coverage is $|\eta| < 2.58$ [56].

polyethylene fibers in the barrel and polypropylene foils in the endcap. This interleaving of materials gives rise to the transition radiation of particles as the particle passes through materials with different indices of refraction. As transition radiation is more prominent in low mass particles with a high$-\gamma$ factor [58], electrons passing through the TRT will produce a large amount of photons. The TRT, therefore, plays a major role in electron identification.

The TRT barrel consists of 73 layers of straws, oriented collinear to beam-axis, organized in three concentric sections extending from R = 554 mm to R = 1082 mm. The endcap modules contain 160 straw planes, aligned perpendicular to the beam axis, and cover the range $|z| \approx 820$ mm - 2750 mm. The TRT straws used two gas mixtures during Run-2, with one composed of 70% Xe, 27% CO_2, 3% O_2 while the alternative mixture replaces the xenon gas with argon [59].

By tracking particles through all of the ID, the ATLAS design track momentum resolution is given by

$$\frac{\sigma_{p_{\mathrm{T}}}}{p_{\mathrm{T}}} = 0.05\% p_{\mathrm{T}} \oplus 1\% \tag{69}$$

where the and the units for p_{T} are given in GeV. Therefore, the tracking system has a precise momentum resolution for low-p_{T} objects and the resolution degrades as a function of p_{T}.

2.3.3 Electromagnetic and Hadronic Calorimeters

The energy of charged and neutral particles are measured by the ATLAS calorimeters. The calorimeters are designed to stop the particles coming out of the IP from the proton collisions, such that these particles deposit most of their energy in these detectors. ATLAS uses two types of detectors, electromagnetic and hadronic calorimeters, to measure the properties of electrons, photons, tau leptons, and jets.

All of the ATLAS calorimeters act as sampling calorimeters where alternating sections of absorbing material, which initiates the showering of particles, and active medium providing energy measurements are interleaved to "sample" the shower as a particle passes through the calorimeter. This sampling provides energy measurements by summing all of the signals from the different layers and can also determine the shape of the shower by comparing the amount of energy deposited in each layer. This shape information provides useful input for particle identification in the ATLAS detector.

Electromagnetic Calorimeters: The ATLAS electromagnetic calorimeter consists of a high-granularity lead-liquid argon (LAr) electromagnetic (EM) sampling

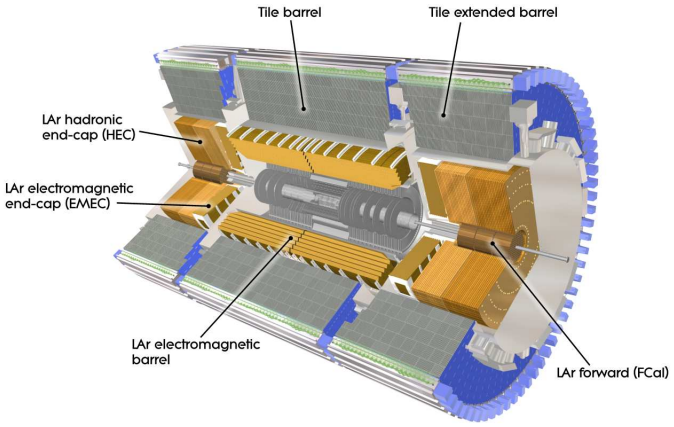

Figure 11: A diagram of the ATLAS electromagnetic and hadronic calorimeters. The forward calorimeter (FCAL) acts as both an electromagnetic and hadronic calorimeter, while the barrel and endcap calorimeters use two separate subdetectors to measure the electromagnetic and hadronic particle energies.

calorimeter that covers the $|\eta| < 3.2$ region. At the energy scales associated with LHC collisions, the energy deposits in the EM calorimeter mostly come from charged particles in jets, bremsstrahlung radiation produced by electrons, and pair production of electrons produced by photon as these particles travel through the calorimeter. Further, the electromagnetic calorimeters are arranged in an accordion geometrical design in order to provide complete coverage of the azimuthal (ϕ) direction.

The electromagnetic calorimeter is composed of several sections, with two half-barrel calorimeter (EMB) sections covering the region $|\eta| < 1.475$ and two endcap (EMEC) calorimeters covering the region $1.375 < |\eta| < 3.2$. Three cryostats provides the cooling capability necessary to keep the liquid argon below its boiling point, and the endcap cryostat services cause a loss of detection material in the region $1.37 < |\eta| < 1.52$. The readout cells in the calorimeters are segmented in both η and ϕ, providing highly granular measurements in the lateral and longitudinal directions for particle showers. The segmentation of an EMB module is shown in Figure 12. The first layer is composed of narrow strips that provides precision measurements of the initial energy. The second layer is made up of square cells and extends for 24 X_0 such that photons and electrons of moderate energies ($E < 50$ GeV) deposit the majority of their energy inside this layer. The back layer is less granular and is meant to add additional stopping power for higher momentum electrons and photons. In addition

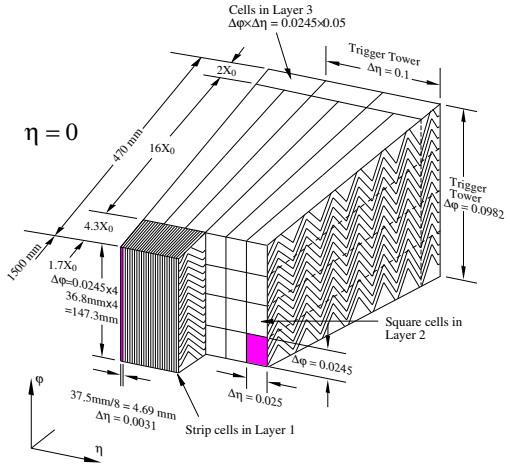

Figure 12: An illustration of segmentation and layer granularity for a module in the electromagnetic calorimeter barrel.

to these main layers, a pre-sampling layer consisting only the liquid argon layer without absorbers covers the region $|\eta| < 1.8$ to provide additional measurements of the energy loss from the particle interactions in the ID and solenoidal field. The EMEC segmentation follows that of the EMB calorimeter up to the edge of the tracking acceptance ($\eta < 2.5$), after which the granularity becomes coarser, having only two sample layers.

Hadronic Calorimeters: The hadronic calorimeters are comprised of two different technologies, depending on the η region of the detector. In the region $|\eta| < 1.7$, behind the LAr EM barrel calorimeter, the hadronic calorimeter is composed of alternating layers of plastic scintillator tiles for the active material and steel for the absorbing material. As hadronic particles pass through the material, a cascade of particles is produced via strong and electromagnetic interactions providing a similar shower to the EM showers, except that hadron to hadron decays also occur. When charged hadrons pass through the scintillator material, light is emitted and these signal are detected by photomultiplier tubes in order to provide the energy measurement. The tile calorimeter is separated into three sections, with a barrel section covering the region $|\eta| < 1.0$ and two extended barrel sections covering the region $0.8 < |\eta| < 1.7$.

The hadronic end-cap (HEC) calorimeter, which covers the region $1.5 < |\eta| < 3.2$, uses liquid argon as the active material and copper plates as the absorber and is arranged in a parallel-plate design. The HEC is composed of two wheels behind each

end-cap cryostats, with each wheel being divided into two longitudinal segments, totaling four detection layers. The granularity of each HEC readout cell is much coarser compared to the segmentation found in the EM calorimeter, being comprised of $(\Delta\eta, \Delta\phi) = (0.1 \times 0.1 \text{ or } 0.2 \times 0.2)$ depending on the η region of the detector.

Forward Calorimeter: The forward calorimeter (FCal) provides coverage from 3.1 $< |\eta| <$4.9 and measures both electromagnetic and hadronic energy in this region. The FCal is composed of three cylindrical modules which are arranged similarly to the barrel and endcap calorimeters. These three modules are setup as a series of concentric rods and tubes, with liquid argon in the gaps between the rod and tubes as the active medium. The first module uses copper as an absorber and is used primarily for electromagnetic energy measurement. The second and third modules use tungsten as the absorbing medium and are used for hadronic energy measurement.

The energy resolution for calorimeters is determined via the equation

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \tag{70}$$

where the stochastic, or sampling, term a depends on calorimeter design as it provides a measure of the particle fluctuations in the shower, the noise term b is determined from electronics and pileup noise in the detector, and the constant term c depends on the non-linearities in the pulse shapes. Here, the operator $\oplus$ is used to indicate that each term is quadratically summed to give the final resolution.

For the ATLAS calorimeters, the energy resolution was measured in test beam data under several different setups. For the electromagnetic calorimeters, electrons beams with energy $E \leq 245$ GeV were used to determine the response of the detectors [60, 61]. For the LAr hadronic calorimeters, pion beams with energy $E \leq 200$ GeV were used [61, 62], while pion beams with $E \leq 350$ GeV were used to determine the response of the hadronic barrel calorimeter [63]. In these beam tests, only the sampling and constant noise terms were measured and these values are given for each subdetector in Table 7.

2.3.4 Muon Spectrometer and Toroidal Magnets

As the muon is ≈ 200 times more massive than the electron, the muon can escape detection in the calorimeter absorbers. In order to identify muons, the outermost layer of ATLAS is dedicated to specialized muon detectors know as the muon spectrometer (MS).

The MS is comprised of several different detector technologies which emphasize two different uses: chambers dedicated to the precision tracking of each muon

Subdetector	Sampling Term, a	Constant Term, c
EM Barrel	$(10.1 \pm 0.1)\%$	$(0.17 \pm 0.04)\%$
EM End-cap (low-η)	$(12.1 \pm 0.2)\%$	$(0.4 \pm 0.1)\%$
EM End-cap (high-η)	$(13.5 \pm 0.5)\%$	$(0.7 \pm 0.1)\%$
EM Forward	$(29.3 \pm 0.7)\%$	$(3.0 \pm 0.1)\%$
Hadronic Barrel	$(52.9 \pm 0.9)\%$	$(5.7 \pm 0.2)\%$
Hadronic Endcap (low-η)	$(84.1 \pm 0.3)\%$	$(0 \pm 0)\%$
Hadronic Endcap (high-η)	$(88 \pm 5)\%$	$(6.8 \pm 0.4)\%$
Hadronic Forward	$(98.5 \pm 4.0)\%$	$(6.4 \pm 0.4)\%$

Table 7: The measured resolution values for the ATLAS calorimeters, derived from test beam data, where the impact points for each subdetector correspond to $\eta = 0 - 0.7$ (EM/Hadronic Barrel), $\eta = 1.6 - 1.8$ (EMEC/HEC low-η), $\eta = 2.8$ (EMEC/HEC high-η), and $\eta = 3.65$ for the FCAL.

candidate and chambers dedicated to fast readout for the ATLAS trigger system (Section 2.3.5). These chambers cover the pseudorapidity range $|\eta| < 2.7$ for the tracking chambers, but the trigger chambers are only instrumented in the region $|\eta| < 2.4$.

The MS is split into three common regions: a barrel section covering $|\eta| < 1.05$ and two endcaps covering the $1.05 < |\eta| < 2.7$ region. The barrel MS is comprised of three cylindrical layers, concentric to the beam axis, with radii of approximately 5m, 7.5m, and 10m from the beam-pipe. These barrel layers are composed of eight large chambers and eight small chambers, differing only in lateral length, that follow the ϕ symmetry of the ATLAS toroidal magnet. These large and small chambers overlap slightly in the ϕ plane to provide full azimuthal coverage.

The endcap MS is arranged perpendicular to the beam axis and is arranged in wheel structures which are located at distances in $|z|$ of approximately 7.4m (Small Wheels, the endcap inner layers), 10.8m (the "extra" layers of the endcap situated around the toroidal magnets), 14m (first big wheel, endcap middle layers), and 21.5m (last big wheel, endcap outer layers) from the ATLAS IP. Details of the geometry and positioning of the MS are given in Figure 13.

The magnetic field for the barrel MS is provided by eight coils of superconducting magnets, referred to as the barrel toroid. The barrel toroid provides a magnetic field of approximately 0.5 T and the eight coils are symmetric in the azimuthal angle and exclusively covers the region $|\eta| < 1.4$. Each endcap MS section contains a smaller toroidal magnet consisting of eight square coils and eight wedge coils that provide a

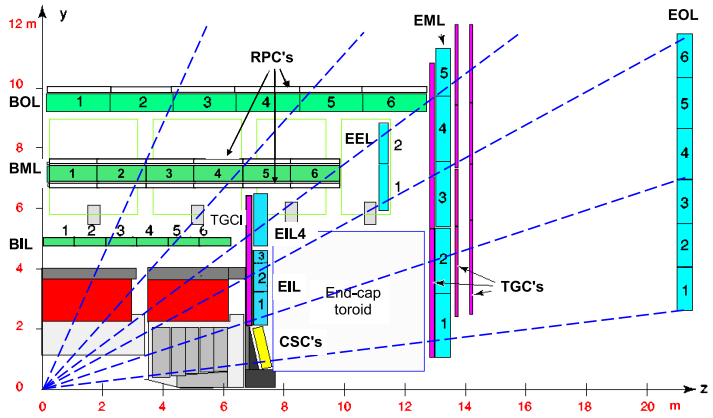

Figure 13: A cross-section of the ATLAS muon spectrometer system in the R-z plane (bending plane) with positions of the various subdetectors. The naming scheme of the MDTs is as follows: (B,E) refers to the barrel or endcap positioning, (I,M,O,E) refers to the inner, middle, outer, or "extra" layers positions, and the L stands for large as the small chambers are not shown [55].

magnetic field of 1.0 T. These endcap toroids exclusively provide the magnetic field for the region $1.6 < |\eta| < 2.7$. In the barrel-endcap transition region, $1.4 < |\eta| < 1.6$, the two magnetic fields from the barrel toroid and endcap magnets overlap and provide a combined magnetic field.

The four detector technologies in the MS are the monitored drift tubes (MDTs), resistive plate chambers (RPCs), cathode stripe chambers (CSCs) and thin-gap chambers (TGCs). Each detector type is setup in a complementary fashion, focusing on precision tracking measurements in the MDTs and CSCs and fast track readout for the trigger in the RPCs and TGCs. The MDTs and CSCs are aligned to precisely measure the primary coordinate of muon tracks in the bending (R-z) plane while the RPCs and TGCs measure the track coordinates in both the bending and non-bending (ϕ) planes. The precision and non-bending coordinate measurements are later matched (Section 3.3.1) in order to provide the muon momentum measurement. Descriptions of the detector technologies are given below.

Monitored Drift Tubes: The MDT chambers are used for precision tracking in the MS, providing an average resolution of 80 μm for a track in the bending plane. The MS consists of 3 layers of MDT chambers, where each chamber is arranged in 3-8 layers of drift tubes. The outer two layers of MDTs cover the entire MS detection region $|\eta| < 2.7$, but the innermost layer only covers the region $|\eta| < 2.0$.

Each chamber consists of drift tubes with a diameter of 30 mm filled with a Ar/CO_2 mixture held at 3 atmospheres. After ionization of the gas by a muon, the electrons drift towards a 50μ m W-Rh anode wire in the center of the tube. The anode wires are held at a voltage of ≈ 3 kV [64].

Resistive plate chambers: The RPC consists of a parallel plate detector separated by 2mm using an insulating space and this volume is filled with a gas mixture of predominantly tetrafluoroethane ($C_2H_2F_4$) for charge deposition [65]. It consists of three concentric cylindrical layers of this parallel plate design, with each layer known as a "muon station". The large distance between the innermost RPC station and outermost RPC station allows for determination of muon candidates in the p_{T} = 9-35 GeV range, while a low-p_{T} (4-9 GeV) identification is determined by using coincidences between the two layers of the inner chambers.

Cathode Strip Chambers: The CSC detector replaces the MDTs in the first (innermost) layer in the region $2 \leq |\eta| \leq 2.7$, where a high interaction rate during proton collisions would cause degraded performance in the MDTs. The CSC detector is organized into two disks of 16 chambers each, with small and large chambers overlapping in the ϕ coordinate. Each CSC chamber consists of a multi-wire proportional chamber and two segmented cathode strips. One cathode is oriented perpendicular to the wires in order to provide precision coordinate tracking and the other is oriented parallel in order to capture the transverse component of the muon tracks.

Thin Gap chambers: The TGCs are similar to multi-wire proportional chambers and consist of of 50 μm gold-plated tungsten anode wires separated by a distance of 1.8 mm and graphite planes acting as cathodes with a cathode-anode distance of 1.4 mm. The TGCs operate with a mixture of CO_2 and n-pentane acting as the gas medium and a 2.9 kV potential allows for fast measurements of the track coordinates in the MS endcaps. In addition to the triggering capabilities, the TGCs also provide the azimuthal coordinate of the tracks to complement the MDTs. The TGCs consist of seven layers of wire chambers behind the MDT middle layer and two layers behind the MDT inner layer. The TGCs cover the region $1.0 < |\eta| < 2.4$.

2.3.5 Trigger System

The ATLAS trigger system is dedicated to identifying the bunch crossings filled with proton collisions that contain interesting events and must these select these events for further offline analysis by rejecting the large QCD mulit-jet background. For a sense of scale, the LHC collision rate occurs at 40 MHz, but only about 1-2 KHz of events can be recorded into permanent storage for data analysis.

Since the beginning of Run-2, ATLAS uses a two-level trigger system to select events for further offline analysis. These two levels consist of the hardware based Level-1 (L1) trigger and the software based high-level trigger (HLT). These two systems are described below:

Level-1: The first level is a hardware-based trigger, known as the Level-1 (L1) trigger, which consists of three main systems: the L1 Muon Trigger (L1Muon), the L1 calorimeter trigger (L1Calo), and the L1 Central trigger processor (CTP). The L1Muon trigger processors receive input from the dedicated muon trigger chambers described in Section 2.3.4. In contrast to the L1Muon system, the L1Calo trigger relies on a reduced granularity calorimeter readout called calorimeter trigger towers. Trigger towers provide energy and timing sums for detector regions of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for the barrel calorimeters and become less granular (up to $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$) in the forward η calorimeter regions.

During Run-2, a new hardware based trigger known as the L1 topological trigger (L1Topo) was commissioned. L1Topo has several unique properties, namely that it can calculate geometric and kinematic relationships between different trigger objects and it can receive inputs from both the L1Muon and L1Calo trigger paths. The kinematic quantities used during Run-2 included topological measurements of $\Delta\phi$ and $\Delta\eta$ between objects and some algorithms are dedicated to calculating the invariant mass of input objects for specific trigger chains [66].

The CTP receives the trigger information from the L1Muon, L1Calo, and L1Topo processors and synchronizes the trigger inputs to the LHC collision clock in order to align all signal with the correct bunch crossing [67]. After synchronization, all of the trigger inputs are compared within CTP to a series of trigger thresholds for each type of physics being targeted. If the processed information passes at least one of the L1 thresholds, then the CTP issues a L1 accept (L1A) decision and propagates this decision to all ATLAS subdetectors. Upon receiving the L1A, all of the data from each subdetector for that given event is read out and propagated to the next step in the trigger process.

The choice of thresholds for the L1 trigger are determined by the maximum L1 output rate of approximately 100 KHz, which represents a factor of 400 reduction from the 40 MHz input rate. This rate reduction is typically accomplished by cutting on high transverse momentum signals from the L1Muon chambers and high transverse momentum electrons, taus decaying to hadronic final states, jet candidates, and large missing transverse energy (see Section 3.8) in the L1Calo trigger system.

High-Level Trigger: Any event selected by the L1 trigger decision is then passed on to the high-level-trigger (HLT) which is a software based trigger that uses the full

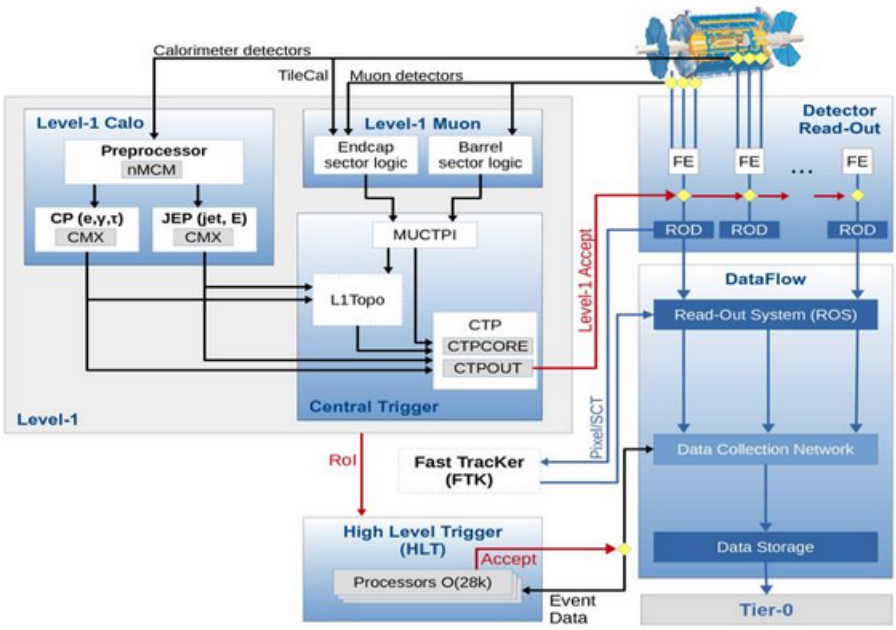

Figure 14: A schematic overview of the ATLAS trigger system during Run-2. Here, the inputs from the calorimeter and muon detector paths are separated into their trigger readout paths (left) and the full event readout paths (right) which lead to offline storage. Figure adapted from Ref. [68].

granularity of the ATLAS detector to make more refined decisions on which events to save into the ATLAS output data. For example, the inner detector trackers are only readout at HLT step and the track quantities are able to more accurately identify the primary vertex and enhance particle identification. Via these more complex algorithms, the HLT reduces the data rate from 100 KHz to approximately 1 KHz of output. Being software based, the HLT trigger is limited by the processing time of an event; in order to prevent CPU overloads, most trigger decisions for an event must be made in $\mathcal{O}$(100 ms - 1 s). In order to make these quick decisions, CPU intensive steps such as track reconstruction are only performed on the ROIs coming from the L1 trigger.

A schematic overview of the ATLAS trigger system and data-acquisition chain can be seen in Figure 14.

2.3.6 Forward Detectors

The ATLAS experiment also has several specialized detectors in the forward region. These detectors include:

LUCID: LUCID (LUminosity Cherenkov Integrating Detector) is the primary instantaneous luminosity monitor for the ATLAS experiment. It provides measurements of the inelastic proton-proton scattering in the forward region in order to constrain the integrated luminosity. LUCID consists of photomultiplier tubes (PMTs) with quartz windows that are arranged symmetrically around the beam pipe 17 m away from the ATLAS IP on both sides of the detector. When particles of sufficient energy pass through the quartz window of a PMT, Cherenkov light is emitted and detected by the PMT. This design differs from the Run-1 LUCID detector, which previously had a gas radiator providing the Cherenkov thresholds for particle counting [69]. The LUCID detector measures the relative luminosity μ_{vis} via several methods: event counting, hit counting, or charge integration. In order to calibrate these values, an absolute luminosity measurement is required.

ALFA: The ALFA (Absolute Luminosity for ATLAS) detector is comprised of four scintillating fiber trackers, two per side, that use the Roman-pot technique to measure elastic proton-proton scattering. ALFA is located ± 240 m away from the ATLAS interaction point and the two detectors per side are inserted close to the proton beams in the vertical plane. The absolute luminosity can be measured by the elastic scattering cross section, as this cross-section is related to the elastic scattering amplitude f_{el} via the optical theorem

$$\sigma_{\text{tot}} = 4\pi \text{Im}[f_{el}(t=0)] \tag{71}$$

where $t = 0$ defines the momentum transfer scale [70]. As this very low momentum transfer can only be measured at angles that are below the nominal LHC beam divergence [55], the ALFA detector can only measure this cross section during dedicated LHC fills. These fills are required to have high-β^* optics and reduced beam emittance and help to provide the general luminosity calibration of the ATLAS physics program.

AFP: The AFP (ATLAS Forward Proton) detector is similar to the ALFA detector, using Roman pots that contain a silicon tracker and are inserted close to the LHC beam. The AFP detector is inserted in the horizontal plane of the LHC beam in order to measure diffractive proton-proton collisions. The AFP detector consists of two near and two far stations, located ± 205 m and ± 217 m from the ATLAS interaction

point respectively. In addition to the silicon trackers, the far stations also contain quartz time-of-flight detectors in order to improve the tracking resolution in the z-direction for these forward events.

ZDC: The Zero Degree Calorimeter (ZDC) measures the energies of neutral particles during low-μ proton-proton collisions and heavy-ion collisions. The ZDC is located on either side of the ATLAS experiment, 140 m away from the interaction point.

3 ATLAS Object Reconstruction and Particle Identification

During proton-proton collisions, a large number of particles are produced and detected by the ATLAS detectors. These particles include both electrically charged objects, such as electrons, muons, and charged hadrons (π^- for example), and electrically neutral objects such as photons, neutrinos and neutral hadrons (π^0,etc). In order to accomplish the diverse physics goals of the ATLAS experiment, all of these particles must be reconstructed and identified with a high efficiency. This chapter discusses the object reconstruction and identification techniques for the objects used in the $H \to aa \to bb\mu\mu$ search described in Sections 4-6.

3.1 Track Reconstruction

The reconstruction of charged particles in the ATLAS detector is important for almost all physics object identification and reconstruction. The charged particle trajectories, referred to as *tracks*, are used explicitly as inputs to the reconstruction of muons (Section 3.3.1) and electrons (Section 3.5.1) and aid in isolating these leptons from the hadronic background at the LHC. Further, tracks are used to distinguish jets originating from b-quark decays (Section 3.7), calibrate the jet energy scale, determine the soft (low momentum) contributions to the transverse missing energy (Section 3.8), and reduce the jet background coming from pileup (Section 3.6.1).

Charged particles reconstruction begins with the measurement of the particles positions at multiple points (hits) in the ID tracking system (Section 2.3.2), particularly in the pixel (including IBL) and SCT detectors. The measurements from these subdetectors are first clustered by matching hits above threshold if the energy deposits share common edges or corners. These clusters are then assigned three-dimensional measurement values, called space-points. In order to cut down on combinatorial fakes, preliminary track seeds are formed on groups of three space-points; these track seeds are formed by starting with space-points only found in the SCT, then those found only in the pixel detectors, and finally those which cross multiple detectors. This prioritization helps to maximize the purity of the track seeds. These track seeds then undergo an ambiguity resolving step, where bad quality tracks candidates are rejected if they fail the following criteria [71]

- Each track must have transverse momentum $p_T > 500$ MeV
- Each track must have pseudorapidity within the tracking volume $|\eta| < 2.5$

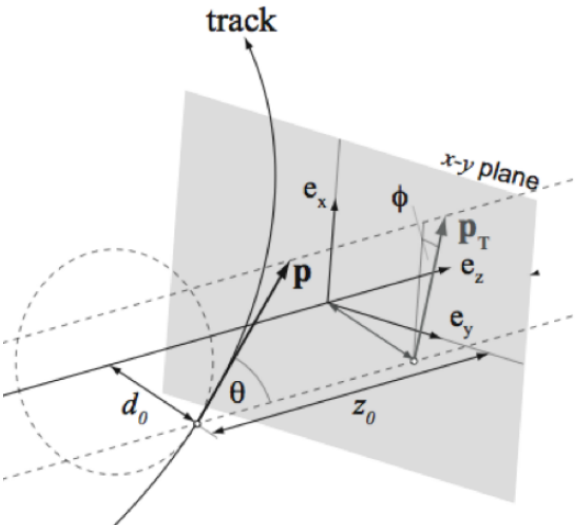

Figure 15: An illustration of the track parametrization in that ATLAS coordinate system. The transverse and longitudinal impact parameters d_0 and z_0 respectively, are defined as the distance of closest approach for the track in the x-y and z planes. ϕ defines the azimuthal angle at the closest approach and θ is the polar angle of the track. Along with the charge of the track, these parameters form the perigee parametrization.

- Number of silicon hits (SCT+Pixel) $N_{\text{Si}} \geq 7$
- Number of shared modules $(N_{\text{Pix}}^{\text{sh}} + N_{\text{SCT}}^{\text{sh}}/2)$, $N_{\text{mod}}^{\text{sh}} \leq 1$
- The number of holes in the silicon (SCT+Pixel)$N_{\text{Si}}^{\text{hole}} \leq 2$
- The number of holes in the Pixel detectors $N_{\text{Pixel}}^{\text{hole}} \leq 1$
- $|d_0^{\text{BL}}| < 2.0$ mm
- $|z_0^{\text{BL}}\sin\theta| < 3.0$ mm

where d_0^{BL} and z_0^{BL} refer to the transverse and longitudinal impact parameters of the track, as measured with respect to the LHC beam-line. The impact parameters are illustrated, along with the other perigee parameters ϕ and θ in Figure 15.

The hole requirement removes tracks if their passage through the Pixel or SCT detectors was expected to yield a hit in a specific area and the track failed to do so. This hole assigning accounts for inactive regions of the detector [57] such that

high-quality tracks in bad detector regions are not removed. The addition of shared module requirements helps to resolve ambiguities, as each cluster can only seed a maximum of two tracks. This requirement cuts down on the number of mis-assigned clusters to tracks [71]. These requirements correspond to the "Loose" ATLAS track criteria and are used as inputs to the other objects described in this Section.

Alternative tracking methods also include the TRT outside-in tracking, where hits in TRT are extended inward to try and match hits from the SCT to form tracks which can recover some inefficiency for long-lived particles which do not decay in the silicon trackers or recover low p_{T} electrons with a high energy loss. Additionally, the inside-out tracking can be extended to look for tracks matched to hits in the TRT. This TRT track extension [72] can prove useful for electron identification.

3.2 Vertex Reconstruction

In order to determine the hard-scatter portion of a proton-proton collision, every proton-proton interaction (including pileup interactions) must be reconstructed. These interactions are determined by grouping together reconstructed tracks which are consistent with comping from the same initial position. The position resulting from this grouping is called a vertex.

Vertex reconstruction begins by associating tracks to an initial vertex seed position. The initial x-y positions are determined as the center of the LHC beam spot in the x-y plane, while the z-coordinate starts from the mode of all track z-coordinates associated with the beam spot position [73]. Next, an iterative process attempts to minimize the χ^2 value for the best fit vertex from the tracks associated to the seed. The track compatibility with the fitted vertex is calculated using a χ^2 fit which assigns weights to each track of the form

$$w(\chi^2) = \frac{1}{1 + \exp\left(\frac{\chi^2 - \chi^2_{\mathrm{cutoff}}}{2T}\right)} \tag{72}$$

where χ^2_{cutoff} defines the cutoff where an individual track weight becomes equal to 0.5 and T is a parameter that iterates with the fit in order to smooth out the weight distributions from the fit [73]. After performing this fit, any tracks that are 7σ away from the best-fit vertex position are treated as incompatible with that vertex and are removed from the fit. After the first vertex fit is completed, the tracks not associated to the first vertex are then considered in another vertex fit. This approach continues until all tracks have been associated to a vertex or there are no tracks which are compatible enough to form a vertex ($N_{\mathrm{trk}} \geq 2$ is required).

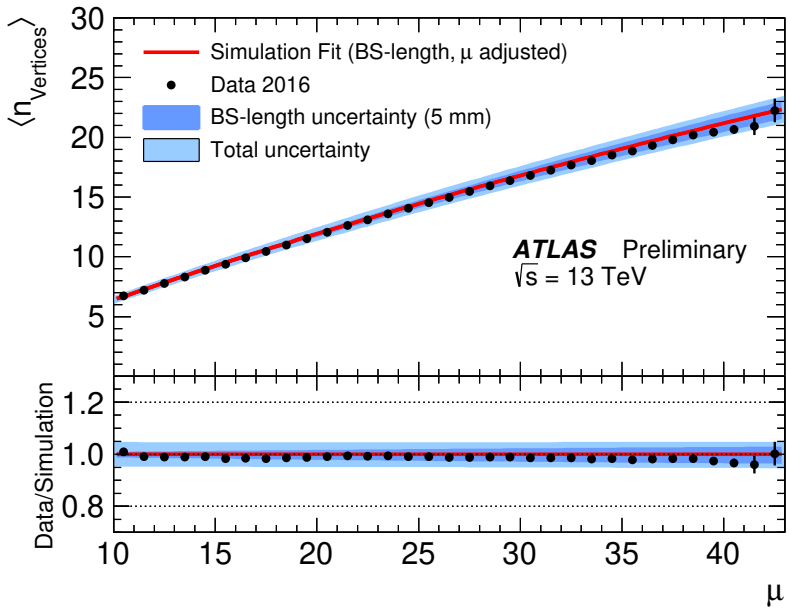

Figure 16: Distribution of the average number of reconstructed vertices as a function of $\langle\mu\rangle$ in 2016 ATLAS data. The red curve corresponds to the best fit from a minimum-bias simulation events, while the data points correspond to zero-bias data collected in 2016 [74].

After all vertices are reconstructed, the primary vertex is identified by selecting the vertex with the highest sum of squared transverse momenta of all associated tracks, $\sum_{\text{trk}} p_{\text{T}}^2$. The primary vertex acts as a proxy for the hard-scatter of the proton collisions, while the other reconstructed vertices in an event are considered to come from pileup interactions. The number of vertices identified in an events as a function of the $\langle\mu\rangle$ is shown in Figure 16 [74].

3.3 Muons

3.3.1 Reconstruction

In ATLAS, muons can be reconstructed using tracking information from the ID, MS, and even from energy deposits in the calorimeters. Muon candidate reconstruction in the ID uses the same reconstruction techniques as discussed for the general charged particle case in Section 3.1, but the quality requirements for muons are discussed in

Section 3.3.2. The reconstruction in the MS, however, is complicated by the magnetic field and detector orientation discussed in Section 2.3.4.

For MS muon reconstruction, hit patterns are first identified in the MDT and trigger chambers to determine track segments along a bending plane trajectory. These trajectories are taken to be straight-lines in order to correlate hits between different MDT layers. The ϕ coordinate of these segments can also be constrained using hits from the RPC and TGC detectors. The CSC, due to its unique capability to measure the η and ϕ trajectories, uses a combinatorial search to build segments.

These straight line segments from the MDT layers are then matched together, starting from the middle layer and attempts to match segments in the outer and inner layers using a χ^2 fit. After matching, these segments now form a muon track candidate in the MS. Each muon candidate is required to have at least two matching segments, except in the barrel-endcap transition region where a quality requirement allows single segment tracks to pass [75].

A single segment can be used for multiple muon track fits, which results in some ambiguity for geometrically close muons. To mitigate this ambiguity and maintain high selection efficiency for these muons, track segment sharing is allows in the inner MS layers. The ambiguity is resolved as long as the muons separate and share no hits in the outermost layer.

Given the ability to reconstruct muon tracks in the ID and MS, along with detect some energy in the calorimeters, several categories of reconstructed muons are used in ATLAS. These include:

Combined Muons: Combined muons combine the track fits performed separately in the MS and ID into a single track using hits from both detectors. The track fit can proceed along two pathways, the inside-out and the outside-in reconstruction. The inside-out approach extrapolates the track from the ID and looks for matching MS tracks, while the outside-in approach extrapolates the muons from the MS inward to try and match ID tracks. The majority of combined muons are reconstructed via the outside-in reconstruction [75].

Extrapolated Muons: Extrapolated muons only use track candidates from the MS in order to form reconstructed muons. In addition to the nominal tracking requirements, extrapolated muons tracks must be compatible with an origin in the LHC beam spot envelope and the track must cross at least two layers of the MS detectors or three layers in the CSC region. These muons provide unique sensitivity in the region $2.5 < |\eta| < 2.7$, which is outside of the ID coverage. These muons are also referred to as stand-alone muons [76].

Segment-tagged Muons: Segment-tagged muons start with track candidates reconstructed in the ID and classify these tracks as muons if the extrapolated track position matches at least one segment in the MDT or CSC. These muons cover the phase space where the MS has limited coverage or when the muon p_{T} is too low to traverse multiple detector layers.

Calorimeter-tagged muons: Muons traverse through the ATLAS calorimeter and deposit energy in the EM and hadronic calorimeters due to ionization and radiative energy losses [77]. Calorimeter-tagged muons are reconstructed from any ID tracks matched to a calorimeter energy deposit consistent with a minimum ionizing particle. As this signature is hard to disentangle from the jet background, calorimeter tagged muons have the lowest purity. Despite this issue, these muons provide additional muon reconstruction efficiency for the $|\eta| < 0.1$ region where the MS has very limited acceptance.

Each of the muon reconstruction techniques are performed independently on every event. As muon candidates can be reconstructed by multiple techniques, overlapping muons are removed by prioritizing combined muons, then segment tagged muons, and finally calorimeter tagged muons. The extrapolated muons do not enter this prioritization as they are not initially reconstructed in the ID. Any overlaps between the extrapolated muons and the ID originated muons is resolved by choosing the track with the better fit quality and higher number of tracks [75].

3.3.2 Identification

Muon identification aims to suppress hadronic background by applying quality requirements that select prompt muons with a high efficiency and provide a good understanding of the muon momentum measurement.

Five muon identification selections are employed by ATLAS to target the different needs of physics analyses, include the Loose, Medium,Tight, Low-p_T and High-p_T selections. The Loose, Medium, and Tight selections are inclusive and simply increase purity at the cost of reconstruction efficiency. The Loose category uses all of the muon reconstruction techniques to identify muon candidates, while the Medium and Tight selections only use muons coming from the Combined and Extrapolated techniques.

The identification criteria rely on several variables related to the compatibility of tracks produced in the ID and those produced in the MS. These variables include [75]:

- The q/p significance, $\frac{q/p}{\sigma_{q/p}}$, which measures the difference between the ratio of the charge and momentum of the muons measured in the ID and MS and normalizes the difference by the corresponding tracking uncertainties

- $\rho' = \frac{|p_{\mathrm{T}}^{\mathrm{ID}} - p_{\mathrm{T}}^{\mathrm{MS}}|}{p_{\mathrm{T}}^{\mu}}$, which measures the difference between the transverse momentum the absolute value of the difference between the transverse momentum measurements in the ID and MS divided by the p_{T} of the combined track
- The normalized χ^2 of the combined track fit corresponding to a muon candidate

In addition to these quantities, the track requirements for combined muons require ID tracks with $N_{\mathrm{Pix}} \geq 1$ hit, $N_{\mathrm{SCT}} \geq 5$ hits, $N_{\mathrm{Si}}^{\mathrm{hole}} < 3$, and that $\geq 10\%$ of the TRT hits originally assigned to the ID track are included in the global track fit [75].

As an illustration of the additional quality requirements, Medium muons require that the q/p significance be less than 7 and the number of hits in at least 2 MDT layers must be greater than or equal to 3 for all regions except $|\eta| < 0.1$ where the requirement is loosened to 1 MDT layer. Tight muons must follow these requirements and are further required to pass a two dimensional selection on ρ' and q/p significance made as a function of p_{T} of the muon. Further, Tight muons require the χ^2 of the muon track to be less than 8. Further details on the identification requirements can be found in Ref. [75].

3.4 Topological Cluster Reconstruction

In the LAr and Tile calorimeters described in Section 2.3.3, the lateral and longitudinal segmentation allows particle showers to be reconstructed as three-dimensional objects. To accomplish this reconstruction, ATLAS uses a (4,2,0) growing volume algorithm to cluster of cells into topological clusters [54]. In this algorithm, energetic cell are given a signal significance defined by

$$\zeta_{\mathrm{cell}}^{\mathrm{EM}} = \frac{E_{\mathrm{cell}}^{\mathrm{EM}}}{\sigma_{\mathrm{noise,cell}}^{\mathrm{EM}}} \tag{73}$$

where $E_{\mathrm{cell}}^{\mathrm{EM}}$ defines the energy deposited in the EM calorimeter cell under test and $\sigma_{\mathrm{noise,cell}}^{\mathrm{EM}}$ is a noise term that depends on both electronics noise and noise coming from the pileup background. The $\langle\mu\rangle$ induced noise is the dominant term for the LHC Run-2 conditions. After all cells are assigned a significance, every cell that has $|\zeta| > 4$ is used as a primary seed for clustering. The nearest cells in (η, ϕ) around this primary cell, including those in different subdetectors, are then added into the primary cell cluster if they have $|\zeta| > 2$. Overlaps of cells with $|\zeta| > 4$ are resolved via merging these clusters and the algorithm progresses iteratively. The boundary condition for the last cells in the clusters are set as $|\zeta| > 0$, maintaining any cells near the noise thresholds.

Of particular note is the use of negative energy seeds in the iterative process as well. Typically, negative energy signals arise from pileup interactions as the ATLAS calorimeter shapers assign negative energy values to any interaction that comes from out-of-time pileup. The final physics objects for electrons, jets, and the transverse missing energy are only build from topological clusters with $E^{\mathrm{EM}} > 0$, despite the use of cell energies $E^{\mathrm{EM}}_{\mathrm{cell}} < 0$. Although not associated to signal processes, any negative energy clusters can provide useful input for pileup cancellations in the $E^{\mathrm{miss}}_{\mathrm{T}}$, electron, or jet objects [78].

3.5 Electrons

3.5.1 Reconstruction

Electrons from proton collisions in the ATLAS detector fall into two broad categories: central and forward electrons. A central electron is defined as any electron within the inner tracking detector and precision EM calorimeter acceptance ($|\eta| < 2.47$); these electrons are reconstructed using a combination of ID tracks and EM energy deposits. Forward electrons are any electrons falling outside of this acceptance ($|\eta| > 2.47$) and are only reconstructed using EM energy deposits. The central electron reconstruction algorithm is detailed in the following section and the forward electron reconstruction follows the same algorithm without relying on any track inputs.

Electron reconstruction begins by building clusters out of energy deposited in the EM calorimeter. The energy deposited in "towers" of size $\Delta\eta \times \Delta\Phi = 0.025 \times 0.025$ from each EM layer, including the pre-sampler for $|\eta| < 1.8$, is summed to form the energy of the tower. The seed clusters are then formed by using a sliding-window algorithm [79] to identify local transverse energy (E_{T}) maxima in clusters of 3×5 towers. Each cluster is required to have $E_{\mathrm{T}} > 2.5$ GeV and ambiguities are resolved by taking the highest E_{T} cluster in any 5×9 tower ($\Delta\mathrm{R} \approx 0.26$) region [80].

After the EM deposit clustering, ID tracks are reconstructed following the standard track reconstruction detailed in Sec. 3.1 which relies on a pion hypothesis. A special caveat exists for electron tracks, as track seeds with $p_{\mathrm{T}} > 1$ GeV matched to loose EM cluster ROIs are considered even if these seeds cannot form a track that meets the requirement of at least seven silicon hits. In this case, the track fitting is changed to allow for electron energy loss due to bremsstrahlung if the global Chi^2 fit fails [80].

After the initial track reconstruction and fitting, a generalized fitter that better models electron bremsstrahlung is applied to the initial track candidates. This fitter, by utilizing the Gaussian Sum Filter (GSF) [81], uses track candidates with $N_{\mathrm{Si}} \geq 4$ and are matched to the EM ROIs via geometrical matching requirements:

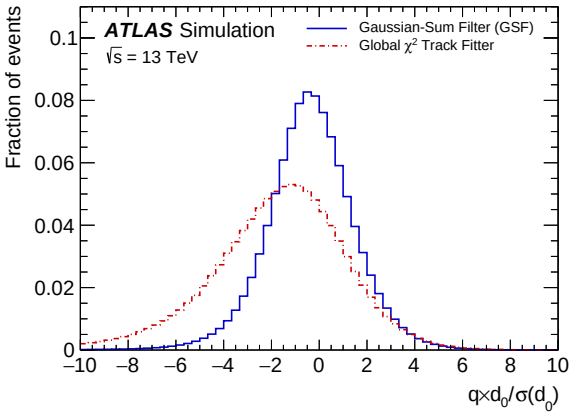

Figure 17: Distribution of the reconstructed electric charge of electrons candidates multiplied by the transverse impact parameter significance, $q \times d_0/\sigma_(d_0)$. The distributions show the differences in simulated event candidates using tracks fitted with the Global χ^2 Track Fitter (dashed red lines) and for tracks fitted with the GSF (solid blue line). Figure adapted from Ref. [80].

$$\begin{aligned} |\eta_{\text{EMcluster}} - \eta_{\text{track}}| &< 0.05 \\ -0.20 < q \times [\Delta(\phi_{\text{EMcluster}}, \phi_{\text{track}})] &< 0.05 \\ \text{or} -0.1 < \Delta\phi_{\text{res}} &< 0.05 \end{aligned} \tag{74}$$

where q accounts for the charge of the track and $\Delta\phi_{\text{res}}$ is the separation between the cluster position in the second layer of the EM calorimeter and the GSF-track ϕ [80]. This fit allows for better resolution in the tracking quantities, with the example of the transverse impact parameter being shown in Fig. 17.

The final reconstruction of the electron candidate identifies the primary track via tightening the track-cluster matching criteria and using the number of silicon hits to associate the electron as track qualities. Afterwards, the initial EM seed cluster is also expanded to 3×7 ($|\eta| < 1.37$) or 5×5 ($1.52 < |\eta| < 2.47$) or both ($1.37 < |\eta| < 1.52$) tower sizes. These extended windows were optimized for electron energy calibration and account for the energy distribution differences between the barrel and endcap calorimeters [82].

3.5.2 Identification

The current ATLAS method for electron identification is known as the likelihood-based (LH) identification. In this identification, the likelihood takes the form are input into a series of probability density functions of the form

$$L_{\mathrm{S(B)}} = \prod_{i=1}^{n} P_{\mathrm{S(B),i}} x(i) \tag{75}$$

where $\mathbf{x}$ is a vector of discriminating variables, $P_{S,i}$ is the value of the signal pdf for a given variable i at the value x_i and the $P_{B,i}$ defines the same value but for the background pdf. A full list of discriminating variables for the electron likelihood can be found in Ref. [82] and these variables generally relate to the shower shape and energy deposits in different EM calorimeter layers, hadronic leakage values, tracking quantities, and track-cluster matching variables. The signal PDF defines electrons arising from prompt decays, such as $Z \to ee$ decays, while the background PDF defines electrons coming from several sources including jet fakes, electrons from photon conversions, and electrons from semi-leptonic decays of heavy flavor hadrons.

Each electron candidate is assigned a discriminant of the form

$$d_L = \frac{L_S}{L_S + L_B} \tag{76}$$

which is transformed using an inverse sigmoid function

$$d'_L = -\frac{1}{15}\ln(d_L^{-1} - 1) \tag{77}$$

where d'_L provides the actual discrimination and allows for easier binning of the likelihood to define various operating points for signal identification. ATLAS defines several operating points, including VeryLoose, Loose, Medium, and Tight, with each step indicating an increase in the threshold required for the discriminant value.

3.6 Jet Reconstruction

Due to color confinement, isolated quarks do not propagate from the proton collisions outward to the detector; instead, the quarks undergo hadronization which results in a shower of stable particles. Decay products of short lived hadrons are embedded inside of these showers. In order to reconstruct these particle showers, an infrared and collinear safe jet reconstruction algorithm type known as sequential recombination

is used. In this reconstruction, a distance between input objects d_{ij} and the distance between and object and the beam (B) d_{iB} [83]

$$d_{ij} = \min(p_{\mathrm{T},i}^{2n}, p_{\mathrm{T},j}^{2n}) R_{ij}^2 \quad (78)$$

$$d_{iB} = p_{\mathrm{T,i}}^{2n} R_0^2 \quad (79)$$

where $p_{\mathrm{T,i}}$ is the transverse momentum of a particular object i, $R_i j$ is the ΔR separation between the input objects. The two parameters, R_0 and n, define the radius of the final jet size and the relative power of the momentum scale versus the geometrical scale respectively. The algorithm proceeds iteratively by combining objects (cluster or proto-jets) i and j if $d_{ij} < d_{iB}$; otherwise, if d_{iB} is the smaller of the two distances, the object i is identified as a jet and is removed from the list of objects queued for reclustering.

Several values of n have been used in high-energy particle physics experiments. The values of $n = 1$ corresponds to the k_t clustering algorithm [84], $n = 0$ corresponds to the Cambridge-Aachen algorithm [85, 86], and $n = -1$ corresponds to the anti-k_t algorithm [87]. These schemes prioritize different energy clusters for recombination, with Anti-k_t prioritizing the highest p_{T} clusters first, the Cambridge-Aachen algorithm having no preference on the p_{T} of the clusters, and the k_t algorithm prioritizing the lowest p_{T} cluster first. In this thesis, all jets are reclustered using the anti-k$_t$ algorithm .

For the physics analysis described in this thesis, there are three main algorithms used for jet reconstruction. These algorithms differ in the fundamental treatment of the tracking and calorimeter signals used to reconstruct hadronic jets.

Calorimeter jets: Calorimeter jet reconstruction begins with the clustering of adjacent calorimeter readout cells into topo-clusters, following the algorithm discussed in Section 3.4. All clusters with $E_{\mathrm{T}} > 0$ are kept as inputs for the jet reconstruction. The topo-cluster seeds are typically taken from the electromagnetic calorimeter scale and calibrated at that scale; because of this origin of topo-clusters, jets built using these inputs are typically referred to as *EMTopo* jets. After these clustering step, these jets are typically calibrated by associating tracks to the full jet cone and using this information to calibrate the response and remove pileup interactions.

Particle Flow jets: Particle flow jets are built out of "particle flow objects" (PFOs), which are objects that combine tracks and topo-cluster signals. PFOs are built by attempting to geometrically match a track to a topo-cluster deposit. The energy of this track is then subtracted from the topo-cluster signal via subtracting the readout cell energies that are associated to the track. If the topo-cluster signal remains positive after all associated tracks have removed their energy, the topo-cluster is

kept. The topo-cluster and tracks are then used in the recombination algorithm in order to form jets.

Track jets: Track jets, which are described further in Section 3.6.2, are reconstructed by inputing only the tracks into the reclustering algorithm and are independent of the topo-clusters and other calorimeter signals. The jets are not subject to the pileup jet rejection algorithm detailed in Section 3.6.1, as they are composed entirely of tracks associated to the primary vertex.

3.6.1 Pileup Jet Rejection

To suppress jets arising from pileup interactions in an event, a two variable likelihood called the jet-vertex-tagger (JVT) [88] is used. The JVT likelihood is build from the charged p_{T} fraction of a jet, i.e. the ratio of jet p_{T} associated with the primary vertex, $R_{p_{\mathrm{T}}}$ and the pileup corrected jet vertex fraction (corrJVF). The inputs to the JVT are given by [89]

$$\mathrm{corrJVF} = \frac{\sum_i p_{\mathrm{T}}^{trk_i}(\mathrm{PV}_0)}{\sum_j p_{\mathrm{T}}^{trk_j}(PV0) + \frac{\sum_n \sum_k p_{\mathrm{T}}^{trk_k}(PV_n)}{(k \times n_{trk}^{PU})}} \tag{80}$$

$$R_{p_{\mathrm{T}}} = \frac{\sum_i p_{\mathrm{T}}^{trk_i}(PV0)}{p_{\mathrm{T}}^{\mathrm{jet}}} \tag{81}$$

where $\sum_i p_{\mathrm{T}}^{trk_i}(PV0)$ is the scalar sum of the tracks in a jet associated with the primary vertex, $\sum_n \sum_k p_{\mathrm{T}}^{trk_k}(PV_n)$ is the scalar sum of the tracks in a jet associated with pileup vertices in the event, and the k factor is a weighting constant that corrects for the total number of pileup tracks in an event. The JVT distribution is given in Figure 18, where the JVT values range from 0 to 1, with JVT scores of 1 indicating that the jet is strongly associated with the primary vertex. A negative value of the likelihood is assigned if there are no tracks in the jet associated to the primary vertex.

3.6.2 Track Jets

Track jets are constructed from ID tracks that satisfy the 'Loose' tracking criteria given in Sec. 3.1. Further, the tracks must be associated to the PV in each event, which has been found to suppress track jets arising from pileup interactions.

In order to recluster these PV-associated tracks into a jet, a variable radius (VR) reclustering algorithm [83] is used. In VR reclustering, a non-trivial extension of the distance measure described in Eq. 79 is defined, which exchanges R_0 for a fixed cone

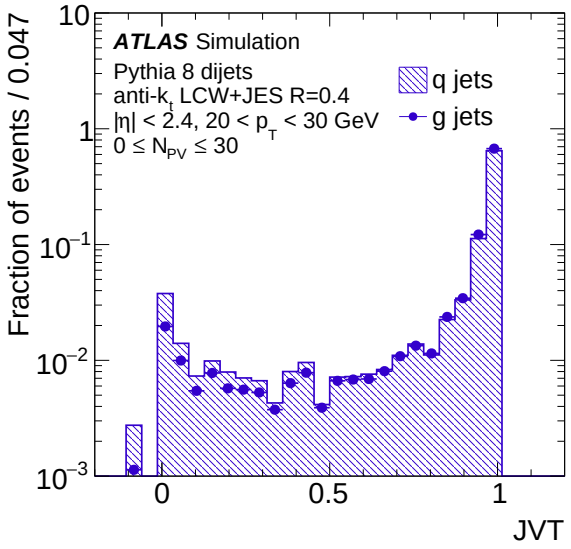

Figure 18: The Jet Vertex Tagger (JVT) distribution for Anti-k_t jets with $R = 0.4$ for dijet events where the jets have $20 \leq p_{\mathrm{T}} \leq 30$ GeV after calibration [89].

algorithm with an effective radius $R_{\mathrm{eff}}(p_{\mathrm{T,i}})$. This effective radius is calculated for each d_{iB} and only clusters four-momenta (only for $n \leq 0$) if the following condition is satisfied

$$\frac{d_{ij}}{d_{iB}} = \frac{R_{ij}^2}{R_{\mathrm{eff}}}(p_{\mathrm{T,i}})^2 < 1 \tag{82}$$

For the track jet implementation in ATLAS, the anti-k$_t$ algorithm is used and the effective radius is defined as

$$R_{\mathrm{eff}} = \max\left[0.02, \min\left(0.4, \frac{\rho}{p_{\mathrm{T}}}\right)\right] \tag{83}$$

where $\rho = 30$ GeV is chosen. This choice of ρ ensures that the track jets form a $R = 0.4$ jet for track jet $p_{\mathrm{T}} < 75$ GeV and the choice of this minimum value also enforces that extremely low-p_{T} jets do not form large jets. The additional absolute minimum $R = 0.02$ is chosen as a cut-off to ensure that extremely high-p_{T} jets remain infrared safe.

Track jets do not contain the full jet energy information, as only charged particles deposit tracks as they pass through the ATLAS tracking system. This lack of neutral

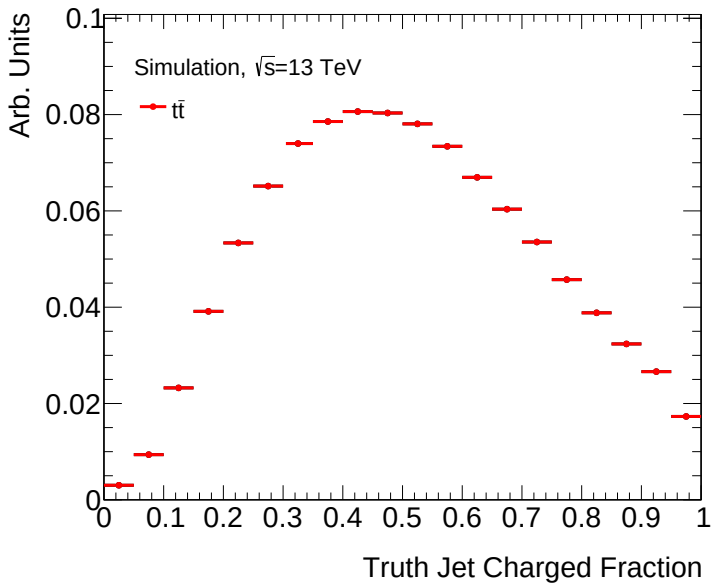

Figure 19: The truth charged fraction of b-jets matched to track-jets with $7 \leq p_{\mathrm{T}} \leq 13$ GeV for simulated $t\bar{t}$ events. The width of this distribution affects any correction factor from the tracking scale to the truth scale and can result in large uncertainties on the final track-jet energy scale.

hadronic response means that these jets are difficult to calibrate to the true jet energy scale, as the neutral component varies by a large amount. As an example of this variation, the charged fraction for b-jets in simulated top quark pair production ($t\bar{t}$) events is shown in Figure 19.

The main benefit of track jets is the resistance to pileup interactions, brought on by the choice of tracks as inputs for the clustering. Therefore, lower momentum decays can be reconstructed when compared to calorimeter or particle flow jets, which are restricted to $p_{\mathrm{T}} > 20$ GeV. Further details on the importance of this low momentum identification are discussed in Section 6.1.

3.7 B-Jet Identification

In order to improve the identification of the $H \to bb\mu\mu$ signal from other background processes with jets, the jets from the Higgs system can be identified as originating from b-hadron decays. This identification is called *b-tagging*. B-tagging relies on several properties of b-quark hadronization: the relatively long lifetime ($c\tau \sim 10^{-12}$

s) which gives rise to large impact parameters (IP) of the decay products and a secondary vertex (SV) in the jet; the higher p_{T} of the heavy hadron decay products compared to light hadrons leading to higher fragmentation from a heavy hadron; and the relatively large branching fraction to semi-leptonic decays through the weak interaction which produces low p_{T} leptons. To target these properties, ATLAS has developed several B-hadron identification algorithms that rely on SV identification, IP determination in 2 and 3 dimensions, and an algorithm known as JetFitter which relies on identification of B→C transitions [90].

All of these algorithms are used as inputs to a BDT-based, multi-variate (MV) tagging algorithm known as MV2. The MV2 algorithm is trained on MC simulation, where simulated events containing a large fraction of B-hadron initiated jets are treated as the signal in the BDT and simulated events containing jets coming light-flavor quarks and gluons (light-flavor jets) are treated as background. In addition to the light-flavor jets, a small percentage of jets originating from prompt charm hadrons are added to the background. This addition of charm jets improves the b-tagging rejection of charm jets, as C-hadron lifetimes, impact parameters, and decay vertices are quite similar to the B-hadron properties. The nominal choice of parameters used by ATLAS, called MV2c10, is trained on a sample with $\approx 7\%$ charm jet composition in the background sample [91].

3.8 Missing Transverse Energy Reconstruction

The missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$) in ATLAS is calculated from two distinct inputs denoted as the hard-object term and the soft-object term [78]. The generalized form of the hard-object contribution is calculated by taking the negative vectorial sum of all reconstructed and calibrated electrons, photons, τ-leptons, muons, and jets in a given event. The soft-object contribution is calculated by taking all of the ID tracks associated with the primary vertex in the event, but not associated to any object in the hard-object term. These two contributions are described by

$$\begin{aligned} \overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}(\mathrm{Hard}) &= -\sum \overrightarrow{p}_{\mathrm{T}}^{e} - \sum \overrightarrow{p}_{\mathrm{T}}^{\mu} - \sum \overrightarrow{p}_{\mathrm{T}}^{\tau} - \sum \overrightarrow{p}_{\mathrm{T}}^{\gamma} - \sum \overrightarrow{p}_{\mathrm{T}}^{\mathrm{j}} \\ \overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}(\mathrm{Soft}) &= -\sum \overrightarrow{p}_{\mathrm{T}}^{\mathrm{trk}} \\ E_{\mathrm{T}}^{\mathrm{miss}} &= |\overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}(\mathrm{Hard}) + \overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}(\mathrm{Soft})| \end{aligned} \tag{84}$$

where the summations are taken as vector sums over all identified and calibrated particles in each category for a given event. The first $\overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}$ contribution is meant to account for loss of energy from invisible particles (neutrinos for example) by looking for the imbalance of energy with respect to the visible particles identified in the

event. The soft-object term is important to improve the $\overrightarrow{p}_{\mathrm{T}}^{\mathrm{miss}}$ scale and resolution by accounting for at least some of the energy from particles which fail to meet the reconstruction criteria for a hard object [78].

To prevent double-counting of the same calorimeter energy associated to multiple objects, an overlap removal procedure (Section 4.5.5) is performed based on a priority scheme. For the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation, a sequential ordering of objects is performed. In general, the sequence prioritizes electrons, then photons, followed by hadronically decaying τ-leptons, and then finally jets for the calorimeter signal contributions.

As a variety of identification criteria are used in ATLAS analyses, the $E_{\mathrm{T}}^{\mathrm{miss}}$ given by Eq. 84 may not be the form used by most analyses. In the analysis detailed in Section 4, for example, the hard-object term is only calculated with respect to the electrons, muons, and jets.

4 Searching for exotic Higgs decays

4.1 General Analysis Strategy

The exotic Higgs decay mode considered in this analysis is the 2b2μ final state arising when one a decays to a b-quark pair and the other a decays to a muon pair. The mass range for the new a bosons considered here is 18 - 62 GeV, so the signal is characterized by two oppositely charged muons, two b-tagged jets, a relatively small amount of missing transverse energy, and the dimuon and di-bjet masses being below the Z mass.

The Standard Model background processes with the same or similar experimental signature at the LHC are the following:

- **Drell-Yan + jets background**: The Drell–Yan (DY) + jets background can have the same experimental signature if $\mathrm{Z}^{(*)}/\gamma^* \to \mu\mu$ or $\mathrm{Z}/\gamma \to \tau\tau$ and the leptonic decays of the taus do not produce a significant amount of $E_{\mathrm{T}}^{\mathrm{miss}}$. The jets are produced via radiative processes, which predominantly yield light flavor jets. A significant cross-section, however, allows for the production with two b-jets, yielding a difficult background. This background contribution can be reduced via b-jet identification techniques and also by vetoing the $m_{\mu\mu} > 75$ GeV, which removes the majority of the Z production cross section.

- **Top-quark backgrounds**: Another prominent background arises from the dileptonic decays of top-quark pairs $t\bar{t}$, where t→Wb and each W decays leptonically, $W \to \ell\nu$. In addition, a similar signature exists for a top quark produced in association with a W boson, Wt. These decays contain the entire bb$\mu\mu$ signature and typically also have a significant amount of $E_{\mathrm{T}}^{\mathrm{miss}}$ from the unmeasured neutrino energy.

- **Diboson backgrounds**: The dibosonic processes WW, WZ, and ZZ give rise to the same experimental signature for certain decays. The ZZ process, $ZZ \to qq\ell\ell$, can give rise to $bb\mu\mu$ directly when the quarks are b-quark decays; alternatively, when the Z decays to light or charm quarks, these jets can be misidentified as b-jets to enter the analysis selections. The WZ decays $WZ \to qq\ell\ell$ or $WZ \to qq\ell\nu$ can enter into the analysis selections either through misidentification of light jets as b-jets or via a fake muon being identified. Finally, the WW process $WW \to \mu\nu\mu\nu$ can enter into the selections if it is produced in association with two extra jets. The cross sections for all of these processes are quite small and the misidentification of two charm or light jets lowers the probability of detection by an additional factor of 40 - 10^4 based

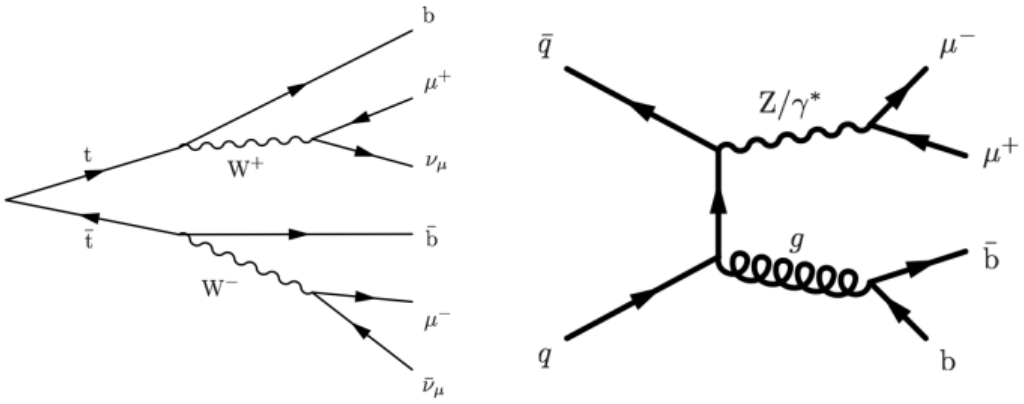

Figure 20: The leading order Feynman diagrams for top quark pair production (left) and the Drell–Yan production in associated with b-jets (right).

on the MV2c10 mistag rates. These factors, coupled with the muon isolation and Z mass veto, effectively suppress these background contributions.

- **W boson background**: Due to the relatively low-p_{T} scale of the leptons in the $H \to bb\mu\mu$ signature, the W+jets process can contribute to the signal region selections if $W \to \mu\nu$ and an additional jet is misidentified as a muon. In order to veto this background, selections based on lepton isolation are applied. Further selections requiring b-jet identification and low $E_{\mathrm{T}}^{\mathrm{miss}}$ also suppress the contribution of this process.

- **tt+V background**: Although the top quark pair production with a vector boson $t\bar{t} + X$ where $X = (\mathrm{H,Z,W,}\gamma)$ does provide a similar experimental signature, it is characterized by a larger number of final state objects. In addition, the cross sections for these processes are quite small and any analysis selection aiming to veto $t\bar{t}$ will render this background negligible.

The background process cross sections are given in Table 8. Based on these cross sections and the baseline selections discussed in Section 4.6, the dominant backgrounds come from the Drell–Yan + jets and $t\bar{t}$ backgrounds. The Feynman diagrams for these two dominant backgrounds are shown in Figure 20.

Process	$\sigma_X \times$ BR [pb]	Order of Calculation	Reference
Dileptonic $t\bar{t}$	87.6	NNLL	[92, 93]
$Z/\gamma^* \to \ell\ell$, 10 GeV$< m_{\ell\ell} <$ 40 GeV			
without b-partons	2319.2	NNLO+NNLL	[94]
only b-partons	87.1	NNLO+NNLL	[94]
$Z/\gamma^* \to \ell\ell$, $p_T(V) \leq$ 500 GeV			
with only light partons, $m_{\ell\ell} > 40$ GeV	1689.9	NNLO+NNLL	[94]
with only charm partons, $m_{\ell\ell} > 40$ GeV	251.7	NNLO	[94]
with only b partons, $m_{\ell\ell} > 40$ GeV	147.3	NNLO	[94]
$Z/\gamma^* \to \ell\ell$, $p_T(V) \geq$ 500GeV, inclusive jet	1.9	NNLO	[94]
W$\to \ell\nu$, $p_T(V) \leq$ 500 GeV with light jets	16172.5	NNLO	[94]
W$\to \ell\nu$, $p_T(V) \leq$ 500 GeV with charm jets	2777.9	NNLO	[94]
W$\to \ell\nu$, $p_T(V) \leq$ 500 GeV with b jets	951.5	NNLO	[94]
W$\to \ell\nu$,$p_T(V) \geq$ 500 GeV, inclusive jet	15.8	NNLO	[94]
top + W boson dileptonic	7.6	NNLO	[95, 96]
top + Z boson, t-channel, >= 1 lepton	0.24	NLO	[97]
$t\bar{t} + W$, inclusive	0.6	NLO	[98, 99]
$t\bar{t} + Z(\nu\nu)$	0.2	NLO	[98, 99]
$t\bar{t} + Z(qq)$	0.6	NLO	[98, 99]
$t\bar{t} + Z(ee)$	0.04	NLO	[98, 99]
$t\bar{t} + Z(\mu\mu)$	0.04	NLO	[98, 99]
$t\bar{t} + Z(\tau\tau)$	0.04	NLO	[98, 99]
ZZ$\to qq\nu\nu$	4.3	NLO	[34, 100, 101]
ZZ$\to qq\ell\ell$	2.2	NLO	[34, 100, 101]
ZZ$\to 4\ell$	1.3	NLO	[34, 100, 101]
ZZ$\to 4\nu$	0.6	NLO	[34, 100, 101]
WZ,ZZ$\to 2\ell 2\nu$	12.5	NLO	[34, 100, 101]
WZ$\to \ell\ell\ell\nu$	4.6	NLO	[34, 100, 101]
WZ$\to \ell\nu\nu\nu$	3.2	NLO	[34, 100, 101]
WZ$\to qq\nu\nu$	6.8	NLO	[34, 100, 101]
WZ$\to qq\ell\ell$	3.4	NLO	[34, 100, 101]
WZ$\to \ell\nu qq$	11.4	NLO	[34, 100, 101]
WW$\to qq\ell\nu$	49.4	NLO	[34, 100, 101]

Table 8: Background process cross-sections used in this analysis for a subset of the backgrounds simulated. The production process, order of calculation, production cross section, and references for calculation are given. The W and Drell–Yan predictions are binned as a function of jet flavor and vector boson p_T and given as per lepton flavor cross sections.

4.2 Theoretical Cross Sections and Modeling

4.2.1 Signal Cross Section

As our signal assumes that the Higgs boson at 125 GeV is SM-like, the theoretical cross section for the SM Higgs boson is used for the signal production cross section. The full Higgs boson cross section is calculated at varying order for the different production processes. A further separation is made between the calculation accuracy depending on the addition of gluonic loops and electroweak processes. Table 9 gives the order of each calculation and the central value of the cross section for each Higgs production mode used in the analysis.

Process	σ_H [pb]	$\Delta_\sigma^{\text{scale}}$ [%]	$\Delta_\sigma^{\text{PDF}+\alpha_s}$ [%]	Reference
Gluon Fusion	48.52	± 3.9	± 3.2	[16, 102–104]
Vector Boson Fusion	3.776	$^{+0.4}_{-0.3}$	± 2.1	[105–108]
pp→WH	1.369	$^{+0.5}_{-0.7}$	± 1.9	[105, 109–112]
pp→ZH	0.8824	$^{+3.8}_{-3.0}$	± 1.6	

Table 9: Higgs signal cross-sections used in this analysis. The production process, order of calculation, production cross section, and references for calculation are given.

4.2.2 Signal and Background Modeling

The signal and background processes in this analysis are modeled using Monte Carlo (MC) generation techniques as generally described in Section 1.4. Due to the challenges of modeling all physics processes with a single MC technique, the specific details for the choice of generator, parton shower, and PDF choices for the signal and background processes are given below.

The gluon fusion Higgs signal process is modeled at next-to-leading order (NLO) using POWHEG-BOX v2 [43, 44, 113] interfaced with Pythia 8.212 [45] using the AZNLO set of tuned parameters [114] for the simulation, parton showering, and hadronization of the $H \to aa \to bb\mu\mu$ decay. The vector boson fusion production mode for the signal process is also modeled with the same generator and shower parameters. The MC samples from these processes are then weighted to the higher order cross-sections given in Section 4.2.1. Five mass points are simulated in the range $m_a = 20 - 60$ GeV in steps of 10 GeV for both the gluon-fusion and VBF production modes.

The top background processes are modeled using several generators. The $t\bar{t}$, single-top processes are generated using POWHEG-BOX v2 [43] using the CT10

PDF set [35] interfaced with PYTHIA v6.428 [115] and the Perugia 2012 set of tuned parameters [116] for the parton shower. In these samples, the mass of the top quark (m_t) is set to 172.5 GeV. In order to regulate the high-p_{T} radiation in the top decays, the h_{damp} parameter in POWHEG is set to m_t which has been shown to improve the agreement between ATLAS data and simulation in the high p_{T}^t region [117]. The $t\bar{t}$ cross-section was calculated at NNLO in QCD including resummation of next-to-next-to-leading logarithmic (NNLL) soft gluon terms [92, 93]. The cross-section for the single top-quark production processes was calculated using the prescriptions described in Refs. [95, 96]. The production of $t\bar{t}$ in association with W or Z bosons (denoted by $t\bar{t}+V$) was modeled with samples generated at leading order using MADGRAPH5_AMC@NLO v2.2.2 [42] and showered with PYTHIA v8.186. These samples are then normalized to the NLO cross-sections computed in Refs. [98, 99].

SHERPA 2.2.1 [34] with the NNPDF3.0 [37] PDF set was used for the generation of Drell–Yan (Z/γ^*)+jets, W+jets and diboson (WW, WZ, ZZ) processes. The MC samples for the $Z^{(*)}/\gamma^*$+jets and W+jets samples are explicitly generated at NLO accuracy for events with 0-2 additional partons in the matrix element and at LO accuracy for events with up to 4 additional partons [118]. These samples additional select only the leptonic decays of the $Z^{(*)}/\gamma^*$ and W. The samples are weighted to cross-sections calculated at NNLO QCD accuracy for $Z^{(*)}/\gamma^*$+jets and W+jets production [94] and at NLO including LO contributions with two additional partons for the diboson processes [34, 100, 101].

All MC samples included in the nominal background estimation are processed through the full ATLAS detector simulation [119] based on GEANT4 [120]. Every sample also has additional proton-proton collisions overlaid in each event to simulate the effect of pile-up interactions; these collisions are generated at leading order with generated with PYTHIA v8.186, using the A2 set of tuned parameters [121] and the MSTW2008LO PDF set [122]. Finally, the reconstruction and analysis selections applied on the data are the same selections that are applied to the MC samples.

A summary of the Monte Carlo generators, cross section, and parton shower configurations for the processes consider in this analysis are given in Table 10.

4.2.3 Monte Carlo Correction Factors

Despite using the same reconstruction and selection techniques in the simulation and data, there are several known deficiencies that must be corrected in the simulation. Correction factors to address these deficiencies are derived in a series of dedicated control samples. These correction factors can address the response differences in the calorimeter energy determination that affect jet energy calibration scales (Sec-

Process	MC Generator	Parton Shower	PDF
$H \to aa \to bb\mu\mu$	POWHEG-BOX v2	PYTHIA 8.212	CTEQ6L1
$t\bar{t}$,tWb	POWHEG-BOX v2	PYTHIA 6.428	CT10
$tq + Z/\gamma^*$	MadGraph5_aMC@NLO 2.2.3	PYTHIA 6.428	CTEQ6L1
$t\bar{t}$ + W/Z	MadGraph5_aMC@NLO 2.2.1	PYTHIA 8.186	NNPDF 2.3
(Z/γ^*)+jets, W+jets WZ/ZZ/WW	Sherpa 2.2.1	–	NNPDF 3.0

Table 10: List of MC generators, parton shower programs, and PDFs used to simulate the signal and background processes in this analysis. For the Sherpa generated processes, Sherpa uses its own internal parton showering program.

tion 3.6) and identification. Other examples include the lepton identification, isolation, association to the primary vertex, and trigger efficiencies (Section 4.5.2 for muon examples); these efficiencies can differ due to changes in the response or the differences in the alignment of the ID and MS between data and MC simulation.

As the average number of interactions per bunch crossing $\langle\mu\rangle$ is not known a priori, the values used for simulation must be corrected during analysis to reflect the actual running conditions in the data. This correction is performed by reweighting the generated $\langle\mu\rangle$ profile in the simulation to the distribution found in the data as a function of data-taking period (and taking into account the good running conditions described in Section 4.3). This "pile-up reweighting" correction is performed on each sample individually.

4.3 LHC Dataset and Data Quality Requirements

This analysis uses the partial Run-2 dataset, corresponding to the ATLAS data collected in 2015 and 2016 at a center of mass energy of $\sqrt{s} = 13$ TeV. The integrated luminosity associated with this dataset after data quality selections is 36.1 fb^{-1} with an uncertainty of 2.1%. In order to ensure good data-taking conditions, the following data quality requirements must be met:

Good Run List (GRL): Under typical data-taking conditions, each ATLAS run is associated to an LHC fill and divided into units of data-taking time called lumiblocks. Each lumiblock lasts for approximately one minute. The GRL is a list of lumiblocks for each ATLAS run that is prepared globally (for all/most analyses) and requires that all data meets (sub)detector-level data quality requirements to be classified as "good for physics". These detector level requirements typically disallow subdetector

defects, propagated via a software flag, that are given in further detail below. Further, the GRL requires that the ATLAS detector has certain nominal running conditions such as "stable beams" being declared by the LHC and that the ATLAS magnets are running at the nominal currents. Finally, the GRL requires that the data come only from LHC collisions with 25 ns bunch spacing, as approximately 0.1 fb^{-1} of 50 ns data was collected in 2015.

Detector Data Quality Defects: Events within a lumiblock can be affected by problems associated with specific subdetectors of the ATLAS experiment. These events are removed if an "intolerable" defect is assigned. For example, noise bursts in the LAr calorimeter during data-taking may be flagged as intolerable if the noise is non-transient; in such a case, the lumiblock could be labeled with a severe noise burst defect and all events in that lumiblock would be removed from the GRL. Additionally, non-nominal voltages in either calorimeter can also lead to defects being assigned to specific lumiblocks, via the larError and tileError flags.

The largest subdetector fault affecting this analysis was the loss of the IBL for data-period H1 during 2015, when the IBL was turned off. This loss was on the order of 0.2 fb^{-1}. The largest loss overall came from the ATLAS toroid being off during several data-taking periods in 2016. This loss accounted for 0.6 fb^{-1}. These two losses are already accounted for and were removed in order to obtain the 36.1 fb^{-1} dataset.

In certain cases, a subdetector data buffer may fill up due to a fault in the subdetector which could lead to a halt in data-taking. In these cases, recovery of the data-taking conditions can occur by restarting the trigger, timing, and control (TTC) system of ATLAS. This TTC restart procedure prevents a large amount of data-loss, as the TTC restart time is often much shorter in duration than starting an entire new run of the ATLAS detector. As data-taking has halted, however, event fragments from some detectors must be deleted and the events removed from the GRL. An incomplete event flag is used to veto such events in the analysis.

4.4 Trigger Selections

For this analysis, only the combination of lowest unprescaled[8] single muon triggers are used. The exact ATLAS trigger names and types are listed in Table 11. The triggers with isolation are discussed in the next paragraph and the OR of the isolated and non-isolated triggers is used in order to maximize efficiency for higher p_{T} leptons.

[8]Unprescaled means that all events which pass the requirements set in the trigger are accepted. Prescaling a trigger with a factor of 10, for example, would mean that only 1 event in every 10 that pass the requirements are accepted for further processing.

Trigger Type	ATLAS Trigger Name
Single Muon	HLT_mu{20,24}_iloose HLT_mu{24,26}_ivarmedium HLT_mu{40,50}

Table 11: The single lepton HLT trigger names for the triggers used in the 36 fb^{-1} analysis. The analysis selection involves a logical "OR" of all triggers in the category to account for isolation and identification inefficiencies at different momentum scales. The braced terms refer to different thresholds used to control trigger rates as the LHC delivered higher luminosities to the ATLAS experiment during different run periods. The "iX" names refer to the isolation details referred to in Section 4.4. Further details on the triggers can be in Refs. [123, 124].

During the data-taking period from 2015-2016, the number of pileup interactions increased significantly from a peak $\langle\mu\rangle = 15$ in 2015 to a peak $\langle\mu\rangle = 40$ by the end of 2016. As such, the trigger selections for the single muon trigger were tightened to reduce the total HLT output rate to a manageable level of 1-2 KHz. In order to prevent high losses in the single muon triggers, track isolation (Section 4.5.1) was applied to the muon identified by the single muon trigger. This track isolation rejects events coming from non-prompt muon decays and helped keep the p_{T} threshold from increase dramatically. This isolation definition, therefore, requires at least one of the leptons in the final analysis selections to be well isolated.

For the $H \to aa \to bb\mu\mu$ signal, the single lepton trigger efficiency is found to be 67 - 71% with respect to the inclusive simulated sample. By requiring the simulated events to contain at least one muon with transverse momentum on the plateau of the trigger ($p_{\mathrm{T}} > 27$ GeV), the signal efficiency is found to be around 80% which corresponds to the L1 efficiency loss. The increase in signal efficiency by using additional dilepton triggers for the 2015-2016 data-taking period was found to be on the order of 5-10%. The use of these additional triggers increases the complexity of the analysis, but a further discussion can be found in Section 6.2 with regards to the full Run-2 dataset result.

4.5 Object Selections

4.5.1 Muon Selection

Muons are identified using the "Medium" quality criteria discussed in Section 3.3.1, which enforces that the muon momentum measurement uses the combined track

fits from the ID and MS tracking information. The muon momentum is refined by correcting for energy lost in the calorimeters and by applying a momentum energy scale and resolution correction to the simulated MC events. As a baseline selection, all muon candidates are required to be within the acceptance of the ID ($|\eta| < 2.5$) and have $p_T > 7$ GeV.

In addition, the muons are required to be identified with the primary vertex via tighter tracking selections on the longitudinal and transverse impact parameters. The selection requires that the transverse impact parameter significance (d_0/σ_{d_0}) of the primary muon track be less than 3. The selection on the longitudinal impact parameter requires that $|z_0^{\mathrm{BL}} \times \sin(\theta_{\mathrm{trk}})|$ be less than 0.5 mm for the same primary muon track. These additional track selections reduce the presence of non-prompt muon backgrounds coming from pileup or cosmic interactions.

To further remove the hadronic and non-prompt backgrounds, muons must satisfy both calorimeter and track based isolation requirements. The calorimeter isolation requires that any additional calorimeter energy within a fixed cone of size $\Delta R = 0.2$ (referred to as *topoetcone20*) around the muon, excluding the energy contained in the muon, is required to be less than a specified (p_T, η) dependent value. These values are tuned to be 95% efficient for prompt muons and are correspond to the ATLAS "Tight" isolation [125].

The track based isolation uses a similar cone isolation requirement as the calorimeter isolation, but the track isolation uses a variable radius cone which requires $\Delta R = \max(0.3, \frac{10 \text{ GeV}}{\mathrm{p_T}})$ rather than the fixed cone size used in the calorimeter isolation. This quantity is known as the $pTvarcone30/p_{\mathrm{T}}$. The exact values are again (p_T, η) dependent and are tuned to be 99% efficient for prompt muons from Z$\rightarrow \mu\mu$ and J$/\psi \rightarrow \mu\mu$ events. The (p_{T}, η) values of the isolation cuts used in this analysis can be found in Appendix A.

As described in Section 4.4, the isolation on the muon triggers requires that the trigger matched muon must additionally satisfy the trigger isolation requirement. This tracking isolation requirement varies as a function of data-taking period. During 2015, the trigger track isolation requires there to be $< 12\%$ of p_{T} not associated with muon within a fixed radius of $R = 0.2$. During 2016, this trigger isolation requires instead the $pTvarcone30/p_{\mathrm{T}}$ to be < 0.06. The trigger-matched muon is therefore typically required to be more isolated than the Tight isolation.

A summary of the analysis selections for baseline and signal muons can be found in Table 12. These requirements select signal muons with an identification efficiency of $\sim$94% and the isolation efficiency ranges from 91% for $m_a = 20$ GeV to 95% for $m_a = 60$ GeV.

Selection	Baseline Requirement	Signal Requirement
Reconstruction	Combined	Combined
Reconstruction Quality	Medium	Medium
Momentum Acceptance	p_{T} >7 GeV	$p_{\mathrm{T}} > \{27, 7\}$ GeV
ID Acceptance	$\lvert\eta\rvert < 2.5$	$\lvert\eta\rvert < 2.5$
Calorimeter Isolation	(-)	Tight, $\varepsilon = 95\%$
Track Isolation	(-)	Tight, $\varepsilon = 99\%$
Trigger Matching	Required for ≥ 1 muon	Required for ≥ 1 muon
d_0/σ_{d_0}	< 3	< 3
$\lvert z_0^{\mathrm{BL}} \times \sin(\theta_{trk})\rvert$	< 0.5 mm	< 0.5 mm

Table 12: Summary of muon selection Requirements for the $bb\mu\mu$ analysis. (-) denotes that there is no requirement applied for this quantity.

4.5.2 Muon efficiencies and scale factors

The muons used in this analysis have scale factors derived for a variety of quantities, including the lepton track-to-vertex association (TTVA), the muon isolation, the muon reconstruction, and the muon trigger scale factor. The muon efficiency can be factorized into its individual components

$$\varepsilon_{\mathrm{Total}} = \varepsilon_{\mathrm{reco}} \times \varepsilon_{\mathrm{TTVA}} \times \varepsilon_{\mathrm{trig}} \times \varepsilon_{\mathrm{Iso}} \tag{85}$$

where the total efficiency is a product of the reconstruction and identification (reco), trigger (trig), track-to-vertex association (TTVA), and isolation (iso) efficiencies. Each of these efficiencies is typically determined sequentially with respect to one another. The reconstruction efficiencies are determined first in the data, then the trigger efficiencies, and finally the isolation efficiencies. The final MC simulation events are corrected for each muon using the combination given above.

The muon efficiencies are measured using tag-and-probe methods on $Z \to \mu\mu$ which are used to correct muons with $p_{\mathrm{T}} > 15$ GeV while $J/\psi \to \mu\mu$ events are used to extract the corrections for muons with $p_{\mathrm{T}} > 7$ GeV. These measurements are described in detail in Ref [75]. In the tag-and-probe method, muon pairs from the Z or J/ψ decay are selected and one muon is required to follow a specific quality requirement and the other is "probed" in order to determine its efficiency for a given reconstruction quality. As an example for the reconstruction scale factor, one muon is required to be reconstructed as a Medium quality muon and the other is allowed to be a Loose, Medium, Tight, LowPt, or HighPt Muon in order to determine the efficiency for selecting each of these quality muons.

Selection	Requirement
Reconstruction	Combined
Reconstruction Quality	Medium LH
Momentum Acceptance	p_{T} >7 GeV
ID Acceptance	$\|\eta\| < 2.47$, excluding 1.37 $\|\eta\| < 1.52$
Calorimeter Isolation	(-)
Track Isolation	(-)
Trigger Matching	(-)
d_0/σ_{d_0}	< 5
$\|z_0^{\mathrm{BL}} \times \sin(\theta_{trk})\|$	< 0.5 mm

Table 13: Summary of the electron selection requirements for the $bb\mu\mu$ analysis. (-) denotes that there is no requirement applied for this quantity.

4.5.3 Electron Selection

While electrons are not explicitly used during the event selection, the electron candidates are used to calculate quantities such as the $E_{\mathrm{T}}^{\mathrm{miss}}$ and are considered during the object overlap removal. As a baseline, the electron candidates are required to be within the geometrical acceptance of high-granularity portion of the EM calorimeter ($|\eta| < 2.47$), outside of the cryostat crack region ($1.37 < |\eta| < 1.52$). Further, all electrons must be identified as passing the "Medium" likelihood value, which removes the hadronic jets reconstructed as electrons and removes electrons from photon conversions. The electron energy scale correction is derived using data CRs that select the decays $Z \to ee$ and $J/\psi \to ee$ and these corrections smear the energy of the electrons in the MC simulation. After this energy calibration and resolution smearing, all electron candidates are requires to have $p_{\mathrm{T}} > 7$ GeV.

In addition, the electrons are required to be identified with the primary vertex in a similar way to the muons. The requirements on the longitudinal impact parameter are the exact same as for the muon case, but the transverse impact parameter significance of the primary electron track is only required to be less than 5. No further selections for the electrons are required, as the isolation is not considered during the overlap removal procedure. A summary of the analysis selections for baseline electrons are given in Table 13. Additionally, no efficiency scale factors are applied to events with electrons as the analysis is insensitive to these effects.

4.5.4 Jet Selection

In the 2015-2016 analysis, jets are reconstructed from topological calorimeter deposits that are clustered using the topo-cluster algorithm discussed in Section 3.4 and reconstructed using the anti-k_t algorithm with a radius parameter of R = 0.4 using the implementation in the FASTJET [126] package. Jets are calibrated using jet energy scale corrections and jet resolution smearing derived in data CRs. This calibration procedure is detailed in Ref. [127]. After this calibration, jets are required to have $p_T > 20$ GeV and must be within the tracking acceptance, $|\eta| < 2.5$.

Jets that arise from non-collision background, i.e. from beam induced background or cosmic-ray interactions with the ATLAS detector, are rejected via a series of cleaning requirements. These cleaning requirements rely on the ratio of jet energy deposited in single layers of the electromagnetic and hadronic calorimeters. The amount of charged particle content in a jet measured by the ID tracker is also used for further discrimination power. The exact quality requirements for discarding a jet ("BadLoose" requirement) can be found in Ref. [128]. These cleaning requirements are found to have a jet selection efficiency on hard-scatter jets of greater than 99%.

The JVT selection, detailed in Section 3.6.1, is applied to all jets with $20 \leq p_T \leq 60$ GeV and contained mostly within the ID tracking volume $|\eta| < 2.4$ in order to suppress jets coming from pileup interactions. The JVT selection requires that the jets have JVT > 0.59. This cut results in an efficiency of signal jets of close to 100% and has a fake rate on the order of 1% for pileup jets in this kinematic region [89]. Any jet with $p_T \geq 60$ GeV is not subject to the JVT requirement.

An event cleaning requirement is also applied in this analysis, due to the use of E_T^{miss} in the analysis selections. It is recommended to veto any event if a jet in the event passes the JVT selection, but is labeled as "BadLoose" jet coming from other detector effects. This requirement prevents large mismeasurements of the E_T^{miss} due to the effects from beam backgrounds.

Jets consistent with decay of b-hadrons (b-jets) are identified using the standard ATLAS multivariate MV2 discriminant trained against a background sample containing 7% charm-initiated jets as described in Section 3.7. The MV2 algorithm is calibrated for several standard working points corresponding to different b-tagging efficiencies. This analysis uses the 77% b-jet identification efficiency working point as it gives the highest signal significance over the background. For the 77% efficiency working point, the mis-identification probability of tagging a jet containing charm quarks as a b-jet is determined from simulated $t\bar{t}$ events and found to be approximately 16%. A similar mis-identification of light jets under the same conditions is found to be negligible (< 1%) [91]. The identification of b-jets following the baseline analysis selections is found to have a purity of approximately 95%.

Selection	Requirement
Kinematic Acceptance	$p_{\mathrm{T}} > 20$ GeV
Geometric Acceptance	$\lvert\eta\rvert < 2.5$
Pileup Suppression	JVT > 0.59 for jets with $p_{\mathrm{T}} \leq 60$ GeV, $\lvert\eta\rvert < 2.4$
B-jet Identification	77% WP, MV2c10_weight > 0.65

Table 14: A summary of the jet selection requirements in this analysis. The b-jet identification is applied to all jets, but is used as a classifier rather than a strict selection requirement. Therefore, the jets which do not pass the identification are labeled as "light-jets" and the passing jets are labeled as "b-jets" in the analysis.

A summary of the jet selection requirements can be found in Table 14. It should be noted that the b-tagging requirement is used to classify jets as b-jets or not b-jets; therefore, all jets in an event are not required to pass this identification.

4.5.5 Overlap Removal Criteria

For the objects which pass the baseline selections given in Sections 4.5.1-4.5.4, an overlap removal procedure removes the ambiguity between objects which get identified by more than one object criterion. The overlaps are performed sequentially using the distance measure ΔR_y, as defined in Section 2.3.1, and proceed as follows: electron-muon overlaps are performed, then jet-electron overlaps, and finally jet-muon overlaps are determined. After each step in this overlap procedure, the overlapping object that is not retained is removed from the event. As an example, if a muon is rejected in the electron-muon overlap, it is not considered in the jet-muon overlap. In order to maximize the lepton and b-jet identification efficiencies, a b-jet specific overlap procedure is used. This procedure keeps jets from being removed if a low p_{T} lepton is identified inside the jet cone, thereby saving the jets with semi-leptonic B hadron decays.

For the electron-muon pairs, if the muon is identified as a calorimeter-tagged muon and shares an ID track with the electron, the muon is removed from the event. Otherwise, if the muon is not calo-tagged and the pair share an ID track, the electron is removed from the event. For electron-jet pairs, if $\Delta R_y(\text{electron, jet}) \leq 0.2$ and the jet is not b-tagged, the electron is kept and the jet is removed from the event. Otherwise, if $\Delta R_y(\text{electron, jet}) \leq \min(0.4,\ 0.04 + 10\ \text{GeV}/p_{\mathrm{T}}^{e})$, the jet is retained and the electron is removed from the event. Finally, for muon-jet pairs, if the jet has fewer than three tracks with $p_{\mathrm{T}} > 500$ MeV or $\Delta R_y(\text{muon, jet}) \leq 0.2$ and the jet is

not b-tagged, the muon is retained and the jet is removed from the event. Otherwise, if $\Delta R_y(\text{muon, jet}) \leq \min(0.4, 0.04 + 10 \text{ GeV}/p_{\text{T}}^{\mu})$, the jet is retained and the muon is removed from the event.

4.5.6 Missing Transverse Momentum Selection

The missing transverse momentum ($E_{\text{T}}^{\text{miss}}$) used in the 2015-2016 analysis is calculated as the magnitude of the negative vector sum of all calibrated and selected objects which pass the overlap removal criteria in the events. Therefore, the $E_{\text{T}}^{\text{miss}}$ is determined using all of the jets, muons, and electrons in each event. In addition, the soft track term described in Section 3.8 is also included in the $E_{\text{T}}^{\text{miss}}$ calculation for each event. In order to calculate the soft track term, the track inputs must meet the requirements given in Section 3.1, including the impact parameter requirements to be matched to the primary vertex.

4.6 Event Selection

The production of the $H \to aa \to 2b2\mu$ process gives a final state characterized by two jets identified as b-jets, two oppositely signed muons with an invariant mass less than $m_H/2$, and a negligible amount $E_{\text{T}}^{\text{miss}}$. Further, the bb and $\mu\mu$ decays arise from the same massive particle, and so these events can be constrained to have equivalent masses of the two systems. These kinematic properties are targeted by the analysis selections given in the following sections.

4.6.1 Preliminary Selections

The preliminary selections, referred to as the *preselection*, of the $2b2\mu$ analysis require the following:

1. Events must pass the GRL selections, detector data quality flags, and must pass the jet event cleaning described in Section 4.5.4.

2. All events must pass the logical OR of the lowest threshold single muon triggers, which includes triggers require muon isolation and triggers which have higher p_{T} thresholds that do not require muon isolation. Additionally, one of the selected muons in the event is required to match the trigger based muon that caused the event to be accepted in the data. This trigger matching requirement cuts down on the amount of hadronic fakes firing the trigger and allows for the muon trigger efficiency scale factors to be applied appropriately.

3. Each event must contain at least one reconstructed vertex, with the primary vertex (PV) defined as the vertex with the largest sum of the track momentum squared $\sum_{\text{trk}} p_{\text{T}}^2$. The PV is required to contain ≥ 3 ID tracks that are consistent with the LHC beam envelope.

4. All muons, jets, and electrons are selected according to the selections given in Section 4.5 and the overlap removal selections described in Section 4.5.5 are applied.

5. After overlap removal, each event must have exactly 2 reconstructed muons with opposite charge. The leading muon in the event is further required to have $p_T > 27$ GeV, which allows for the highest threshold, lowest unprescaled single muon trigger, mu26_ivarmedium, to be maximally-efficient with respect to the trigger threshold. The subleading muon is simply required to pass the baseline selections given in Table 12.

6. Every event must contain exactly 2 muons based on the requirements above to reject potential dibosonic backgrounds. Further, at least 2 jets meeting the selection criteria are required for every event.

7. The mass of the dimuon pair is then restricted to the range $16 < m_{\mu\mu} < 64$ GeV. The upper bound assumes that the 125 GeV Higgs boson decays into two on-shell particles of equal masses (62.5 GeV with a dimuon mass resolution of 2 GeV) while the lower bound is motivated by the kinematics of the simulated a-boson decays. For lower values of m_a, many of the subleading signal jets have p_{T} lower than the reconstruction threshold; further, the jets coming from a-decay in this mass regime tend to overlap geometrically in the detector, so the selection efficiency of requiring ≥ 2 jets lowers the sensitivity of the analysis to the $H \to aa$ signal. The dimuon mass restrictions also vetos the on-shell Z mass window, removing many background events associated with the DY+jets background.

After these selection cuts, the overall background composition plots are shown in Figures 21-22 for several kinematic variables. The background composition after these selections is dominated by Drell–Yan+jet events which accounts for over 80% of the background yield. The overall composition of the background and signal processes in this inclusive region is given in Table 15.

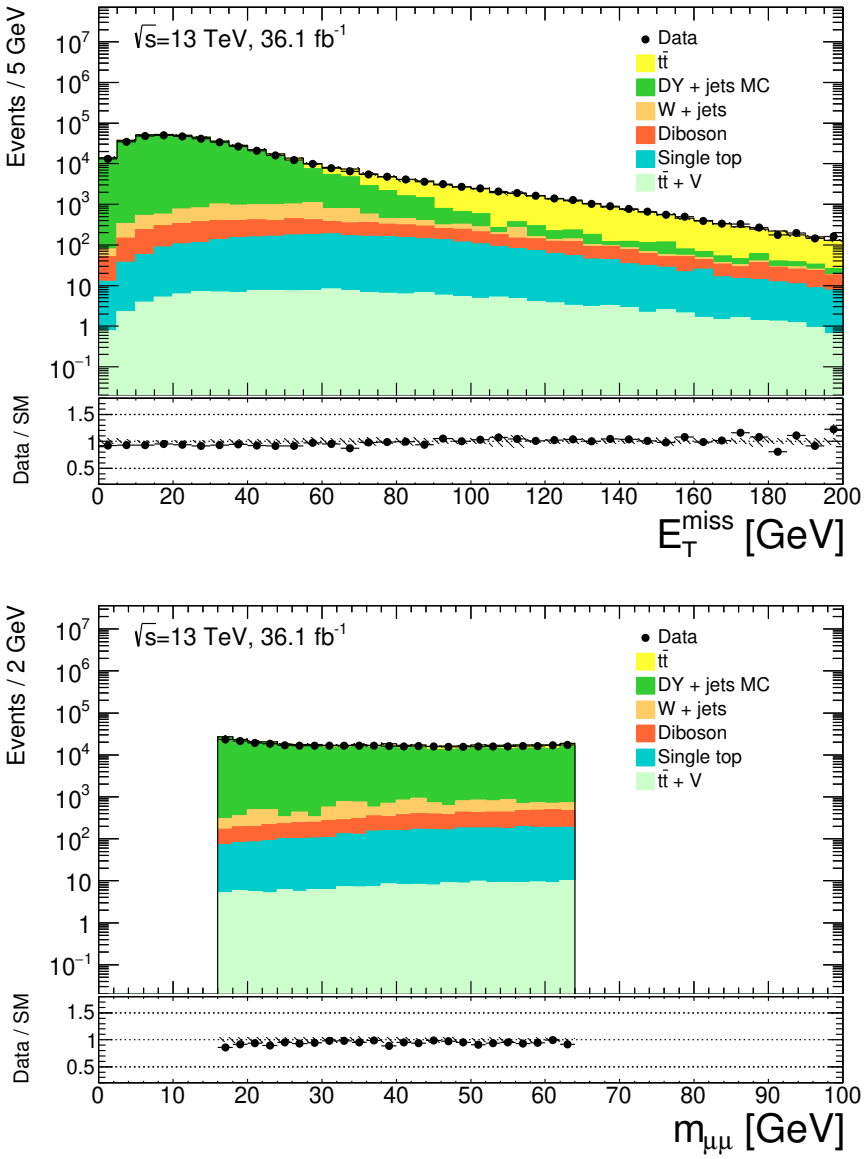

Figure 21: $E_{\mathrm{T}}^{\mathrm{miss}}$ and $M_{\mu\mu}$ distributions after the analysis preselection. The uncertainty bands are shown only for the MC statistical uncertainties. The ratio plots display the ratio of the data over the Standard Model (SM) background predictions.

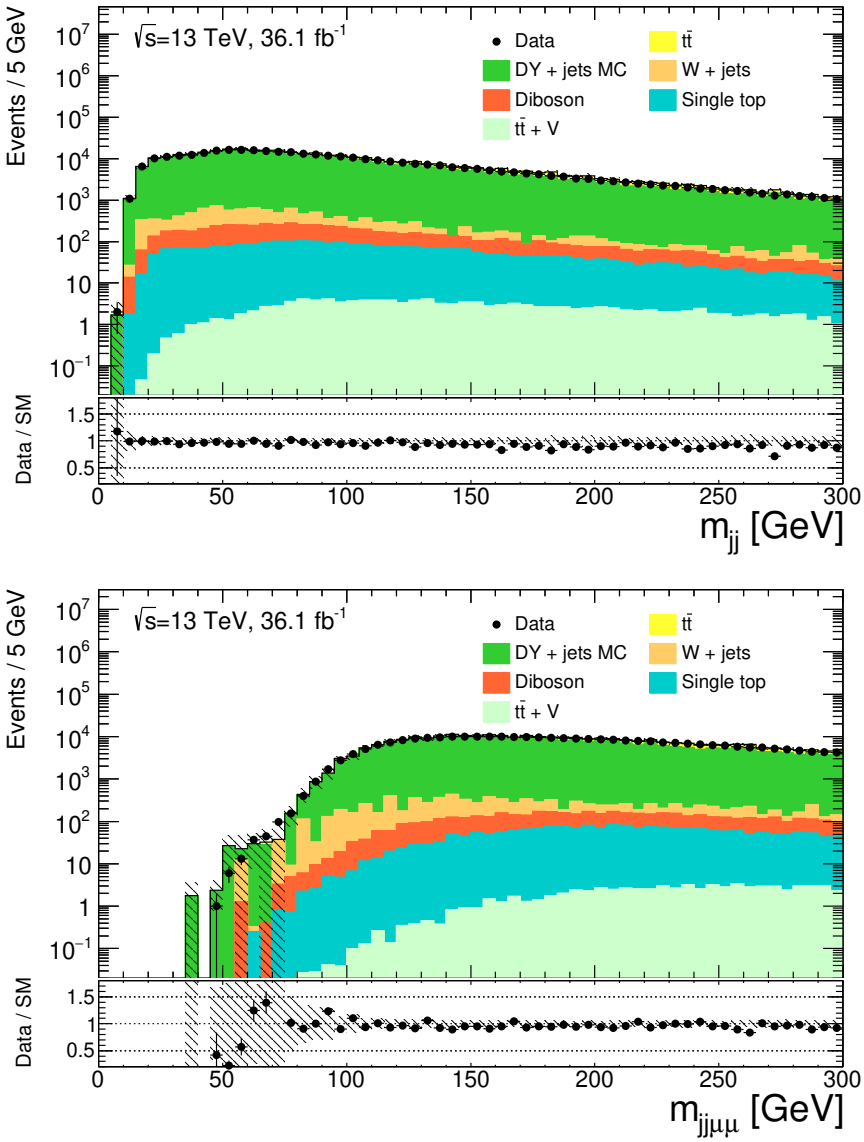

Figure 22: M_{jj} and $M_{jj\mu\mu}$ distributions after the analysis preselection. In these plots, the two highest p_{T} jets are chosen to build the invariant mass of the dijet system. The uncertainty bands are shown only for the MC statistical uncertainties. The ratio plots display the ratio of the data over the Standard Model (SM) background predictions.

Physics Process	Fraction of Events [%]
DY+jets	83.9
$t\bar{t}$	12.7
W+jets	1.6
Diboson	1.1
Single top	0.7
tt + V	< 0.1
Signal	0.1

Table 15: Composition of the background and signal processes, estimated from MC simulation samples, after the preselection requirements have been applied. The Drell–Yan and $t\bar{t}$ backgrounds are found to dominate in this region, with minor contributions from other backgrounds.

4.6.2 Signal Region Selections

Further analysis selections aim to exploit the kinematics of the $H \to aa \to bb\mu\mu$ decays. These selections boost the significance of the signal by rejecting the large DY and $t\bar{t}$ backgrounds, while maintaining a high efficiency for the signal sample. The signal region (SR) selections include the following requirements:

b-jet Identification The DY+jets process is typically characterized by the presence of light quark or gluon jets, as these are the largest contributing diagrams at the LHC. In order to suppress this background, 2 jets in each event are required to be identified as b-jets. The selection explicitly requires the b-jet multiplicity, $n_{\mathrm{b-jets}} = 2$, with the b-jets identified using the 77% efficiency working point of the MV2c10 tagger discussed in Section 3.7. The rejection power of this selection on the DY background is illustrated in Figures 23-26, where the $n_{\mathrm{b-jets}} = 0$ and $n_{\mathrm{b-jets}} = 2$ selections are separated. The $n_{\mathrm{b-jets}} = 0$ selection will be discussed in Section 4.8, where it provides a region used for the estimation of the DY background.

Missing Transverse Energy selections After suppression of the Drell-Yan background seen in Figures 25-26, the dominant background is coming from the $t\bar{t}$ process which preferentially has 2 b-jets in the event. The analysis selection then relies on the fact that neutrinos coming from leptonic decays of the top quark carry a significant amount of energy. Therefore, the top background events are characterized by having a significant amount of $E_{\mathrm{T}}^{\mathrm{miss}}$. The $E_{\mathrm{T}}^{\mathrm{miss}}$ distributions for the signal and background processes after the preselection and b-jet identification selection can be seen in Figure. 27. In order to maximize the selection efficiency of the signal and suppress the top quark initiated backgrounds, an upper cut of $E_{\mathrm{T}}^{\mathrm{miss}} < 60$ GeV is placed. The

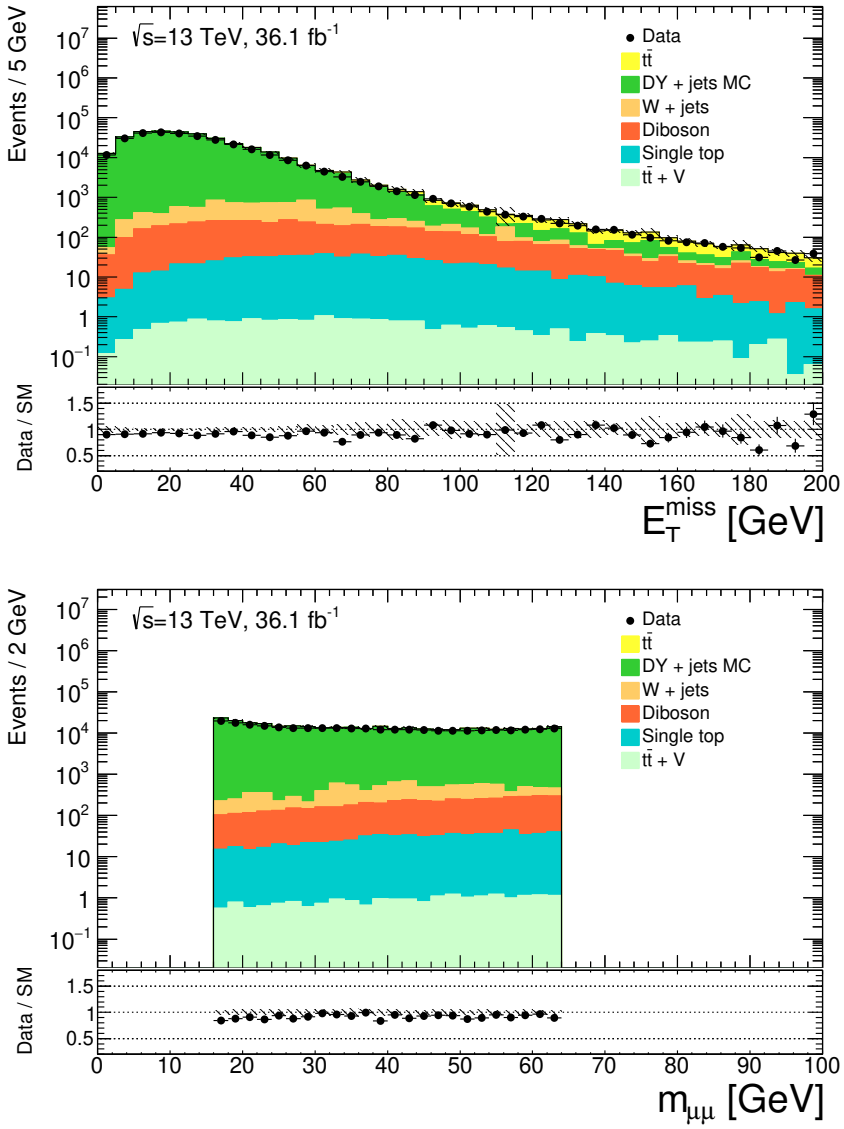

Figure 23: $E_{\mathrm{T}}^{\mathrm{miss}}$ and $M_{\mu\mu}$, distributions after the analysis preselection requirements and an additional selection requiring $n_{\mathrm{b-jet}} = 0$. The uncertainty bands are shown only for the MC statistical uncertainties.

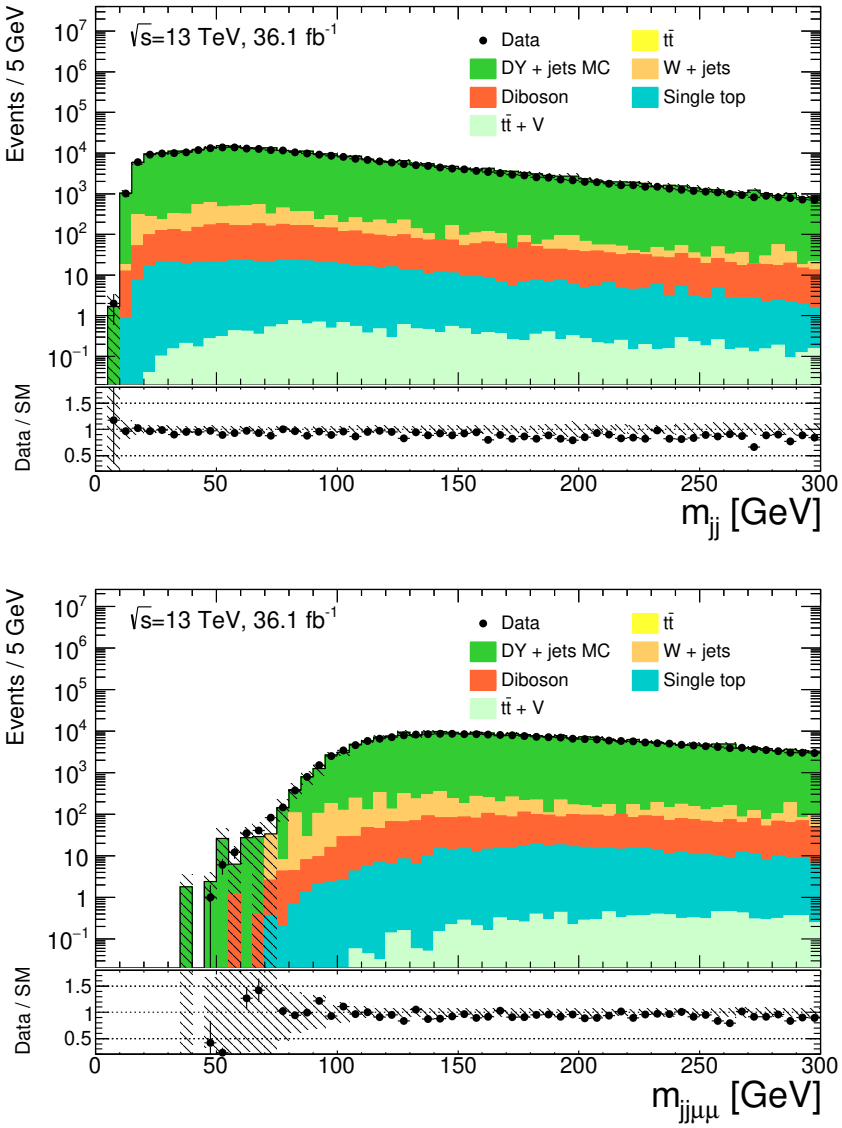

Figure 24: M_{jj} and $M_{jj\mu\mu}$ distributions after the analysis preselection requirements and an additional selection requiring $n_{\mathrm{b-jet}} = 0$. The uncertainty bands are shown only for the MC statistical uncertainties.

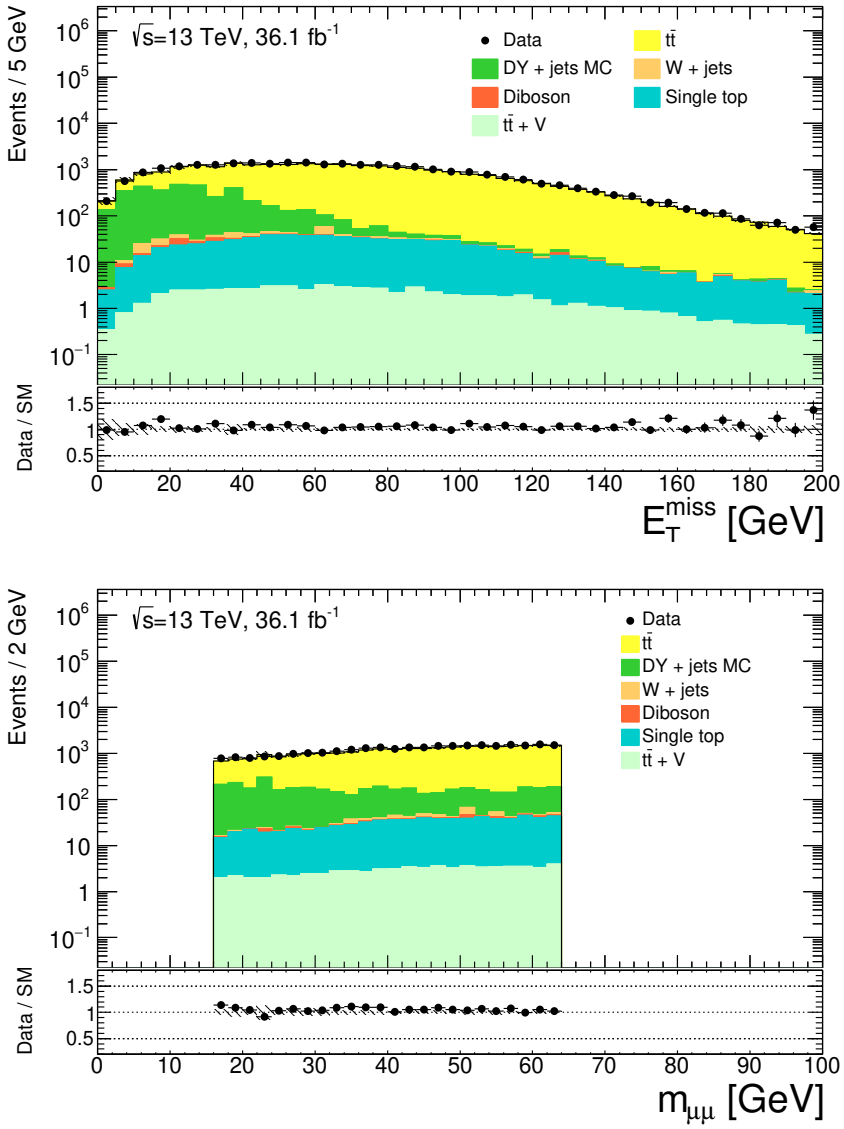

Figure 25: $E_{\mathrm{T}}^{\mathrm{miss}}$ and $M_{\mu\mu}$, distributions after the analysis preselection requirements and an additional selection requiring $n_{\mathrm{b-jet}} = 2$. The uncertainty bands are shown only for the MC statistical uncertainties.

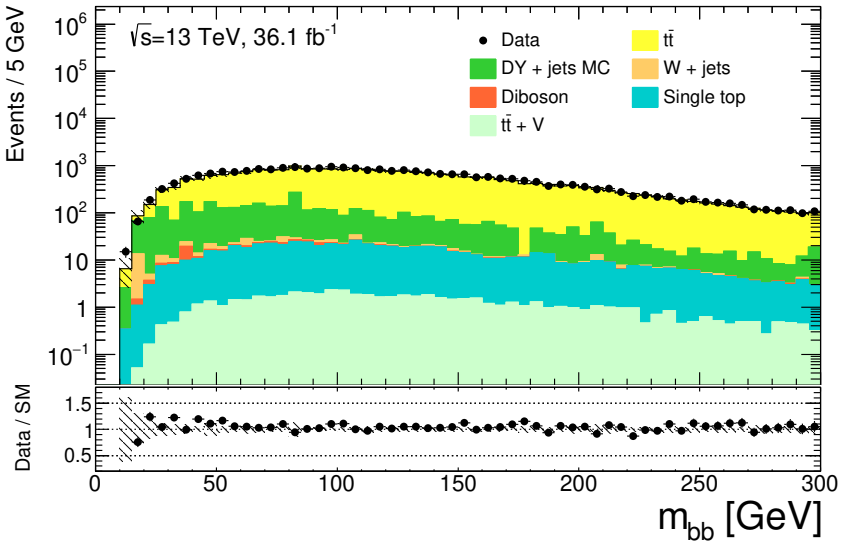

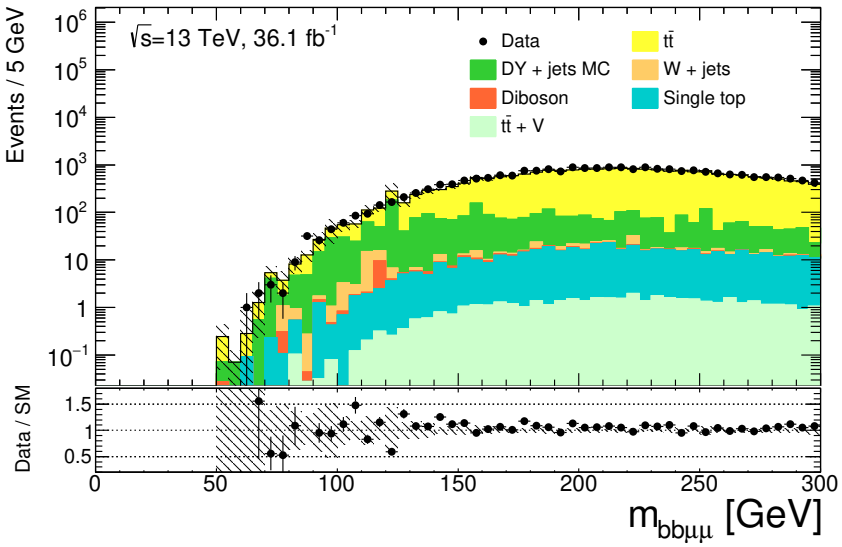

Figure 26: M_{bb} and $M_{bb\mu\mu}$ distributions after the analysis preselection requirements and an additional selection requiring $n_{\mathrm{b-jet}} = 2$. The uncertainty bands are shown only for the MC statistical uncertainties.

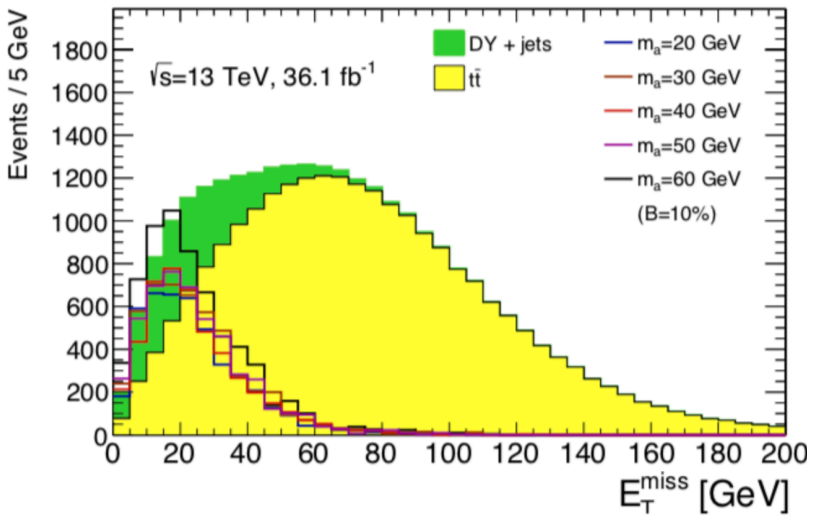

Figure 27: $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution after preselection and 2 b-tag multiplicity requirements, using MC simulation only. The distribution illustrates the low $E_{\mathrm{T}}^{\mathrm{miss}}$ of the signal and Drell–Yan + jets background. Here, the exotic Higgs signal is normalized to 10% of the total Higgs cross section for illustration purposes.

resulting distributions after the $E_{\mathrm{T}}^{\mathrm{miss}}$ are given in Figures 28-29. After this cut, the Drell–Yan MC statistics begin to show signs of poor modeling, indicating the need for the template background estimation discussed in Section 4.8.1.

Kinematic Selections Targeting the Higgs Mass After including the b-jet and $E_{\mathrm{T}}^{\mathrm{miss}}$ selections, the analysis selections exploit the kinematics of the exotic Higgs decays by using the masses of the di-bjet, dimuon and $bb\mu\mu$ systems.

The dimuon side of the H$\rightarrow aa$ decay ($a \rightarrow \mu\mu$) is measured with a more precise resolution compared to the di-bjet side of the decay (H$\rightarrow bb$). This resolution difference arises from the calorimeter response given in Section 2.3.3 and also the fragmentation of the jets into a large area of the calorimeter. As an example, for a 20 GeV jet in the calorimeter, the energy resolution is $\sigma_{\mathrm{E}} \approx 11\%$, while a 20 GeV muon has a resolution of $\sigma_{\mathrm{p_T}} \approx 2\%$. The resolution differences for the masses are illustrated in Fig. 30

As the majority of 125 GeV Higgs bosons produced via gluon-fusion are produced with a small amount of transverse momentum [16], it is expected that the b-jets and muons will carry $\mathcal{O}(30)$ GeV of energy (approximately $m_H/4$). This low jet momentum scale results in a significantly broader reconstructed m_{bb} compared to the true m_{bb} of the $a \rightarrow bb$ decay. As the di-bjet resolution becomes broader, the

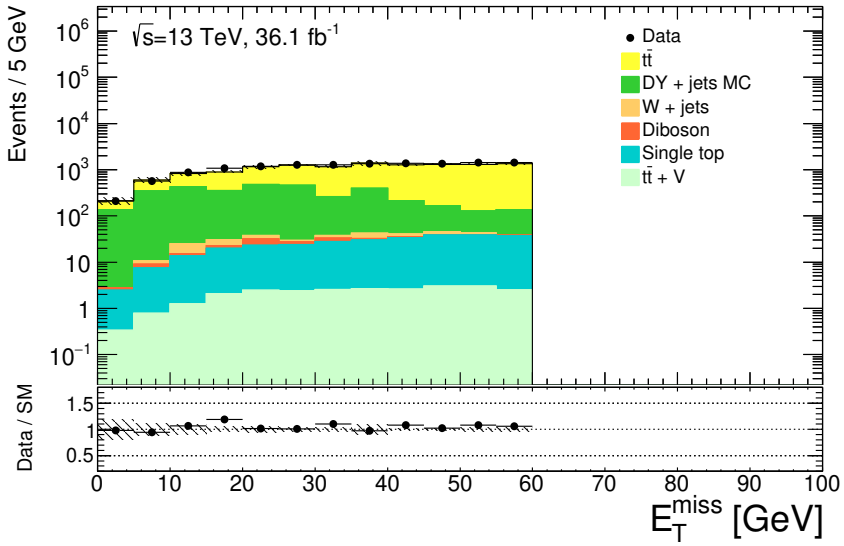

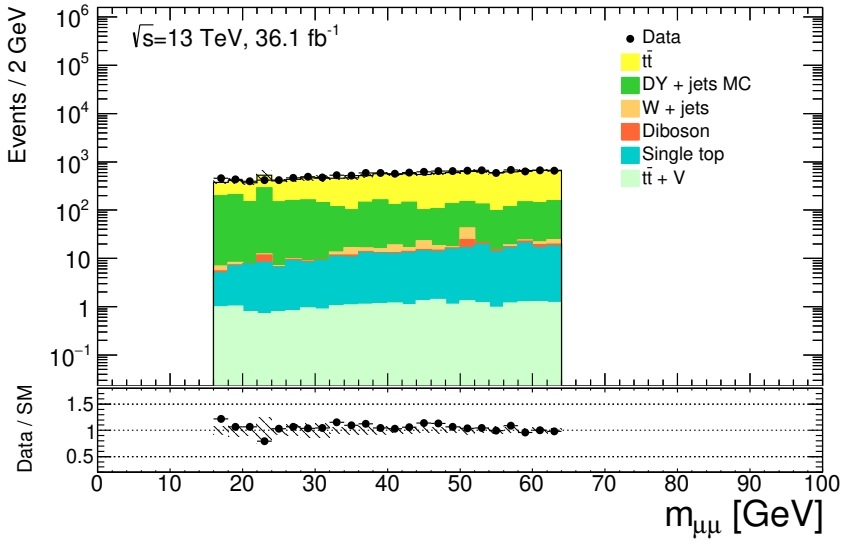

Figure 28: $E_{\mathrm{T}}^{\mathrm{miss}}$ and $M_{\mu\mu}$ distributions after the analysis pre-selection and requiring $n_{\mathrm{b-jet}} = 2$, $E_{\mathrm{T}}^{\mathrm{miss}} < 60$. The uncertainty bands are shown only for the MC statistical uncertainties.

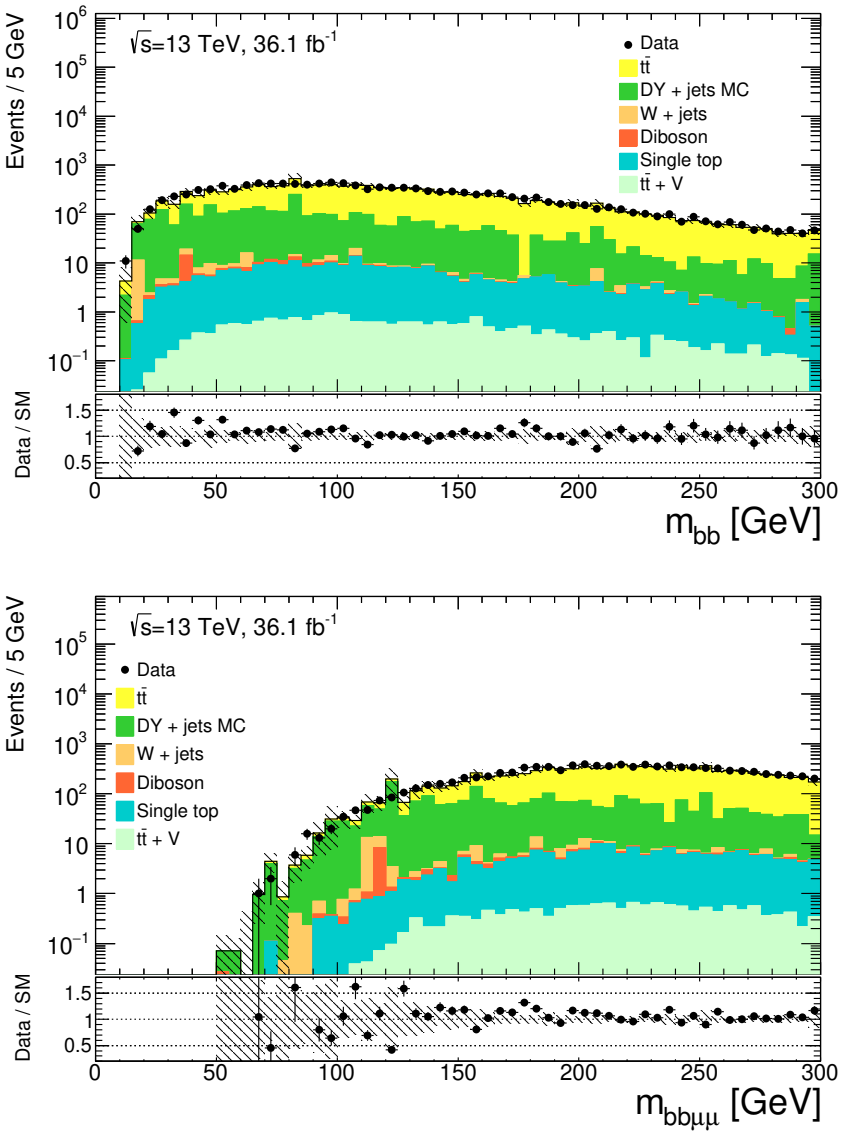

Figure 29: M_{bb} and $M_{bb\mu\mu}$ distributions after the analysis pre-selection and requiring $n_{\mathrm{b-jet}} = 2$, $E_{\mathrm{T}}^{\mathrm{miss}} < 60$. The uncertainty bands are shown only for the MC statistical uncertainties.

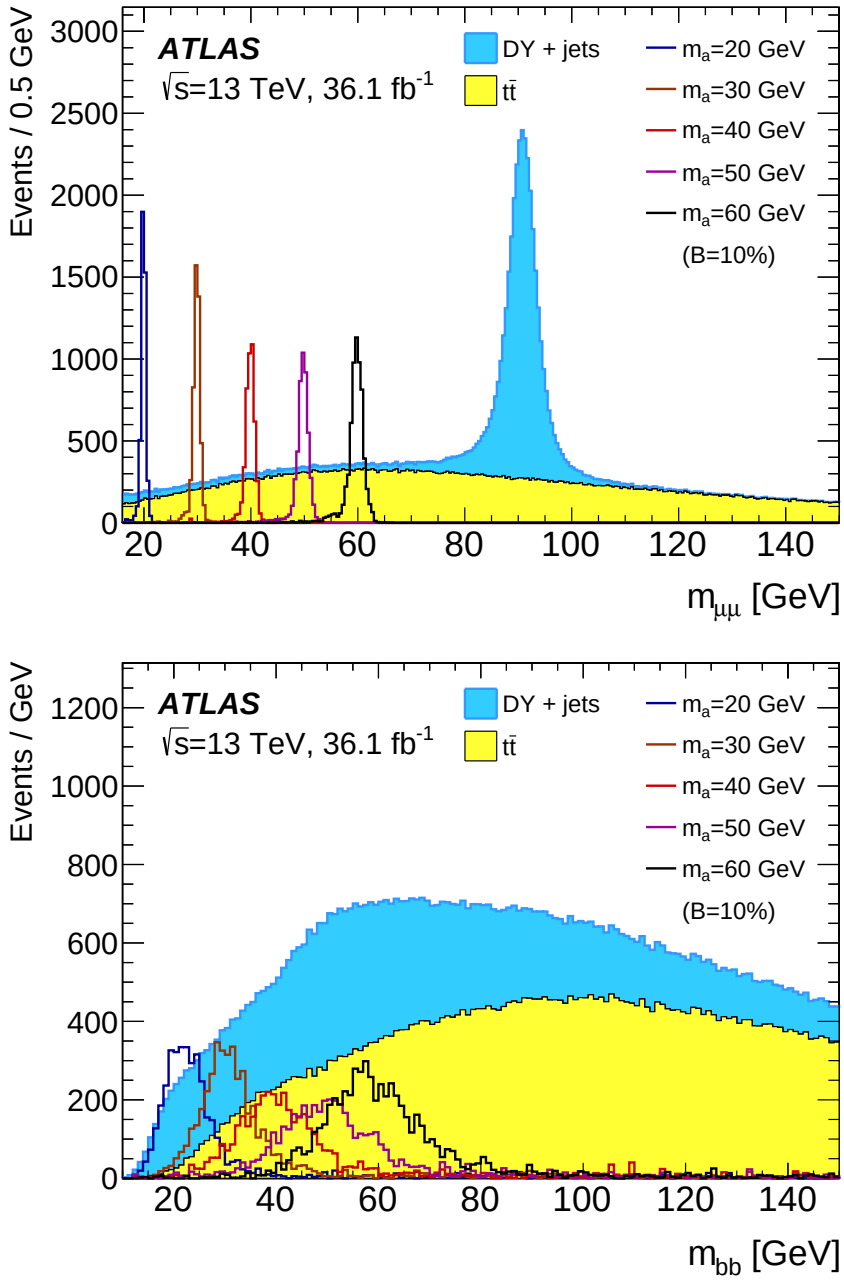

Figure 30: The mass distributions for the Drell–Yan, $t\bar{t}$, and $H \to aa \to bb\mu\mu$ signals for the dimuon (top) and di-bjet (bottom) decays.

resolution of the reconstructed Higgs boson mass m_H will also become broader and the phase space cuts cannot target the Higgs mass with a high efficiency. The distributions for the $m_{bb\mu\mu}$ mass for the signal and background Monte Carlo samples show this trend in Figure 32

To better target the $H \to aa$ decays, a kinematic likelihood (KL) fit implemented in the KLFitter package [129] is used. This fit aims to exploit the equivalence of the two pseudoscalar masses by constraining m_{bb} to the dimuon system by testing the compatibility of the hypothesis that $m_{bb} = m_{\mu\mu}$.

The likelihood takes the form

$$L = W(\hat{E}_{b_1}, E_{b_1}) \cdot W(\hat{E}_{b_2}, E_{b_2}) \cdot F_{\mathrm{BW}}(m_{bb}^{KL}, m_{\mu\mu}) \tag{86}$$

where $W(\hat{E}_i, E_i)$ denotes transfer functions that take the measured energy E_i of the leading or subleading jet energy and adjust these values to obtain $\hat{E}_i$ which correspond to the jet energies at the maximum value of the likelihood. The $\hat{E}_i$ in the maximum likelihood case act as a proxy for the true jet energy and the varied energies are used to compute the adjusted di-bjet mass, m_{bb}^{KL}. $F_{\mathrm{BW}}(m_{bb}^{\mathrm{KL}}, m_{\mu\mu})$ denotes a relativistic Breit-Wigner function centered at the value of $m_{\mu\mu}$ in the event under test with a width parameter $\Gamma = 0.5$ GeV. This Γ value is chosen to be small compared to the intrinsic resolution of the dijet mass. The transfer functions are parametrized as double Gaussian probability density functions that are derived from MC simulated events. The transfer functions are derived as a function of reconstructed jet p_{T} and are binned in $|\eta|$ to account for different detector energy responses. Further details about the exact form and values used in the double Gaussian parameterization can be found in Appendix B. After the KLFit is performed, the resolution of m_{bb} is matched to the dimuon resolution as shown in Figure 31 and the resolution of $m_{bb\mu\mu}$ improves by up to a factor of 2 depending on the value of m_a as shown in Figure 32.

After these improvements from the KLFit, several selections are made on the derived quantities to further target the $H \to aa \to bb\mu\mu$ signal. These consist of the following selections:

- Higgs Mass constraint - As the signal arises from the 125 GeV Higgs boson, a selection is placed on the 4-object mass in order to reject background outside of this mass region. The selection explicitly requires events to be within 15 GeV of the Higgs mass after the KLFit is performed on the b-jets, $|m_{bb\mu\mu}^{KL} - 125| < 15$ GeV.

- aa hypothesis: As the a bosons are identical in mass, a further requirement on the $m_{bb} \approx m_{\mu\mu}$ hypothesis can further reject the background. As the KLFitter

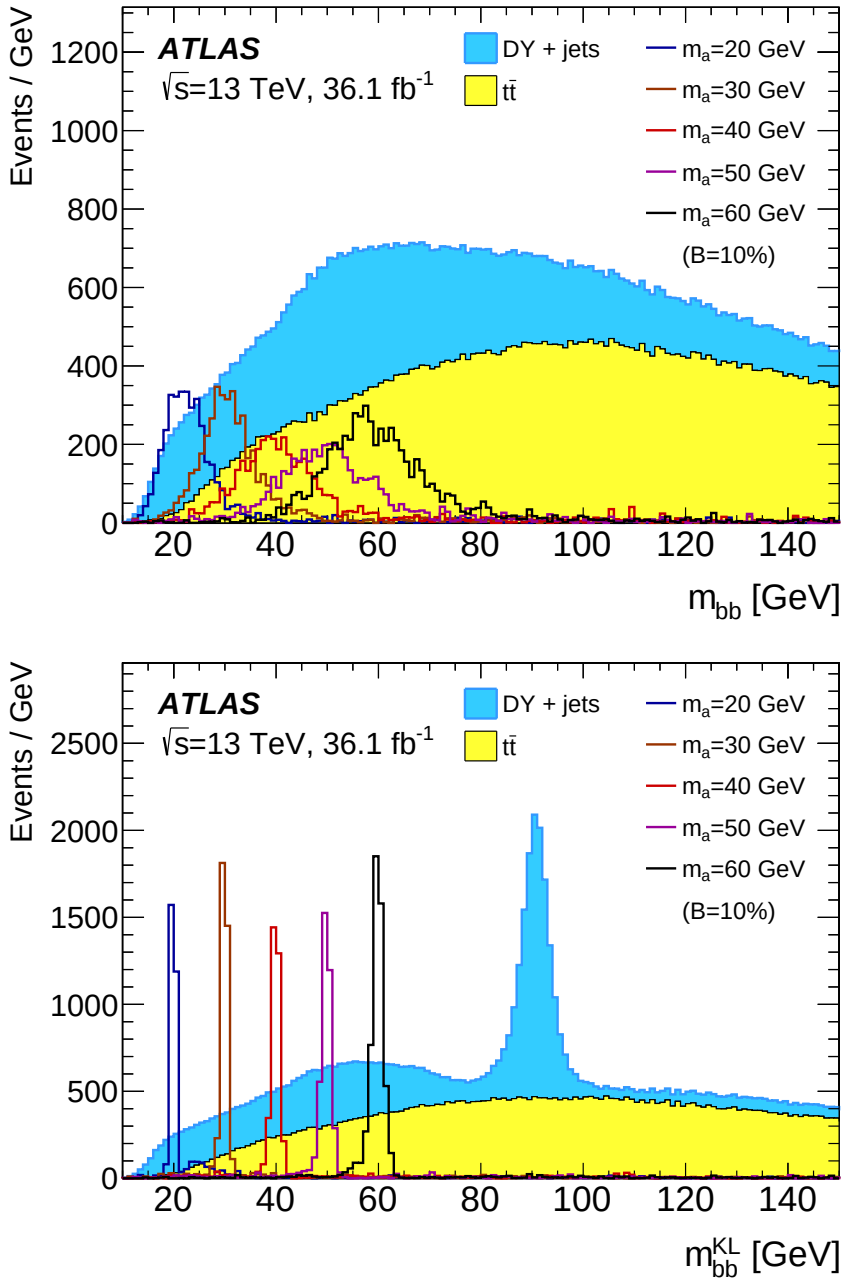

Figure 31: The mass distributions for the Drell–Yan, $t\bar{t}$, and $H \to aa \to bb\mu\mu$ signals for the di-bjet before (top) and after (bottom) the jets undergo the KLFit.

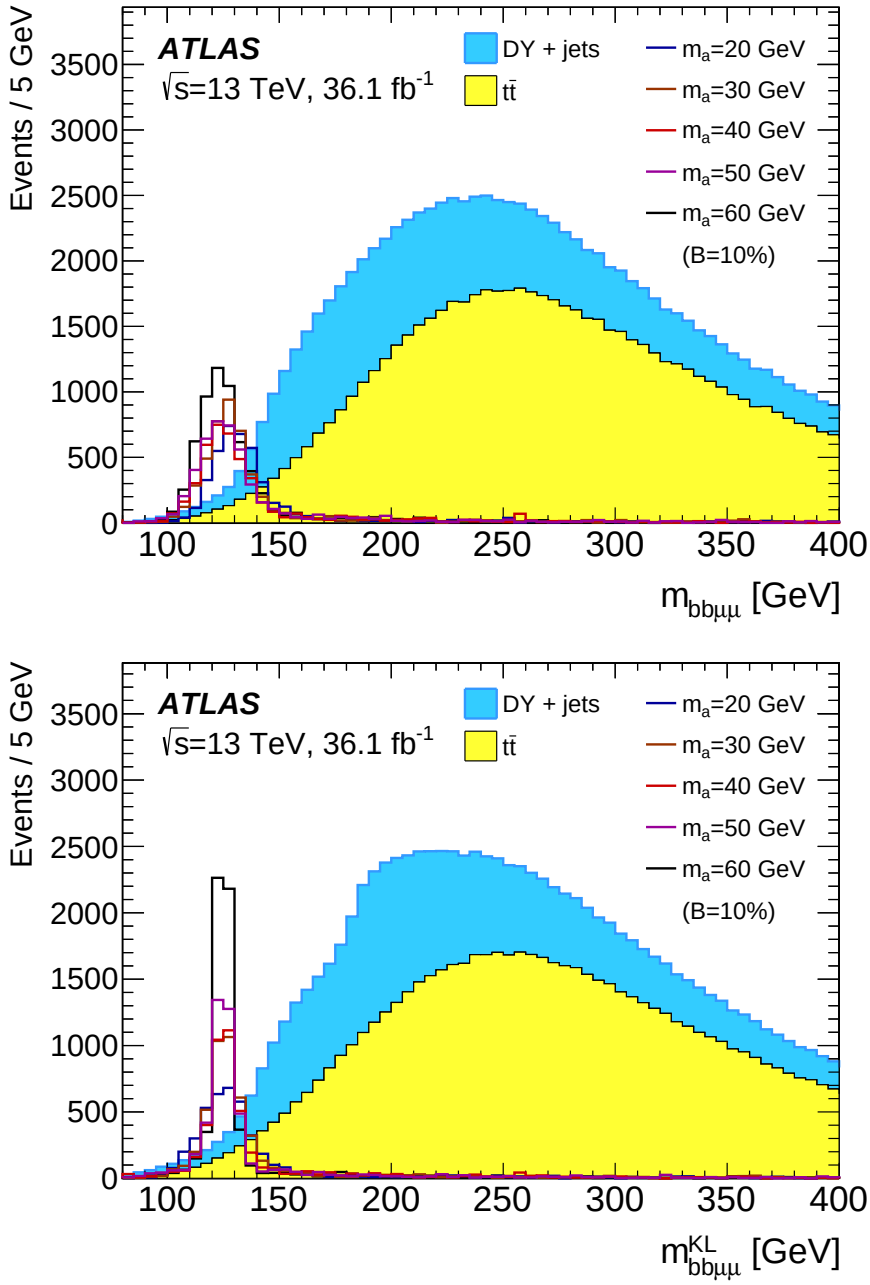

Figure 32: The $m_{bb\mu\mu}$ distributions for the Drell–Yan, $t\bar{t}$, and $H \to aa \to bb\mu\mu$ signals for the before (top) and after the KLFit (bottom).

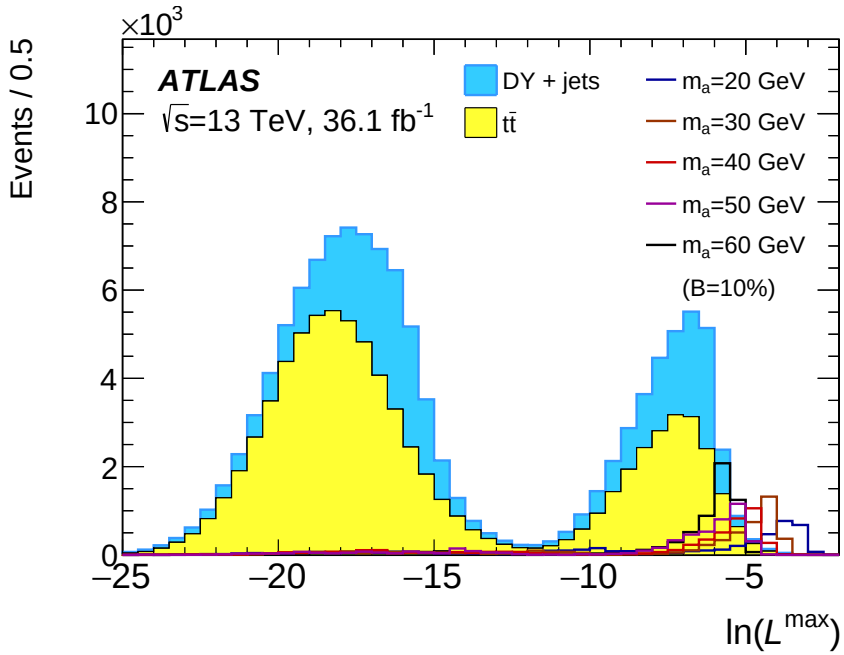

Figure 33: The KLFit maximum log-likelihood after the preselection and 2 b-tag requirement for the Drell–Yan, $t\bar{t}$, and signal processes.

likelihood profiles the level of equivalence between the two masses, a cut of output likelihood score can accomplish this selection. The Ln($L^{\max}$) distributions for the signal and major background processes is shown in Figure 33 after the preselection + 2 b-jet requirement. A maximal efficiency signal selection of Ln($L^{\max}$) > -8 is chosen as the event requirement.

After all of these selection cuts, a search for the a boson is performed by scanning the $m_{\mu\mu}$ distribution in multiple bins centered around a hypothesized m_a. To maximize the sensitivity of the analysis, a optimization is performed to determine the most sensitive $m_{\mu\mu}$ binning. This optimization is discussed in Section 5.6 as the optimization is done after performing the full background estimation. The optimized bin widths vary as a function of dimuon mass, with a width of 2 GeV chosen for $16 \leq m_{\mu\mu} \leq 40$ GeV, a width of 3 GeV for $40 < m_{\mu\mu} < 50$ GeV, and 4 GeV bin widths chosen for for $50 \leq m_{\mu\mu} < 64$ GeV.

The signal acceptance as a function of these cuts is given be in Table 16 for the $m_a = 30$ GeV signal sample. The other generated mass points explored in this analysis have similar acceptances and the analysis cutflows for these mass points are given in Appendix C.

Selection Criterion	MC Events	Event Yield	Eff./Step [%]	Tot. Eff. [%]
$N_\mu = 2$ & $N_{jets} > 0$	41670	1212.643	1.000	1.000
Pass Trigger	28367	809.654	0.681	0.681
$p_T^{\mu_1} > 27$, $p_T^{\mu_2} > 7$	27754	790.545	0.978	0.666
$16 < M_{\mu\mu} < 64$	27713	789.542	0.999	0.665
OS Muons	27625	786.969	0.997	0.663
$N_{b-jets} = 2$	2553	76.677	0.092	0.061
MET < 60	2478	74.286	0.971	0.059
$\mid M_{bb\mu\mu}^{KL} - M_h \mid \leq 15$	1979	57.260	0.799	0.047
Log(L) > -8	1830	52.480	0.925	0.044

Table 16: Analysis cutflow for $m_a = 30$ GeV, with event yields weighted to Br(H→aa)=100% and the BR(a→bb)BR(a→ $\mu\mu$) = 1.6×10^{-3} taken from the Type-III 2HDM+S model described in Ref. [33].

4.7 Systematic Uncertainties

4.7.1 Luminosity Uncertainties

The uncertainty on the integrated luminosity for the 36.1 fb^{-1} 2015+2016 datasets was determined to be 2.1%. This uncertainty is derived from the calibration of the luminosity scale using dedicated van der Meer scans, i.e. x-y beam-separation scans, performed in August 2015 and May 2016 in order to calibrate to the absolute luminosity. Additional checks on the luminosity measurement from the LUCID detector are done by comparing following a methodology Further details can be found in Ref. [130].

The luminosity uncertainty is applied to all backgrounds estimated using only simulation and is also applied to the generated $H \to aa$ signal samples.

4.7.2 Lepton Experimental Uncertainties

In order to fully describe the modeling of muons, the scale factors discussed in Section 4.5.2 are applied to the MC simulation samples in the nominal selections. As these scale factors rely on the composition of tag and probe events, each scale factor comes with a corresponding uncertainty. The final impact on the analysis is calculated by varying the the central value within its uncertainty (SF$\pm\sigma_{\text{SF}}$) and determining the impact on the final event yields. As lower momentum muons suffer larger fractional energy losses when traversing the calorimeter material, an additional

scale uncertainty is applied to the muon reconstruction efficiency scale factor for any muon with momentum in the range $7 < p_{\mathrm{T}} < 15$ GeV.

The energy and momentum scale of the muons used in this analysis are measured by comparing mass measurements of the Z-boson and J/ψ dimuon decays. The uncertainties on these energy scales are split between the MS and ID subdetectors, and are binned as a function of η to account for different material effects, magnetic field effects, and subdetector technologies. These momentum scale uncertainties are parametrized as

$$\frac{\sigma(p_{\mathrm{T}})}{p_{\mathrm{T}}} = r_0/p_{\mathrm{T}} + r_1 + r_2 \cdot p_{\mathrm{T}} \tag{87}$$

where the r_0 term accounts for the fluctuations in energy loss due to the calorimeter material, the r_1 term accounts for multiple scattering effects, magnetic field inhomogeneities affecting the track sagitta measurement and radial displacement of muon track hits, while the r_2 term describes the spatial resolution fluctuations and misalignment affects of the MS [75]. Each of these effects is accounted for by independently shifting all muon p_{T} values by $\pm 1\sigma$ and assessing the impact on the final yields.

Finally, as the muons are also treated as inputs for the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation, the momentum varied muons are re-run through this calculation. The correlations for the $E_{\mathrm{T}}^{\mathrm{miss}}$ and the muon momentum scale are then explicitly determined in this way.

Although the electrons are only used in the overlap removal and $E_{\mathrm{T}}^{\mathrm{miss}}$ calculations, the momentum scale and resolution uncertainties are determined for every electron. These uncertainties are calculated from measurements of the dielectron mass in $Z \to ee$ and $J/\psi \to ee$ decays, where the difference between the MC and data responses is used to derived both scale and resolution uncertainties simultaneously. The scaling factors required in order to correct the MC to data can then be extrapolated to determine the uncertainties on the energy scale [131]. These uncertainties, like the muon uncertainties, are propagated through the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation in order to derive the correct correlation terms.

4.7.3 Jet Experimental Uncertainties

The largest contributions to the systematic uncertainties affecting jet come from the understanding of the jet energy scale (JES) and resolution (JER) measurements. Both of these uncertainties are measured in the data (in-situ) by calculating the response between MC and data events which can depend on the p_{T}, η region of the detectors. In general, there are several categories for the jet energy scale systematic uncertainties:

In-situ Uncertainties In-situ techniques assess the ability of the MC simulation to accurately model the jet energy spectra seen in specific data regions. Two types of uncertainties can be assessed, one dealing with an absolute energy scale measurement and a separate one dealing with relative energy scale measurements.

The absolute JES measurement takes events with well known reconstructed objects, such as $Z \to \ell\ell$+jet or γ+jet events, and uses the one of two techniques to calibrate the jet: the direct balance (DB) or the missing-E_T projection fraction (MPF) technique. The DB measures the ratio of the recoiling jet p_T to the p_T of the well-measured system in order determine the jet energy scale. As multiple radiation can lead to significant changes in this balance, events are required not to have a subleading jet with large momentum. Complementary to this single jet approach, the MPF technique takes all jets in the event and attempts to calibrate the entire hadronic response by also accounting for the E_T^{miss} present in the event. As the leptons and photons are typically restricted to the central calorimeter region for precision concerns, the Z/γ+jet corrections are used in the region $20 < p_T < 950$ GeV and $|\eta| < 0.8$.

The relative JES uncertainty, also called the $|\eta|$-intercalibration uncertainty, arises from the measurement of the jet response in the $0.8 < |\eta| < 4.5$. Two methods are used to determine the jet energy scale in this region, the central jet reference method and the matrix method. In the central jet reference method, a forward jet is calibrated by comparing its p_T response to that of a central jet that is calibrated using the Z/γ+jet calibration. These events are required to have two jets and veto events with a third jet with substantial p_T. The complementary matrix method divides the entire forward region into several regions and calibrates jets by measuring the response in one forward jet bin to all other reference regions simultaneously.

For the relative JES uncertainty, the largest uncertainty comes from the differences in jet modeling seen between different MC generators, with the predictions from POWHEG interfaced with PYTHIA6 being compared to those from Sherpa 2.2.1 [127].

Finally, multi-jet balance method targets events with jets with $p_T > 300$ GeV and attempts to calibrate jet momenta up to 2 TeV. These events balance the high-p_T jets with the jets calibrated via the Z/γ+jet measurements. The largest uncertainties in this method come from the flavor composition uncertainties.

Pileup The pileup uncertainty coming from the modeling of pileup contributions to the jet energy. The pileup contributions to the jet are subtracted using an area-based p_T subtraction known as ρ-subtraction. This subtraction depends on the number of matched charged particles and the number of vertices reconstructed in the event in order to correctly model the contributions from pileup. As the forward

calorimeter does not have tracking information available, an additional MC based residual correction is used to correct for the pileup in this region.

Jet Flavor Composition The jet flavor composition uncertainties come from the difference in response seen in the MC modeling of the fragmentation for gluon-initiated jets and for quark-initiated jets. These uncertainties can be analysis dependent, as the exact composition of the quark jet to gluon jet can be determined via other analysis methods and this choice can affect the uncertainty derived on fragmentation. A further uncertainty can come from the difference in fragmentation between light jets and b-jets, as the heavy flavor jets tend to have higher constituent multiplicities. These are important as the global JES does not rely on flavor tagging due to analysis regions often having different flavor contributions.

Punch-through Punch-through uncertainties arise from high-p_{T} jets ($p_{\mathrm{T}} > 1$ TeV) which are energetic enough to pass through the hadronic calorimeter and deposit energy within the MS. Due to the low p_{T} scale involved in most objects in this analysis, this uncertainty source is negligible.

Non-closure MC non-closure uncertainties describe the response differences that come from the choice of calorimeter modeling in ATLAS, either using a full detector simulation or a look-up parametrization (ATLFAST-II [132]). As this analysis uses full simulation for all nominal signal and background samples, the MC non-closure is non-existent.

All of these uncertainty sources contribute several nuisances parameters, leading to around 80 nuisance parameters[9] needed to fully describe the JES uncertainty and the corresponding correlations. These nuisances parameters consist of 75 nuisance parameters from the in-situ Z/γ+jet and multi-jet JES measurements, 3 uncertainties from the $|\eta|$-intercalibration, 4 uncertainties from the pileup modeling of $N_{\mathrm{PV}}, \mu, \rho$ and the residual MC correction, 3 uncertainties from the JES response differences between b-jets, light jets, and gluon initiated jets, and finally the nuisance parameters from the specialized uncertainties dealing with detector simulation, punch-through, and a high-$p_{\mathrm{T}} > 2$ TeV extrapolation uncertainty due to the limitations of the calibration methods.

As many analyses can be insensitive to correlations between the various parameterizations, the number of nuisance parameters can be reduced by forming a covariance matrix if the nuisance parameters are based on the same kinematic variables. For example, the 67 absolute in-situ measurement nuisance parameters can be re-

[9]Nuisance parameters are used to characterize the systematic uncertainty, typically a variation, of a measured parameter. In the case of jets, the energy is the measured and the effects described above can cause the measurement to vary within a specific uncertainty range.

duced via diagonalization and removal of sub-dominant eigenvalues in a residual term. A side effect of this diagonalization, however, is that the new eigenvalues no longer have a physical meaning. In this analysis, the ATLAS strongly reduced set of 4 nuisance parameters is used, which describe the flavor composition and in-situ measurements in 3 combined nuisance parameters and 1 parameter describing the η-intercalibration term.

The JER is determined using the same function form as Equation. 70, exchanging E values for p_{T} in that equation. This resolution is measured using similar in-situ techniques to the JES, which can measure the width of the p_{T} balance. The noise term, N/p_{T}, however, is determined in zero-bias data using either a random cone or in pileup data using a ρ-subtraction technique [133]. In the random cone approach, a cone of R=0.4 is projected into a random η, ϕ area in the detector and all topo-clusters in this cone are summed to form a proxy for jets. Another cone is placed at new $\phi + \pi$ and a new random η and the difference in the random cone p_{T} is taken as the noise term. This noise term is then propagated to higher pileup by adding it to the combining the random cone noise with the ρ determined noise term at higher pileups.

The uncertainties for the JER can be described by the same uncertainty parameters used in JES parameterization, with the caveat that the modeling of the noise term induces new nuisance parameters. A similar eigenvalue decomposition is again used on these uncertainties and the final analysis uses a single nuisance parameter to describe the JER.

An additional uncertainty based on the data-MC scale factor for the JVT is also applied to all jets with $p_{\mathrm{T}} < 60$ and $|\eta| < 2.4$. This uncertainty accounts for the difference in hard-scatter efficiency for a fixed JVT cut observed between difference MC generators. The uncertainty, given as a function of p_{T} of the jets and is found to be between 1-2% for all bins [88].

4.7.4 Flavor Tagging Uncertainties

Efficiency scale factors are applied to simulated event samples to correct for the differences between data and simulation in the b-tagging efficiency and mis-tag rates for b, charm, and light-jets. These scale factors take the form

$$\kappa_b = \frac{\varepsilon_b^{\mathrm{data}}}{\varepsilon_b^{\mathrm{MC}}} \tag{88}$$

where the efficiencies in data and MC are measured for a given tagging efficiency point, i.e. 77% as in this analysis. These correction factors are derived using data

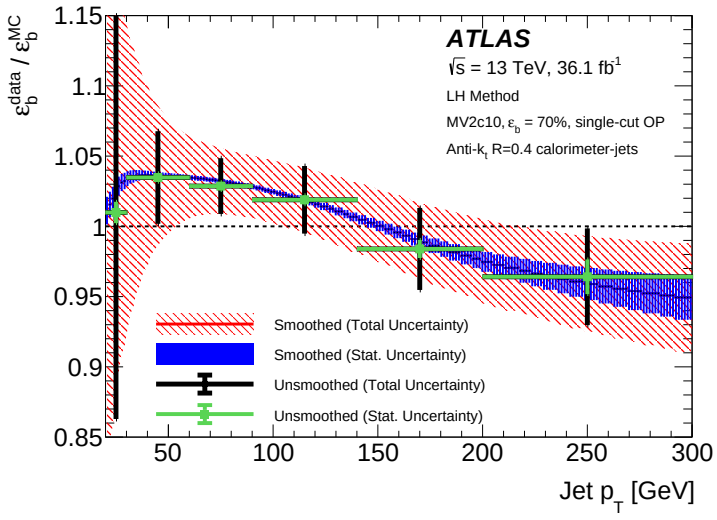

Figure 34: An illustration of the b-tagging simulation scale factors for the 70% WP of the MV2c10 b-tagging algorithm on R=0.4 calorimeter jets. The uncertainties are compared for a smoothed uncertainty and an un-smoothed uncertainty for both the statistical and total uncertainty values. Figure adapted from Ref. [91].

events selecting $t\bar{t}$ events with two leptons and are binned as a function of p_{T} and $\langle\mu\rangle$ in order to remove any kinematic selection or pileup biases. The uncertainties on these scale factors come from a variety of sources, including the choice of generator, the parton shower fragmentation model, the B-hadron lifetime modeling, the background cross-section uncertainties and other factor affecting the data purity of the sample, and the jet energy and resolutions in the data sample [91]. These uncertainties are again decomposed into several eigenvalues that provide the dominant uncertainties as a function of the jet flavor, with 26 distinct nuisance parameters.

An example correction factor for the 70% WP and its corresponding uncertainties are shown in Figure 34, where the scale factors are found to be close to unity and the large uncertainties in the low-p_{T} region are especially relevant for this analysis.

4.7.5 $E_{\mathrm{T}}^{\mathrm{miss}}$ Uncertainties

The systematic uncertainty on the $E_{\mathrm{T}}^{\mathrm{miss}}$ arises due to the propagation of the experimental systematics from the input objects, namely the jets, muons, and electrons. This variance determines the overall scale uncertainty of the $E_{\mathrm{T}}^{\mathrm{miss}}$.

In addition to this source of uncertainty, the soft-term uncertainty from the underlying tracks is determined using Z$\rightarrow \mu\mu$ events with no additional jets. In this

event, the leptons are not expected to recoil and the total sum of the $E_{\mathrm{T}}^{\mathrm{miss}}$ hard and soft terms should equal to 0 within experimental resolution. In order to determine the systematic, the $p_{\mathrm{T}}^{\mathrm{miss}}$ soft term is projected onto the $p_{\mathrm{T}}^{\mathrm{miss}}$ hard term and the variance (RMS) in the longitudinal and transverse projections of this quantity are compared between data and several MC simulation samples to extract the systematic uncertainties. These uncertainties account for two additional nuisance parameters, referred to as the parallel (MET_para) and perpendicular (MET_perp) uncertainties [78].

4.7.6 Theoretical Uncertainties

Given the various techniques used to calculate cross-sections and MC simulation procedures used to model the interaction of the protons in LHC collisions, theoretical uncertainties on the signal and background modeling are assessed in the analysis. Three types of modeling uncertainties in the MC generators are considered, relating to the choice of hard-scatter matrix element calculator and parton shower program, the choice of PDF set, and the fundamental choice of renormalization and factorization scales for the perturbative calculations of the hard scatter. These uncertainties are discussed in general terms below, while specific uncertainties can be found in Section 4.8 for the background theory uncertainties and Section 4.7.7 for the signal theory uncertainties.

Hard Scatter Generator/Parton Shower Choice MC simulation programs often implement physics processes at different orders and choose different models for the hadronization and radiation effects discussed in Section 1.4. These choices are considered to be theoretical systematic uncertainties, as the choice of MC generator or PS program can lead to significant differences in the yields and kinematic distribution shapes when modeling the LHC data. To assess this modeling uncertainty, a MC sample is produced with a different generator or parton shower configuration and compared to the nominal MC sample.

PDF Uncertainties An uncertainty on the choice of PDF set can be determined in several ways, which include comparing the results obtain from two different PDF sets and taking the difference as an uncertainty or by varying the results by the fit uncertainties within a single PDF set. In this analysis, the PDF uncertainties are determined via PDF set comparisons using the PDF4LHC recommendations [134], which recommend the comparison of a nominal PDF set to the PDF4LHC combined results which incorporate the CT14, MMHT14, and NNPDF3.0 PDF sets.

In practice, these comparisons are performed by reweighting events based on shifts in the parton kinematics for the generated events by comparing using event

weights of the form

$$w = \frac{\mathrm{PDF}_{\mathrm{alt.}}(x_1, f_1, Q) \times \mathrm{PDF}_{\mathrm{alt.}}(x_2, f_2, Q)}{\mathrm{PDF}_{\mathrm{nom.}}(x_1, f_1, Q) \times \mathrm{PDF}_{\mathrm{nom.}}(x_2, f_2, Q)} \tag{89}$$

where $\mathrm{PDF}_{\mathrm{alt.}}$ refers to an alternative PDF sample and $\mathrm{PDF}_{\mathrm{nom.}}$ refers to the PDF used during the initial event generation. Additionally, the PDF sets used for comparison must be the same order in α_s, as comparing a LO to an NLO PDF prediction can inflate the uncertainties due to inconsistencies. The background and signal samples with non-negligible PDF uncertainties include the $H \to aa$ signal samples and the $t\bar{t}$ samples.

Scale Variations After the renormalization of matrix elements calculations in perturbative calculations, the coupling obtains a scale dependence that depends on the renormalization group equations. This scale typically refers to the momentum scale where a physics process is measured and is denoted as μ_R. The factorization scale, denoted as μ_F provides a corollary evolution scale as it defines the point at which the QCD behavior changes from the perturbative (UV) scale to the confined/non-perturbative (IR) scale. The factorization scale is necessary in order to calculate the perturbative regions of QCD, otherwise the hard scatter calculation would not be possible due to large divergences. It is typically chosen to be on the order of the hard-process in an event, being given by $\mu_F = m_H$ for Higgs events of $\mu_F \sim m_t$ for events involving the top quark. In order to avoid large logarithms in fixed order calculations from MC generators, the two scales are typically set to be equivalent. The uncertainty on the choice of scales is estimated by varying μ_R and μ_F by a factors of 1/2 or 2 and taking the largest difference in the MC predictions to be the uncertainty. An uncertainty on the scale choice is determined and applied for the signal and $t\bar{t}$ background.

As an additional complication in the choice of generator and parton showering choices comes from the amount of ISR or FSR generated. As ISR/FSR can affect the recoil of a system, the amount of radiation can affect observable quantities such as the p_{T} of the Higgs or top-quark systems. To account for these effects, an uncertainty can be derived by altering the MC simulation parameters for ISR and FSR using different MC generator tunes. These tunes can change the p_{T} cutoff for radiation, the factorization and renormalization scales, and also can affect the value of the QCD coupling α_s used to model the shower evolution.

4.7.7 Signal Uncertainties

The $H \to aa \to bb\mu\mu$ signal is subject to all experimental uncertainties detailed in the previous sections. Additional theoretical modeling uncertainties are derived on

the signal MC samples as described below.

For the $H \to aa$ signal samples, the PDF uncertainty is determined by computing the the weights in Equation 89 using LHAPDF6 [135] for several different PDF sets. The nominal CTEQ6 PDF set is compared with the 101 eigenvector variations of the NNPDF3.0 PDF and the PDF4LHC15_nlo_mc PDF sets. The envelope of these variations with respect to the nominal configuration is taken as the final uncertainty.

For the factorization μ_F and renormalization μ_R scale uncertainties, the 3 mass points of the $H \to aa \to bb\mu\mu$ signal samples are generated with the following pairs of $\{\mu_F,\mu_R\}$: {0.5,0.5},{0.5,1},{1,0.5},{1,1},{1,2},{2,1},{2,2}. The deviations of each pair from the nominal prediction ({1,1}) are computed and the largest deviation found is taken as the systematic uncertainty. Additional off-diagonal elements, {0.5,2} and {2,0.5}, were generated, but not included in the systematic determination due to complications with large logarithms inherent in these matrix element calculations.

For the parton shower uncertainties, the PYTHIA internal parameters for the ISR and FSR are used to determine the influence these variations have on the invariant mass distributions of each signal mass point. The ISR and FSR parameters are varied up and down by 1σ in order to increase and decrease the amount of radiation arising from the Higgs decay products.

The effect of these theoretical modeling uncertainties are given in Tables 17-19 for the $m_a = 20, 40, 60$ GeV signal mass points. The largest uncertainty comes from the scale variations, which can change the final signal yields by up to 6%.

Uncertainties in the calculation of the gluon-fusion and vector-boson fusion Higgs production cross-section are assigned following the recommendations of the LHC Higgs Cross Section working group given in Ref. [136]. These uncertainties amount to 3.9% for the gluon-fusion production and 2.1% for the VBF production process.

Given the analysis selections, a small signal contribution comes from the associated production of the Higgs boson with a vector boson (VH). As this contribution is not explicitly modeled using MC simulation samples, it is accounted for by increasing the cross section weights of the gluon fusion and VBF processes by the estimated 3% VH contribution. As the VH production mode kinematics differ from those of the gluon-fusion and VBF processes and the additional object vetos can potentially change as a function of the experimental uncertainties, a conservative 100% uncertainty is applied to this weighting procedure.

As a final systematic uncertainty and scaling factor, the signal samples are compared to MC samples of Higgs boson the p_{T} and rapidity spectra generated using the NNLOPS approach [30]. The $H \to aa$ signal samples are found to accurately model the Higgs rapidity distributions when compared to the NNLOPS sample, but the

Theoretical Variation	Effect on Signal Yields (%)
μ_F: 0.5, μ_R: 0.5	5.0 $\pm$ 0.25
μ_F: 0.5, μ_R: 1	4.9 $\pm$ 0.24
μ_F: 0.5, μ_R: 2	1.0 $\pm$ 0.05
μ_F: 1, μ_R: 0.5	0.63 $\pm$ 0.03
μ_F: 1, μ_R: 2	5.5 $\pm$ 0.27
μ_F: 2, μ_R: 0.5	2.0 $\pm$ 0.10
μ_F: 2, μ_R: 1	0.13 $\pm$ 0.01
μ_F: 2, μ_R: 2	2.4 $\pm$ 0.12
ISR Variation Up	4.1 $\pm$ 0.21
ISR Variation Down	0.88 $\pm$ 0.04
Maximum PDF	0.30

Table 17: The breakdown of theoretical systematics for the gluon-fusion Higgs signal sample, with $m_a = 20$ GeV. The largest uncertainty is found to come from the factorization and renormalization scale variations. The uncertainty on the yields from the PDF variations are negligible and therefore omitted.

Theoretical Variation	Effect on Signal Yields (%)
μ_F: 0.5, μ_R: 0.5	2.5 $\pm$ 0.10
μ_F: 0.5, μ_R: 1	5.5 $\pm$ 0.23
μ_F: 0.5, μ_R: 2	3.1 $\pm$ 0.13
μ_F: 1, μ_R: 0.5	0.34 $\pm$ 0.01
μ_F: 1, μ_R: 2	5.6 $\pm$ 0.23
μ_F: 2, μ_R: 0.5	5.4 $\pm$ 0.23
μ_F: 2, μ_R: 1	2.6 $\pm$ 0.11
μ_F: 2, μ_R: 2	4.5 $\pm$ 0.20
ISR Variation Up	1.5 $\pm$ 0.06
ISR Variation Down	0.25 $\pm$ 0.01
Maximum PDF	0.35

Table 18: The breakdown of theoretical systematics for the gluon-fusion Higgs signal sample, with $m_a = 40$ GeV. The largest uncertainty is found to come from the factorization and renormalization scale variations. The uncertainty on the yields from the PDF variations are negligible and therefore omitted.

Theoretical Variation	Effect on Signal Yields (%)
μ_F: 0.5, μ_R: 0.5	0.52 ± 0.02
μ_F: 0.5, μ_R: 1	1.1 ± 0.03
μ_F: 0.5, μ_R: 2	1.1 ± 0.03
μ_F: 1, μ_R: 0.5	0.39 ± 0.01
μ_F: 1, μ_R: 2	3.2 ± 0.10
μ_F: 2, μ_R: 0.5	0.65 ± 0.02
μ_F: 2, μ_R: 1	2.1 ± 0.06
μ_F: 2, μ_R: 2	2.6 ± 0.08
ISR Variation Up	1.7 ± 0.05
ISR Variation Down	1.6 ± 0.05
Maximum PDF	0.11

Table 19: The breakdown of theoretical systematics for the gluon-fusion Higgs signal sample, with $m_a = 60$ GeV. The largest uncertainty is found to come from the factorization and renormalization scale variations. The uncertainty on the yields from the PDF variations are negligible and therefore omitted.

Higgs p_{T} is found to be harder than the distribution from the higher order generator. A comparison of the nominal and NNLOPS p_{T} and rapidity distributions is shown in Figure 35.

In order to correct the MC simulation samples for this effect, a reweighting is derived as a function of the Higgs p_{T} by fitting the ratio of the generator p_{T} distributions using continuous functional fit. The ratio and fit are shown in Figure 36. This function is used to apply event weights to the final signal selection and a 2% difference in the event yield after all signal selections is observed between the weighted and unweighted signal samples. This difference is taken as a systematic uncertainty on the signal yields in the final fit.

4.8 Background Estimation

The following sections describe how each background yield in the $H \to bb\mu\mu$ analysis is determined. The DY + jets background estimate is determined using a data-driven template method described in Section 4.8.1. The $t\bar{t}$ background is estimated using a hybrid approach, where the shapes of the kinematic distribution are taken from MC predictions and the overall yield is normalized to the data yield in a top-enriched data control region. This control region and method is further described in Section 4.8.2. All other backgrounds, including the W+jet, diboson, single top, and

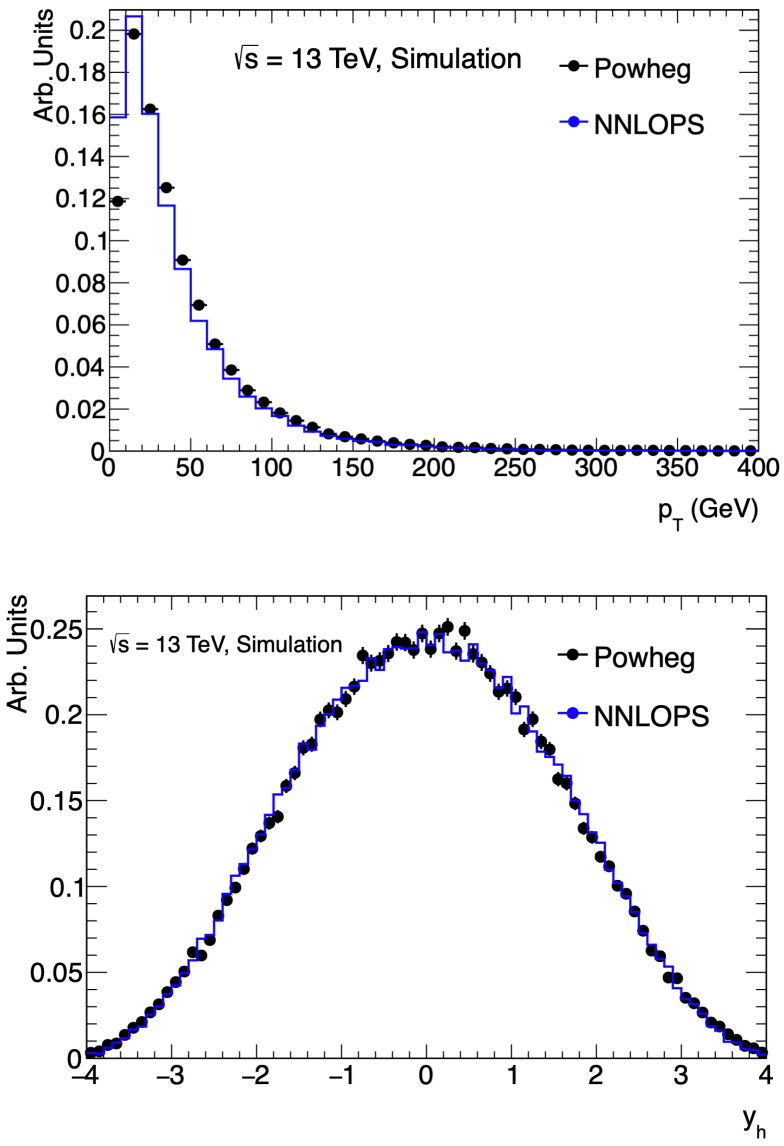

Figure 35: Comparison of Higgs p_{T} (top) and rapidity (bottom) spectra as modeled by the NLO Powheg and NNLOPS Powheg generators. The rapidity distributions are consistent between the generators, but the NLO Powheg prediction indicates a harder p_{T} spectrum compared to the higher order generator.

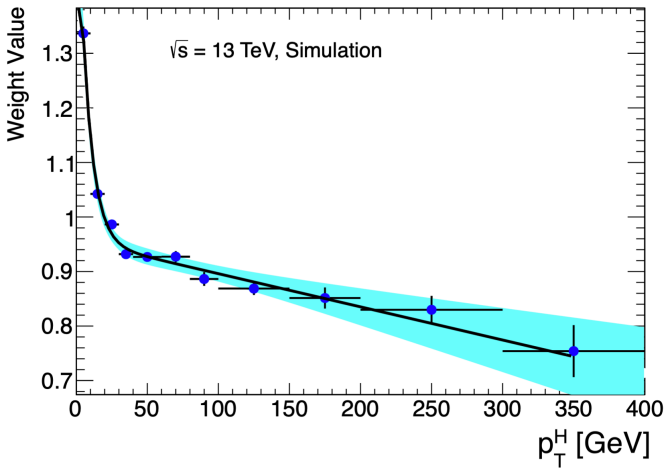

Figure 36: Distribution of the ratio between the p_{T} spectra generated by NLO and NNLOPS Powheg event generators. The fit of this ratio, which is used to derive an event reweighting on the signal samples, and the 95% confidence interval are also shown.

$t\bar{t}$ + X backgrounds are determined solely from MC predictions. These backgrounds are labeled as the "minor backgrounds" in this analysis and the estimation of these backgrounds is discussed in Section 4.8.3.

The control region definitions for the DY + jets and $t\bar{t}$ processes are shown in Figure 37, along with the signal and template region definitions.

4.8.1 Drell–Yan + Jets Background Estimation

The Drell–Yan + jets background is modeled using a data-driven template method, rather than a MC simulation estimate, due to a statistically limited MC sample. To illustrate this statistical issue, events from the Drell–Yan and the $t\bar{t}$ MC samples are shown after the signal region selections in Figure 38. In this region, the MC statistical error on the DY background is found to be on the order of 100% for certain $m_{\mu\mu}$ bins. This modeling issue is further exacerbated by large negative event weights in the generated sample.

In order to model the Drell-Yan background with better precision and circumvent these issues, a data-driven template approach is used. In this approach, templates are defined by applying all kinematic selections corresponding to a given control, validation, and signal region to a data sample, with the exception that the b-tagged

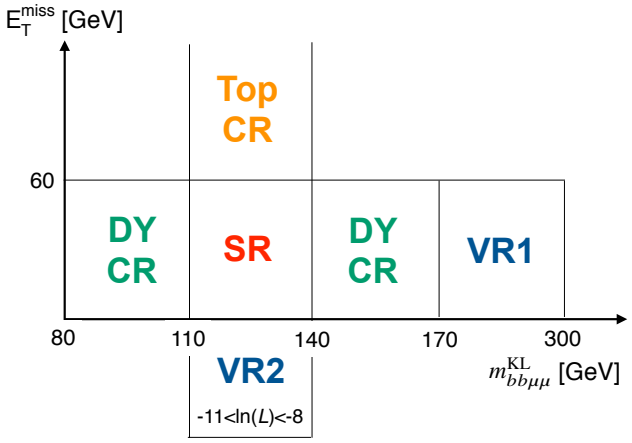

Figure 37: A pictorial representation of the control and validation regions used in the H→bb$\mu\mu$ analysis. The Drell-Yan (DY) and $t\bar{t}$ (Top) control regions (CR) are chosen to be close to the signal region (SR) in order to reduce kinematic biases. The validation regions (VR) are independent ranges used to validate closure of the background estimation methods performed in the analysis.

jet multiplicity is required to be exactly zero. This selection, which denotes a data template region (TR), is enriched in Drell–Yan+jet events due to the rejection of the $t\bar{t}$ background because of the inverted b-tagging requirement.

The shape of the Drell–Yan background distribution is then extracted by subtracting off the contributions from the $t\bar{t}$ and minor backgrounds in the TR using the MC simulated samples; the shape from the TR is then applied to the corresponding control region, validation region, or signal region with the correct $n_{\mathrm{b-jet}}$. The final normalizations for the Drell–Yan background are then determined in a series of fits to the corresponding data yield in the control and signal regions (Section 5.4).

An illustration of the comparison between the template region and the nominal b-tag region for the loose analysis selections with the additional $E_{\mathrm{T}}^{\mathrm{miss}}$ cut are shown in Figure 39. The data derived templates are compared to the template shapes found in the MC simulation and the trends are found to be similar for the same selections.

In the data templates, however, the b-jet spectrum is found to have a more boosted p_{T} spectrum compared to that coming from light jets. The spectra for the leading and subleading jets between the 0-tag template and the 2 b-tag loose selection region are illustrated in Figures 40-42, where the yields are normalized by subtracting off the contribution from the MC predicted yields of the $t\bar{t}$ background and the minor backgrounds.

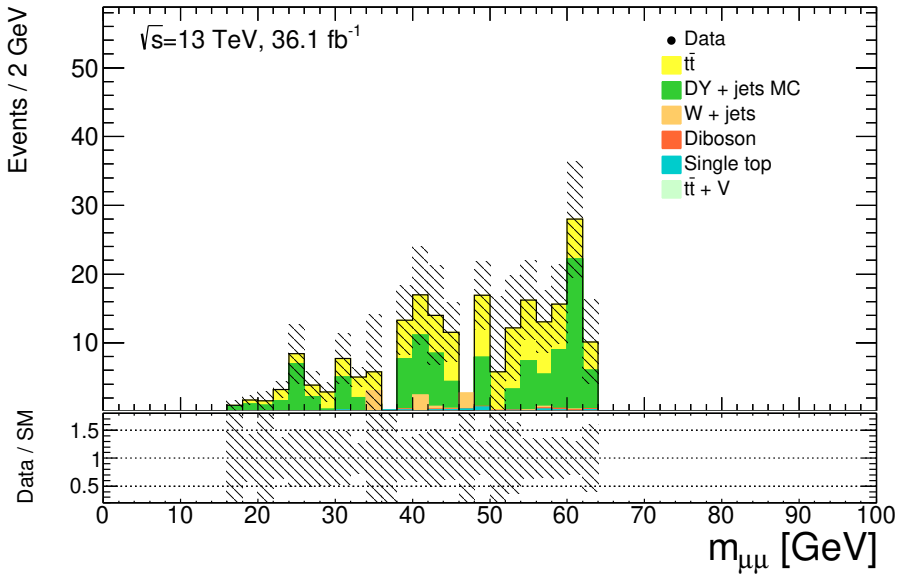

Figure 38: The dimuon mass distribution after signal region selections have been applied to the Drell-Yan + jets and $t\bar{t}$ Monte Carlo simulation sample. The statistical modeling of the Drell-Yan background is found to be quite poor, as can be seen in the large yield fluctuations and large uncertainties. Uncertainties shown here reflect only the statistical uncertainty on the MC simulation samples.

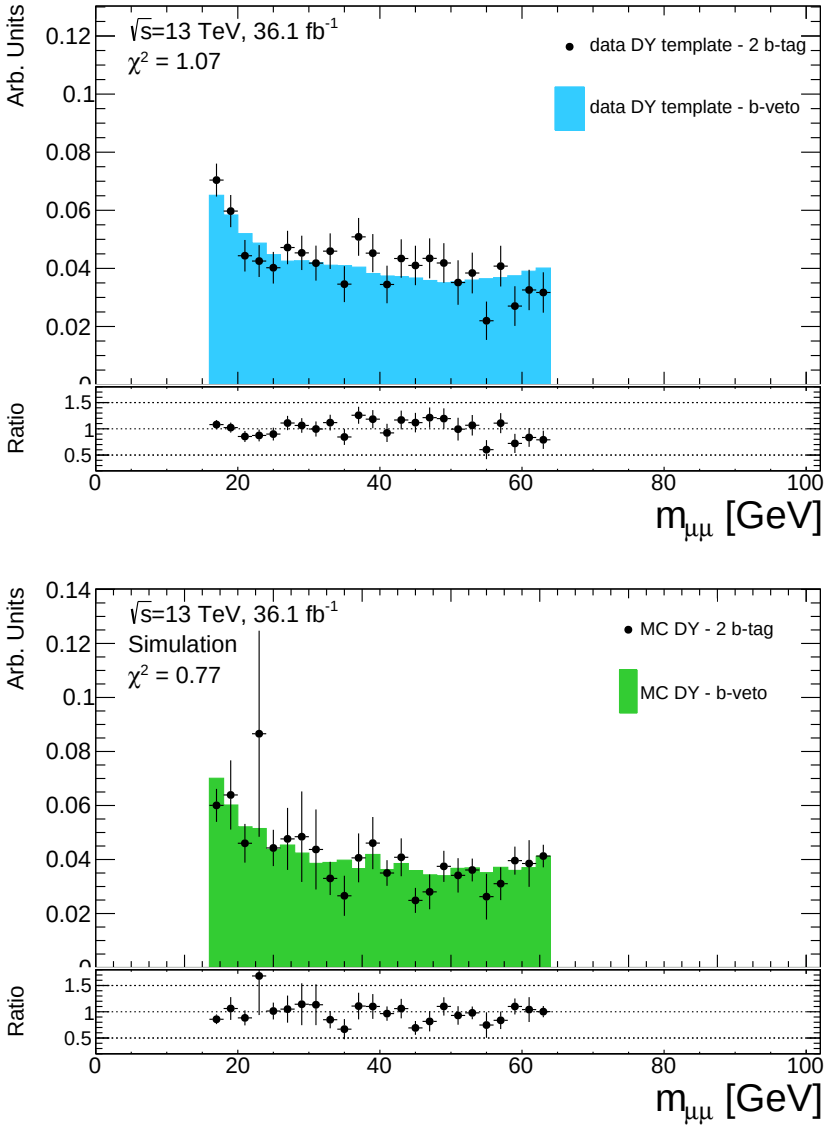

Figure 39: The template region shapes in the $m_{\mu\mu}$ distribution, derived from the data templates (top) and from MC simulation (bottom). The two distributions show the same trend and the 0-tag template models the 2-tag nominal region remarkably well. The output of a χ^2 fit comparing the 0-tag and 2-tag shapes is given in the upper left of each plot.

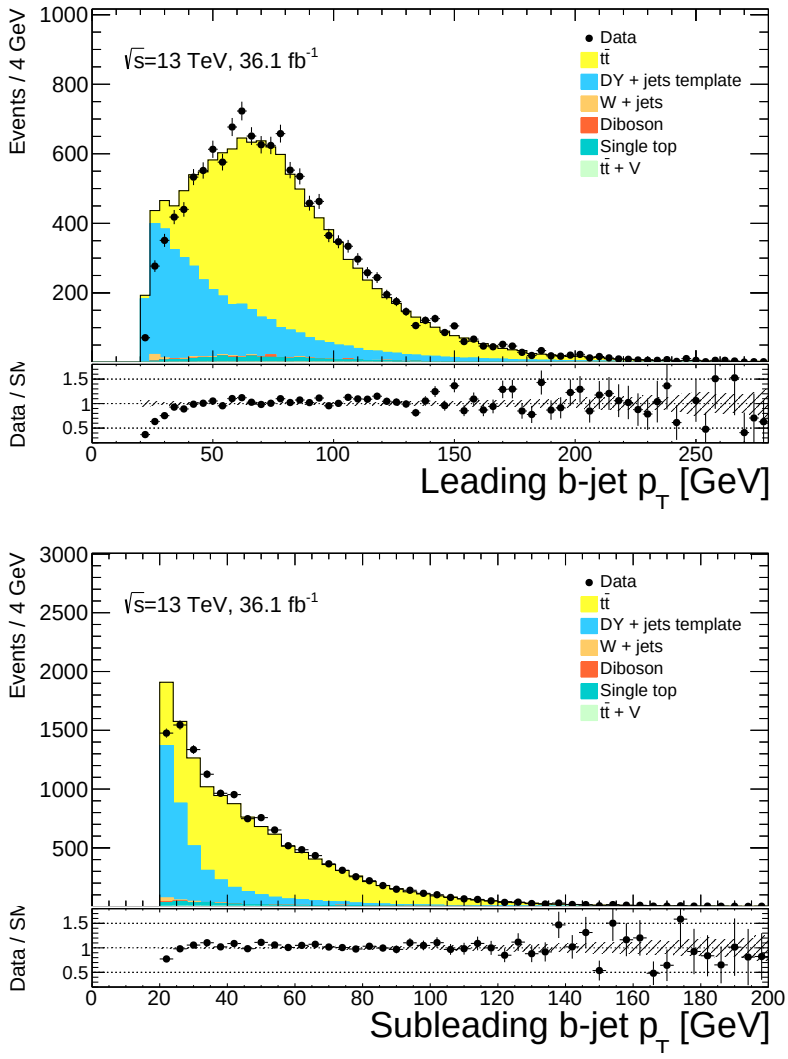

Figure 40: Jet p_{T} variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) before the jet reweighting procedure is applied. The p_{T} distribution of the leading jet illustrates the overshoot of the template predictions at low p_{T}.

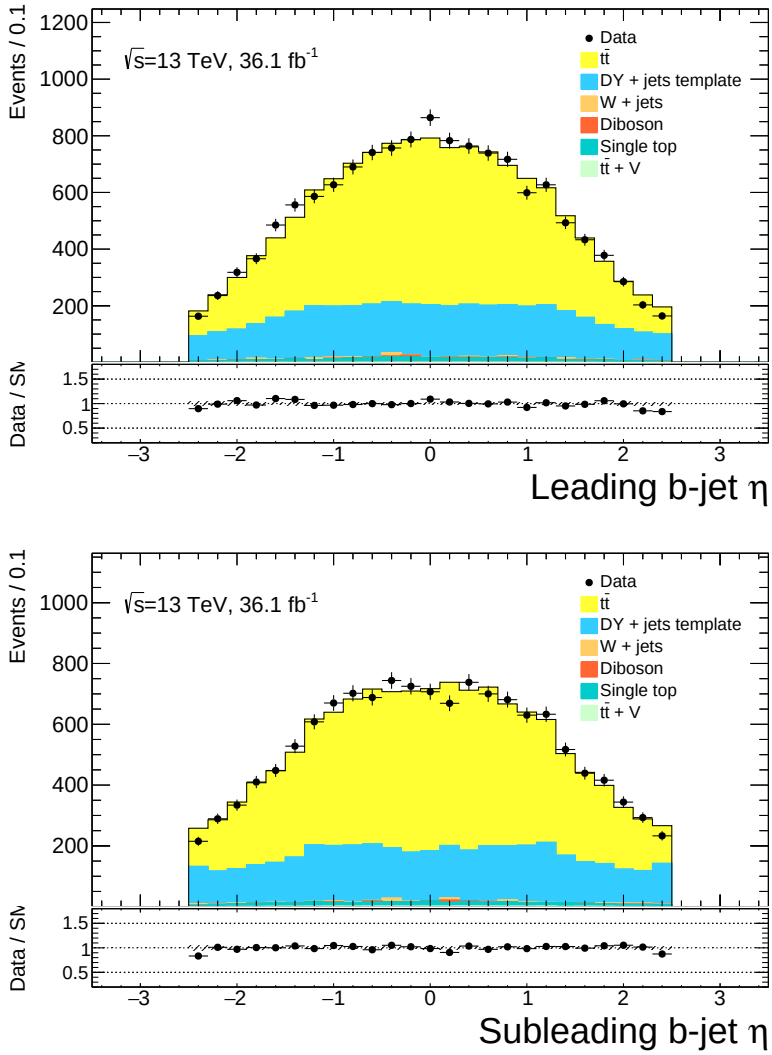

Figure 41: Jet η variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) before the jet reweighting procedure is applied. The agreement in the η distributions illustrates that the template prediction issues are localized only to low p_{T}.

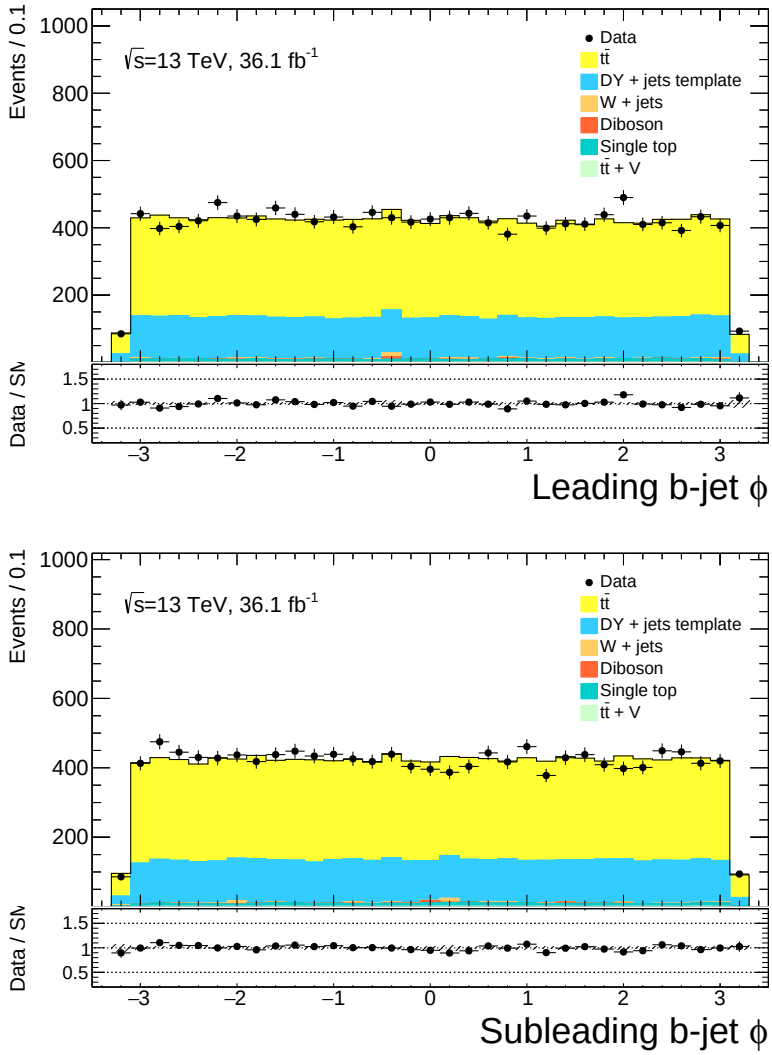

Figure 42: Jet ϕ variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) before the jet reweighting procedure is applied. The agreement in the ϕ distributions illustrates that the template prediction issues are localized only to low p_{T}.

In order to correct for the slope at low-p_{T}, an event weight based on the ratios between the template and the 2 b-tag shape is derived by fitting the ratio between the leading jet p_{T} in the 0-tag and 2-tag region. The jet spectra after this reweighting are shown in Figures 43-45 and the improvement this reweighting has on the modeling of the m_{bb}^{KL}, $m_{bb\mu\mu}^{\mathrm{KL}}$, and $\mathrm{Ln}(L^{\mathrm{max}})$ variables is illustrated in Figures 46-48.

With this data-driven approach, several new uncertainties are assessed on the Drell–Yan template method. These uncertainties include the effect of flavor composition on the data template prediction, the effect of background subtraction on the template shape, and the effect of data statistics in each $m_{\mu\mu}$ bin. The uncertainty related to the flavor composition is assessed by requiring the selection of exactly one b-jet in the template region definition rather than exactly 0 b-jets. This increase in the b-jet fraction of the template which brings the template kinematically closer to the signal region than the nominal 0 b-jets template. Given the jet reweighting that occurs, this flavor uncertainty is expect to impact the shape of the template region rather than the final normalization. The comparison of the shapes in the $m_{\mu\mu}$ distribution between the 1 b-tag and 0 b-tag templates for the preselection plus $E_{\mathrm{T}}^{\mathrm{miss}}$ selection and the inverted SR selection is shown in Figure 49. The average per-bin difference between the two templates in the shape of the $m_{\mu\mu}$ distribution is taken as a conservative estimate of the uncertainty; this uncertainty is found to be $\sim 14\%$ in the signal region before the final analysis fits.

The MC subtraction uncertainty on the template is estimated by comparing the nominal 0 b-tag template with all of the non-DY backgrounds subtracted from the template yield in the data to the case where no MC yields have been subtracted from the nominal template data. This conservative systematic affects the DY yields in the signal region by up to 4%.

The flavor composition uncertainty also accounts for the jet p_{T} reweighting, so an additional uncertainty based on the modeling of the fit from the not assessed.

4.8.2 $t\bar{t}$ Background Estimation

The background contribution arising from the dileptonic $t\bar{t}$ process is partially estimated using a statistically precise MC sample; in the estimation methods, the kinematic shapes are taken from MC simulation samples and the overall normalization is determined in a dedicated control region. The $t\bar{t}$ control region (TCR) is constructed by requiring all of the final analysis selections, except that the requirement on the transverse missing energy is inverted, $E_{\mathrm{T}}^{\mathrm{miss}} > 60$ GeV. The shapes of the kinematic distributions, in particular $m_{\mu\mu}$, are found to match the shapes of the data in the TCR. The MC yields in this region, however, are smaller than the data yields by

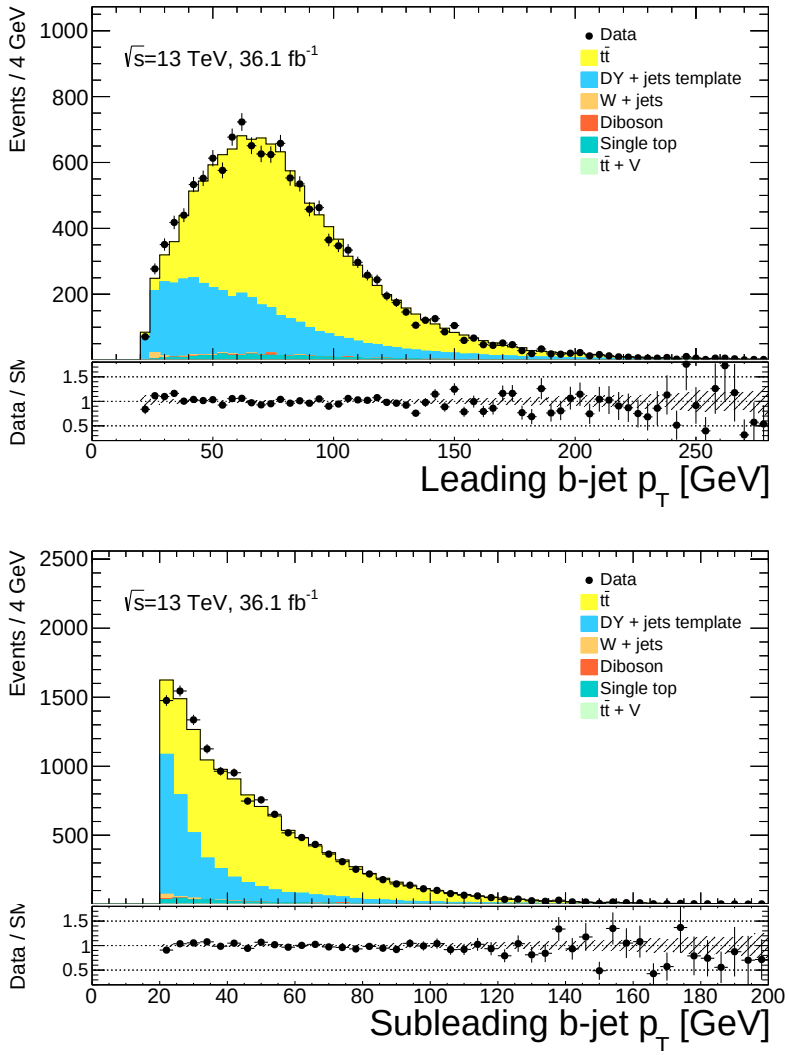

Figure 43: Jet p_{T} variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) after the jet reweighting procedure is applied. The overshoot in the p_{T} distribution of the leading jet is corrected and the agreement between the data and subleading jet p_{T} from the reweighted template also improves.

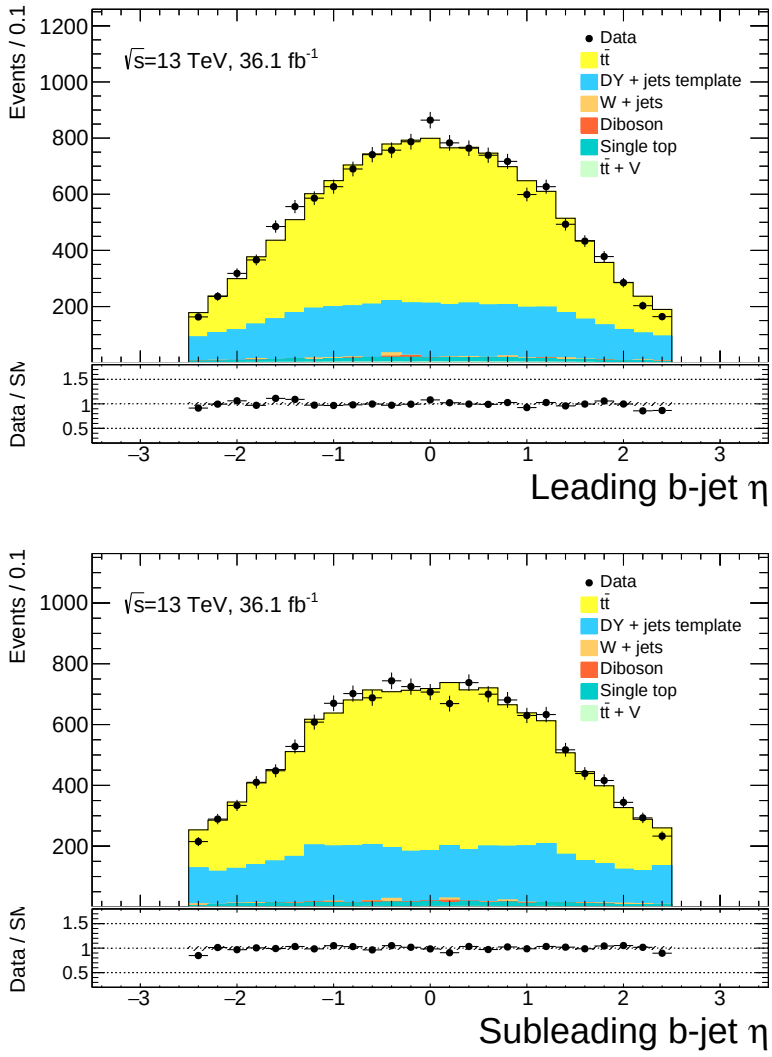

Figure 44: Jet η variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) after the jet reweighting procedure is applied.

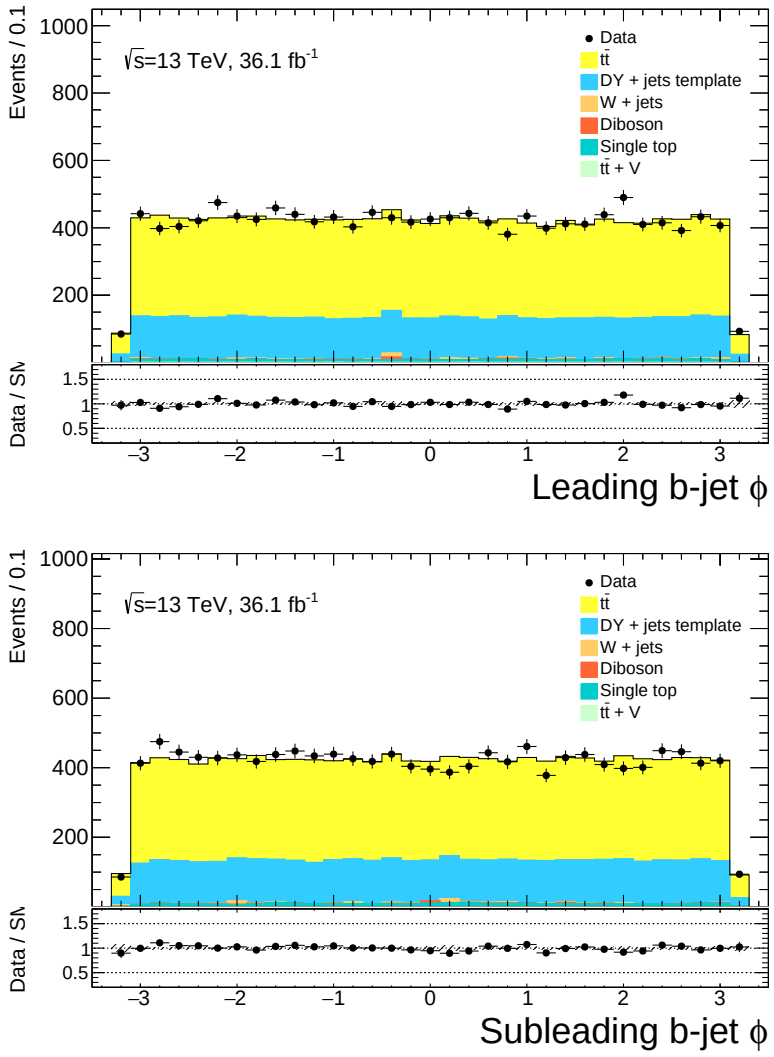

Figure 45: Jet ϕ variables modeled in a data region defined by the preselection and 2 b-jet requirements for the leading jet (top) and subleading jet (bottom) after the jet reweighting procedure is applied.

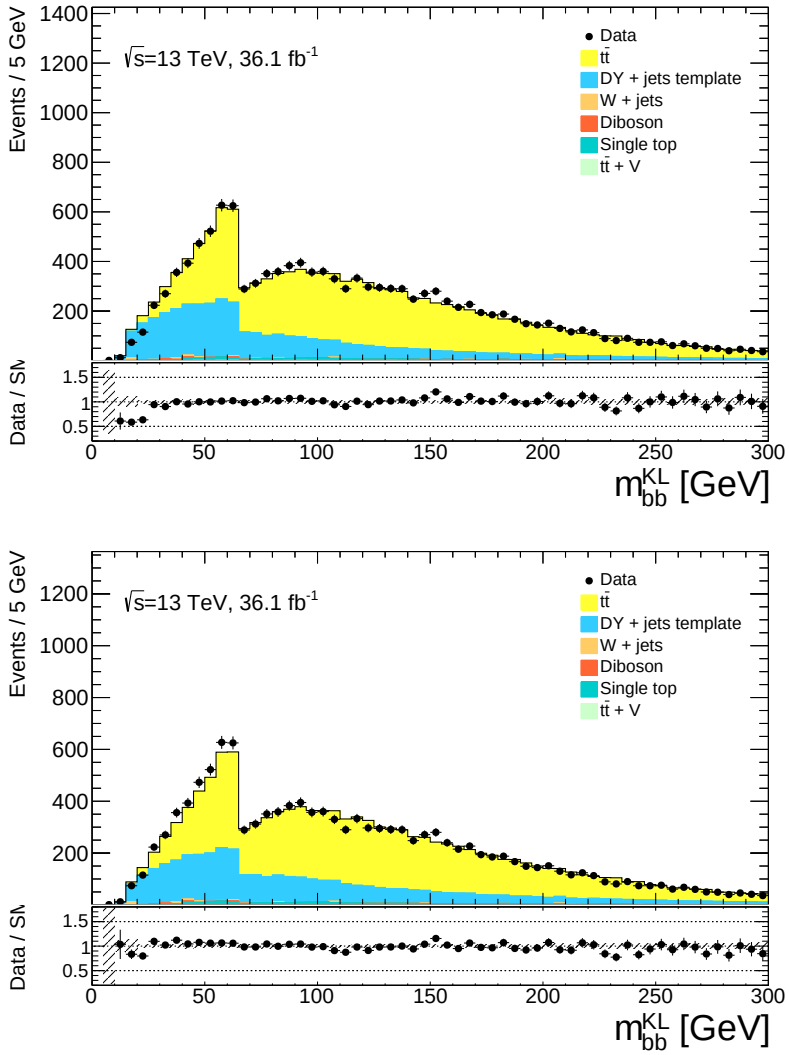

Figure 46: Effect of jet reweighting on the Drell–Yan template prediction for the di-bjet mass after the KLFit, m_{bb}^{KL}, with the distributions before (top) and after (bottom) jet reweighting is applied.

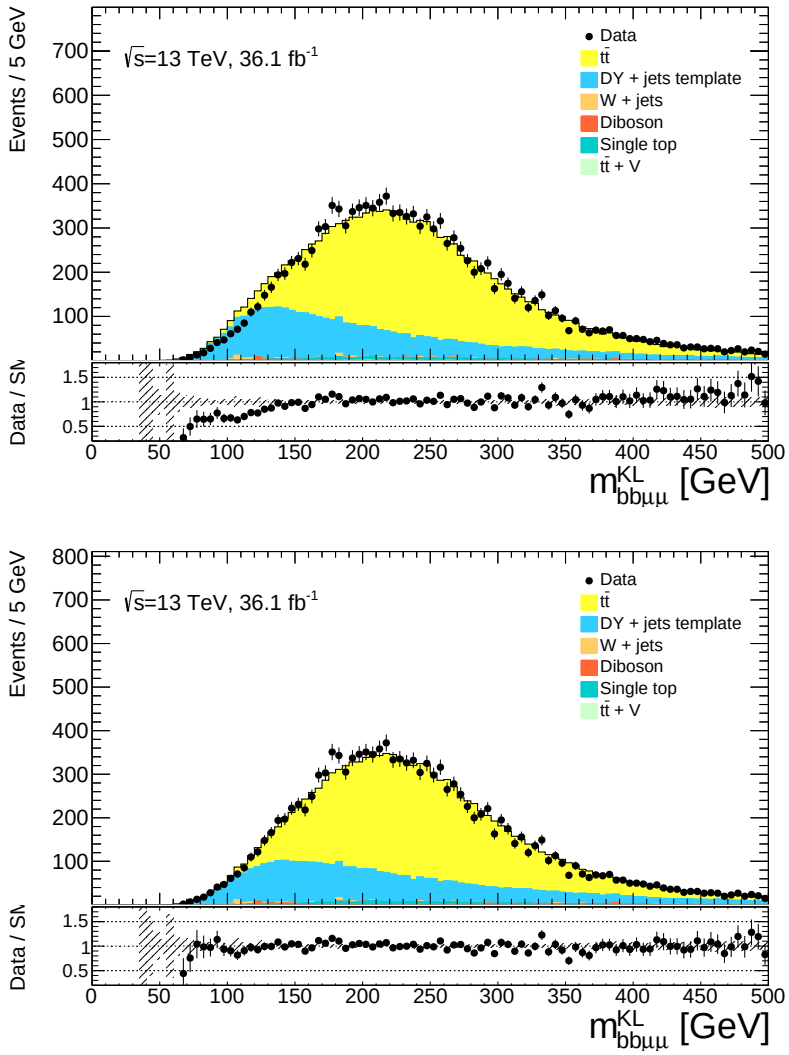

Figure 47: Effect of jet reweighting on the Drell–Yan Template Prediction for the $m^{\mathrm{KL}}_{bb\mu\mu}$ variable, with the distributions before (top) and after (bottom) jet reweighting is applied.

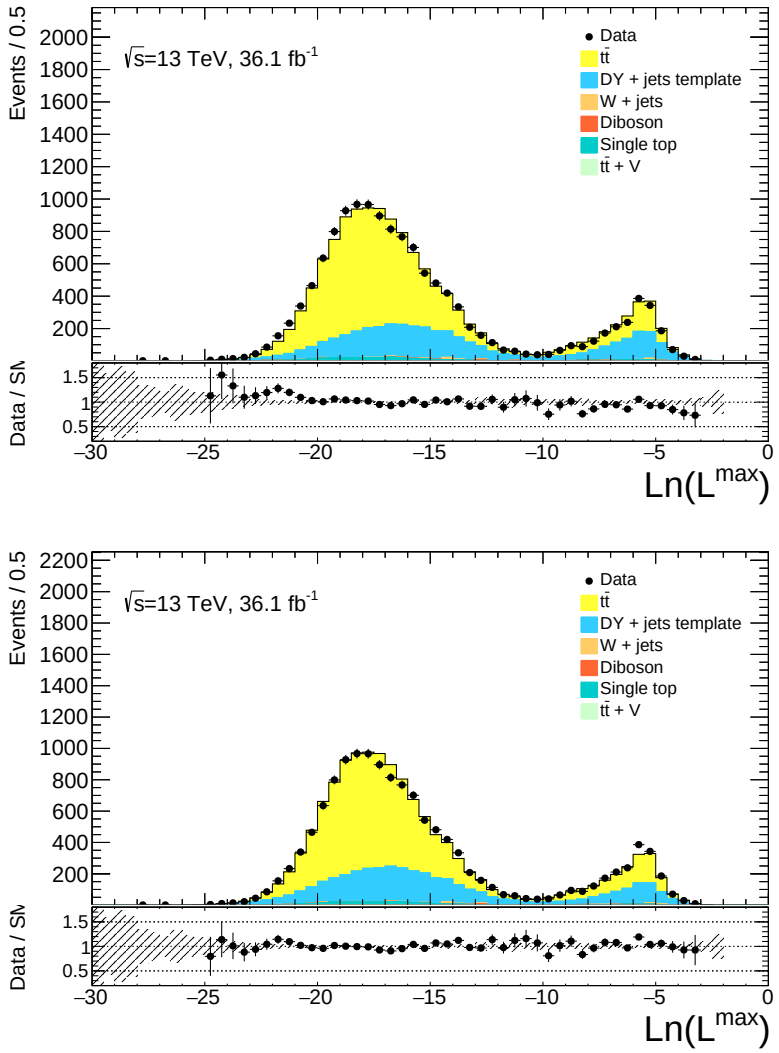

Figure 48: Effect of jet reweighting on the KLFit Log-Likelihood output, Ln($L^{\max}$), with the distributions before (top) and after (bottom) jet reweighting is applied.

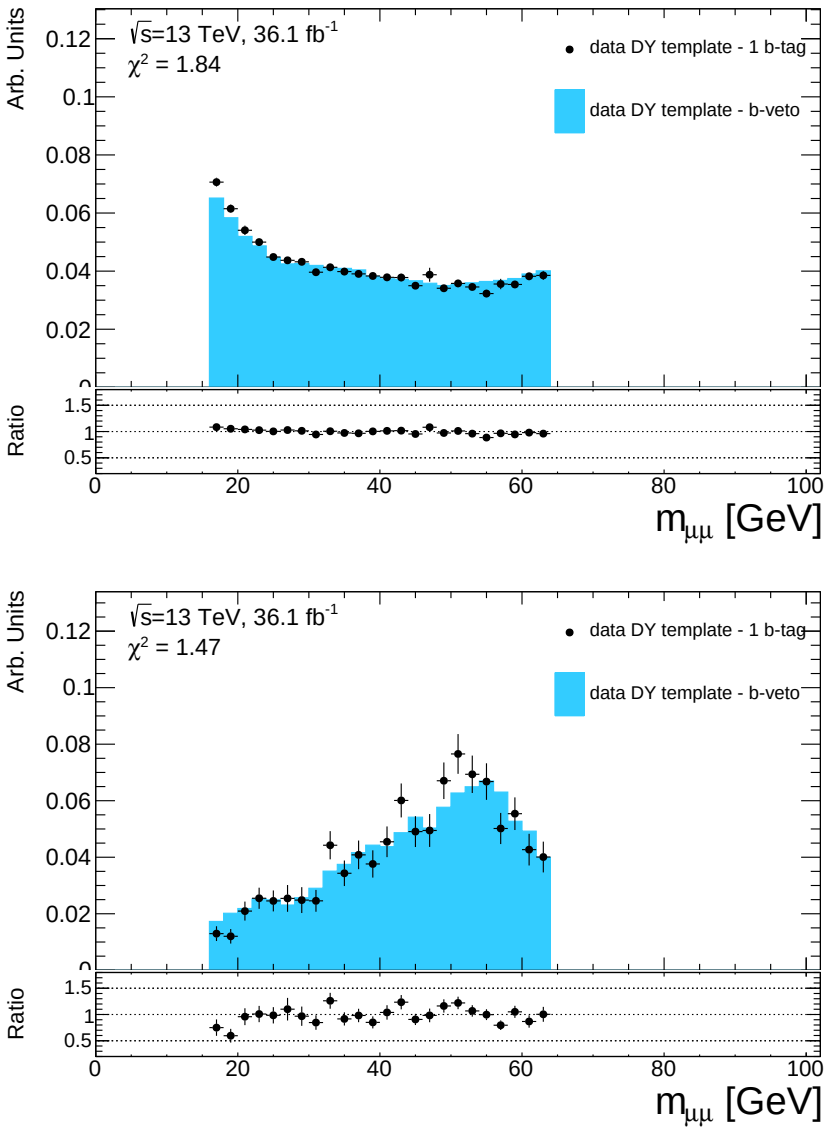

Figure 49: A comparison of the 0 b-jet and 1 b-jet template used to describe the Drell–Yan background in the $m_{\mu\mu}$ distribution. The distribution of the template is shown for the preselection plus $E_{\mathrm{T}}^{\mathrm{miss}}$ (top) and inverted signal region (bottom) selections. The ratio plot depicts the ratio of the 1-tag template prediction to the 0-tag prediction. The output of a χ^2 fit comparing the 0-tag and 1-tag template shapes is given in the upper left corner of each plot.

about 30%. This discrepancy is seen not only for the nominal MC simulation sample used in this analysis, but is also found at a similar value for an alternative generator using PYTHIA8 for the parton shower development. The shapes of several kinematic distributions in the TCR are shown in Figures 50-51, where a normalization of 1.3 has been applied to the $t\bar{t}$ predictions.

For the theoretical uncertainties on the $t\bar{t}$ background, several comparisons are made between different MC generators and parton showering programs. For the hard scatter uncertainty estimation, the event yields from a sample generated using POWHEG are compared to yields obtained from an alternative MC sample generated with aMC@NLO. Both of the samples use HERWIG++ v2.7.1 for parton showering and hadronization. An additional uncertainty on hadronization and parton shower is determined by comparing the nominal POWHEG sample interfaced with PYTHIA6 to the same POWHEG sample interfaced with HERWIG++. These theoretical uncertainties are found to be the largest uncertainties on the $t\bar{t}$ background predictions, affecting the yields in the signal region by 18% for the hard-scatter uncertainty and 16% for the parton-showering uncertainty.

The scale and radiation uncertainties for the $t\bar{t}$ background are determined by comparing the nominal POWHEG interfaced with the PYTHIA v6.248 sample to a sample using two alternative settings in PYTHIA v6.428. The first alternative setting uses a different shower tune, PERUGIA2012radHi tune, and the renormalization and factorization scales are both set to twice their nominal values. Additionally, the h_{damp} parameter in POWHEG is set to $2m_t$. These settings result in more radiation in final generated $t\bar{t}$ state. The second alternative settings uses the PERUGIA2012radLo tune, h_{damp} is set to m_t, and the renormalization and factorization scales are set to $0.5\times$ their nominal values. These changes result in less radiation in the final generated $t\bar{t}$ event.

Finally, the $t\bar{t}$ PDF uncertainty is estimated using an alternative sample generated using aMC@NLO using the CT10 PDF set interfaced with Herwig++ is used to make PDF comparison. With this sample, the nominal CT10 PDF set is compared to the PDF4LHC15 error set to obtain an uncertainty value and then also compared to the predictions from the CT14 PDF sets. The event weights used for reweighting are obtained using LHAPDF6 and the maximal difference between these PDF sets is assigned as an uncertainty. The PDF uncertainty is found to affect the event yields of $t\bar{t}$ in the signal region by up to 2%.

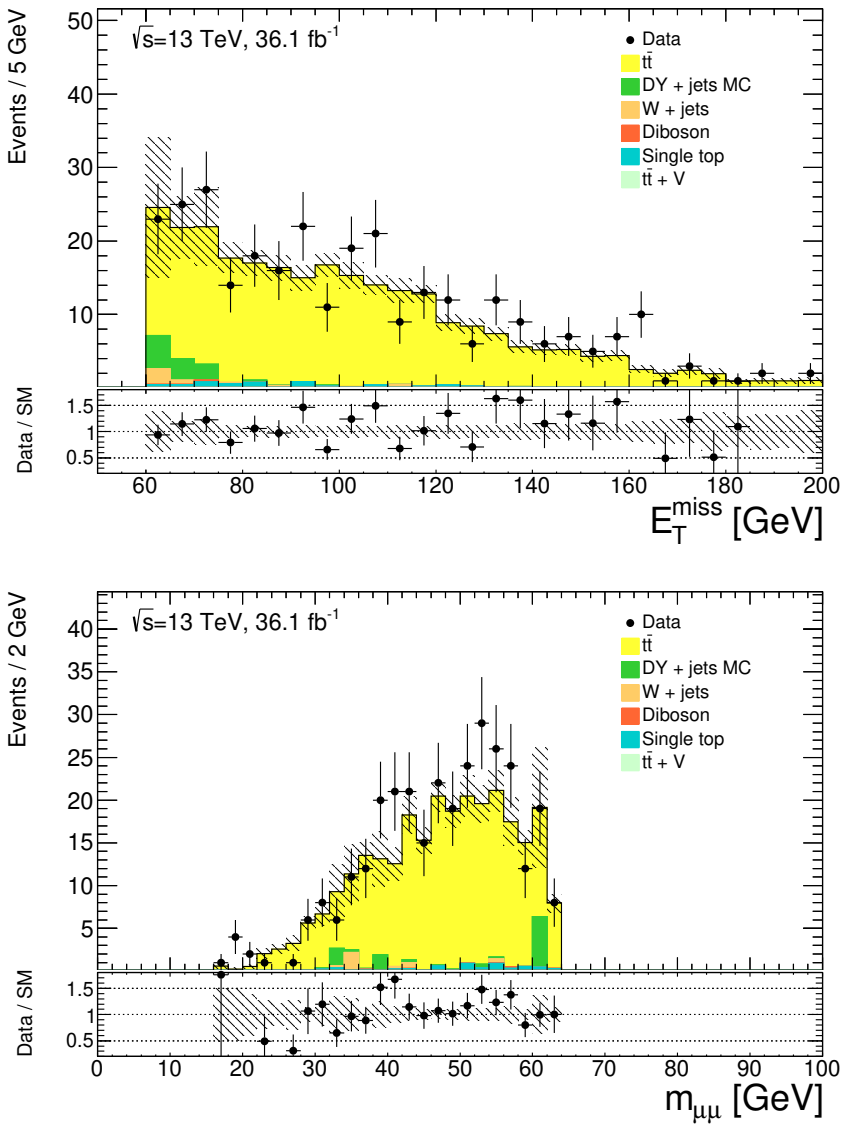

Figure 50: Kinematic Distributions of $E_{\mathrm{T}}^{\mathrm{miss}}$ and $m_{\mu\mu}$ in the Top Control Region (TCR), with data-MC simulation comparisons. The MC simulation samples have been normalized to account for the difference seen between data-MC in the $t\bar{t}$ predictions. The uncertainties include the theoretical and experimental uncertainties that affect $t\bar{t}$ production and the shape uncertainties from the small template background.

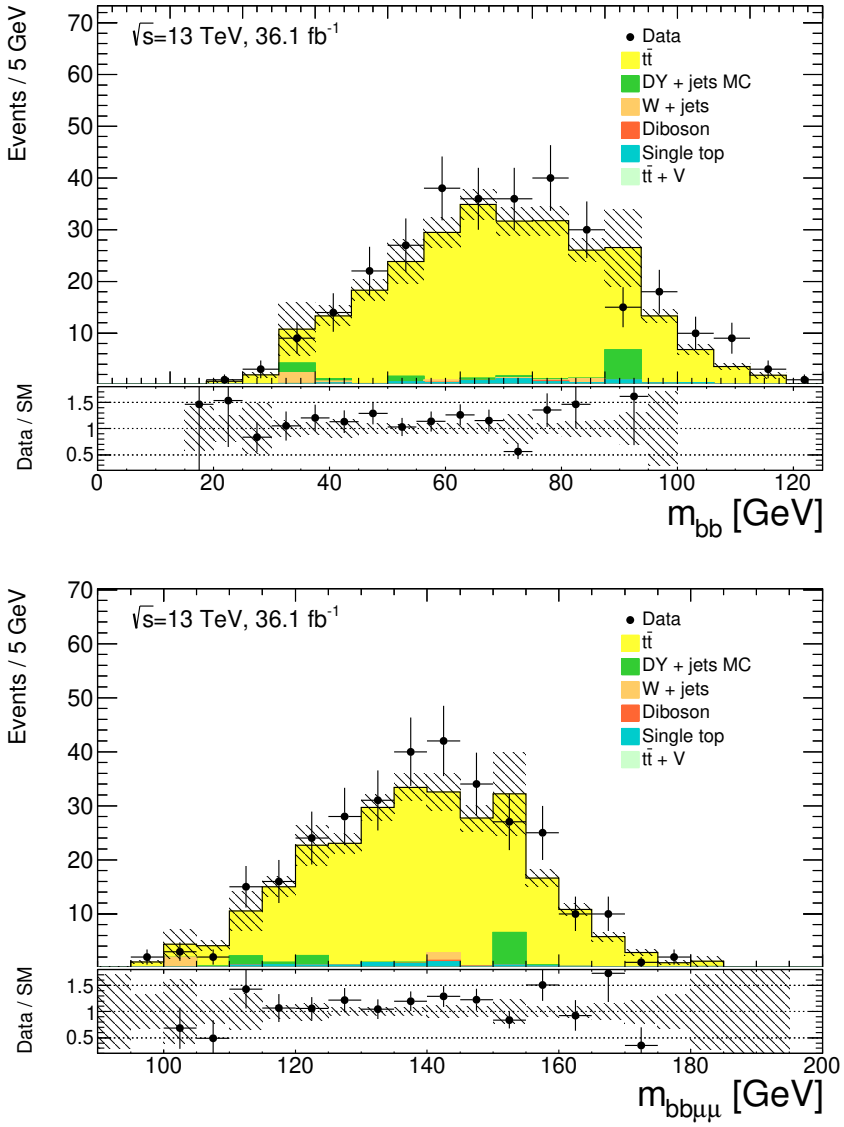

Figure 51: Kinematic Distributions of m_{bb} and $m_{bb\mu]mu}$ in the Top Control Region (TCR), with data-MC simulation comparisons. The MC simulation samples have been normalized to account for the difference seen between data-MC in the $t\bar{t}$ predictions. The uncertainties include the theoretical and experimental uncertainties that affect $t\bar{t}$ production and the shape uncertainties from the small template background.

4.8.3 Estimation of Minor Backgrounds

The minor backgrounds included in the analysis selections contribute a small number of events in the final analysis selections. Given the small effect on the analysis based on these yields, all of the yields and shapes of these backgrounds are taken solely from MC predictions.

The single-top background can enter into the final analysis selections in several ways. For the production in association with a W boson Wt, the (mis)identification of additional ISR/FSR radiation jets as a b-jet can lead to a contribution from this background. In the s-channel and t-channel production, the identification of an additional muon arising from semi-leptonic decays B-hadron decays or a QCD jet faking a muon can also pass the signal selections. The associated production with a Z boson, tZq process, can enter into the analysis selections if a fake muon is identified or via the misidentification of an additional b-jet depending on the Z and top decays. All of these effects are found to be small in the final signal selections. For the Wt and tZq, a cross-section uncertainty of 5% is applied to the MC predictions to account for the PDF and scale uncertainties on the NNLO cross-section calculation [95].

Several dibosonic processes can enter into the analysis selections. The ZZ$\rightarrow \mu\mu qq$ and WZ$\rightarrow \mu\mu qq$ backgrounds enter in the analysis selections when the dilepton mass is below $m_{\mu\mu} < 64$ GeV and the two jets are (mis)identified as b-jets. The Z mass veto removes a large portion of the ZZ background and the b-tagging requirements remove the WZ background. The WZ$\rightarrow \mu\nu qq$ process can also enter the analysis selections if a non-prompt muon passes the isolation criteria. More exotic combinations, such as decays involving electrons being identified as the b-jets, are found to be negligible. For the dibosonic MC samples, a normalization uncertainty of 10% is applied to account for the effects of higher order QCD corrections that affect the cross section of these processes [137].

W boson production in associated with 2 b-jets can enter the analysis selections if a non-prompt muon is reconstructed and identified as isolated. The W boson production is modeled in Sherpa 2.2.1 yield and shapes are taken from the MC predictions. The nominal theoretical uncertainty on the cross section for these samples is 5% [118]. A conservative uncertainty of 100% uncertainty is assigned to the W+jets MC yield due to the lack of statistical precision in the MC sample for modeling the non-prompt muon isolation.

In addition to these specialized cross section uncertainties, all experimental uncertainties are determined for these minor backgrounds, though the effect of these uncertainties is found to be negligible effect on the final background yields.

5 Results and Sensitivity to Exotic Higgs Decays

The observation of a new particle in the exotic decays of the 125 GeV Higgs boson, or the lack of evidence for this new state, relies on the use of a dedicated statistical procedure to quantify the background and signal process event yields after the final analysis selections. The statistical methods are presented in Section 5.1 and the different versions of the fit are discussed in Section 5.2. The exact treatment of systematic uncertainties as nuisance parameters is discussed in Section 5.3, while the results of the background-only fit, validation of the fits, and final results are given in Sections 5.4-5.6.

5.1 Statistical Methods

To estimate the background and potential signal contributions in each signal region defined by the dimuon mass windows, a series of profile likelihood fits are performed [138]. In each fit, a probability density function (PDF) for the signal and background processes is built from binned histograms of the background contribution using the MC simulation and the Drell–Yan template described in Section 4.8.1.

To determine the probability for observing an exotic Higgs signal in the signal regions, a model must be built that addresses both the signal and background contributions. To begin, the probability for observing N events for the signal region of interest can be expressed as a Poisson distribution

$$P(N|\lambda) = \frac{e^{-\lambda}\lambda^{N}}{N!} \tag{90}$$

where λ is the number of expected or predicted events, compared to N events observed in the signal region. For the sum of the signal plus background yields, a PDF can be built to determine the number of predicted events. The PDF for a signal model and several backgrounds for a single signal region bin is given by

$$F(x) = \mu_{\text{sig}} S(x) + \sum_{i} \mu_i B_i(x) \tag{91}$$

where S(x) describes the PDF of the signal model, μ_{sig} defines the signal strength, each μ_i defines the normalization parameter for a given background, and $B(x)$ defines the background PDF. The μ values are treated as free parameters that are adjusted for both the signal and background in order to best match the data. To accomplish this adjustment for the background, measurements are made in the dedicated data CR regions discussed in Section 4.8, where the signal contribution is expected to be minimal.

Although the final results present μ_{bkg} for the Drell–Yan and $t\bar{t}$, a redefinition of $\mu_{\text{bkg}} \equiv \beta_i$ and $\mu_{\text{sig}} = \mu$ is made to simplify notation.

To extract these normalizations, a binned profile log likelihood ratio is applied to each region of interest, either a signal or control region. The likelihood, $L(\mu, \theta)$, used in these fits takes the form [138]

$$\begin{aligned} L(\boldsymbol{n}, \boldsymbol{\theta^0} | \mu, \boldsymbol{\beta}, \boldsymbol{\theta}) &= P_{\text{SR}} \times P_{\text{CR}} \times C_{\text{Syst}} \\ &= P(n_{\text{S}} | \lambda_{\text{S}}(\mu, \boldsymbol{\beta}, \boldsymbol{\theta})) \times \prod_{i \in \text{CR}} P(n_i | \lambda_i(\mu, \boldsymbol{\beta}, \boldsymbol{\theta})) \\ &\times C_{\text{Syst}}(\boldsymbol{\theta^0}, \boldsymbol{\theta}) \\ &= \prod_{i=1}^{n_S} \frac{e^{-\lambda_i^S} \lambda^{N_i^S}}{N_i^S!} \times \prod_{j=1}^{n_B} \frac{e^{-\lambda_j^S} \lambda^{N_j^B}}{N_j^B!} \times C_{\text{Syst}}(\boldsymbol{\theta^0}, \boldsymbol{\theta}) \end{aligned} \tag{92}$$

where each i denotes a signal region measurement, each j describes a background control region measurement, each λ accounts for the Poisson expectations for the signal and background processes, n_S and n_B are the number of bins used for measurement, $\boldsymbol{\beta}$ is a vector of background normalizations, and $\boldsymbol{\theta}$ defines a vector of nuisance parameters that account for the systematic uncertainties. $C_{\text{Syst}}(\boldsymbol{\theta^0}, \boldsymbol{\theta})$ describes a product of PDFs for each nuisance parameter, with θ_i^0 being the nominal value before each systematic is allowed to act on the yield in a signal or control region. In this analysis, the PDFs for most of the nuisance parameters is described as a Gaussian distribution of the form

$$f(\theta_0 | \theta) = e^{-(\theta - \theta^0)^2/2} \tag{93}$$

An exception is made for counting uncertainties, such as MC simulation statistical uncertainties, which are treated with Poisson probability terms (Eq. 90). The nominal value for each nuisance parameter, θ_i^0, is set to 0 and each θ_i is allowed to vary around this central value. The values of $\theta_i = \pm 1$ then correspond to one sigma deviation values of each systematic. The exact values for θ_i are obtained via the methods described in Section 4.7 and all systematic uncertainties are put into the likelihood. The fit, therefore, can account for the correlations between different uncertainties such as the correlations between the lepton momenta variations and the $E_{\text{T}}^{\text{miss}}$ uncertainties.

Each Poisson expectation term λ for the signal and background depend inherently on the normalization of the backgrounds in the n_B control regions, the signal strength μ, and the systematic uncertainty nuisance parameters due to their effect on the normalizations. In order to extrapolate the background predictions from the control

regions to the signal region, the Poisson expectations can be defined in terms of transfer factors c,

$$\lambda_i(\mu, \boldsymbol{\beta}, \boldsymbol{\theta}) = c_{s,i}(\boldsymbol{\theta}) \times \mu N^S + \sum_j c_{j,i}(\boldsymbol{\theta}) \times \beta_j N_j^B \tag{94}$$

and each transfer factor is defined as

$$c_{j,i}(\boldsymbol{\theta}) = \frac{N_i^{B(S)}}{N_j^{B(S)}} \times \left(1 + \sum_k \Delta_{j,i,k}\theta_k\right) \tag{95}$$

where the index i denotes the signal region, index j runs overall all of the CRs and backgrounds, $N^{B(S)}$ denote the unnormalized predictions of the background (signal) in a given region, and $\Delta_{j,i,k}$ defines the magnitude for a given uncertainty θ_k on the λ_k in region i for the physics process j. As an example, the 14% uncertainty on the $t\bar{t}$ background process due to the X systematic gives rise to $\Delta_{j,i,k} = 0.06$.

In this picture, if there are uncertainties that affect the control region and signal region equally, then the transfer factors cancel out the effect of this uncertainty. These cancellations are optimized by choosing control regions that are as similar to the signal region kinematics as possible, by inverting as few selections as possible. This motivation provided the minimal selections for the CRs in the $H \to aa \to bb\mu\mu$ analysis selections.

With all of these uncertainties and background and signal expectations accounted for, the data is profiled in the full fit. In this analysis, the compatibility of the data and the signal strength μ, is accomplished using a profile likelihood ratio. The profile likelihood ratio is commonly expressed as [139]

$$\Lambda(\mu) = \frac{L(\mu, \hat{\hat{\boldsymbol{\theta}}}(\mu))}{L(\hat{\mu}, \hat{\boldsymbol{\theta}})} \tag{96}$$

where $\hat{\hat{\boldsymbol{\theta}}}_{\boldsymbol{i}}$ defines the value of each nuisance parameter that maximizes the likelihood for a specific value of μ under test. Similarly, $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the values that maximize the unconditional likelihood function. The ratio of these likelihood values is defined on the range $0 \leq L(\mu)/L(\hat{\mu}) \leq 1$, where values closer to 1 indicate agreement between the data and the value of μ under test.

In the fits for this analysis, it is assumed that the signal yield must be greater than or equal to 0, with any deficit seen simply providing a stronger upper limit on the signal production process. In the fits, however, $\hat{\mu}$ is modeled as a Gaussian

distributed variable, necessitating an effective profile likelihood ratio of the form

$$\tilde{\Lambda}(\mu) = \begin{cases} \frac{L(\mu,\hat{\hat{\boldsymbol{\theta}}}(\mu))}{L(\hat{\mu},\hat{\boldsymbol{\theta}})}, & \hat{\mu} \geq 0 \\ \\ \frac{L(\mu,\hat{\hat{\boldsymbol{\theta}}}(\mu))}{L(0,\hat{\boldsymbol{\theta}}(0))}, & \hat{\mu} < 0 \end{cases} \tag{97}$$

where the value of the likelihood for $\mu = 0$ is taken as the best estimation for any value when $\hat{\mu} < 0$.

The test-statistic for the given μ hypothesis, denoted as t_μ is related to $\Lambda(\mu)$ via

$$t_\mu = -2 \text{ Ln } \Lambda(\mu) \tag{98}$$

and the appropriately biased terms are denoted as $\tilde{t}_\mu$. Given the constraint on $0 \leq \Lambda(\mu) \leq 1$, large values of t_μ indicate an incompatibility between the data and the μ value under test. In order to calculate the level of disagreement between the observed data and the hypothesized yield being tests, the p-value is defined

$$p_\mu = \int_{t_\mu,\text{obs.}}^{\infty} f(t_\mu|\mu) dt_\mu \tag{99}$$

where $t_{\mu,\text{obs.}}$ is the value of t_μ observed in the data and $f(t_\mu|\mu)$ is the PDF of t_μ for a particular signal strength μ.

Equation 97 is modified in the special cases of testing the background-only hypothesis when $\mu = 0$, where rejecting the background-only hypothesis leads to the discovery of a new physics signal, and the case of setting upper limits on the signal strength. In the case where $\mu = 0$ the biased test-statistic $\tilde{t}_0 \equiv q_0$ is given by

$$q_0 \equiv \begin{cases} -2 \text{ Ln } \frac{\text{L}(0,\hat{\hat{\boldsymbol{\theta}}}(0)}{\text{L}(\hat{\mu},\hat{\boldsymbol{\theta}})} & \hat{\mu} \geq 0 \\ \\ 0 & \hat{\mu} < 0 \end{cases} \tag{100}$$

where q_0 is used rather than t_0, as t_0 allows for the positive (increasing yield) fluctuations in the background data to reject the background only hypothesis. In the case of testing the signal from $H \rightarrow aa \rightarrow bb\mu\mu$, a large negative $\hat{\mu}$ should only indicate mis-modeling, not the discovery of a new physics signal. This compatibility test statistic q_0 also comes with a corresponding p_0 which denotes the p-value in the case of the background only fit.

Alternatively, the upper limit test statistic is defined by

$$q_\mu \equiv \begin{cases} -2 \text{ Ln } \dfrac{\mathrm{L}(\mu, \hat{\hat{\boldsymbol{\theta}}}(\mu)}{\mathrm{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \tag{101}$$

where q_μ goes to zero in the presence of upward data fluctuations ($\hat{\mu} > \mu$). This fact contrasts with Equation 100, where the downward fluctuations are set to 0. When the number of events in the data sample is sufficiently large, Equations 100-101 follow the asymptotic form, where the discovery test statistic is given by [140, 141]

$$q_0 \equiv \begin{cases} \dfrac{\hat{\mu}^2}{\sigma^2} & \hat{\mu} \geq 0 \\ 0 & \hat{\mu} \leq 0 \end{cases} \tag{102}$$

and the upper limit test statistic takes the form

$$q_\mu \equiv \begin{cases} \dfrac{(\mu - \hat{\mu})^2}{\sigma^2} & \hat{\mu} \geq 0 \\ 0 & \hat{\mu} \geq \mu \end{cases} \tag{103}$$

Each of these test statistics assumes that $\hat{\mu}$ follows a Gaussian distribution with a standard deviation of σ.

To derive upper limits on new-physics models, the CL_s method [139, 142] is used. The CL_s value defines a confidence level given by

$$\begin{aligned} \mathrm{CL}_s &= \frac{P(q \geq q_{\mathrm{obs}} | \mathrm{sig} + \mathrm{bkg})}{P(q \geq q_{\mathrm{obs}} | \mathrm{bkg})} \\ &= \frac{p_{\mathrm{sig+bkg}}}{1 - p_{\mathrm{bkg}}} \end{aligned} \tag{104}$$

where p_{s+b} denotes the p-value for a signal and background hypothesis, while p_b is the p-value for the background only. In high-energy physics, exclusions are typically given at the 95% confidence level, which requires $\mathrm{CL}_s < 0.05$. For a constraint of 0.05, the likelihood that the hypothesis was incorrectly excluded is then required to be below 5%. Additionally, the use of CL_s, due to the explicit requirement on $s + b$, ensures that only models to which the experiment is sensitive are excluded.

5.2 Fit Configurations

In this analysis, three configurations of the profile likelihood fits are performed:

Background-only fit: In the background only fit, the top control region and Drell-Yan control regions are used to extract the background predictions for these two backgrounds. In this version of the fit, μ_{sig} is set to 0 and the minor background processes are fixed to yields derived from the MC simulation samples. $\mu_{\text{t}\bar{\text{t}}}$ and μ_{DY} are allowed to float and the best fit value for both parameters is used to transfer the result of the fit to the signal and validation regions.

After estimating the background contributions in the signal region with the background fits, a signal optimization can occur in the dimuon mass spectrum. The signal can be overlaid to estimate the significance, characterized as the number of signal events over the square root of the number of background events, $N_{\text{sig}}/\sqrt{N_{\text{bkg}}}$. The number of signal and background events are then constrained within dimuon mass windows in order to determine the highest discovery potential, or in the absence of signal, the best upper limits on physics process production.

Model Dependent Fit: In this version of the fit, the TCR, DYCR, and SR are used to extract the DY, $t\bar{t}$, and signal yield predictions. The signal sample corresponding to the SR of interest is included, with all of its nuisance parameters, in addition to the background samples. The μ_{sig},$\mu_{t\bar{t}}$, and μ_{DY} are all simultaneously fit to the data in all three regions. As the signal contamination is only approximately 0 in both of the control regions, the fit values of $\mu_{t\bar{t}}$, and μ_{DY} may differ from the values derived in the background-only fit.

Model Independent Fit: In this version of the fit, the TCR, DYCR, and SR are again used to extract the DY and $t\bar{t}$ yields, but the signal yield and uncertainties are not determined from the MC simulation samples. Instead, the systematic uncertainties for the signal are removed from the fit and a test signal yield is set to exactly 1 event. $\mu_{t\bar{t}}$, and μ_{DY} are still accounted for in the fit, but μ_{sig} is exchanged for μ_{95} which corresponds to the highest normalization allowable by the 95% CL. In the absence of a definitive signal model, μ_{95} corresponds to a limit on the number of new physics events allowed in the final signal regions.

5.3 Profiling Systematic Uncertainties

The systematic uncertainties described in Section 4.7 affect the overall expected event yields in the DYCR, TCR, and the SR. The change in these event yields

affects the outcome of the best fit to the normalization parameters $\mu_{t\bar{t}}$,μ_{DY} and μ_{sig} and so the fit must account for the changes each systematic has on the underlying distributions. Each uncertainty is classified in one of two ways, depending on the effect the systematic has during the fits:

Shape Uncertainties: Shape uncertainties are assumed to affect only the shapes of the nominal input distribution (in this case, the dimuon mass distribution). These uncertainties are assigned to affect only the transfer factors and do not affect the normalization parameters μ_i. For the Drell-Yan and $t\bar{t}$ backgrounds and $bb\mu\mu$ signal, almost all of the experimental and theoretical uncertainties are classified as shape uncertainties.

Normalization Uncertainties: Normalization uncertainties affect the μ_i parameters for the background and signal processes. The uncertainties on the minor backgrounds (including all experimental, theoretical, and cross section uncertainties) account for most of the normalization uncertainties as the normalization of these samples is not explicitly derived from data control regions.

During the likelihood fit, the optimal values and errors on the nuisances parameters are determined simultaneously; therefore, the fit accounts for all correlations between these nuisance parameters.

All of the background uncertainties discussed in Sections 4.7 and 4.8 are put into the fit. For the experimental uncertainties, any uncertainty found to have a negligible effect on both the signal and background is removed from the fit in order to improve fit stability and computation time. The criterion for being labeled as negligible is if the systematic is associated with a minor background and has an impact on the minor background yield of $< 10\%$.

For the theoretical uncertainties, a similar pruning approach is taken. For the minor backgrounds, only the uncertainty on the cross section is taken as a nuisance parameter in the fits. For the $t\bar{t}$ background, all of the theoretical uncertainties stemming from variations in the PS, the PDF, the ME choice, and the hard-scatter generator choice are profiled in the fit. Similarly, the signal uncertainties on the cross section, hadronization, factorization, and PDF choices are accounted for in the model-dependent fits.

A summary of all of the uncertainties that are profiled in the fits and the background estimations that these uncertainties affect are given in Tables 20-21.

Source	Description	Systematics name	Parameter Name in fit
Pile Up	Data Pile up SF	PRW_DATASF	PRW
Trigger	Efficiency SF		
Muons	p_{T} resolution MS	MUONS_MS	MU_MS
Muons	p_{T} resolution ID	MUONS_ID	MU_ID
Muons	p_{T} resolution	MUON_SAGITTA_RESBIAS	MU_SagRes
Muons	p_{T} resolution	MUON_SAGITTA_RHO	MU_SagRho
Muons	p_{T} scale	MUONS_SCALE	MU_Scale
Muons	Isolation efficiency SF	MUON_ISO_SYS	MU_ISO_SYS
Muons	Isolation efficiency SF	MUON_ISO_STAT	MU_ISO_STAT
Muons	Identification efficiency	MUON_TTVA_SYS	MU_TTVA_SYS
Muons	Identification efficiency	MUON_TTVA_STAT	MU_TTVA_STAT
Muons	Identification efficiency SF	MUON_EFF_SYS	MUON_EFF_SYS
Muons	Identification efficiency SF	MUON_EFF_STAT	MUON_EFF_STAT
Muons	ID and Reco Efficiency	MUON_EFF_STAT_LOWPT	MUON_EFF_STAT_LOWPT
Muons	ID and Reco Efficiency	MUON_EFF_SYS_LOWPT	MUON_EFF_SYS_LOWPT
MET	Soft term	MET_SoftTrk_ResoPerp	MET_ResPerp
MET	Soft term	MET_SoftTrk_ResoPara	MET_ResPara
MET	Soft term	MET_SoftTrk_ScaleUp	MET_Scale
Jets	JES strongly reduced	JET_GroupedNP_1	JES_SR_NP1
Jets	JES strongly reduced	JET_GroupedNP_2	JES_SR_NP2
Jets	JES strongly reduced	JET_GroupedNP_3	JES_SR_NP3
Jets.	JES strongly reduced	JET_SR1_JET_EtaIntercalibration_NonClosure	JES_SR_NP4
Jets	Energy resolution	JET_JER_SINGLE_NP	JER
Jets	JVT efficiency SF	JET_JvtEfficiency	JVT
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_Light0	btag_L0
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_Light1	btag_L1
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_Light2	btag_L2
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_Light3	btag_L3
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_Light4	btag_L4
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_B0	btag_B0
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_B1	btag_B1
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_B2	btag_B2
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_C0	btag_C0
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_C1	btag_C1
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_C2	btag_C2
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_C3	btag_C3
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_extrapolation	btag_extrap
B-tagging	Flavor tagging scale factors	FT_EFF_Eigen_extrapolation_from_charm	btag_extrapC
e/γ	e/γ resolution smearing	EG_RESOLUTION_ALL	EG_Res
e/γ	e/γ scale smearing	EG_SCALE_ALL	EG_Scale
Modeling	$t\bar{t}$ Hard Scatter	TopMCaNLOtt	top_PowMCnlo
Modeling	$t\bar{t}$ Radiation	ISR/FSR	top_Rad
Modeling	$t\bar{t}$ Parton Shower	Parton Shower	top_pythHerw
Modeling	$t\bar{t}$ PDF	–	
Modeling	DY template shape	DDZ_shape	DDZ_shape
Modeling	DY template: MC subtraction	DDZ_subtraction	DDZ_subtraction
Modeling	DY template jet reweighting	DDZ_reweighting	DDZ_reweighting
Luminosity	Luminosity measurement	LUMI	
Modeling	Cross Section, simulated bkgs	W+jets, db, st, $t\bar{t}$V	*_Xsec
Modeling	Cross Section	signal	sig_Xsec
Modeling	$\{\mu_R, \mu_F\}$ scales	signal	sig_RenFact
Modeling	Radiation	signal	sig_ISR
Modeling	PDF	signal	sig_PDF
Modeling	H_{p_T}	signal	sig_HpTrew
Modeling	VH contribution	signal	sig_VH
Modeling	signal acceptance interpolation	interpolated signal points only	sig_interp

Table 20: A summary and glossary of all systematic uncertainties affecting the signal and background estimation in this analysis.

Uncertainty Name	Sample			
	Signal MC	$t\bar{t}$	DY	Other
PRW_DATASF	✓	✓		✓
LUMI	✓			✓
MUONS_MS	✓	✓		✓
MUONS_ID	✓	✓		✓
MUON_SAGITTA_RESBIAS	✓	✓		✓
MUON_SAGITTA_RHO	✓	✓		✓
MUONS_SCALE	✓	✓		✓
MUON_ISO_SYS	✓	✓		✓
MUON_ISO_STAT	✓	✓		✓
MUON_TTVA_SYS	✓	✓		✓
MUON_TTVA_STAT	✓	✓		✓
MUON_EFF_SYS	✓	✓		✓
MUON_EFF_STAT	✓	✓		✓
MUON_EFF_STAT_LOWPT	✓	✓		✓
MUON_EFF_SYS_LOWPT	✓	✓		✓
MET_SoftTrk_ResoPerp	✓	✓		✓
MET_SoftTrk_ResoPara	✓	✓		✓
MET_SoftTrk_ScaleUp	✓	✓		✓
JET_GroupedNP_1	✓	✓		✓
JET_GroupedNP_2	✓	✓		✓
JET_GroupedNP_3	✓	✓		✓
JET_SR1_JET_EtaIntercalibration_NonClosure	✓	✓		✓
JET_JER_SINGLE_NP	✓	✓		✓
JET_JvtEfficiency	✓	✓		✓
FT_EFF_Eigen_Light0	✓	✓		✓
FT_EFF_Eigen_Light1	✓	✓		✓
FT_EFF_Eigen_Light2	✓	✓		✓
FT_EFF_Eigen_Light3	✓	✓		✓
FT_EFF_Eigen_Light4	✓	✓		✓
FT_EFF_Eigen_B0	✓	✓		✓
FT_EFF_Eigen_B1	✓	✓		✓
FT_EFF_Eigen_B2	✓	✓		✓
FT_EFF_Eigen_C0	✓	✓		✓
FT_EFF_Eigen_C1	✓	✓		✓
FT_EFF_Eigen_C2	✓	✓		✓
FT_EFF_Eigen_C3	✓	✓		✓
FT_EFF_Eigen_extrapolation	✓	✓		✓
FT_EFF_Eigen_extrapolation_from_charm	✓	✓		✓
EG_RESOLUTION_ALL	✓	✓		✓
EG_SCALE_ALL	✓	✓		✓
TopMCaNLOtt		✓		
TopRad		✓		
TopHerwig		✓		
TopPDF		✓		
DDZ_shape			✓	
DDZ_reweighting			✓	
DDZ_subtraction			✓	
*_XSEC				✓
sig_Xsec	✓			
sig_RenFact	✓			
sig_ISR	✓			
sig_PDF	✓			
sig_HpTrew	✓			

Table 21: Categorization of systematic uncertainties by the affected background and signal estimates.

5.4 Background Only Fits

The first results shown are the results from the background only fit described in Section 5.2, where the values for $\mu_{t\bar{t}}$ and μ_{DY} are extracted from simultaneous fits to the TCR and DYCR. The fits are run using the statical implementation of the methods described above using the HistFitter [138] package, which takes as inputs the predicted and observed total yields in the TCR and DYCR. Each of the control regions is treated as a single bin, with the normalization of each background adjusted to best fit the total yield of each control region.

The DYTR distributions, after applying the leading jet reweighting to account for the kinematic bias, have yields which are arbitrary when considering the DY contributions to the DYCR and SR[10]. As such, the unscaled DYTR predictions overshoot the data by a large factor before being constrained by the profile likelihood fits. This effect is well illustrated in Table 22 where the event yields are a factor of 40 higher than the total data yield in the DYCR before the fit. All of the other background yields are taken from MC simulation and are only allowed to vary within their systematic uncertainties. The normalization factors determined by the background only fit are found in Table 23.

The $m_{bb}^{\mathrm{KL}}, m_{\mu\mu}, m_{bb\mu\mu}^{\mathrm{KL}}$,and $E_{\mathrm{T}}^{\mathrm{miss}}$ distributions for the TCR and DYCR are shown in Figures 52-55 after the profile likelihood fit. From these distributions, the fit is found to model the control region data well.

Before finalizing the background-only fit, a signal optimization is performed to identify the most sensitive dimuon mass bins for the m_a signals. The background only fit is run with three setups, each setup differing in the treatment dimuon mass width corresponding to 2, 3, 4 GeV around the signal mass hypotheses. The lowest dimuon width is selected due to the minimum ATLAS experimental dimuon resolution in the range $20 \leq m_a \leq 60$. The resulting expected limits are shown in Figure 56, where only the gluon-fusion simulated signals are used for simplicity. From these results, the best limits are obtained using 2 GeV dimuon mass window widths in the range $20 \leq m_{\mu\mu} \leq 40$ GeV, 3 GeV windows for $40 < m_a < 50$ GeV, and 4 GeV windows for $50 \leq m_a \leq 60$ GeV. These same mass windows are also used in the model-independent limits as these optimizations align with the dimuon mass experimental resolution.

[10]The yield in the DYCR corresponds approximately to the jet-by-jet convolution of the ratio between the b-jet to light jet cross-sections of the Drell–Yan production, where full p_{T}, b-tagging efficiency and mis-tag rates, and other experimental uncertainties have been propagated.

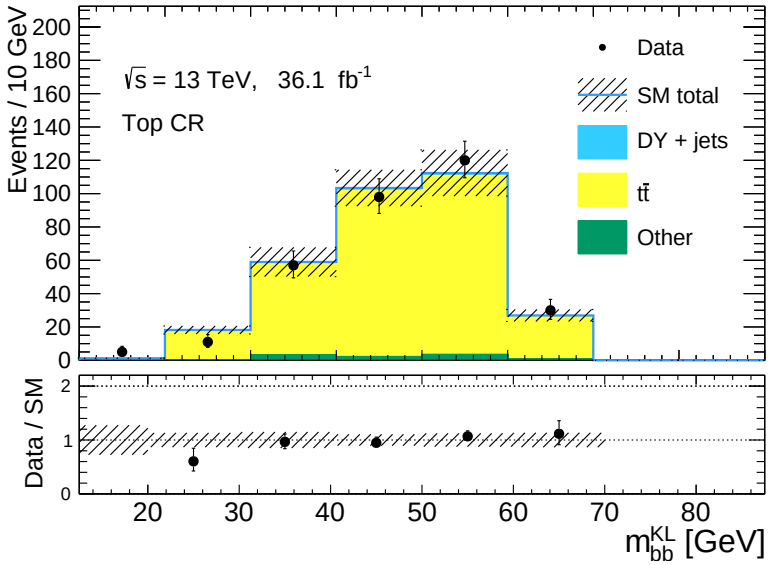

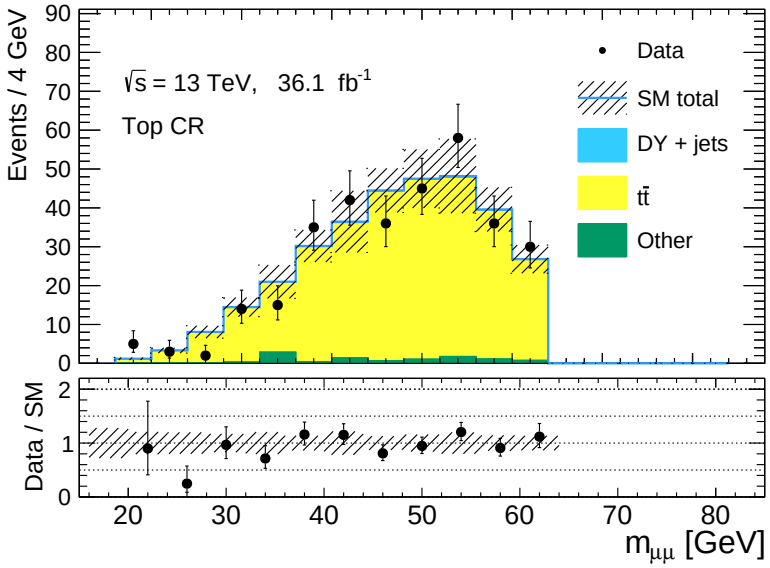

Figure 52: The m_{bb}^{KL} and $m_{\mu\mu}$ distributions in the top control region (TCR) after the profile likelihood fits in the background-only configuration. All distributions are found to be well modeled.

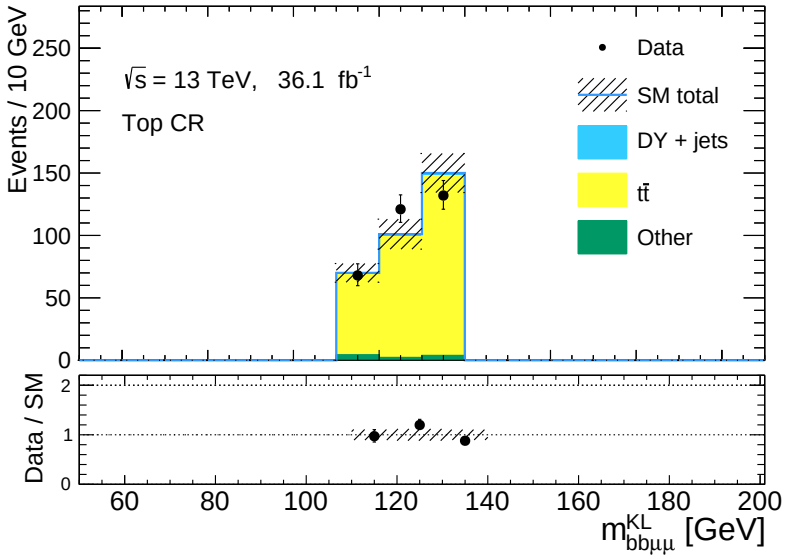

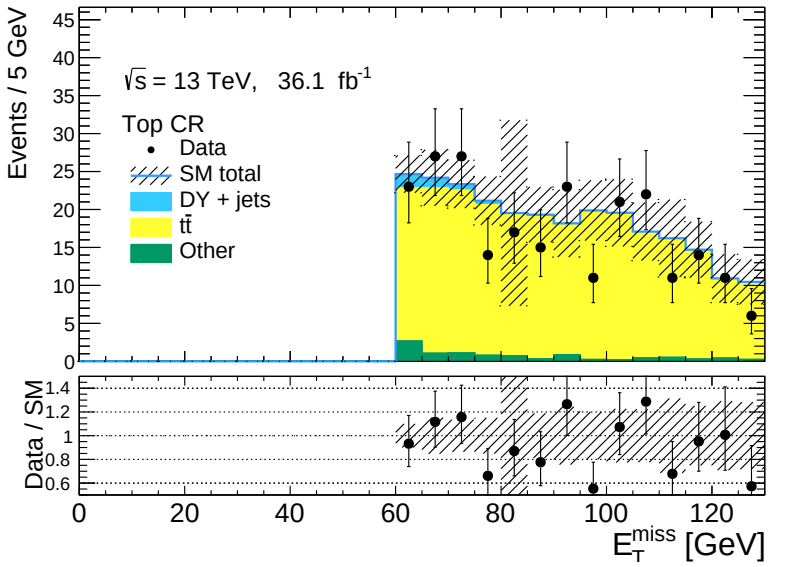

Figure 53: The $m_{bb\mu\mu}^{\mathrm{KL}}$ and $E_{\mathrm{T}}^{\mathrm{miss}}$ distributions in the top control region (TCR) after the profile likelihood fits in the background-only configuration. All distributions are found to be well modeled.

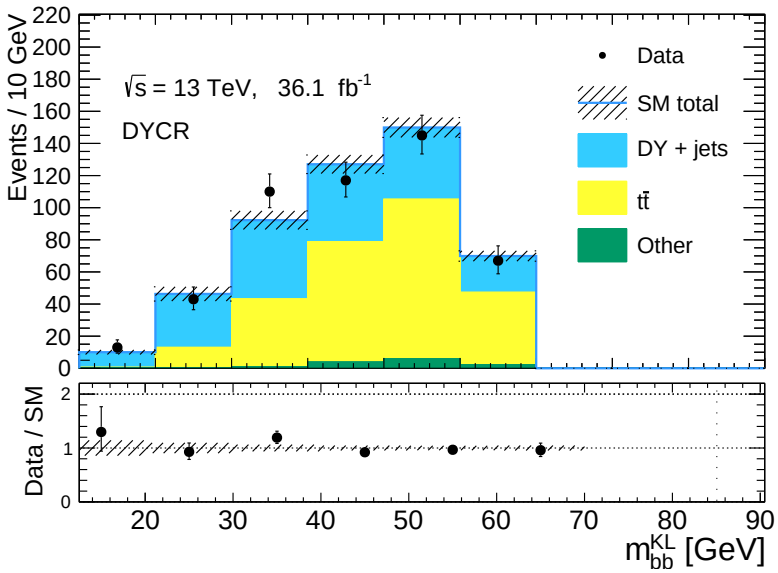

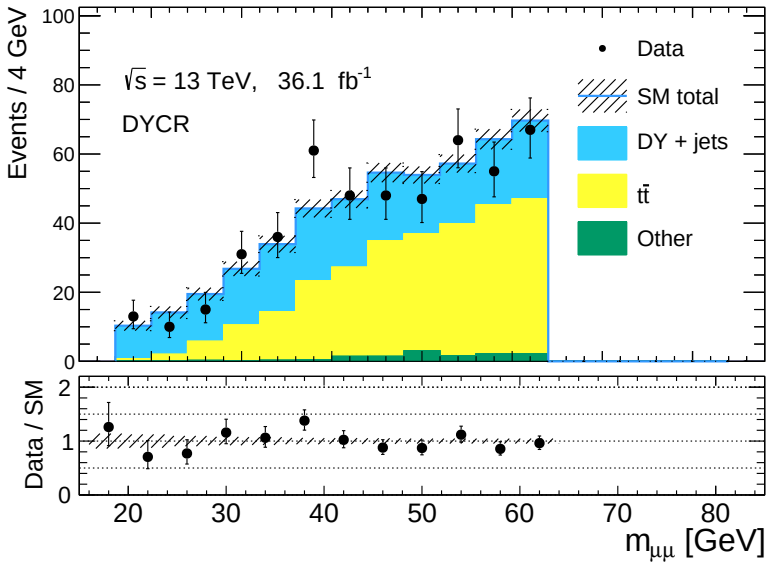

Figure 54: The m_{bb}^{KL} and $m_{\mu\mu}$ distributions in the Drell–Yan control region (DYCR) after the profile likelihood fits in the background-only configuration. All distributions are found to be well modeled after the profile likelihood fits.

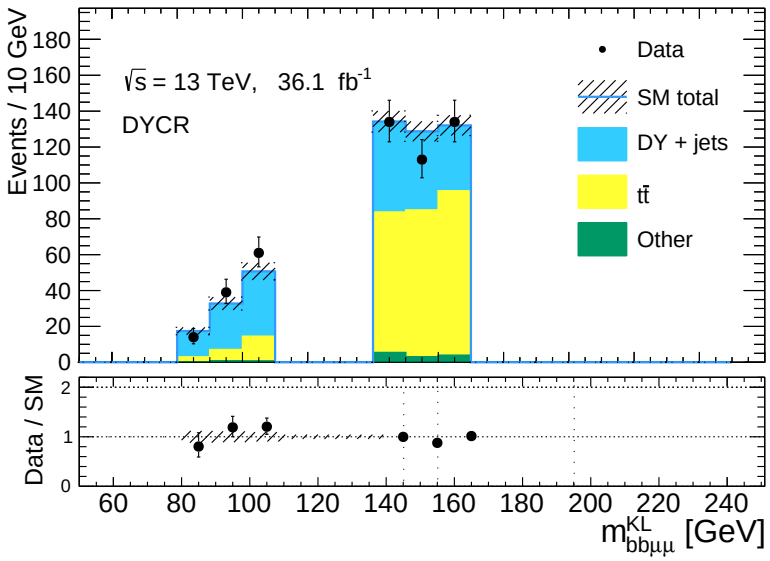

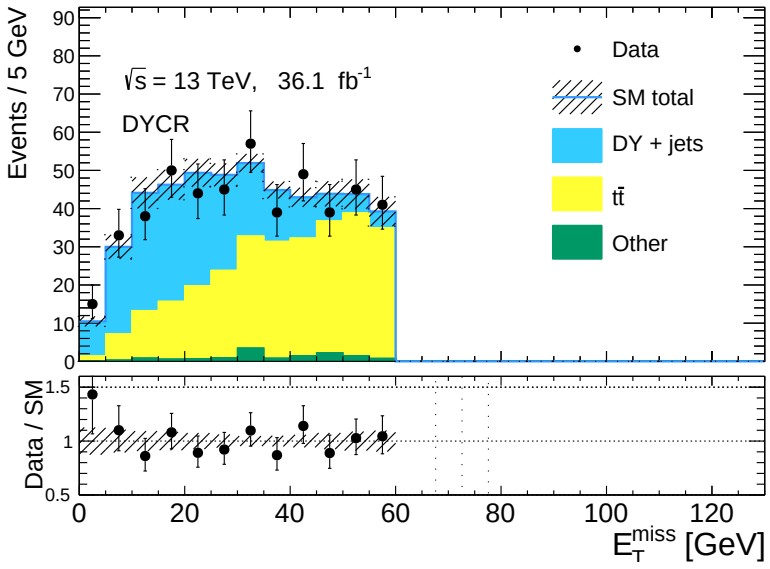

Figure 55: The $m^{\mathrm{KL}}_{bb\mu\mu}$ and $E^{\mathrm{miss}}_{\mathrm{T}}$ distributions in the Drell–Yan control region (DYCR) after the profile likelihood fits in the background-only configuration. All distributions are found to be well modeled after the profile likelihood fits.

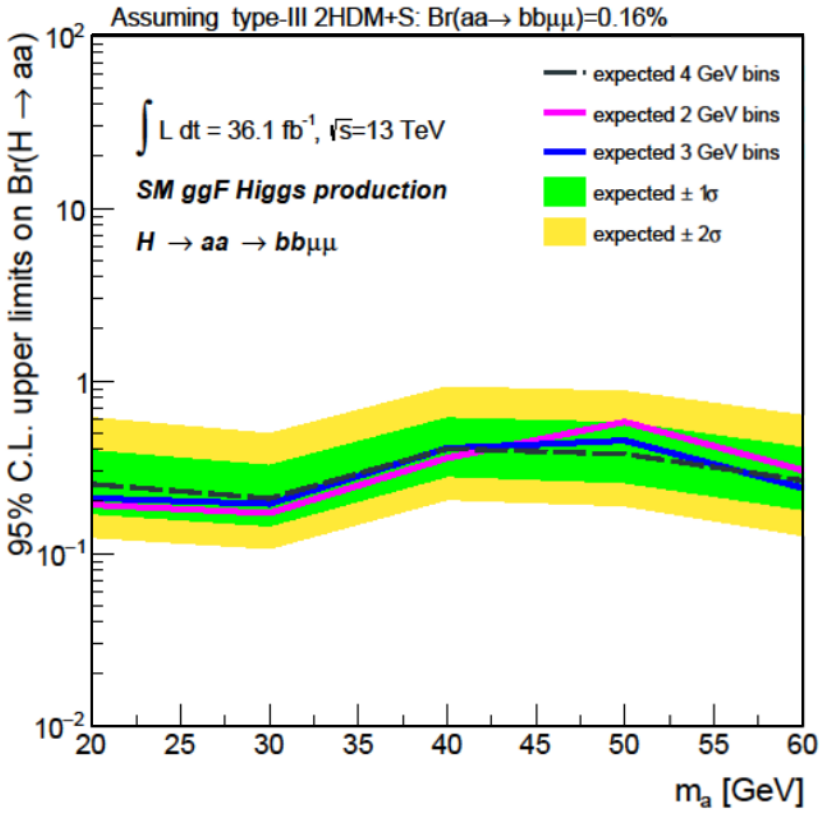

Figure 56: The resulting limits from the background only fit of the $H \to aa \to bb\mu\mu$ analysis with varying dimuon mass windows as the final selection criteria. The mass window widths of 2, 3, and 4 GeV are compared for the 5 generated signal mass points. The limits are presented in terms of the Type-III 2HDM+S, tan $\beta = 2$ model, where Br(aa→bb$\mu\mu$)=1.6×10^{-3}. Here, only the gluon-fusion (ggF) signal samples are used to calculate the upper limits.

	TCR	SB	VR1	VR2
Observed events	321	495	896	83
Fitted bkg events	320.99 ± 24.41	494.99 ± 30.50	907.15 ± 114.62	87.42 ± 27.62
Fitted ttbar events	305.28 ± 23.18	274.67 ± 19.66	701.44 ± 112.39	50.84 ± 27.20
Fitted DDZ events	5.37 ± 3.66	207.37 ± 22.48	180.36 ± 21.30	33.31 ± 4.48
Fitted db events	0.53 ± 0.37	1.74 ± 0.40	2.77 ± 0.82	0.49 ± 0.22
Fitted st events	5.63 ± 0.35	7.33 ± 0.37	18.62 ± 0.94	2.71 ± 1.91
Fitted ttV events	0.71 ± 0.10	0.67 ± 0.09	1.23 ± 0.16	0.07 ± 0.05
Fitted Wjets events	$3.47^{+4.78}_{-3.47}$	$3.22^{+4.62}_{-3.22}$	$2.73^{+3.61}_{-2.73}$	0.00 ± 0.00
MC exp. SM events	728.75	18698.53	16640.42	3009.36
MC exp. ttbar events	240.45	216.39	552.00	39.94
MC exp. DDZ events	477.98	18469.20	16063.09	2966.16
MC exp. db events	0.53	1.74	2.77	0.49
MC exp. st events	5.63	7.33	18.62	2.71
MC exp. ttV events	0.71	0.66	1.23	0.07
MC exp. Wjets events	3.46	3.21	2.72	0.00

Table 22: Results of the background only fit, with the yields in the TCR, DYCR and DYVR and VR determined using only the TCR and DYCR yields as fit inputs. The uncertainties shown are full uncertainties, with statistical and systematic uncertainties.

Normalization Parameter	Value
$\mu_{t\bar{t}}$	1.27 ± 0.06
μ_{DY}	0.011 ± 0.001

Table 23: The normalization factors for the $t\bar{t}$ and Drell–Yan backgrounds, extracted from the background only fit. The Drell–Yan background normalization is small as the normalization is derived from the template yields and the Drell–Yan + light jets cross section is much larger than the cross section for Drell–Yan + b jets. When using the MC based Drell-Yan sample, the corresponding normalization is found to be $\mu_{\mathrm{DY}} = 0.97 \pm 0.20$. This factor was determined by removing events with large, negative weight from the final analysis selections.

5.5 Validation Regions

Prior to unblinding the signal region, two regions kinematically similar to the signal region are chosen to test the modeling of the background fits. These regions are defined as follows

- High-mass validation region: A 4-object mass of $170 < m_{bb\mu\mu} < 300$ GeV is constructed in order to test both the Drell–Yan and $t\bar{t}$ yields outside of each respective control region. As the background composition in this region is approximately equivalent between the two backgrounds, it provides a test of both major background normalizations. This validation region is denoted as VR1.

- Log Likelihood Validation region: In order to test the modeling across the kinematic likelihood, a validation region in the Higgs mass window but with a $-11 < \mathrm{Ln}(L^{\mathrm{max}}) < -8$ is chosen. This validation region is denoted as VR2.

The validation regions are explicitly not included as part of the fit and therefore do not constrain the background predictions; the values of $\mu_{t\bar{t}}$ and μ_{DY} derived in the background-only fit are simply applied to the MC simulation prediction and DYTR yield under the new kinematic selections. The modeling of these validation regions under the background-only hypothesis is shown for the kinematic variables of interest in Figure 57 and the predictions are found to be in good agreement with the data.

5.6 Results

5.6.1 Model Dependent Results

The background normalization and signal contributions are determined by a series of profile likelihood fits which determine the compatibility between the observed and predicted yields for each m_a a series of test masses.

The fits are run using the statistical implementation of the methods described above using the HistFitter package, which takes as inputs the predicted and observed total yields in the TCR and DYCR and a mass window centered at a target m_a. A total of 36 profile likelihood fits are applied to the data in order to test the different m_a hypotheses in steps half the size of the optimized dimuon mass window width. The fit parameters are an overall top background correction factor$\mu_{t\bar{t}}$, an overall DY normalization factor μ_{DY} and the signal strength μ defined as the signal yield determined by the fit divided by the predicted yield. The center of a given window

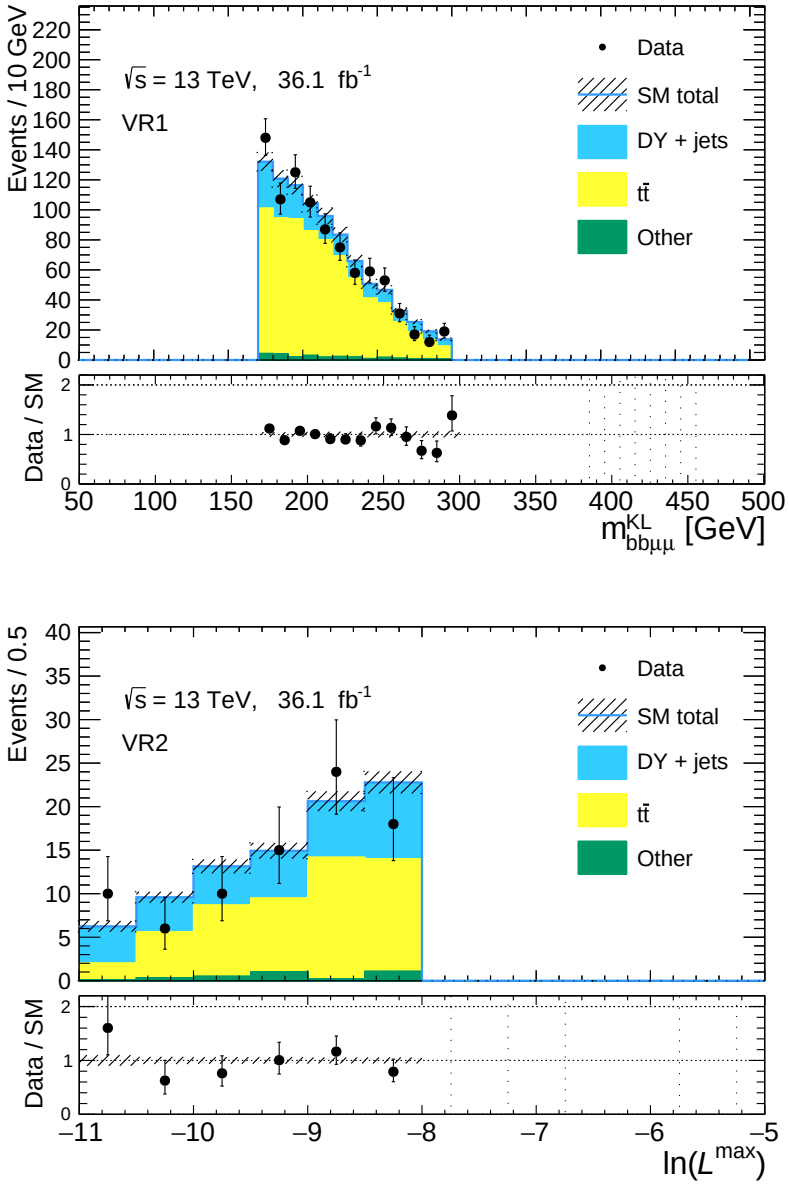

Figure 57: (Top) The distribution of $m_{bb\mu\mu}$ for the validation region defined by the selection 170 GeV $< m_{bb\mu\mu} <$ 300 GeV, denoted here as a validation region (VR1). (Bottom) The distribution of the KLFit log-likelihood for the validation region defined by the selection $-11 < \mathrm{Ln}(L^{\mathrm{max}}) < -8$, denoted here as an additional validation region (VR2).

$m_{\mu\mu}$ bin [GeV]	[19–21]	[29–31]	[39–41]	[48–52]	[58–62]
Observed events	6	6	16	48	29
Total background	4.84 ± 0.97	7.8 ± 1.2	13.7 ± 2.2	37.9 ± 5.1	30.8 ± 4.2
$t\bar{t}$	0.96 ± 0.29	3.08 ± 0.74	6.6 ± 1.5	18.1 ± 4.3	14.8 ± 3.3
DY	3.88 ± 0.92	4.5 ± 1.1	7.1 ± 1.7	19.0 ± 4.5	15.5 ± 3.6
Diboson	< 0.01	< 0.01	$0.02^{+0.04}_{-0.02}$	0.26 ± 0.16	0.3 ± 0.1
Single top	< 0.01	0.2 ± 0.2	< 0.01	$0.65^{+0.97}_{-0.65}$	$0.09^{+0.19}_{-0.09}$
$t\bar{t}V$	< 0.01	< 0.01	< 0.01	$0.01^{+0.02}_{-0.01}$	0.05 ± 0.03

Table 24: Background and observed data yields in the five simulated mass m_a signal regions. The background yields are extracted from the background-only fit. The yield uncertainties include all systematic and statistical effects. As the W+jets contribution in these signal regions is found to be negligible, it is not given in the table.

corresponds to a given target m_a value. The experimental and theoretical systematic uncertainties are incorporated into each fit as nuisance parameters.

The model dependent fit is performed using the predicted yields from the simulated MC samples generated for m_a = 20, 30, 40, 50, and 60 GeV. As the signal is being explicitly fit, the background strengths $\mu_{t\bar{t}}, \mu_{\mathrm{DY}}$ are determined with a specific signal strength for each mass window μ_{sig}. The unblinded dimuon mass distribution and other kinematic variables in the SR (or superset) can be found in Figures 58-59; the yields for each SR centered at the simulated m_a values are given in Table 24. An example breakdown of the uncertainties affecting the background yields in the $m_a = 50$ GeV signal region is given in Table 25 and a comparison of the signal and $t\bar{t}$ uncertainties in same region are given in Table 26. The uncertainties and yields for the other simulated m_a values can be found in Appendix D.

The limits from this fit are interpreted as limits on the branching ratios of the 125 GeV Higgs to the $aa \to bb\mu\mu$ final state. These limits, however, would only cover 5 mass points in the wide $m_{\mu\mu}$ spectrum for which the analysis selections have sensitivity. To provide full coverage of the mass range, an interpolation technique is used to predict the signal yields.

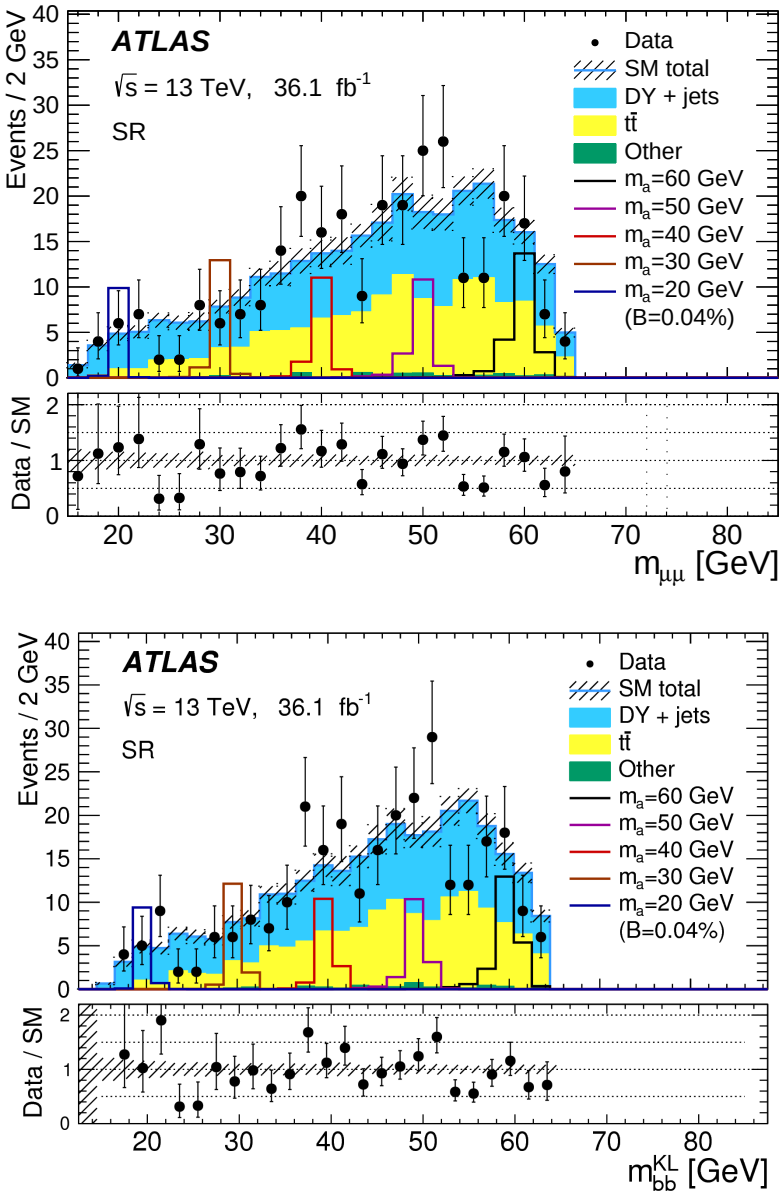

Figure 58: The dimuon mass spectrum in the inclusive signal region of the $H \to aa \to bb\mu\mu$ analysis (top) and the KL fit di-bjet mass spectrum (bottom) after the exclusion fits. The signal contribution for $m_a = 20 - 60$ are normalized to 0.04% of all Higgs decays for illustration purposes. No significant deviations from the SM expectations are found.

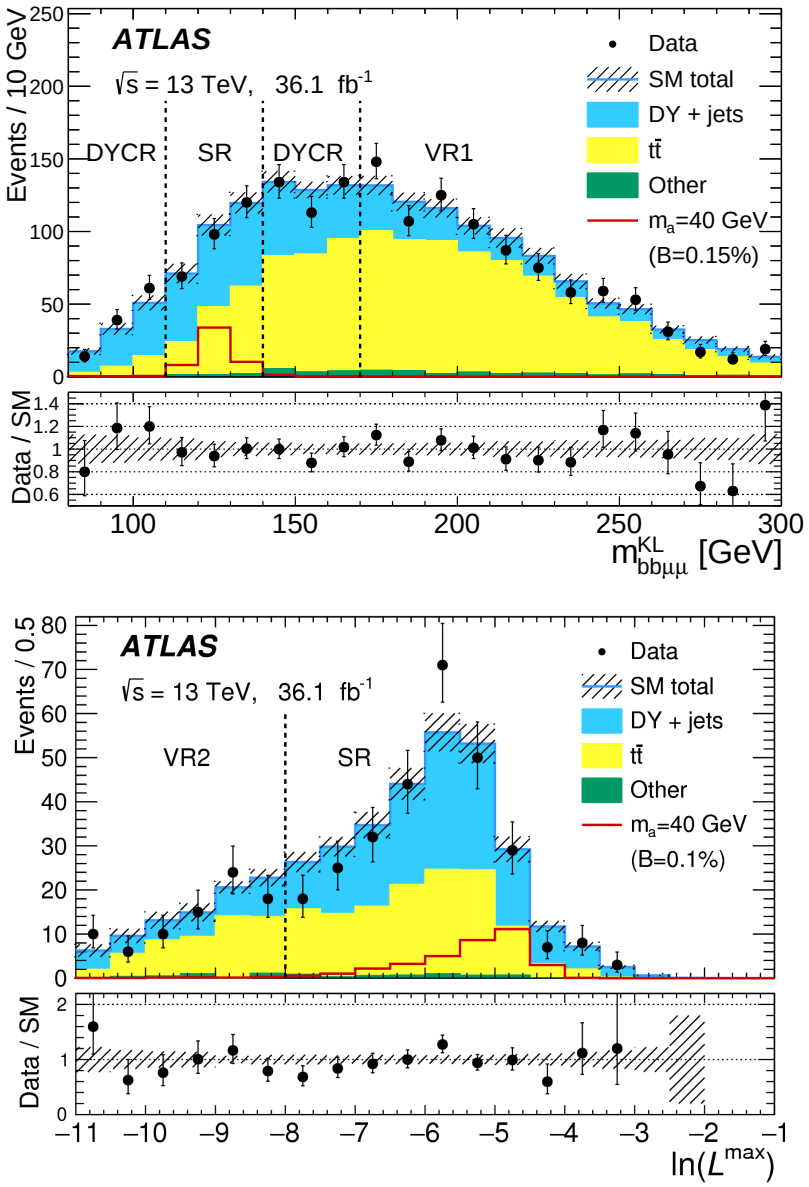

Figure 59: The $M_{bb\mu\mu}$ spectrum showing the unblinded signal region, control regions, and validation regions in the $H \to aa \to bb\mu\mu$ analysis (top) and the KLFit maximum log-likelihood (bottom) in the signal and VR2 regions after the exclusion fits are performed. The signal contribution for $m_a = 40$ GeV is normalized to 0.15% and 0.1% of all Higgs decays for illustration purposes in the figures. No significant deviations from the SM expectations are found.

Uncertainty of channel	SR50_DDZ	SR50_ttbar	SR50_db	SR50_st	SR50_ttV
Total background expectation	18.94	18.03	0.26	0.64	0.01
Total statistical ($\sqrt{N_{\text{exp}}}$)	±4.35	±4.25	±0.51	±0.80	±0.10
Total background systematic	±4.54 [24.00%]	±5.02 [27.82%]	±0.16 [61.46%]	±0.97 [150.52%]	±0.02 [233.60%]
mu_DDZ	±3.56 [18.8%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_DDZ_shape	±2.65 [14.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_DDZ_reweigting	±0.80 [4.2%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
gamma_stat_SR50_cuts_bin_0	±0.47 [2.5%]	±0.45 [2.5%]	±0.01 [2.5%]	±0.02 [2.5%]	±0.00 [2.5%]
alpha_DDZ_subtraction	±0.24 [1.3%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MET_ResPara	±0.00 [0.00%]	±0.05 [0.30%]	±0.06 [23.0%]	±0.00 [0.00%]	±0.00 [0.00%]
mu_Top	±0.00 [0.00%]	±1.62 [9.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_ISO_SYS	±0.00 [0.00%]	±0.00 [0.01%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [6.8%]
alpha_MU_EFF_STAT	±0.00 [0.00%]	±0.00 [0.01%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.15%]
alpha_MU_ID	±0.00 [0.00%]	±0.42 [2.3%]	±0.02 [7.0%]	±0.08 [11.9%]	±0.01 [57.2%]
alpha_MU_MS	±0.00 [0.00%]	±0.05 [0.26%]	±0.03 [11.3%]	±0.00 [0.00%]	±0.00 [4.9%]
alpha_top_PDF	±0.00 [0.00%]	±0.45 [2.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_Rad	±0.00 [0.00%]	±0.90 [5.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_db_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.03 [10.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_trigSF_Stat	±0.00 [0.00%]	±0.01 [0.03%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_ttV_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [13.0%]
alpha_MET_ResPerp	±0.00 [0.00%]	±0.23 [1.3%]	±0.05 [18.7%]	±0.00 [0.00%]	±0.00 [35.7%]
alpha_MU_SagRho	±0.00 [0.00%]	±0.13 [0.71%]	±0.00 [0.00%]	±0.61 [95.5%]	±0.00 [0.05%]
alpha_MU_TTVA_SYS	±0.00 [0.00%]	±0.03 [0.19%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [1.3%]
alpha_btag_B0	±0.00 [0.00%]	±0.39 [2.1%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [12.5%]
alpha_st_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.03 [5.0%]	±0.00 [0.00%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±3.25 [18.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_pythHerw	±0.00 [0.00%]	±2.89 [16.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_SagRes	±0.00 [0.00%]	±0.04 [0.23%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [9.9%]
alpha_MU_Scale	±0.00 [0.00%]	±0.34 [1.9%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.01 [58.4%]
alpha_JES_SR_NP4	±0.00 [0.00%]	±0.10 [0.58%]	±0.00 [0.00%]	±0.40 [61.8%]	±0.00 [26.7%]
alpha_MET_Scale	±0.00 [0.00%]	±0.20 [1.1%]	±0.00 [1.3%]	±0.00 [0.00%]	±0.00 [0.97%]
alpha_JES_SR_NP2	±0.00 [0.00%]	±0.23 [1.3%]	±0.00 [0.00%]	±0.33 [51.8%]	±0.00 [2.2%]
alpha_JES_SR_NP3	±0.00 [0.00%]	±0.21 [1.2%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.69%]
alpha_JES_SR_NP1	±0.00 [0.00%]	±0.77 [4.3%]	±0.05 [19.9%]	±0.34 [53.6%]	±0.00 [16.8%]
alpha_PRW	±0.00 [0.00%]	±0.00 [0.01%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.01 [60.9%]
alpha_btag_L3	±0.00 [0.00%]	±0.04 [0.23%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [10.6%]
alpha_JER	±0.00 [0.00%]	±0.40 [2.2%]	±0.12 [47.2%]	±0.41 [63.2%]	±0.02 [203.4%]

Table 25: Breakdown of the systematic uncertainties on the background estimates in the $m_a = 50$ GeV signal region. All uncertainties with less than 1% effect on any background yield have been truncated. These uncertainties can be correlated; therefore, the quadratic summation of the individual uncertainties may not describe the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

Uncertainty of channel	SR50_a50	SR50_ttbar
Total background expectation	10.12	18.03
Total statistical ($\sqrt{N_{\text{exp}}}$)	±3.18	±4.25
Total background systematic	±9.47 [93.59%]	±5.02 [27.82%]
mu_SIG	±9.81 [97.0%]	±0.00 [0.00%]
alpha_btag_B0	±1.75 [17.3%]	±0.39 [2.1%]
alpha_JER	±0.63 [6.2%]	±0.40 [2.2%]
alpha_JES_SR_NP1	±0.63 [6.2%]	±0.77 [4.3%]
alpha_sig_RenFact	±0.61 [6.0%]	±0.00 [0.00%]
alpha_sig_ISR	±0.40 [4.0%]	±0.00 [0.00%]
alpha_sig_ggFXsec	±0.36 [3.6%]	±0.00 [0.00%]
alpha_sig_VH	±0.35 [3.5%]	±0.00 [0.00%]
alpha_sig_PDFalphaS	±0.30 [3.0%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.25 [2.5%]	±0.00 [0.00%]
gamma_stat_SR50_cuts_bin_0	±0.25 [2.5%]	±0.45 [2.5%]
alpha_JES_SR_NP2	±0.12 [1.1%]	±0.23 [1.3%]
alpha_JES_SR_NP3	±0.04 [0.41%]	±0.21 [1.2%]
alpha_MU_Scale	±0.03 [0.34%]	±0.34 [1.9%]
alpha_MET_Scale	±0.03 [0.33%]	±0.20 [1.1%]
alpha_MU_ID	±0.03 [0.31%]	±0.42 [2.3%]
alpha_MET_ResPerp	±0.01 [0.07%]	±0.23 [1.3%]
mu_Top	±0.00 [0.00%]	±1.62 [9.0%]
alpha_top_PDF	±0.00 [0.00%]	±0.45 [2.5%]
alpha_top_Rad	±0.00 [0.00%]	±0.90 [5.0%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±3.25 [18.0%]
alpha_top_pythHerw	±0.00 [0.00%]	±2.89 [16.0%]

Table 26: Breakdown of the systematic uncertainties on the signal and $t\bar{t}$ background estimates in the $m_a = 50$ GeV signal region. All uncertainties with less than 1% effect on any background yield have been truncated. These uncertainties can be correlated; therefore, the quadratic summation of the individual uncertainties may not describe the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

5.6.2 Model Dependent Interpolation

To obtain model dependent limits for the mass range 18 - 62 GeV, the signal acceptance values must be determined for masses that were not explicitly generated in MC samples. These acceptance values are obtained via an interpolation fit of the signal acceptance between the simulated mass points and an extrapolation fit at the endpoints. Figure 60 illustrates the signal acceptance after all SR selection cuts, including the variable dimuon mass window cut, normalized assuming Br($H \to aa$)=100% and Br(aa→bb$\mu\mu$)=0.16%. For the truth samples, the SR cuts are slightly modified to emulate the log-likelihood fit of the dijet masses, as these samples are explicitly using truth level quantities and do not need to undergo a kinematic fit. These modifications, therefore, provide a cross check of the reconstruction level acceptances. As can be seen in Figure 60, the shapes between the truth level and reconstructed acceptances show similar trends and the yields differ due to lepton efficiencies and reconstruction inefficiencies not being fully propagated to the truth sample.

To provide a cross check for the initial interpolation fits two additional signal samples were generated at the truth level $m_a = 25, 55$ GeV. These two masses were chosen as these points are located in the region where the acceptance behavior changes rapidly. The acceptance of these two points at the truth level was then scaled to produce a reconstruction level acceptance depending on the value relative to the neighboring point. Explicitly, the acceptance at $m_a = 25$ GeV at the truth level is similar to the value at $m_a = 30$ GeV and was therefore assumed to have the same yield; the acceptance at $m_a = 55$ GeV was placed at $\frac{1}{3}$ of the way between m_a =50 GeV and $m_a = 60$ GeV. A conservative 10% uncertainty on these yield values is applied before the interpolation fit.

After these intermediate mass points are assigned a reconstruction yield, the signal acceptance is interpolated from 20 - 60 GeV by fitting these seven points using cubic splines. The spline fits are performed assuming yields of 2, 3, and 4 GeV dimuon mass windows in order to test the stability of the fit predictions. Independent spline fits are used for the gluon-fusion and VBF production modes. All of the fits combinations for production mode and dimuon mass width are shown in Figure 61.

The signal yield is then extracted from the spline fits corresponding to the $m_{\mu\mu}$ windows size used in the specified mass range and the yields from the gluon-fusion and VBF productions are summed. As an example, the yields at $m_a = 27$ GeV are taken as the sum of the values from the gluon-fusion and VBF splines using a 2 GeV window size. The final yield values as a function of m_a are given in Table 27.

The interpolation fits introduce a new systematic uncertainty, as the yields can differ depending on the exact function or spline used to determine the yields. A conservative interpolation uncertainty is assigned by taking the largest deviation

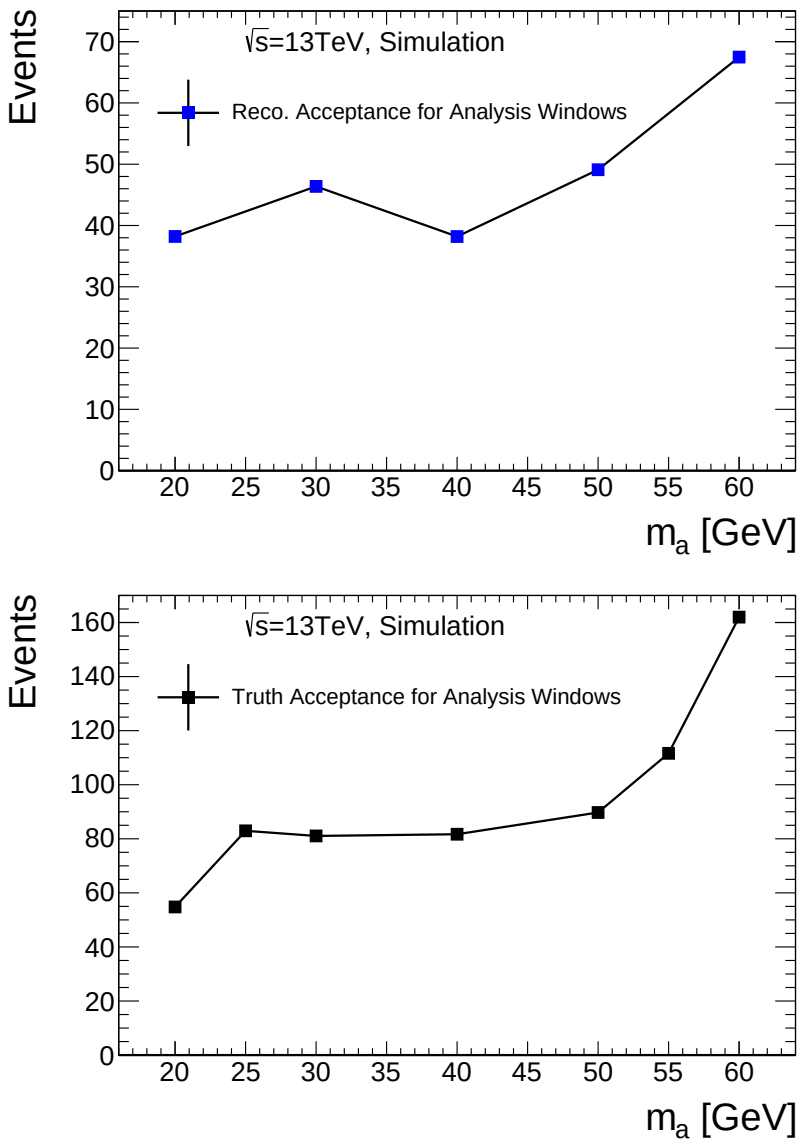

Figure 60: The signal yields after all analysis selection cuts and assuming $\mathrm{Br}(H \to aa) = 100\%$ and $\mathrm{Br}(aa \to bb\mu\mu) = 0.16\%$. at the reconstruction level (top) and the truth level (bottom) with modified constraints to emulate the KL Fit. The truth level yields are higher than the reconstruction level yields as the reconstruction efficiency and trigger scale factors are not taken into account for this plot.

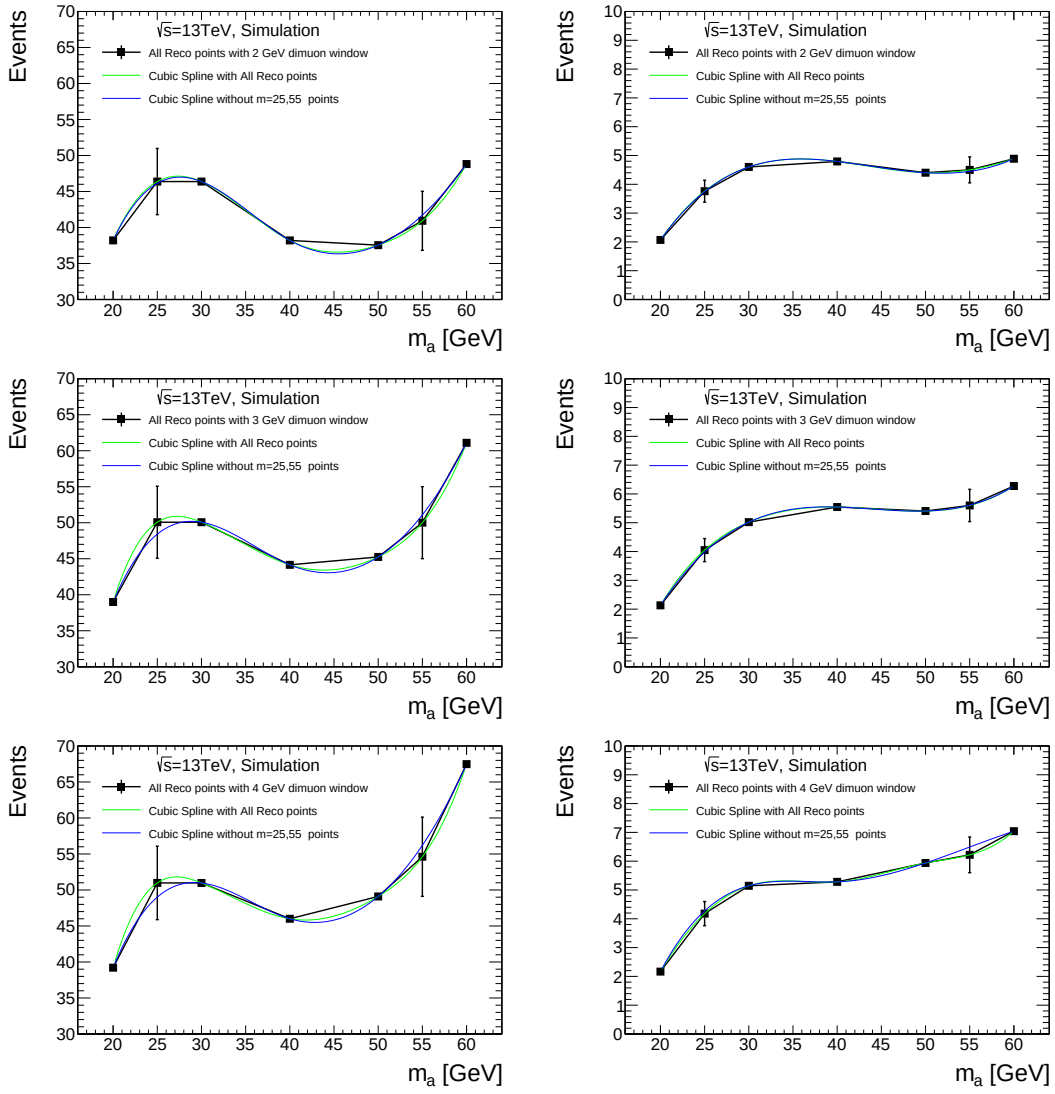

Figure 61: The signal acceptance in the gluon fusion Higgs (left) and VBF (right) production mode after all selection cuts with $m_{\mu\mu}$ window size = 2 GeV (top) = 3 GeV (center) and = 4 GeV (bottom), assuming Br($H \to aa$) = 100% and Br($aa \to bb\mu\mu$) = 0.16%. Two cubic spline and a linear interpolation fits are shown. The cubic splines differ in the inclusion or exclusion of the estimated m_a = 25, 55 GeV acceptances. The 10% errors on these points are for illustration purposes only.

m_a	Gluon-Fusion Yield	VBF Yield
20	38.2	2.1
21	40.7	2.5
22	42.7	2.9
23	44.3	3.2
24	45.5	3.5
25	46.4	3.8
26	46.9	4.0
27	47.1	4.2
28	47.1	4.3
29	46.8	4.5
30	46.4	4.6
31	45.8	4.7
32	45.1	4.8
33	44.2	4.8
34	43.3	4.9
35	42.4	4.9
36	41.5	4.9
37	40.6	4.9
38	39.7	4.9
39	38.9	4.8
40	38.2	4.8
41.5	43.7	5.5
43	43.5	5.5
44.5	43.4	5.5
46	43.6	5.4
47.5	44.1	5.4
50	49.1	5.9
52	50.8	6.0
54	53.1	6.1
56	56.4	6.3
58	61.1	6.6
60	67.5	7.0

Table 27: The expected yields for the $H \to aa \to bb\mu\mu$ model-dependent limits. These yields are extracted from the spline fits in 2, 3, and 4 GeV mass windows and are separated between the gluon-fusion and vector boson fusion simulated sample splines.

Production Mode	Spline $m_{\mu\mu}$ Width	Uncertainty [%]
gluon-fusion	2 GeV	3.0
gluon-fusion	3 GeV	2.7
gluon-fusion	4 GeV	2.0
VBF	2 GeV	4.4
VBF	3 GeV	0.5
VBF	4 GeV	1.8
Inclusive	2 GeV	2.8
Inclusive	3 GeV	2.5
Inclusive	4 GeV	1.9

Table 28: The values of interpolation uncertainty determined for each dimuon mass window spline fit. The uncertainty is taken to be the largest deviation between the cubic spline and a linear interpolation fit. The inclusive uncertainty is derived from the cross-section weighted individual yields added in quadrature. The final uncertainty for the entire mass range is taken to be a conservative 3%.

between the cubic spline fit yield and the yield obtained from a linear interpolation. The deviations from the gluon-fusion and VBF splines are added in quadrature after applying the appropriate cross section weights. The deviation in yields for each fit are given in Table 28. As the inclusive uncertainties for the 2, 3, and 4 GeV splines are found to be similar, a 3% uncertainty is assigned to the full mass range. Using the interpolated signal acceptance, $m_{\mu\mu}$ is scanned in a similar way to the method used for the model-independent case.

For the intermediate masses without full MC simulation, the non-experimental systematic uncertainties are also propagated into the fits. The theoretical uncertainties for the ggH and VBF signals are determined using the same methodology that can be found in Section 4.7.7. These uncertainties include the FSR/ISR uncertainties, the cross section uncertainties, the parton shower uncertainties, and the VH production mode weighting.

Finally, a systematic uncertainty of 19% is assigned for each interpolated signal mass points to account for the average size of the experimental systematic uncertainties on the signal yield (see Tables 43-46). As a cross-check of the method, this simplified systematic setup is applied to the five simulated points. The comparison of the limits between the full and the simplified setup is shown in Figure 62. The limits between the two setups agree within the 1σ uncertainties. The breakdown of the systematic uncertainties using the simplified setup for $m_a = 50$ GeV is show in

Table 29 while an example of the uncertainty breakdown for the interpolated point $m_a = 52$ GeV is given in Table 30. Using this systematic setup, the resulting limits on the cross section times branching ratio $\frac{\sigma_H}{\sigma_{SM}} \times \mathcal{B}(H \to aa \to bb\mu\mu)$ are shown in Figure 63.

5.6.3 Model Independent Results

The model independent scan is performed by removing the Higgs signal inputs and corresponding uncertainties from the model dependent fit; instead, the scans simply look for the number of new physics events allowed after the background fit as described in Section 5.2. The scan predicts the number of new physics events in 36 independent profile likelihood fits, in bins half the size of the $m_{\mu\mu}$ window size used in the model dependent fit. The scan profiles in 1, 1.5, or 2 GeV bins for the dimuon mass widths in the range $20 \leq m_{\mu\mu} \leq 40$ GeV, $40 < m_{\mu\mu} < 50$, and $50 \leq m_{\mu\mu} \leq 60$, analogous to the model-dependent fit. The yields extract from these $m_{\mu\mu}$ bins are then translated into a 95% CL upper limit on the visible production cross-section for new physics times branching ratio into the $bb\mu\mu$ final state, $\sigma_{\text{vis}}(X) \times \mathcal{B}(X \to bb\mu\mu)$. These limits are slightly biased on model choice, as the inclusion of the the KL fit constraint on $m_{bb} \sim m_{\mu\mu}$ requires the new physics to have equal mass decays and the four-object invariant mass constraint means the new resonance X must have a mass in the region 110 - 140 GeV. In this version of the fit, the most significant deviation from the background only hypothesis occurs in the range $37 < m_{\mu\mu} < 39$ GeV and the deviation corresponds to a local significance of 1.57σ. The final results for the expected and observed limits in the model independent interpretation are shown in Figure 64 and in Table 31. The final cross section limits range from 0.1 fb to 0.73 fb, depending on the dimuon mass window.

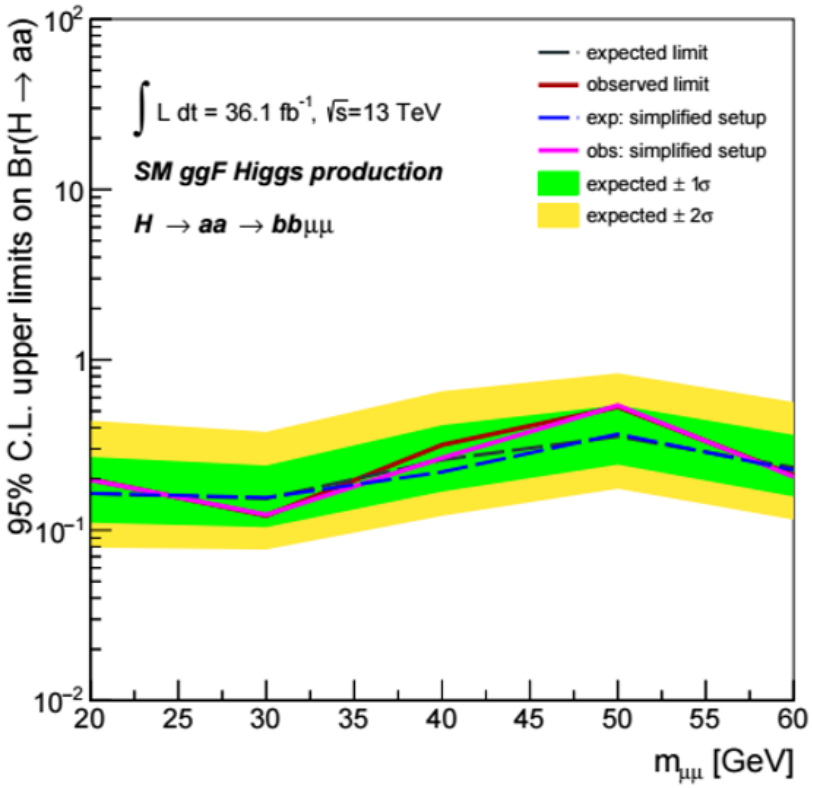

Figure 62: Expected and observed upper limits on $\mathcal{B}(H \to aa$, assuming $\mathcal{B}(aa \to bb\mu\mu)$=0.16% from a type-III 2HDM+S model with tan $\beta = 2$. The limit calculation is performed for the five MC simulated $m_a = 20 - 60$ GeV signal samples. The black (red) lines show the expected (observed) limits with all systematics from Table 21 included in the fit, while the blue (magenta) lines show the expected (observed) limits assuming a simplified experimental uncertainty setup applied to the five generated masses. The one and two sigma bands correspond to the uncertainties on the full systematic setup.

Uncertainty of channel	SR50_a50	SR50_ttbar
Total background expectation	10.16	18.15
Total statistical ($\sqrt{N_{\text{exp}}}$)	±3.19	±4.26
Total background systematic	±9.45 [93.05%]	±5.06 [27.88%]
mu_SIG	±9.72 [95.7%]	±0.00 [0.00%]
alpha_sig_exp	±1.93 [19.0%]	±0.00 [0.00%]
alpha_sig_RenFact	±0.61 [6.0%]	±0.00 [0.00%]
alpha_sig_ISR	±0.41 [4.0%]	±0.00 [0.00%]
alpha_sig_Xsec	±0.38 [3.7%]	±0.00 [0.00%]
alpha_sig_VH	±0.36 [3.5%]	±0.00 [0.00%]
alpha_sig_interp	±0.28 [2.8%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.25 [2.5%]	±0.00 [0.00%]
alpha_sig_PDF	±0.03 [0.30%]	±0.00 [0.00%]
mu_Top	±0.00 [0.00%]	±1.63 [9.0%]
alpha_MU_ID	±0.00 [0.00%]	±0.43 [2.4%]
alpha_top_PDF	±0.00 [0.00%]	±0.45 [2.5%]
alpha_top_Rad	±0.00 [0.00%]	±0.91 [5.0%]
alpha_MET_ResPerp	±0.00 [0.00%]	±0.23 [1.3%]
alpha_btag_B0	±0.00 [0.00%]	±0.39 [2.1%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±3.27 [18.0%]
alpha_top_pythHerw	±0.00 [0.00%]	±2.90 [16.0%]
alpha_MU_Scale	±0.00 [0.00%]	±0.34 [1.9%]
alpha_MET_Scale	±0.00 [0.00%]	±0.21 [1.1%]
alpha_JES_SR_NP2	±0.00 [0.00%]	±0.23 [1.3%]
alpha_JES_SR_NP3	±0.00 [0.00%]	±0.22 [1.2%]
alpha_JES_SR_NP1	±0.00 [0.00%]	±0.78 [4.3%]
gamma_stat_SR50_cuts_bin_0	±0.00 [0.00%]	±0.47 [2.6%]
alpha_JER	±0.00 [0.00%]	±0.41 [2.2%]

Table 29: Breakdown of the model-dependent fit systematic uncertainties and their effect on the signal and background estimates for the $m_a = 50$ GeV using a simplified signal uncertainty and the $t\bar{t}$ background using the full systematic setup. These uncertainties can be correlated; therefore, the quadratic summation of the individual uncertainties may not describe the total background uncertainty. Each percentage shows the size of the uncertainty relative to the total expected event yield of the signal or $t\bar{t}$ background. Uncertainties that are below 1% have been removed from the table.

Uncertainty of channel	SR52_a52	SR52_ttbar
Total background expectation	2.89	17.44
Total statistical ($\sqrt{N_{\text{exp}}}$)	±1.70	±4.18
Total background systematic	±4.50 [155.79%]	±5.17 [29.65%]
mu_SIG	±4.48 [155.2%]	±0.00 [0.00%]
alpha_sig_exp	±0.55 [19.0%]	±0.00 [0.00%]
alpha_sig_RenFact	±0.17 [6.0%]	±0.00 [0.00%]
alpha_sig_ISR	±0.12 [4.0%]	±0.00 [0.00%]
alpha_sig_Xsec	±0.11 [3.7%]	±0.00 [0.00%]
alpha_sig_VH	±0.10 [3.5%]	±0.00 [0.00%]
alpha_sig_interp	±0.08 [2.8%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.07 [2.5%]	±0.00 [0.00%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±3.14 [18.0%]
alpha_top_pythHerw	±0.00 [0.00%]	±2.79 [16.0%]
mu_Top	±0.00 [0.00%]	±1.57 [9.0%]
alpha_JES_SR_NP1	±0.00 [0.00%]	±1.44 [8.3%]
alpha_top_Rad	±0.00 [0.00%]	±0.87 [5.0%]
alpha_JES_SR_NP3	±0.00 [0.00%]	±0.68 [3.9%]
alpha_JER	±0.00 [0.00%]	±0.67 [3.9%]
alpha_MET_ResPerp	±0.00 [0.00%]	±0.50 [2.9%]
alpha_top_PDF	±0.00 [0.00%]	±0.44 [2.5%]
gamma_stat_SR52_cuts_bin_0	±0.00 [0.00%]	±0.42 [2.4%]
alpha_MET_ResPara	±0.00 [0.00%]	±0.37 [2.1%]
alpha_btag_B0	±0.00 [0.00%]	±0.30 [1.7%]
alpha_JES_SR_NP2	±0.00 [0.00%]	±0.27 [1.5%]
alpha_MET_Scale	±0.00 [0.00%]	±0.22 [1.3%]
alpha_MU_SagRho	±0.00 [0.00%]	±0.22 [1.3%]
alpha_JES_SR_NP4	±0.00 [0.00%]	±0.20 [1.1%]

Table 30: Breakdown of the model-dependent fit systematic uncertainties and their effect on the signal and background estimates for the interpolated $m_a = 52$ GeV using a simplified signal uncertainty and the $t\bar{t}$ background using the full systematic setup. These uncertainties can be correlated; therefore, the quadratic summation of the individual uncertainties may not describe the total background uncertainty. Each percentage shows the size of the uncertainty relative to the total expected event yield of the signal or $t\bar{t}$ background. Uncertainties that are below 1% have been removed from the table.

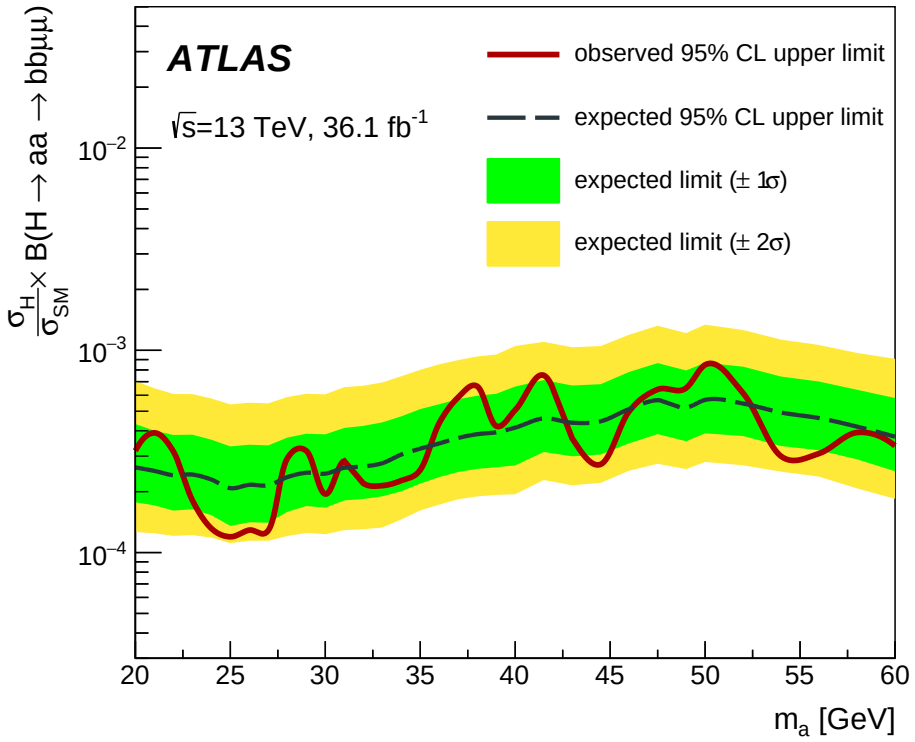

Figure 63: The observed and expected upper limits, presented at the 95% confidence level, on the branching ratio $\mathcal{B}(H \to aa \to bb\mu\mu)$ for $20 < m_a < 60$ GeV. These limits assume the production cross section, σ_H, is given by the SM predictions for the 125 GeV Higgs boson production cross-section in the gluon-fusion, VBF and VH production modes.

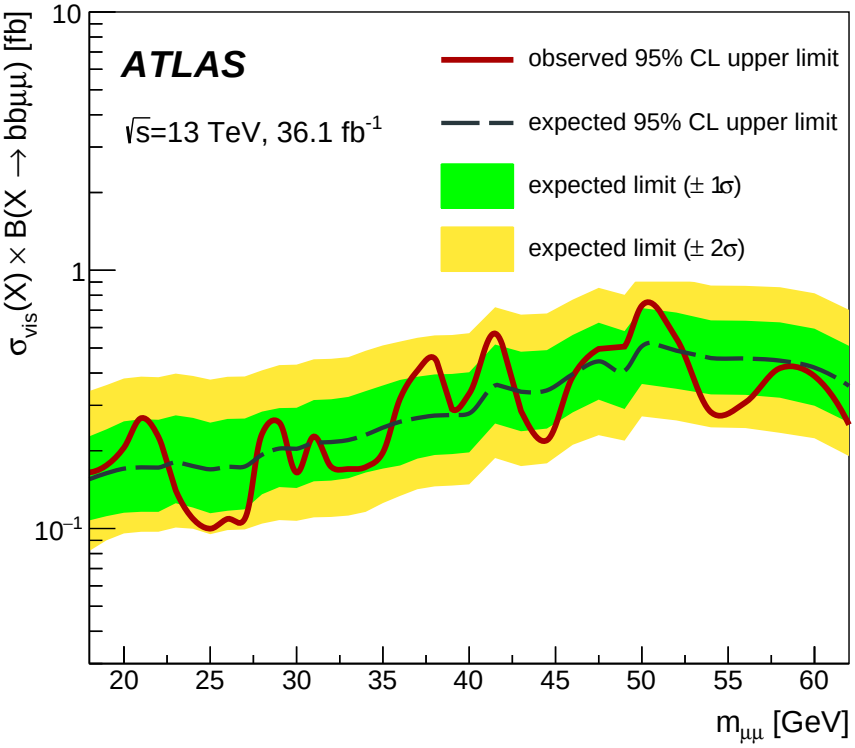

Figure 64: The observed and expected upper limits, presented at the 95% confidence level, on the branching ratio $\mathcal{B}(X \to bb\mu\mu)$ for $18 < m_{\mu\mu} < 62$ GeV from a model independent approach. The limits include all analysis selection, including the $m_{bb} \sim m_{\mu\mu}$ constraint and the narrow width of the dimuon mass in varying $m_{\mu\mu}$ windows.

$m_{\mu\mu}$ bin [GeV]	$\sigma_{\rm vis}(X) \times {\rm Br}(X \to bb\mu\mu)$ [fb]	$S^{\rm obs}$	$S^{\rm exp}$	p_0-value (Z)
$[17-19]$	0.16	5.9	$5.6^{+2.6}_{-1.7}$	0.41 (0.23)
$[18-20]$	0.18	6.3	$5.9^{+2.8}_{-1.9}$	0.38 (0.31)
$[19-21]$	0.21	7.4	$6.2^{+3.2}_{-2.0}$	0.32 (0.46)
$[20-22]$	0.27	9.7	$6.2^{+3.3}_{-2.0}$	0.11 (1.21)
$[21-23]$	0.23	8.2	$6.2^{+3.3}_{-2.0}$	0.23 (0.76)
$[22-24]$	0.14	5.1	$6.5^{+3.4}_{-2.0}$	0.50 (0.00)
$[23-25]$	0.11	3.9	$6.3^{+3.4}_{-1.9}$	0.50 (0.00)
$[24-26]$	0.10	3.6	$6.1^{+3.2}_{-2.0}$	0.50 (0.00)
$[25-27]$	0.11	3.9	$6.3^{+3.3}_{-2.0}$	0.50 (0.00)
$[26-28]$	0.11	3.9	$6.3^{+3.3}_{-2.0}$	0.50 (0.00)
$[27-29]$	0.23	8.4	$7.0^{+3.3}_{-2.1}$	0.27 (0.62)
$[28-30]$	0.26	9.3	$7.4^{+3.2}_{-2.1}$	0.23 (0.74)
$[29-31]$	0.16	5.9	$7.3^{+3.2}_{-2.2}$	0.50 (0.01)
$[30-32]$	0.23	8.2	$7.8^{+3.6}_{-2.3}$	0.42 (0.21)
$[31-33]$	0.17	6.3	$7.8^{+3.6}_{-2.3}$	0.50 (0.00)
$[32-34]$	0.17	6.1	$7.9^{+3.7}_{-2.3}$	0.50 (0.00)
$[33-35]$	0.17	6.2	$8.3^{+4.0}_{-2.4}$	0.50 (0.00)
$[34-36]$	0.20	7.1	$8.9^{+4.0}_{-2.7}$	0.50 (0.00)
$[35-37]$	0.32	11.5	$9.3^{+4.2}_{-3.0}$	0.27 (0.62)
$[36-38]$	0.41	14.7	$9.7^{+4.3}_{-2.9}$	0.09 (1.34)
$[37-39]$	0.46	16.4	$9.9^{+4.4}_{-2.9}$	0.06 (1.57)
$[38-40]$	0.29	10.5	$9.9^{+4.4}_{-2.9}$	0.42 (0.19)
$[39-41]$	0.33	12.0	$10.1^{+4.5}_{-3.0}$	0.31 (0.51)
$[40-43]$	0.57	20.5	$12.9^{+5.7}_{-3.7}$	0.08 (1.40)
$[41.5-44.5]$	0.29	10.3	$12.2^{+5.2}_{-3.6}$	0.50 (0.00)
$[43-46]$	0.22	7.9	$12.4^{+5.3}_{-3.6}$	0.50 (0.00)
$[44.5-47.5]$	0.39	14.0	$14.3^{+6.0}_{-4.1}$	0.50 (0.00)
$[46-49]$	0.50	17.9	$16.0^{+6.6}_{-4.6}$	0.37 (0.34)
$[47.5-50.5]$	0.51	18.3	$14.7^{+6.2}_{-4.3}$	0.27 (0.62)
$[48-52]$	0.73	26.3	$18.3^{+7.4}_{-5.2}$	0.13 (1.12)
$[50-54]$	0.54	19.6	$17.6^{+7.1}_{-5.1}$	0.37 (0.33)
$[52-56]$	0.28	10.2	$16.5^{+6.6}_{-4.6}$	0.50 (0.00)
$[54-58]$	0.31	11.1	$16.4^{+6.6}_{-4.5}$	0.50 (0.00)
$[56-60]$	0.42	15.1	$16.1^{+6.6}_{-4.5}$	0.50 (0.00)
$[58-62]$	0.39	14.0	$15.1^{+6.3}_{-4.4}$	0.50 (0.00)
$[60-64]$	0.25	9.1	$12.9^{+5.5}_{-3.6}$	0.50 (0.00)

Table 31: The results of the model-independent scan of the dimuon mass spectrum after all analysis selections. The first column gives the dimuon mass window, while the second column details the observed 95% CL upper limits on the visible cross section for new physics times branching ratio into the $bb\mu\mu$ final state, $\sigma_{\rm vis}(X) \times \mathcal{B}(X \to bb\mu\mu)$. The visible cross section is defined as the product of the production cross section and acceptance $\times$ efficiency ($\sigma_{\rm vis}(X) = \sigma_{\rm prod}(X) \times \epsilon_{\rm X}$) of any potential signal. When calculating the model-independent limits, no assumptions are made on the values of $\sigma_{\rm prod}(X)$ and $\epsilon_{\rm X}$. The third and fourth columns give the 95% CL upper limits on the observed and expected number of signal events ($S^{\rm obs}$ and $S^{\rm exp}$). The final column gives the discovery p_0-value and the significance of any signal above the background-only hypothesis (Z).

6 Analysis Improvements for the Full Run-2 Dataset

6.1 Improving b-jet Identification

Due to the low-p_{T} of the B-hadron decays in the $a \to bb$ decay, the signal acceptance from requiring two b-tagged jets with $p_T > 20$ GeV is quite low, around 10%. The signal truth jet spectra for the $m_a = 30$ GeV is shown in Figure 65, showing that the spectrum has 50% of subleading jets being below 20 GeV.This loss comes from two sources, namely the loss coming from vetoing jets with $p_{\mathrm{T}} < 20$ GeV and a moderate decrease in the b-tagging efficiency at low reconstructed p_{T}. The majority of the loss comes from the jet momentum cut, as can be seen from the truth leading and subleading truth jets in Figure 66 where the two effects are separated for a sample using b-jets from $t\bar{t}$ decays.

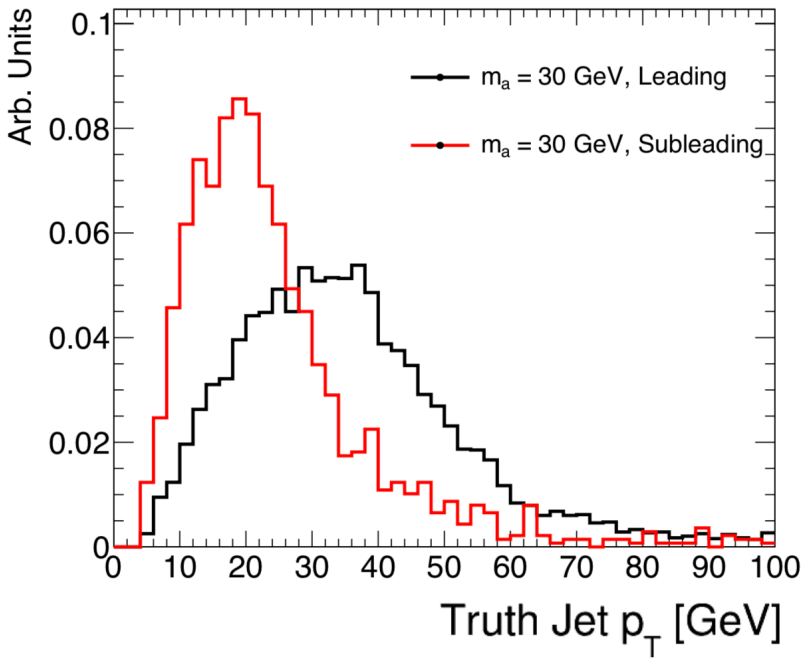

Figure 65: The leading and subleading truth b-jet p_{T} spectra for $m_a = 30$ GeV. The b-jet decays from the a are typically at a low p_{T} scale, with many jets below the nominal ATLAS reconstruction scale of 20 GeV.

Several experimental challenges must be taken into account to target lower jet p_{T} scales. First, as the jet momentum is lowered, the effect of pileup interactions becomes larger due to the natural scale of multi-jet interactions. Therefore, distinguishing signal jets coming from the hard-scatter, or primary vertex, from the background of nearby pileup vertices is crucial. Additionally, the noise term (Eq. 70) in the calorimeter resolution becomes quite significant at lower momentum scales. The response for particle flow reconstruction jets in these low p_{T} regions performs much better than the topological cluster based jets, but the noise term challenges are significant enough to warrant future studies.

Track jets (described in Section 3.6.2) can be used to improve selection efficiency and mitigate pileup contributions while targeting these low p_{T} jet regimes. As the tracks for these jets are explicitly chosen to be associated with the PV in an event, the effects of pileup are naturally constrained. As the b-tagging algorithms of ATLAS run almost exclusively on tracking information, with calorimeter information only being used to determine the jet p_{T} and η, the MV2 algorithms maintain high tagging efficiency for track jets. The reconstruction and tagging efficiency as a function of the truth jet p_{T} comparing track-jet and PFlow jet identification efficiency are shown in Figure 66 where a significant increase in final identification efficiency can be seen for truth jets with $5 < p_{\mathrm{T}} < 20$ GeV.

6.1.1 Track Jet Momentum and Energy Corrections

Track-jets can be explicitly used as b-jet objects for the purpose of constructing invariant masses in the low momentum regions where calorimeter jets cannot be reconstructed. The requirement on the underlying tracks to come from the primary vertex naturally make track-jets resistant to pileup. As an additional benefit, the uncertainty on the track energy scale is much smaller at low p_{T} compared to the calorimeter energy scale. The track scale, however, is missing the energy associated to the neutral particles within a jet. The bias from this lack of neutral energy can be seen in the m_{bb} distributions given in Figure 67. A ratio of the central values $\frac{m_{bb}(\text{particle flow})}{m_{bb}(\text{track})} \approx 1.8$ is used to apply a correction factor to the leading and subleading track-jet energies. The "corrected" di-bjet mass constructed from these jets is shown in Figure 68, where the m_{bb} moves closer to the true m_a value.be

Rather than rely on ad hoc corrections, a true track jet correction factor can be defined as the ratio between the track-jet p_{T} and the geometrically matched truth jet p_{T}.

$$\zeta = \left\langle \frac{\text{truth-jet } p_{\mathrm{T}}}{\text{track-jet } p_{\mathrm{T}}} \right\rangle_{p_{\mathrm{T}}, N_{\mathrm{trks}}} \tag{105}$$

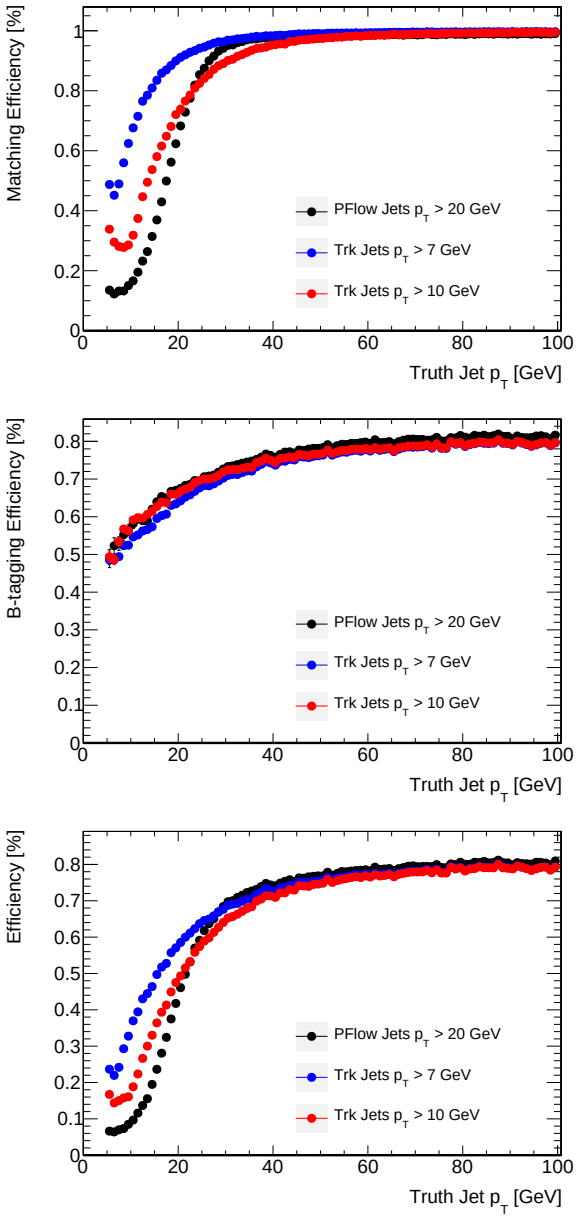

Figure 66: The reconstruction, or matching efficiency (top), b-tagging efficiency (middle), and combined reconstruction and b-tagging efficiency (bottom) for jets in simulated $t\bar{t}$ events as a function of truth jet p_{T}. The b-tagging affects the overall efficiency scale, as the jets must pass the 77% WP of the MV2c10 tagger. The reconstruction level cuts required for each curve are 20 GeV for the particle flow jets (black), 7 GeV for one track jet threshold (blue) and 10 GeV for another choice of track jet threshold (red).

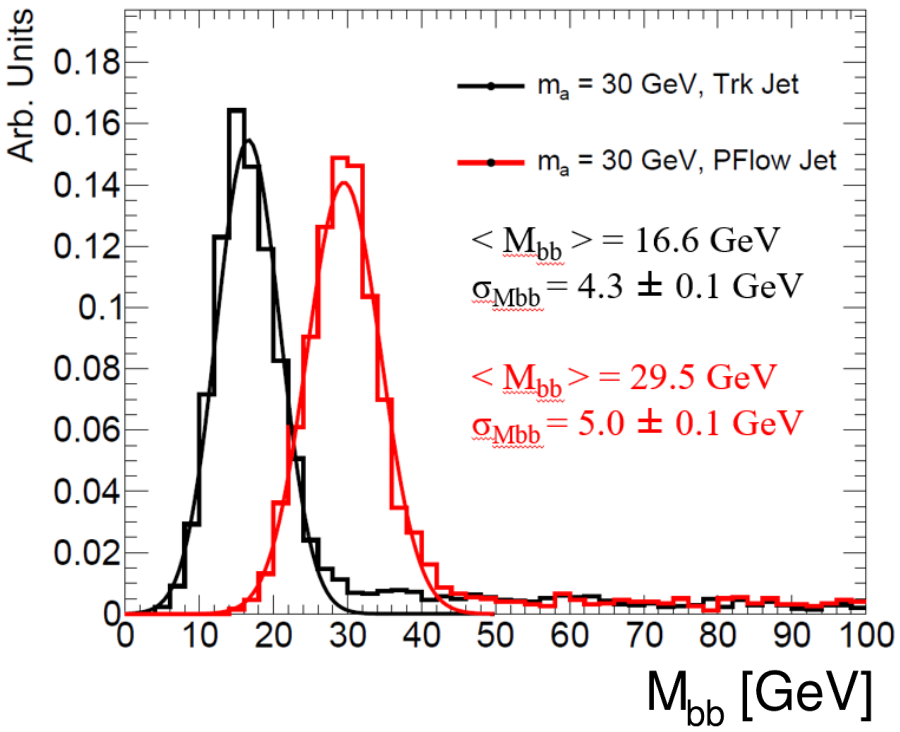

Figure 67: The m_{bb} distributions for 2 b-tagged calorimeter jets (red) and the same distribution build from a category with at least 1 b-tagged calorimeter jet, and 2 b-tagged track jets where the inputs for the m_{bb} calculation are the 2 track jets (black). Here, the explicit bias of the missing neutral fraction in the track jets lowers the overall scale of the M_{bb} distribution.

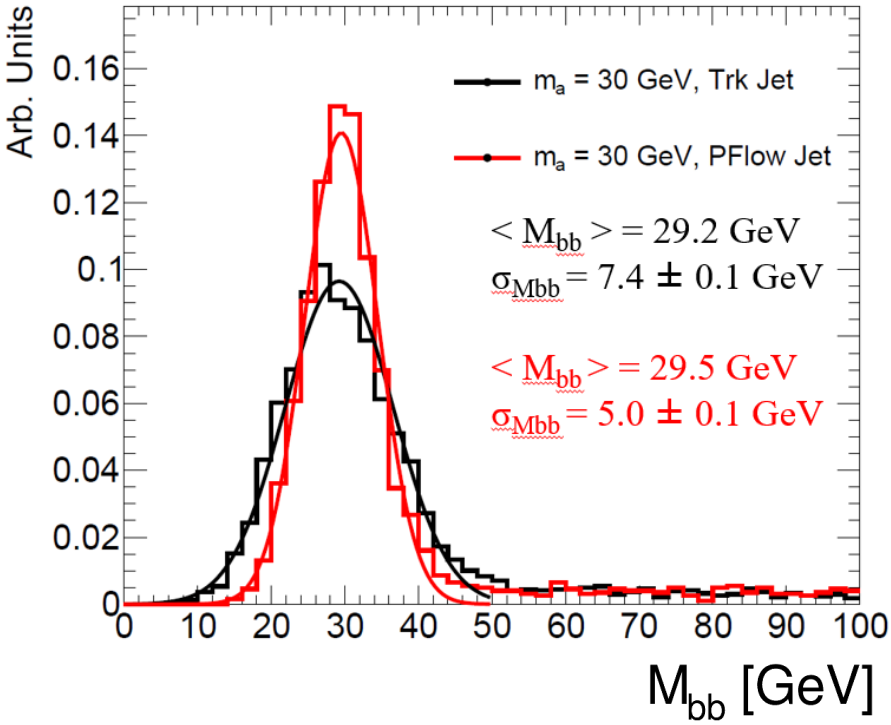

Figure 68: The m_{bb} distributions for 2 b-tagged calorimeter jets (red) and the same distribution build from a category with at least 1 b-tagged calorimeter jet, and 2 b-tagged track jets where the inputs for the m_{bb} calculation are the 2 track jets (black). Here, a pseduo-correction of 1.8, based on the mass scale shift, is applied to every jet and the m_{bb} is recomputed. The new mass scale is centered at approximately the correct value, though the distribution has a larger fractional M_{bb} resolution.

where the angular brackets indicate the mean value of the p_{T} ratio, and the binning in track-jet p_{T} and multiplicity of ID tracks with $p_{\mathrm{T}} > 500$ MeV, N_{trks}, is made explicit. These details are discussed further below in Section 6.1.2.

6.1.2 Event and Object Requirements for Track-Jet Correction Factor

To determine the track-jet correction factor, the dominant backgrounds of $t\bar{t}$ and Drell–Yan + jets are used to calculate ζ. These corrections are derived on MC simulation samples, where the following kinematic selections are applied.

Event Selection Every event for the track-jet correction factor determination is required to have at least two leptons, either muons or electrons, with $p_{\mathrm{T}}^{\ell} \geq 18$ GeV. Any muon candidate is required to be reconstructed as a combined or extrapolated muon, while any electron candidate is required to be identified as passing the medium likelihood discriminant. The medium likelihood discriminant is chosen for electrons as the looser requirements are found to be dominated by jet fakes. These fake jets would impact the final selections, as the overlap removal procedure prioritizes electron over track-jets in order to remove the electron track-jets. All events in the MC simulation are required to pass the emulated single lepton or dileptonic lowest unprescaled triggers. In addition, all events are required to have at least one vertex reconstructed with $N_{\mathrm{trks}} \geq 3$ and the nominal bad jet event veto requirement described in Section 4.5.4.

Track-Jet Momentum: The p_{T} of the track-jet must satisfy the condition that $p_{\mathrm{T}}^{\mathrm{Trkjet}} > 7$ GeV. This selection is chosen in order to prevent low p_{T} leptons and pileup interactions from entering into the selection. Further, this p_{T} cut is also required in the ATLAS flavor-tagging calibration; therefore, the same requirement is used to prevent large uncertainties in the flavor tagging data-MC scale factors.

Geometrical Overlaps: As the track-jets are built from the same tracks that could seed electron and muon objects, both electrons and muons can be reconstructed as track jets. Therefore, an overlap removal procedure is implemented to remove the prompt leptons reconstructed as track jets. The condition is applied at the reconstructed level for all ΔR_y(track jet, lepton) < 0.3, and the lepton is always prioritized. An additional cut applied to the MC simulation uses the truth-level leptons ΔR_y(track jet, truth lepton) < 0.3, where only the prompt truth leptons are used in order to preserve the track jets with non-prompt semi-leptonic B decays.

Track Multiplicity: To further remove low-p_{T} electrons and muons reconstructed as track-jets, a selection on the track multiplicity $N_{\mathrm{trks}} \geq 2$ is applied for all track

jets. This selection has a marginal impact on the b-flavor track-jets from signal decay $a \to bb$, removing between 1%-2% depending on m_a.

Track Jet Reconstruction Overlap: When the constituents in variable radius track-jets reach sufficiently high p_{T}, the possibility of collinear jets arises. In these regimes, the truth-level B-hadron to track-jet association tends to be unreliable, as the track constituents from these decays are distributed between both reconstructed track-jets. In order to prevent these events from entering into the ζ determination, all events which contain a pair of jets that satisfy ΔR_y(track-jet$_i$,track-jet$_j$) < min(Jet radius$_i$, Jet radius$_j$) are vetoed. Here, the jets in the i collection cover all track-jets that are b-labeled, while j runs over all flavors of track-jets meeting the p_{T} and N_{trks} requirements and only cases where $i \neq j$ are considered.

Another geometric overlap condition is crucial for the ζ correction factor determination: the overlap between a reconstructed jet and a truth jet in the MC simulation. Each track jet pair must satisfy ΔR_y(track jet, truth jet) < R_{trk}, where R_{trk} is the track jet radius. R_{trk} is allowed to vary given the condition in Equation 83, so the truth matching criteria must change as a function of the momentum scale of the track jet. This matching explicitly vetoes the case where a truth jet cannot be found within the variable cone radius of the reconstructed track jet. In this case, the event is kept, but a correction factor is not derived for the track jet only the track jet is not considered further for the average correction. The distribution of the minimum R_y between the closest truth jet to a track jet can be found in Figure 69, which indicates that a natural scale cutoff is on the order of R=0.4, as expected.

ζ **Binning**: In order to best correct the track-jet p_{T}, E, several additional kinematic variables were tested to determine their constraining power on ζ. The correction factor is determined primarily in p_{T}, with an additional handle being provided by the track multiplicity N_{trks}. This constraining power is due to the logarithmic relationship between the track-jet p_{T} and N_{trks}. This relationship, along with other kinematic properties of the signal, are illustrated in Figures 70-71.

Example ζ values for different N_{trks} bins as a function of ln(p_{T}) are shown in Figure 72 to illustrate the effect of including track multiplicity in the determination of ζ. The ζ correction factors are binned as a function of the logarithmic p_{T} in order to reduce statistical bin uncertainties across the kinematic range. A more detailed study of the track jet energy correction factor parametrization included binning in several additional kinematic variables including the η, ϕ,track multiplicity of tracks with $p_{\mathrm{T}} > 1$ GeV, and the track width ($\sum(\Delta R(\text{jet, track}) p_{\mathrm{T}}^{\mathrm{trk}}) / \sum p_{\mathrm{T}}^{\mathrm{trk}}$), but these variables did not provide additional constraining information to the energy correction.

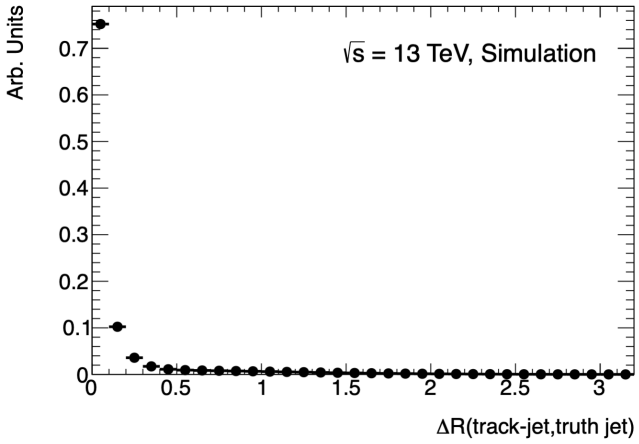

Figure 69: Minimum ΔR_y(track-jet,truth-jet) distribution for each track-jet in simulated $t\bar{t}$ events. This minimum illustrates that the matching criterion $\Delta R_y < R$, where R is the variable radius of the track-jet, provides a natural choice for determining ζ.

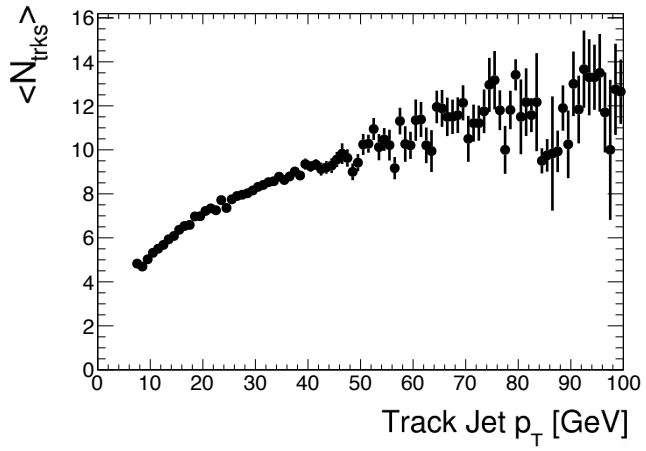

Figure 70: An illustration of the logarithmic relationship between $\langle N_{\text{trks}} \rangle$ and p_{T} for trackjets using an exotic Higgs signal with $m_a = 30$ GeV.

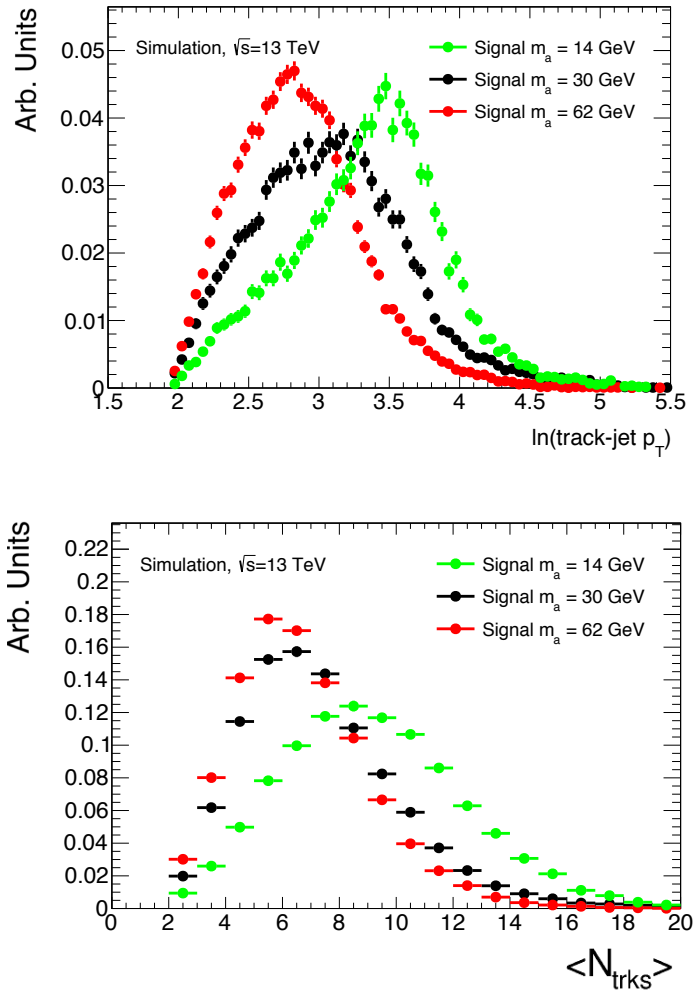

Figure 71: Track-jet $\ln(p_{\mathrm{T}})$, N_{trks} for the $H \to aa \to bb\mu\mu, m_a = 14$ GeV (green), 30 GeV (black), and 62 GeV (red) signal processes.

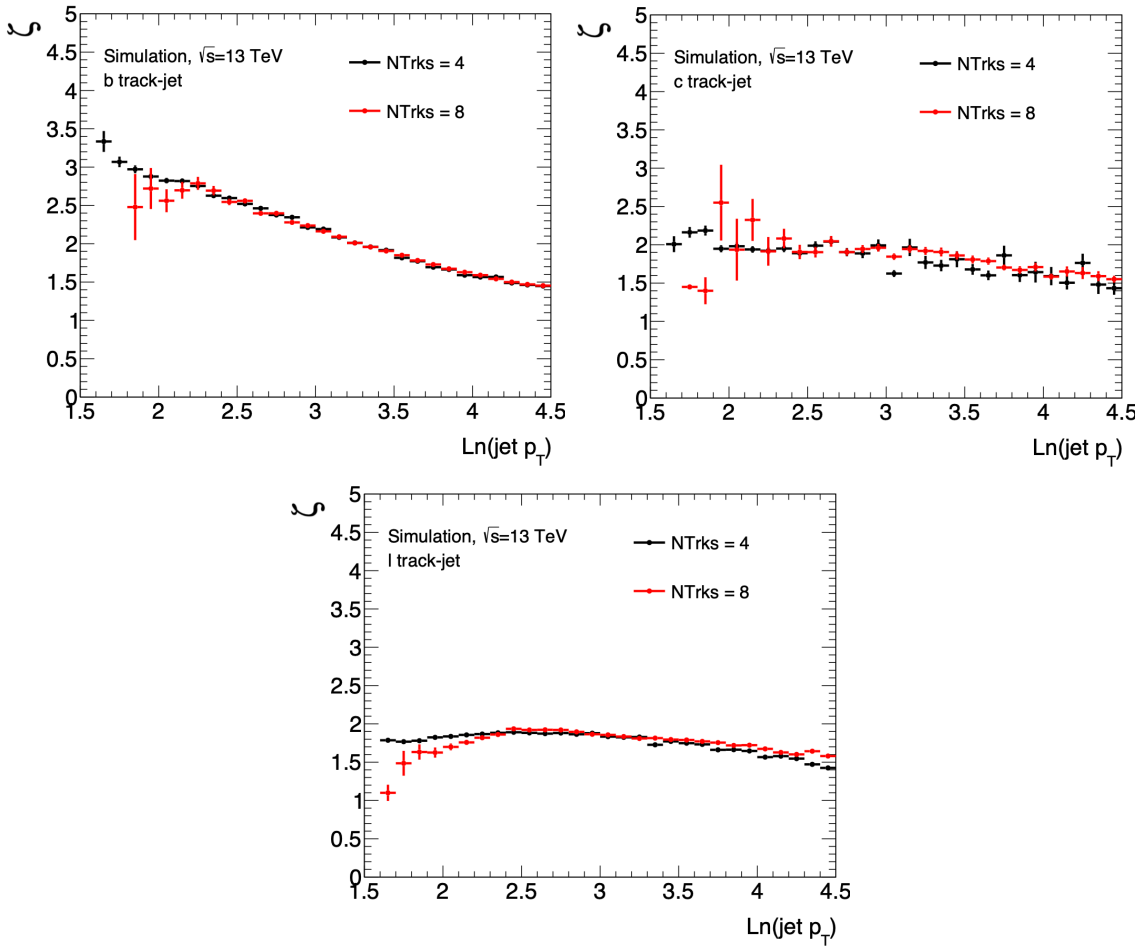

Figure 72: Values of ζ as a function of ln(track-jet p_{T}) for different jet flavors and a varying number of tracks. The ζ values differ at low p_{T} for all flavors of jets and a correction factor in these regions would be biased if the track multiplicity N_{trks} was not taken into account. As the low momentum range is preferred by the $a \to bb$ signal, these small correction are quite significant.

Track Jet Flavor: The correction factor extraction is performed on track-jet distributions categorized by their truth-level identification label. This label is computed by matching reconstructed track jets to truth hadrons in the MC simulation and assigning a number dependent on whether the primary hadron is a light-flavor, charm-flavor, or b-flavor hadron. An example of the differences in ζ based on flavor identification can be seen in Figure 73 for several N_{trks} bins. Based on the similarity between the light and charm quark distributions for the backgrounds, these two correction factors are estimated using the correction factor for the light jets. This grouping also enables b-tagging as a use classifier to extract and apply the ζ correction factor.

6.1.3 Extracting Correction Factors

After the event and object selections, the ζ values are determined for the Drell–Yan and $t\bar{t}$ backgrounds in MC simulation samples. The results for track-jets with $N_{\text{trks}} = 7$, separated into light-flavor and b-flavor categories, are shown in Figure 74. A large discrepancy in ζ is observed between the $t\bar{t}$ and Drell–Yan backgrounds. This discrepancy arises due to reconstruction bias of the track jet p_{T} for these two processes. The truth b-jet p_{T} distribution for the Drell–Yan and $t\bar{t}$ backgrounds is shown in Figure 75 and the truth light jet distribution is illustrated in Figure 76.

The Drell–Yan events tend to contain relatively low p_{T} track-jet, while the $t\bar{t}$ phase space has track jets with a moderate amount of p_{T}; therefore, requiring a low p_{T} track-jet in the $t\bar{t}$ events puts the events in an odd corner of kinematic phase space. In this phase space, most of the $t\bar{t}$ jets preferentially decay to neutral particles, as the amount of energy in the charged particles cannot exceed 7-15 GeV for a 60 GeV jet. The charged fraction, therefore, is biased towards low values in these events by the track-jet reconstruction. This bias is illustrated in Figure 77, where the truth jet charged fraction is plotted for jets which match reconstructed track-jets with $2.0 < \ln(\text{p}_{\text{T}}) < 2.5$; the unbiased truth jet distribution is also shown in Figure 77, where the charged fraction is shown for all truth jets with $2.0 < \ln(\text{p}_{\text{T}}) < 2.5$ and not selecting on the track-jet p_{T}.

To capture the underlying physical scale, the ζ values are derived via a cross-section weighted average. In this weighted average, the Drell–Yan events due to their larger cross-section take precedence over the $t\bar{t}$ events at low p_{T} scales and the $t\bar{t}$ events dominate the average at higher momentum scales. This cross-section weighted average is shown in Figure 78 and the final correction factor corresponds to these binned values. A signal derived correction factor is also given in Figure 78, which shows a trend towards the Drell–Yan correction factors.

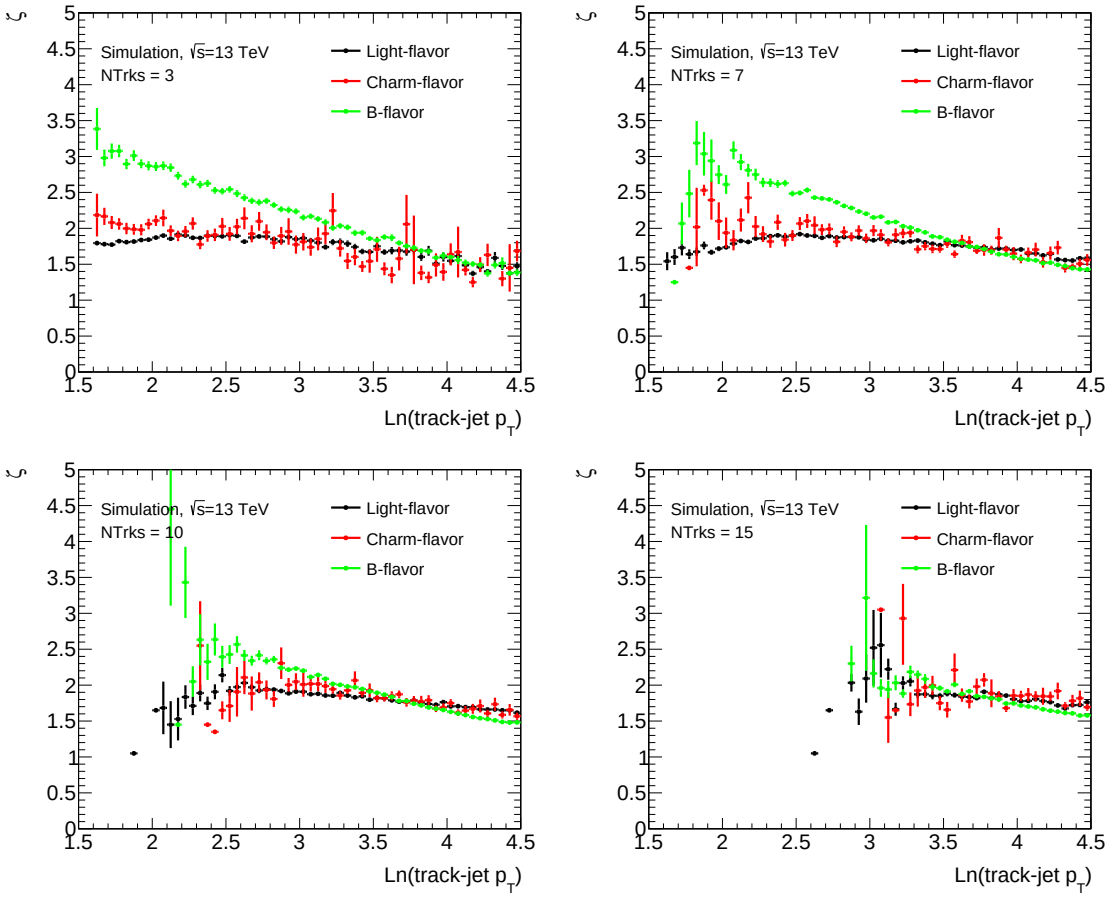

Figure 73: The ζ value determined for track-jets based on the truth flavor identifier for $N_{\text{trks}} = 3$,7,10, and 15 for simulated $t\bar{t}$ events. The similarity of the charm and light flavor ζ values allows for some simplification of these corrections. The depletion of the low p_{T} region for increasing track multiplicities arises due to the logarithmic relationship between these variables.

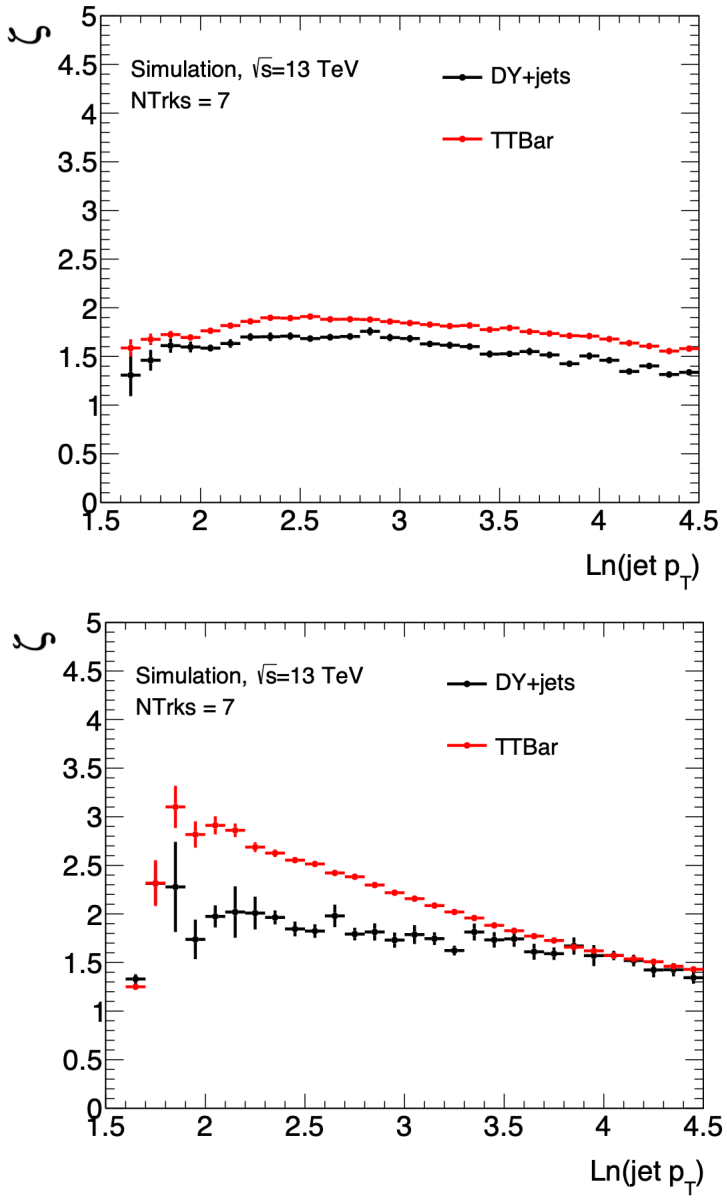

Figure 74: The correction factor ζ derived as a function of ln(track-jet p_{T}) for light (top) and b-flavor (bottom) track-jets with $N_{\mathrm{trks}} = 7$ using simulated Drell–Yan+jet and $t\bar{t}$ events. A large discrepancy can be seen at low p_{T} between the correction factors for the b-labeled track-jets in the two physics processes.

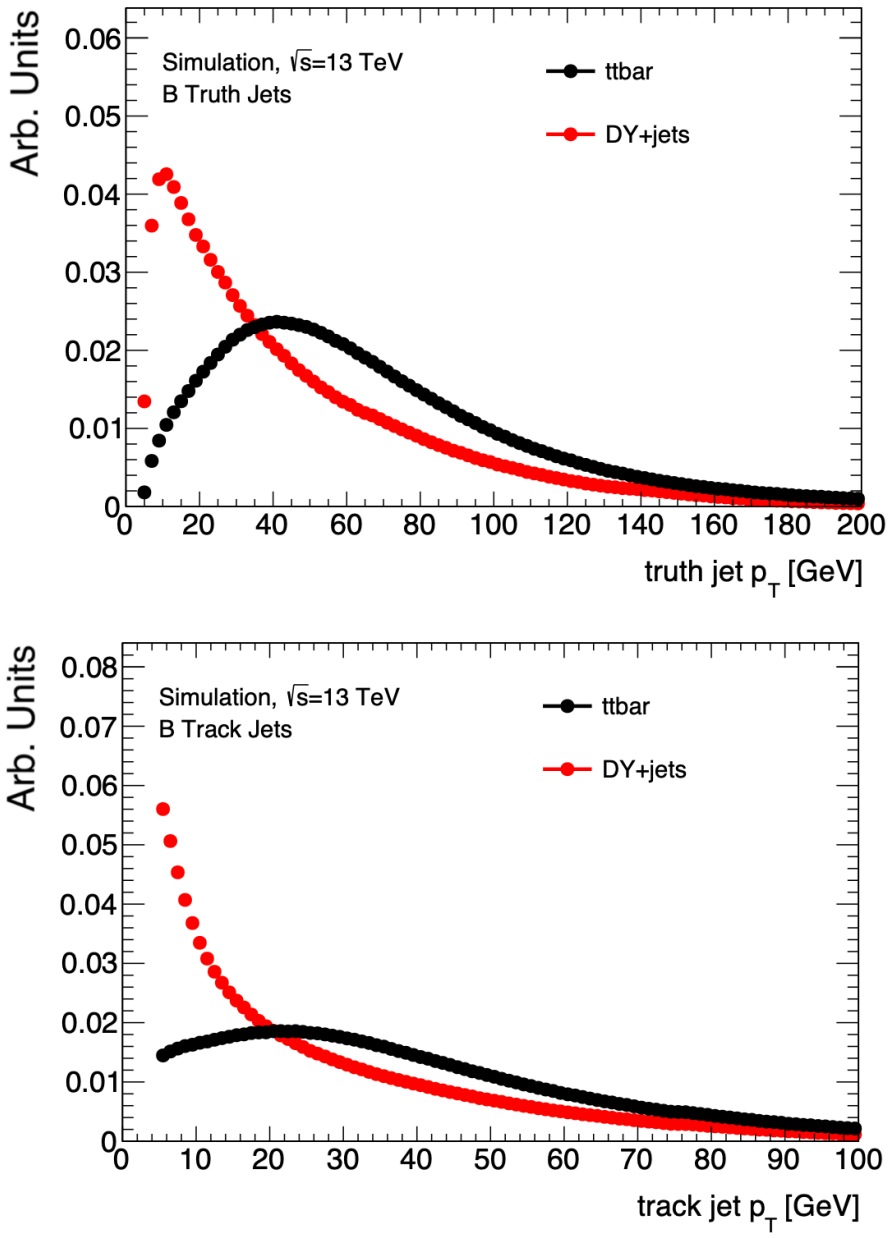

Figure 75: The truth jet (top) and track-jet (bottom) p_{T} distribution for b-jets in the $t\bar{t}$ and Drell–Yan+jet backgrounds. Here, the differences between the two background momenta cause the disagreement in the calculation of ζ.

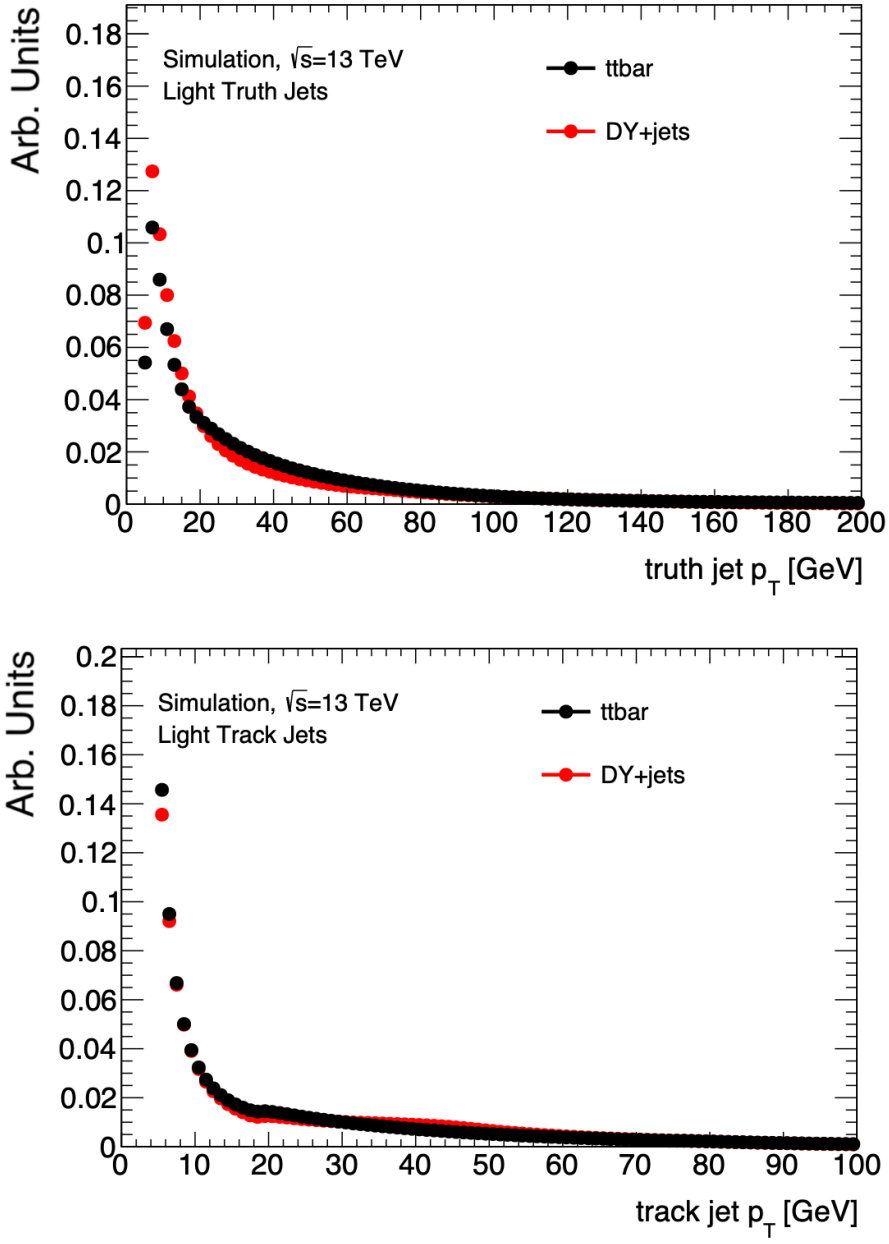

Figure 76: The truth jet (top) and track-jet (bottom) p_{T} distribution for light jets in the $t\bar{t}$ and Drell–Yan+jet backgrounds. Here, the similarity of the two background momenta yields the relative agreement in the calculation of ζ.

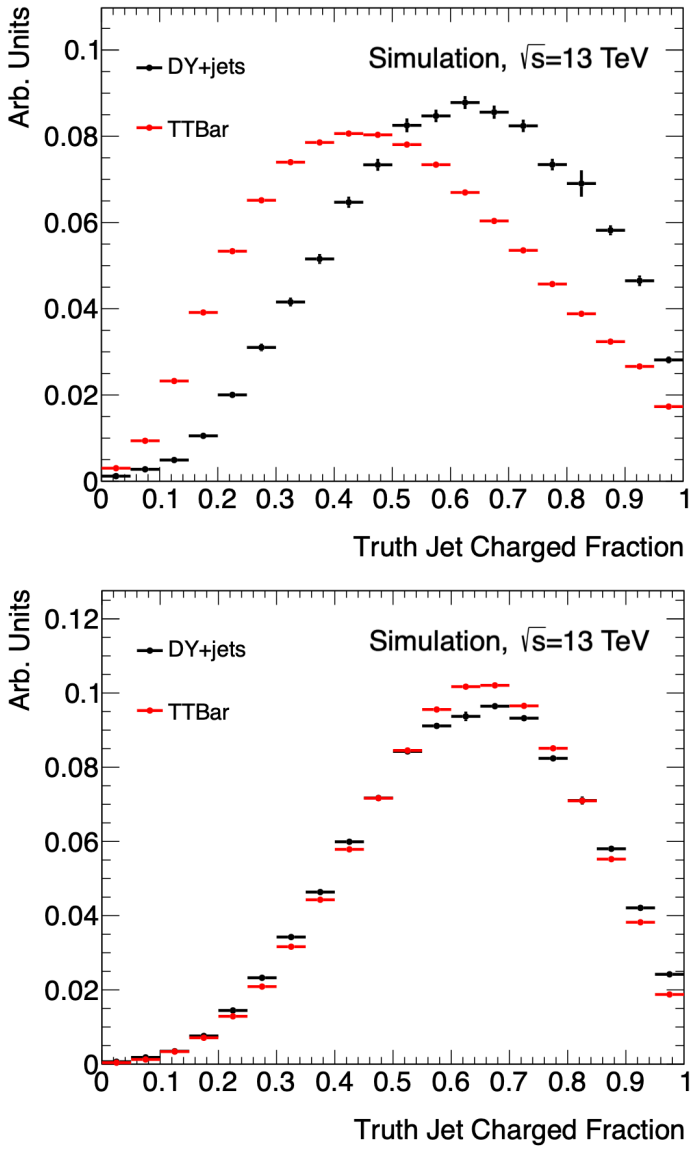

Figure 77: The truth-jet charged fraction distributions for b-flavor truth jets matched to track-jets in the range $2.0 < \ln(p_{\mathrm{T}}^{\mathrm{trk}}) < 2.5$ (top), and b-flavor truth-jets in the range $2.0 < \ln(p_{\mathrm{T}}^{\mathrm{truth}}) < 2.5$ (bottom) without a reconstruction bias using simulated Drell–Yan+jet and $t\bar{t}$ events. A large discrepancy between the physics processes is seen when the track-jet p_{T} is preferentially selected on at low p_{T}.

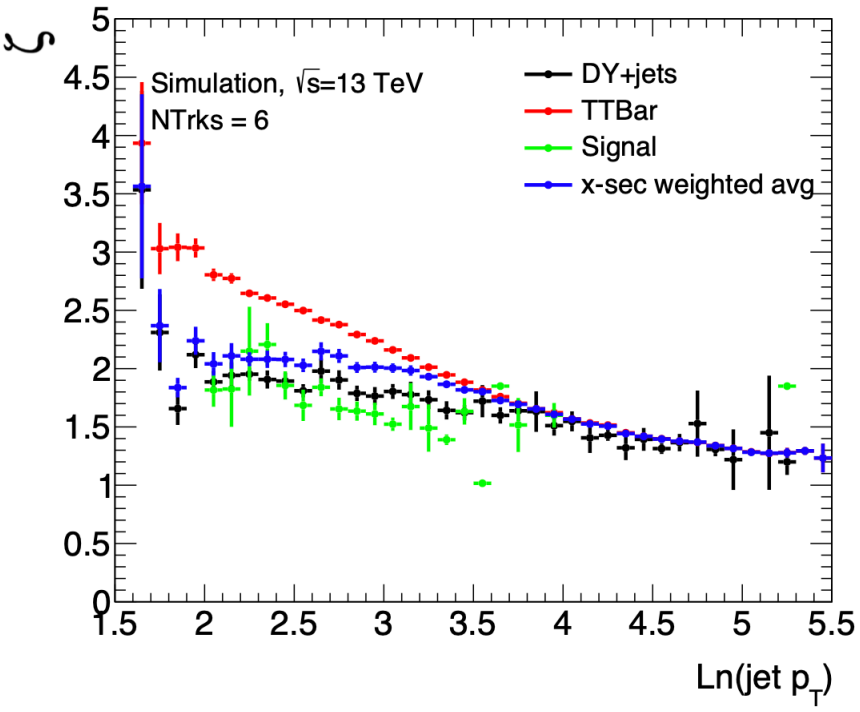

Figure 78: A comparison of the cross-section weighted average ζ values (in blue) for $N_{\text{trk}} = 6$ with the correction factors from the signal and backgrounds given for illustration purposes.

In order to fully correct the track-jet, the energy of the track-jet must also be corrected in addition to its p_{T}. The same procedure for the momentum based correction was repeated for the energy and found not to differ significantly from the momentum based correction. This result is not unexpected given that the energy measurement is completely influenced by the ID tracker determination of the momentum in this phase space.

6.1.4 Additional Acceptance of $bb\mu\mu$ Signal with Track Jets

With the track-jet scale now corrected using the cross-sectional weighted correction factors, there are several categories of event that can be targeted in the $H \to bb\mu\mu$ final state. These include the several categories:

- **Category A**: The full Run-2 data set nominal analysis when 2 particle flow jets with $p_{\text{T}} > 20$ GeV are b-tagged

Signal Mass	Events in Cat. A	Events in Cat. B	Events in Cat. C
14 GeV	14.7	14.1	1.9
30 GeV	47.6	36.9	5.1
62 GeV	137.3	99.7	16.0

Table 32: Weighted Event Yields for $m_a = 14, 30, 62$ GeV for the categories described in Section 6.1.4. These yields quantify the number of signal events passing a modified preselection with 2 muons having $p_{\mathrm{T}}^{\mu} > 18$ GeV, $E_{\mathrm{T}}^{\mathrm{miss}} > 60$ GeV, and modified b-tagging selections based on track-jet and particle flow jets tagged with the MV2c10 77% working point. Categories A-C are mutually exclusive and the event yields are weighted to $\mathrm{Br}(H \to aa \to bb\mu\mu) = 1.6{\times}10^{-3}$ and a luminosity of $\mathcal{L} = 36.1$ fb^{-1}.

- **Category B**: An alternative analysis using 1 particle flow jet with $p_{\mathrm{T}} > 20$ GeV and one unique (not matched to the particle flow jet) track-jet with $p_{\mathrm{T}} > 7$ GeV are b-tagged. This category recover the signal events when the subleading jet falls below the 20 GeV reconstruction threshold.
- **Category C**: Another alternative analysis where two track-jets with $p_{\mathrm{T}} > 7$ GeV are b-tagged. This category would cover the cases where both jets fail the 20 GeV particle flow reconstruction threshold.

In each case, the M_{bb} is calculated from the main analysis objects; for example, in the 1 particle flow, 1 track-jet category, the track-jet (p_{T}, E) are corrected and M_{bb} is constructed from the particle flow jet and corrected track-jet 4-vectors. A breakdown of the event acceptances in dilepton trigger events for these categories is given in Table 32 for three simulated m_a signals. Based on these yields, an analysis targeting both Category A and Category B increases the signal acceptance by up to a factor of 2. Therefore, an alternative analysis using 1 particle flow jet and 1 track-jet is being investigated for the full Run-2 analysis result.

6.2 Muon and Trigger Improvements

During Run-2, the LHC exceeded its design luminosity by a factor of two. This excess resulted in higher than expected pileup and the ATLAS detector raised trigger thresholds to cope with the increased interaction rates. Going forward towards an analysis on the full LHC Run-2 dataset requires a more thorough understanding of the signal acceptance in the dimuon side of the $H \to aa \to bb\mu\mu$ analysis. These

Signal Mass Point	2015 Acceptance [%]	2016 Acceptance [%]
$m_a = 20$ GeV	71.0	64.2
$m_a = 30$ GeV	69.9	61.5
$m_a = 40$ GeV	74.5	64.7
$m_a = 50$ GeV	78.4	68.6
$m_a = 60$ GeV	83.4	74.1

Table 33: The trigger acceptance for the H$\rightarrow aa \rightarrow bb\mu\mu$ signal mass points with respect to the number of events with two medium identification, tight isolation passing muons and $p_{\mathrm{T}} > 15,7$ GeV in the simulated signal sample. The acceptance is defined with respect to the L1 trigger plateau, i.e. events are required to have at least one muon with $p_{\mathrm{T}} > 21$ GeV in the 2015 selection and $p_{\mathrm{T}} > 27$ GeV in the 2016 selection.

include studies of the trigger efficiency, which plays a major role as the first line of identification of the signal.

Because of the increase in the single lepton trigger p_{T} threshold from 20 GeV to 26 GeV, the single lepton selection efficiency for the H$\rightarrow aa \rightarrow bb\mu\mu$ analysis decreased. As an example, the signal acceptance for the $m_a = 60$ GeV simulated signal sample for events with two combined, isolated muons having $p_{\mathrm{T}} > 15,7$ GeV decreased from an acceptance of 72.6% in 2015 to 64.6% in 2016[11]. The acceptance loss between 2015 and 2016 for all mass points can be found in Table 33.

To illustrate the reason for the reduction in signal acceptance, the leading and subleading muon transverse momenta for the H$\rightarrow aa \rightarrow bb\mu\mu$ signals are shown in Figure 79. The leading muons in the signal samples tend to be boosted with a mean value around 30 GeV, while the subleading muons tend to be much softer as the a mass decreases. Therefore, the loss from the trigger increase corresponds to inefficiencies on both of the muon momenta.

To increase the signal acceptance, several new signal regions are constructed by adding in additional trigger requirements. To make each trigger selection orthogonal, the regions are defined with respect to the leading lepton p_{T} threshold. These trigger requirements and the corresponding lepton thresholds are detailed in Table 34. The categories are defined by the following

[11] As these studies include lower p_{T} muons and isolation changes, these numbers differ from the analysis cutflows in Appendix C.

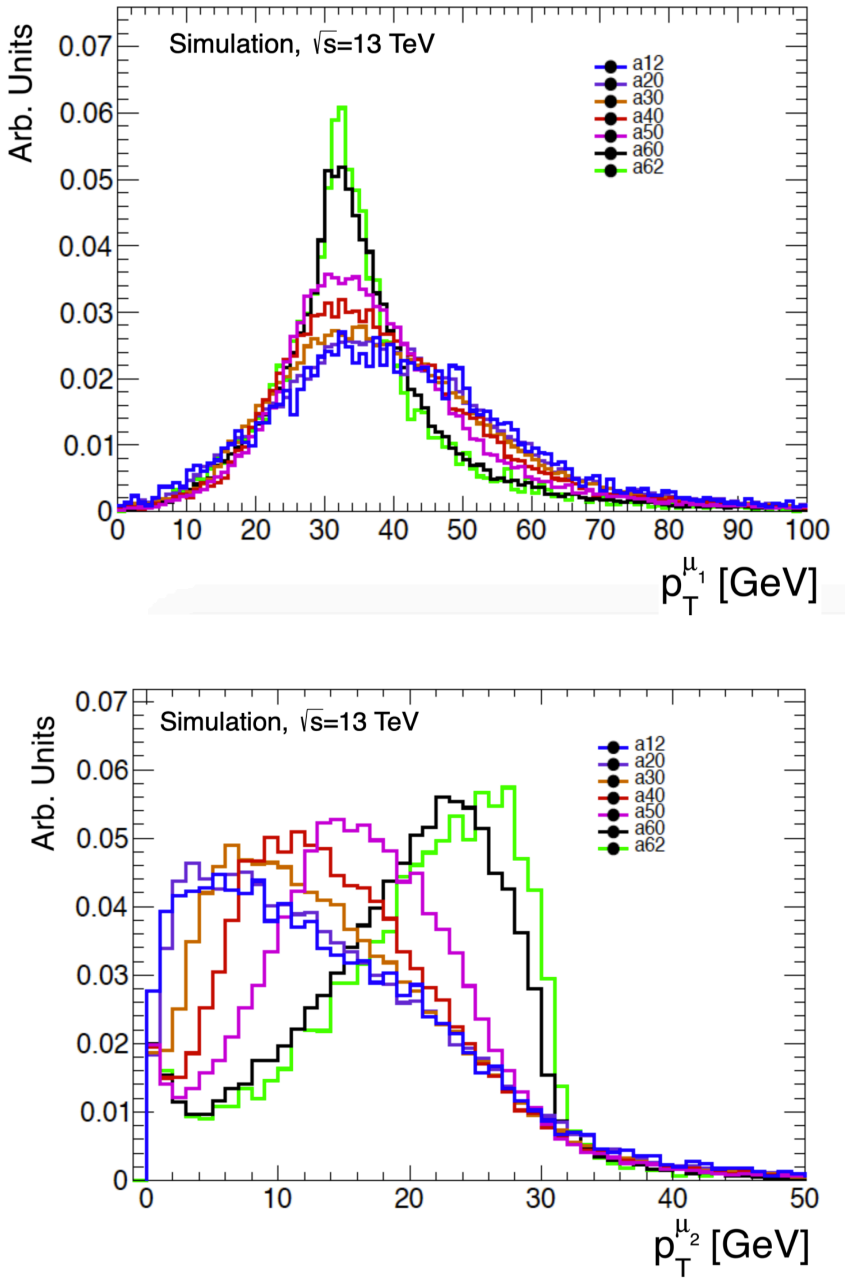

Figure 79: Leading (top) and subleading (bottom) muon p_{T} for the H$\rightarrow aa \rightarrow bb\mu\mu$ signal with 12 GeV $< m_a <$ 62 GeV.

Trigger Type	ATLAS Trigger Name	Analysis Channel
Single Muon, Nominal	mu{20,24}_iloose mu{24,26}_ivarmedium mu{40,50}	$p_{\mathrm{T}}^{\mu_1} > 27$ GeV, $p_{\mathrm{T}}^{\mu_2} > 3$ GeV
Asymmetric Dimuon	mu(18,20,22)_mu8noL1	$p_{\mathrm{T}}^{\mu_1} > 23$ GeV, $p_{\mathrm{T}}^{\mu_2} > 10$ GeV
Symmetric Dimuon	2mu{10,14}	$p_{\mathrm{T}}^{\mu_1} > 15$ GeV, $p_{\mathrm{T}}^{\mu_2} > 15$ GeV

Table 34: The single lepton and dilepton HLT trigger names for the triggers used to optimize signal efficiency. Each trigger type involves a logical "OR" of all triggers in the category to account for isolation and identification inefficiencies at different momentum scales or to account for the lowest thresholds for a given data-taking period. The braced terms refer to different thresholds used to control trigger rates during the LHC luminosity ramp-up. The "iX" names refer to the isolation details given in Section 4.4. Further details on the triggers can be in Refs. [123, 124].

Category 1: In category 1, the single lepton trigger category is expanded by using lower p_{T} muons, as ATLAS has developed a new "LowPt" identification working point to target muons below 7 GeV. This identification loosens the number of MS station coincidences as a low momentum muon is expected to cross less than the standard 3 stations used in the "Medium" requirement. In order to remove non-prompt lepton fakes from this looser selection, additional constraints on the tracking quantities relating the ID muon tracks and the MS muon tracks are applied. The efficiency increase in the barrel region for a general prompt muon signature rises from 70% to 90% while keeping the non-prompt rate below 0.5% [143]. The signal acceptance in this region rises by 3% for the $m_a = 60$ GeV signal mass, but higher gains of up to 16% are seen for lower m_a. These gains comes about due to the phase-space asymmetry coming from the higher p_{T} of the a decays. The gains in signal acceptance are given in Table 35. Further studies are required to optimize the isolation of these muons, as additional fake backgrounds from W+jet and QCD multi-jet events could enter the analysis selection easier as the subleading muon p_{T} is lowered.

Categories 2-3: In categories 2 and 3, the use of dimuon triggers are explored. In these categories, the subleading lepton p_{T} thresholds are raised in order to lower the leading lepton p_{T}. Further gains can be made by additionally requiring isolation on the leading lepton, though these changes are marginal (20 versus 22 GeV) in the ATLAS trigger currently. Category 2 covers the cases with the asymmetric dimuon triggers, where the leading muon is required to have $p_{\mathrm{T}} > 23$ GeV and the subleading

Signal Mass Point	Acceptance Gain ($p_{\mathrm{T}}^{\mu_2} = 7 \rightarrow 3$ GeV) [%]
$m_a = 20$ GeV	15
$m_a = 30$ GeV	16
$m_a = 40$ GeV	10
$m_a = 50$ GeV	5
$m_a = 60$ GeV	3

Table 35: The acceptance change in signal event acceptance for the H$\rightarrow aa \rightarrow bb\mu\mu$ signal mass points by lowering the subleading muon p_{T} from 7 GeV to 3 GeV in the single lepton trigger category. The gains from lowering the momentum threshold tend to affect the lower m_a due to the boost of the a.

muon is required to have $p_{\mathrm{T}} > 10$ GeV. The subleading muon is only selected in the HLT decision and the L1 inefficiency from the muon system is avoided. In category 3, the symmetric dimuon trigger category, 2 L1 decisions are required to be made and the final muon selection requires 2 muons with $p_{\mathrm{T}} > 15$ GeV. The gains from these categories individually and combined are shown in Table 36. By combining both categories, an additional 12-17% of $H \rightarrow bb\mu\mu$ events pass the trigger and preselections.

By taking the OR of all of these trigger and the lower muon p_{T} changes, the increase in signal acceptance after the preselection and $E_{\mathrm{T}}^{\mathrm{miss}}$ selections is given in Table 37. These changes in trigger strategy result in increase of 24% for $m_a = 60$ GeV up to 36% increases for $m_a = 30$ GeV. These improvements, while smaller than the b-tagging gains discussed in Section 6.1, still provide significant gains in sensitivity for the full Run-2 result.

Signal Mass	Gain Category 2 [%]	Gain Category 3 [%]	Gain from 2+3 [%]
$m_a = 20$ GeV	7.4	9.5	11.8
$m_a = 30$ GeV	8.6	12.3	14.8
$m_a = 40$ GeV	11.3	14.1	17.2
$m_a = 50$ GeV	11.0	13.2	16.0
$m_a = 60$ GeV	9.2	11.5	13.7

Table 36: The acceptance gain in signal event acceptance for the H$\rightarrow aa \rightarrow bb\mu\mu$ signal mass points achieved by using the symmetric dimuon trigger, the asymmetric dimuon triggers, and the OR between all triggers. The gains are driven by the asymmetric dimuon trigger due to the relatively high p_{T} of the subleading muon requirement and the L1 inefficiency being applied for the symmetric dimuon trigger. The OR of all triggers, however, remains a viable possibility.

Signal Mass Point	Total Gain [%]
$m_a = 20$ GeV	31
$m_a = 30$ GeV	36
$m_a = 40$ GeV	32
$m_a = 50$ GeV	27
$m_a = 60$ GeV	24

Table 37: The signal event acceptance gains for different trigger configurations in the H$\rightarrow aa \rightarrow bb\mu\mu$ analysis. The trigger optimization categories are described in the text. These include lowering the subleading muon p_{T} from 7 GeV to 3 GeV in the single lepton trigger category and forming two new trigger categories based on the dilepton triggers. All selections have been orthogonalized to emphasize full event gains.

7 Conclusions

A search for exotic decays of the Higgs boson to the $bb\mu\mu$ final states is performed with proton-proton collisions at $\sqrt{s} = 13$ TeV with 36.1 fb^{-1} of data collected by the ATLAS detector at the LHC during the 2015 and 2016 data-taking periods. The search for a narrow dimuon resonance is performed in the range 18 GeV $\leq m_{\mu\mu} \leq$ 62 GeV, using mass windows based on the experimental dimuon resolution. No significant excess is observed in the data above the SM background predictions. In the absence of a signal, upper limits are set on the cross section times branch ratio $(\sigma_H/\sigma_{\text{SM}}) \times \mathcal{B}(H \to aa \to bb\mu\mu)$, which range from 1.2×10^{-4} - 8.4×10^{-4}, depending on m_a. Additional limits on new physics production in the $bb\mu\mu$ final state are determined, with limits on $\sigma_{\text{vis}}(X) \times \mathcal{B}(X \to bb\mu\mu)$ ranging from 0.1 fb to 0.73 fb, depending on the dimuon mass window. Finally, several new paths towards greater sensitivity are explored for muon and b-jet identification and found to contribute significant gains to any future analysis targeting the $H \to aa \to bb\mu\mu$ process.

References

[1] S. L. Glashow. Partial Symmetries of Weak Interactions. *Nucl. Phys.*, 22:579–588, 1961. DOI: 10.1016/0029-5582(61)90469-2.

[2] Abdus Salam. Weak and Electromagnetic Interactions. *Conf. Proc.*, C680519: 367–377, 1968.

[3] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967. DOI: 10.1103/PhysRevLett.19.1264.

[4] Gerard 't Hooft and M. J. G. Veltman. Regularization and Renormalization of Gauge Fields. *Nucl. Phys.*, B44:189–213, 1972. DOI: 10.1016/0550-3213(72)90279-9.

[5] M. Tanabashi et al. Review of Particle Physics. *Phys. Rev.*, D98(3):030001, 2018. DOI: 10.1103/PhysRevD.98.030001.

[6] C. E. Aalseth et al. Search for Neutrinoless Double-Beta Decay in ^{76}Ge with the Majorana Demonstrator. *Phys. Rev. Lett.*, 120(13):132502, 2018. DOI: 10.1103/PhysRevLett.120.132502. arXiv: 1710.11608 [nucl-ex].

[7] Kevork N. Abazajian and Manoj Kaplinghat. Neutrino Physics from the Cosmic Microwave Background and Large-Scale Structure. *Ann. Rev. Nucl. Part. Sci.*, 66(1):401–420, 2016. DOI: 10.1146/annurev-nucl-102014-021908.

[8] Alfred Scharff Goldhaber and Michael Martin Nieto. Photon and Graviton Mass Limits. *Rev. Mod. Phys.*, 82:939–979, 2010. DOI: 10.1103/RevModPhys.82.939. arXiv: 0809.1003 [hep-ph].

[9] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental Test of Parity Conservation in Beta Decay. *Physical Review*, 105:1413–1415, February 1957. DOI: 10.1103/PhysRev.105.1413.

[10] Marco Drewes. The Phenomenology of Right Handed Neutrinos. *Int. J. Mod. Phys.*, E22:1330019, 2013. DOI: 10.1142/S0218301313300191. arXiv: 1303.6912 [hep-ph].

[11] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964. DOI: 10.1103/PhysRevLett.13.508.

[12] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–323, 1964. DOI: 10.1103/PhysRevLett.13.321.

[13] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.*, 13:585–587, 1964. DOI: 10.1103/PhysRevLett.13.585.

[14] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716:1, 2012. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[15] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B*, 716:30, 2012. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].

[16] D de Florian et al. Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector, 2016. arXiv: 1610.07922 [hep-ph].

[17] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters, 2018. arXiv: 1807.06209 [astro-ph.CO].

[18] ATLAS Collaboration. Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment. ATLAS-CONF-2019-005. 2019. URL: https://cds.cern.ch/record/2668375.

[19] CMS Collaboration. Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$. *Eur. Phys. J.*, 2018. arXiv: 1809.10733 [hep-ex].

[20] Martinus Veltman. Radiative corrections to vector boson masses. *Phys. Lett. B*, 91:95–98, 1980. DOI: 10.1016/0370-2693(80)90669-3.

[21] P. Sikivie, Leonard Susskind, Mikhail B. Voloshin, and Valentin I. Zakharov. Isospin Breaking in Technicolor Models. *Nucl. Phys.*, B173:189–207, 1980. DOI: 10.1016/0550-3213(80)90214-X.

[22] Michael E. Peskin and Tatsu Takeuchi. A New constraint on a strongly interacting Higgs sector. *Phys. Rev. Lett.*, 65:964–967, 1990. DOI: 10.1103/PhysRevLett.65.964.

[23] H.E. Haber and G.L. Kane. The search for supersymmetry: probing physics beyond the standard model. *Physics Reports*, 117:75–263, 1985. DOI: 10.1016/0370-1573(85)90051-1.

[24] Ulrich Ellwanger, Cyril Hugonie, and Ana M. Teixeira. The Next-to-Minimal Supersymmetric Standard Model. *Phys. Rept.*, 496:1–77, 2010. DOI: 10.1016/j.physrep.2010.07.001. arXiv: 0910.1785 [hep-ph].

[25] R. Peccei and H.R. Quinn. Cp conservation in the presence of instantons. *Phys.Rev.Lett.*, 38:1440–1443, 1977. DOI: 10.1103/PhysRevLett.38.1440.

[26] John F. Gunion, Howard E. Haber, Gordon L. Kane, and Sally Dawson. The Higgs Hunter's Guide. *Front. Phys.*, 80:1–404, 2000.

[27] David Curtin et al. Exotic decays of the 125 GeV Higgs boson. *Phys. Rev. D*, 90(7):075004, 2014. DOI: 10.1103/PhysRevD.90.075004. arXiv: 1312.4992 [hep-ph].

[28] Marcela Carena, Ian Low, Nausheen R. Shah, and Carlos E. M. Wagner. Impersonating the Standard Model Higgs Boson: Alignment without Decoupling. *JHEP*, 04:015, 2014. DOI: 10.1007/JHEP04(2014)015. arXiv: 1310.2248 [hep-ph].

[29] G.C. Branco et al. Theory and phenomenology of two-Higgs-doublet models. *Phys. Rept.*, 516:1–102, 2012. DOI: 10.1016/j.physrep.2012.02.002. arXiv: 1106.0034 [hep-ph].

[30] Keith Hamilton, Paolo Nason, Emanuele Re, and Giulia Zanderighi. NNLOPS simulation of Higgs boson production. *JHEP*, 10:222, 2013. DOI: 10.1007/JHEP10(2013)222. arXiv: 1309.0017 [hep-ph].

[31] Mirkoantonio Casolino, Trisha Farooque, Aurelio Juste, Tao Liu, and Michael Spannowsky. Probing a light CP-odd scalar in di-top-associated production at the LHC. *Eur. Phys. J.*, C75:498, 2015. DOI: 10.1140/epjc/s10052-015-3708-y. arXiv: 1507.07004 [hep-ph].

[32] Nathaniel Craig, Francesco D'Eramo, Patrick Draper, Scott Thomas, and Hao Zhang. The Hunt for the Rest of the Higgs Bosons. *JHEP*, 06:137, 2015. DOI: 10.1007/JHEP06(2015)137. arXiv: 1504.04630 [hep-ph].

[33] David Curtin, Rouven Essig, and Yi-Ming Zhong. Uncovering light scalars with exotic Higgs decays to $b\bar{b}\mu^+\mu^-$. *JHEP*, 06:025, 2015. DOI: 10.1007/JHEP06(2015)025. arXiv: 1412.4779 [hep-ph].

[34] T. Gleisberg et al. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009. DOI: 10.1088/1126-6708/2009/02/007. arXiv: 0811.4622 [hep-ph].

[35] Hung-Liang Lai, Marco Guzzi, Joey Huston, Zhao Li, Pavel M. Nadolsky, et al. New parton distributions for collider physics. *Phys. Rev. D*, 82:074024, 2010. DOI: 10.1103/PhysRevD.82.074024. arXiv: 1007.2241 [hep-ph].

[36] Sayipjamal Dulat, Tie-Jiun Hou, Jun Gao, Marco Guzzi, Joey Huston, Pavel Nadolsky, Jon Pumplin, Carl Schmidt, Daniel Stump, and C. P. Yuan. New parton distribution functions from a global analysis of quantum chromodynamics. *Phys. Rev.*, D93(3):033006, 2016. DOI: 10.1103/PhysRevD.93.033006. arXiv: 1506.07443 [hep-ph].

[37] Richard D. Ball et al. Parton distributions for the LHC Run II. *JHEP*, 04:040, 2015. DOI: 10.1007/JHEP04(2015)040. arXiv: 1410.8849 [hep-ph].

[38] L. A. Harland-Lang, A. D. Martin, P. Motylinski, and R. S. Thorne. Parton distributions in the LHC era: MMHT 2014 PDFs. *Eur. Phys. J.*, C75(5):204, 2015. DOI: 10.1140/epjc/s10052-015-3397-6. arXiv: 1412.3989 [hep-ph].

[39] ATLAS Collaboration. Measurement of charged-particle distributions sensitive to the underlying event in $\sqrt{s} = 13\,\text{TeV}$ proton–proton collisions with the ATLAS detector at the LHC. *JHEP*, 03:157, 2017. DOI: 10.1007/JHEP03(2017)157. arXiv: 1701.05390 [hep-ex].

[40] Anders Ryd, David Lange, Natalia Kuznetsova, Sophie Versille, Marcello Rotondo, David P. Kirkby, Frank K. Wuerthwein, and Akimasa Ishikawa. EvtGen: A Monte Carlo Generator for B-Physics, 2005.

[41] Stanislaw Jadach, Johann H. Kuhn, and Zbigniew Was. TAUOLA: A Library of Monte Carlo programs to simulate decays of polarized tau leptons. *Comput. Phys. Commun.*, 64:275–299, 1990. DOI: 10.1016/0010-4655(91)90038-M.

[42] J. Alwall et al. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[43] Stefano Frixione, Paolo Nason, and Carlo Oleari. Matching NLO QCD computations with Parton Shower simulations: the POWHEG method. *JHEP*, 11:070, 2007. DOI: 10.1088/1126-6708/2007/11/070. arXiv: 0709.2092 [hep-ph].

[44] Simone Alioli, Paolo Nason, Carlo Oleari, and Emanuele Re. A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX. *JHEP*, 1006:043, 2010. DOI: 10.1007/JHEP06(2010)043. arXiv: 1002.2581 [hep-ph].

[45] T. Sjöstrand et al. An introduction to pythia 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[46] Enrico Bothmann et al. Event Generation with SHERPA 2.2, 2019. arXiv: 1905.09127 [hep-ph].

[47] Bettina Mikulec, Alfred Blas, Christian Carli, Alan Findlay, Klaus Hanke, Giovanni Rumolo, and Jocelyn Tan. LHC Beams from the CERN PS Booster. In *Particle accelerator. Proceedings, 23rd Conference, PAC'09, Vancouver, Canada, May 4-8, 2009*, TU6PFP086, 2010. URL: http://accelconf.web.cern.ch/AccelConf/PAC2009/papers/tu6pfp086.pdf.

[48] I Bejar Alonso and L Rossi. HiLumi LHC Technical Design Report: Deliverable: D1.10. Technical report CERN-ACC-2015-0140, November 2015. URL: https://cds.cern.ch/record/2069130.

[49] Esma Mobs. The CERN accelerator complex. Complexe des accélérateurs du CERN, July 2016. URL: https://cds.cern.ch/record/2197559. General Photo.

[50] P. Grafström and W. Kozanecki. Luminosity determination at proton colliders. *Prog. Part. Nucl. Phys.*, 81:97–148, 2015. DOI: 10.1016/j.ppnp.2014.11.002.

[51] S van der Meer. Calibration of the effective beam height in the ISR. Technical report CERN-ISR-PO-68-31. ISR-PO-68-31, CERN, Geneva, 1968. URL: http://cds.cern.ch/record/296752.

[52] ATLAS Collaboration. Luminosity Public Results Run 2. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2. Accessed: 2019-03-27.

[53] G. Iadarola, L. Mether, and G. Rumolo. Filling schemes and e-cloud contraints for 2017. In *Proceedings, 7th Evian Workshop on LHC beam operation: Evian Les Bains, France, December 13-15, 2016*, pages 239–244, Geneva. CERN, CERN, 2017.

[54] ATLAS Collaboration. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *Eur. Phys. J. C*, 77:490, 2017. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex].

[55] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008. DOI: 10.1088/1748-0221/3/08/S08003.

[56] ATLAS Collaboration. ATLAS Insertable B-Layer Technical Design Report. ATLAS-TDR-19. 2010. URL: https://cds.cern.ch/record/1291633. Addendum: ATLAS-TDR-19-ADD-1. 2012. URL: https://cds.cern.ch/record/1451888.

[57] ATLAS Collaboration. Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV. ATL-PHYS-PUB-2015-018. 2015. URL: https://cds.cern.ch/record/2037683.

[58] A. Andronic and J. P. Wessels. Transition Radiation Detectors. *Nucl. Instrum. Meth.*, A666:130–147, 2012. DOI: 10.1016/j.nima.2011.09.041. arXiv: 1111.4188 [physics.ins-det].

[59] ATLAS Collaboration. Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties. *JINST*, 12:P05002, 2017. DOI: 10.1088/1748-0221/12/05/P05002. arXiv: 1702.06473 [hep-ex].

[60] M. Aharrouche et al. Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam. *Nucl. Instrum. Meth.*, A568:601–623, 2006. DOI: 10.1016/j.nima.2006.07.053. arXiv: physics/0608012 [physics].

[61] Pavol Strizenec and Andrey Minaenko. Performance of the ATLAS liquid argon endcap calorimeter in beam tests. *J. Phys. Conf. Ser.*, 160:012078, 2009. DOI: 10.1088/1742-6596/160/1/012078.

[62] C. Cojocaru et al. Hadronic calibration of the ATLAS liquid argon end-cap calorimeter in the pseudorapidity region $1.6 < |\eta| < 1.8$ in beam tests. *Nucl. Instrum. Meth.*, A531:481–514, 2004. DOI: 10.1016/j.nima.2004.05.133. arXiv: physics/0407009 [physics].

[63] P. Adragna et al. Testbeam studies of production modules of the ATLAS tile calorimeter. *Nucl. Instrum. Meth.*, A606:362–394, 2009. DOI: 10.1016/j.nima.2009.04.009.

[64] N. Amram et al. Streamlined Calibrations of the ATLAS Precision Muon Chambers for Initial LHC Running. *Nucl. Instrum. Meth.*, A671:40–50, 2012. DOI: 10.1016/j.nima.2011.12.086. arXiv: 1103.0797 [physics.ins-det].

[65] ATLAS muon spectrometer: Technical design report, 1997.

[66] Imma Riu. The ATLAS Level-1 Topological Trigger performance in Run 2. *J. Phys. Conf. Ser.*, 898(3):032037, 2017. DOI: 10.1088/1742-6596/898/3/032037.

[67] H. Bertelsen et al. Operation of the upgraded ATLAS Central Trigger Processor during the LHC Run 2. *JINST*, 11(02):C02020, 2016. DOI: 10.1088/1748-0221/11/02/C02020.

[68] ATLAS Collaboration. Performance of the ATLAS Trigger System in 2015. *Eur. Phys. J. C*, 77:317, 2017. DOI: 10.1140/epjc/s10052-017-4852-3. arXiv: 1611.09661 [hep-ex].

[69] G. Avoni et al. The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS. *JINST*, 13(07):P07017, 2018. DOI: 10.1088/1748-0221/13/07/P07017.

[70] S. Abdel Khalek et al. The ALFA Roman Pot Detectors of ATLAS. *JINST*, 11(11):P11013, 2016. DOI: 10.1088/1748-0221/11/11/P11013. arXiv: 1609.00249 [physics.ins-det].

[71] ATLAS Collaboration. Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2. *Eur. Phys. J. C*, 77:673, 2017. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex].

[72] T. Cornelissen, M. Elsing, I. Gavrilenko, W. Liebig, E. Moyse, and Salzburger, A. The new ATLAS track reconstruction (NEWT). *J. Phys. Conf. Ser.*, 119:032014, 2008. DOI: 10.1088/1742-6596/119/3/032014.

[73] ATLAS Collaboration. Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC. *Eur. Phys. J. C*, 77:332, 2017. DOI: 10.1140/epjc/s10052-017-4887-5. arXiv: 1611.10235 [hep-ex].

[74] ATLAS Collaboration. Number of tracks vs. mu with full 2016 data. https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2016-015/. Accessed: 2019-05-27.

[75] ATLAS Collaboration. Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13\,\text{TeV}$. *Eur. Phys. J. C*, 76:292, 2016. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex].

[76] ATLAS Collaboration. Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton–proton collision data. *Eur. Phys. J. C*, 74:3130, 2014. DOI: 10.1140/epjc/s10052-014-3130-x. arXiv: 1407.3935 [hep-ex].

[77] G. Aad et al. Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics, 2009. arXiv: 0901.0512 [hep-ex].

[78] ATLAS Collaboration. Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$. *Eur. Phys. J. C*, 78:903, 2018. DOI: 10.1140/epjc/s10052-018-6288-9. arXiv: 1802.08168 [hep-ex].

[79] Walter Lampl et al. Calorimeter Clustering Algorithms: Description and Performance. ATL-LARG-PUB-2008-002. 2008. URL: https://cds.cern.ch/record/1099735.

[80] ATLAS Collaboration. Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13\,\text{TeV}$. *Eur. Phys. J.*, 2019. arXiv: 1902.04655 [hep-ex].

[81] ATLAS Collaboration. Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung. ATLAS-CONF-2012-047. 2012. URL: https://cds.cern.ch/record/1449796.

[82] ATLAS Collaboration. Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data. *Eur. Phys. J. C*, 77:195, 2017. DOI: 10.1140/epjc/s10052-017-4756-2. arXiv: 1612.01456 [hep-ex].

[83] David Krohn, Jesse Thaler, and Lian-Tao Wang. Jets with Variable R. *JHEP*, 06:059, 2009. DOI: 10.1088/1126-6708/2009/06/059. arXiv: 0903.0392 [hep-ph].

[84] S. Catani, Yuri L. Dokshitzer, M. H. Seymour, and B. R. Webber. Longitudinally invariant K_t clustering algorithms for hadron hadron collisions. *Nucl. Phys.*, B406:187–224, 1993. DOI: 10.1016/0550-3213(93)90166-M.

[85] Yuri L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber. Better jet clustering algorithms. *JHEP*, 08:001, 1997. DOI: 10.1088/1126-6708/1997/08/001. arXiv: hep-ph/9707323 [hep-ph].

[86] M. Wobisch and T. Wengler. Hadronization corrections to jet cross-sections in deep inelastic scattering. In *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*, pages 270–279, 1998. arXiv: hep-ph/9907280 [hep-ph].

[87] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti-k_t jet clustering algorithm. *JHEP*, 04:063, 2008. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[88] ATLAS Collaboration. Tagging and suppression of pileup jets with the ATLAS detector. ATLAS-CONF-2014-018. 2014. URL: https://cds.cern.ch/record/1700870.

[89] ATLAS Collaboration. Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8\,\text{TeV}$ using the ATLAS detector. *Eur. Phys. J. C*, 76:581, 2016. DOI: 10.1140/epjc/s10052-016-4395-z. arXiv: 1510.03823 [hep-ex].

[90] Giacinto Piacquadio and Christian Weiser. A new inclusive secondary vertex algorithm for b-jet tagging in ATLAS. *J. Phys. Conf. Ser.*, 119:032032, 2008. DOI: 10.1088/1742-6596/119/3/032032.

[91] ATLAS Collaboration. Measurements of b-jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events at $\sqrt{s} = 13\,\text{TeV}$. *JHEP*, 08:089, 2018. DOI: 10.1007/JHEP08(2018)089. arXiv: 1805.01845 [hep-ex].

[92] Michal Czakon, Paul Fiedler, and Alexander Mitov. Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_S^4)$. *Phys. Rev. Lett.*, 110:252004, 2013. DOI: 10.1103/PhysRevLett.110.252004. arXiv: 1303.6254 [hep-ph].

[93] Michal Czakon and Alexander Mitov. Top++: A program for the calculation of the top-pair cross-section at hadron colliders. *Comput. Phys. Commun.*, 185:2930, 2014. DOI: 10.1016/j.cpc.2014.06.021. arXiv: 1112.5675 [hep-ph].

[94] C. Anastasiou, L. J. Dixon, K. Melnikov, and F. Petriello. High precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at NNLO. *Phys. Rev. D*, 69:094008, 2004. DOI: 10.1103/PhysRevD.69.094008. arXiv: hep-ph/0312266 [hep-ph].

[95] Nikolaos Kidonakis. Next-to-next-to-leading-order collinear and soft gluon corrections for t-channel single top quark production. *Phys. Rev. D*, 83:091503, 2011. DOI: 10.1103/PhysRevD.83.091503. arXiv: 1103.2792 [hep-ph].

[96] Nikolaos Kidonakis. Two-loop soft anomalous dimensions for single top quark associated production with a W- or H-. *Phys. Rev. D*, 82:054018, 2010. DOI: 10.1103/PhysRevD.82.054018. arXiv: 1005.4451 [hep-ph].

[97] John Campbell, R. Keith Ellis, and Raoul Röntsch. Single top production in association with a Z boson at the LHC. *Phys. Rev.*, D87:114006, 2013. DOI: 10.1103/PhysRevD.87.114006. arXiv: 1302.3856 [hep-ph].

[98] A. Lazopoulos, T. McElmurry, K. Melnikov, and F. Petriello. Next-to-leading order QCD corrections to $t\bar{t}Z$ production at the LHC. *Phys. Lett. B*, 666:62–65, 2008. DOI: 10.1016/j.physletb.2008.06.073. arXiv: 0804.2220 [hep-ph].

[99] J. M. Campbell and R. K. Ellis. $t\bar{t}W$ production and decay at NLO. *JHEP*, 07:052, 2012. DOI: 10.1007/JHEP07(2012)052. arXiv: 1204.5678 [hep-ph].

[100] Stefan Höche, Frank Krauss, Steffen Schumann, and Frank Siegert. NLO matrix elements and truncated showers. *JHEP*, 08:123, 2011. DOI: 10.1088/1126-6708/2009/05/053. arXiv: 1009.1127 [hep-ph].

[101] John M. Campbell, R. Keith Ellis, and Ciaran Williams. Vector boson pair production at the LHC. *JHEP*, 07:018, 2011. DOI: 10.1007/JHEP07(2011)018. arXiv: 1105.0020 [hep-ph].

[102] C. Anastasiou, C. Duhr, F. Dulat, E. Furlan, T. Gehrmann, F. Herzog, A. Lazopoulos, and B. Mistlberger. High precision determination of the gluon fusion Higgs boson cross-section at the LHC. *JHEP*, 05:058, 2016. DOI: 10.1007/JHEP05(2016)058. arXiv: 1602.00695 [hep-ph].

[103] Charalampos Anastasiou, Radja Boughezal, and Frank Petriello. Mixed QCD-electroweak corrections to Higgs boson production in gluon fusion. *JHEP*, 04:003, 2009. DOI: 10.1088/1126-6708/2009/04/003. arXiv: 0811.3458 [hep-ph].

[104] Charalampos Anastasiou, Claude Duhr, Falko Dulat, Franz Herzog, and Bernhard Mistlberger. Higgs Boson Gluon-Fusion Production in QCD at Three Loops. *Phys. Rev. Lett.*, 114:212001, 2015. DOI: 10.1103/PhysRevLett.114.212001. arXiv: 1503.06056 [hep-ph].

[105] J R Andersen et al. Handbook of LHC Higgs Cross Sections: 3. Higgs Properties, 2013. S Heinemeyer, C Mariotti, G Passarino, and R Tanaka, editors. DOI: 10.5170/CERN-2013-004. arXiv: 1307.1347 [hep-ph].

[106] M. Ciccolini, A. Denner, S. Dittmaier. Strong and Electroweak Corrections to the Production of a Higgs Boson + 2 Jets via Weak Interactions at the Large Hadron Collider. *Phys. Rev. Lett.*, 99:161803, 2007. DOI: 10.1103/PhysRevLett.99.161803. arXiv: 0707.0381 [hep-ph].

[107] M. Ciccolini, A. Denner, S. Dittmaier. Electroweak and QCD corrections to Higgs production via vector-boson fusion at the LHC. *Phys. Rev. D*, 77:013002, 2008. DOI: 10.1103/PhysRevD.77.013002. arXiv: 0710.4749 [hep-ph].

[108] P. Bolzoni, F. Maltoni, S.-O. Moch and M. Zaro. Higgs Boson Production via Vector-Boson Fusion at Next-to-Next-to-Leading Order in QCD. *Phys. Rev. Lett.*, 105:011801, 2010. DOI: 10.1103/PhysRevLett.105.011801. arXiv: 1003.4451 [hep-ph].

[109] Robert V. Harlander, Stefan Liebler, and Tom Zirke. Higgs Strahlung at the Large Hadron Collider in the 2-Higgs-Doublet Model. *JHEP*, 02:023, 2014. DOI: 10.1007/JHEP02(2014)023. arXiv: 1307.8122 [hep-ph].

[110] Oliver Brein, Robert V. Harlander, and Tom J. E. Zirke. vh@nnlo - Higgs Strahlung at hadron colliders. *Comput. Phys. Commun.*, 184:998–1003, 2013. DOI: 10.1016/j.cpc.2012.11.002. arXiv: 1210.5347 [hep-ph].

[111] S. Dittmaier et al. Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables, 2011. DOI: 10.5170/CERN-2011-002. arXiv: 1101.0593 [hep-ph].

[112] S. Dittmaier et al. Handbook of LHC Higgs Cross Sections: 2. Differential Distributions, 2012. DOI: 10.5170/CERN-2012-002. arXiv: 1201.3084 [hep-ph].

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

[114] ATLAS Collaboration. Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7\,\text{TeV}$ with the ATLAS detector. *JHEP*, 09:145, 2014. DOI: 10.1007/JHEP09(2014)145. arXiv: 1406.3660 [hep-ex].

[115] Torbjorn Sjöstrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 physics and manual. *JHEP*, 05:026, 2006. DOI: 10.1088/1126-6708/2006/05/026. arXiv: 0603175 [hep-ph].

[116] Peter Zeiler Skands. Tuning Monte Carlo generators: The Perugia tunes. *Phys. Rev. D*, 82:074018, 2010. DOI: 10.1103/PhysRevD.82.074018. arXiv: 1005.3457 [hep-ph].

[117] ATLAS Collaboration. Comparison of Monte Carlo generator predictions for gap fraction and jet multiplicity observables in $t\bar{t}$ events. ATL-PHYS-PUB-2014-005. 2014. URL: https://cds.cern.ch/record/1703034.

[118] ATLAS Collaboration. ATLAS simulation of boson plus jets processes in Run 2. ATL-PHYS-PUB-2017-006. 2017. URL: https://cds.cern.ch/record/2261937.

[119] ATLAS Collaboration. The ATLAS Simulation Infrastructure. *Eur. Phys. J. C*, 70:823, 2010. DOI: 10.1140/epjc/s10052-010-1429-9. arXiv: 1005.4568 [physics.ins-det].

[120] S. Agostinelli et al. GEANT4 - a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003. DOI: 10.1016/S0168-9002(03)01368-8.

[121] ATLAS Collaboration. Summary of ATLAS Pythia 8 tunes. ATL-PHYS-PUB-2012-003. 2012. URL: https://cds.cern.ch/record/1474107.

[122] A.D. Martin, W.J. Stirling, R.S. Thorne, and G. Watt. Parton distributions for the LHC. *Eur. Phys. J. C*, 63:189–285, 2009. DOI: 10.1140/epjc/s10052-009-1072-5. arXiv: 0901.0002 [hep-ph].

[123] Morad Aaboud et al. Performance of the ATLAS Trigger System in 2015. *Eur. Phys. J.*, C77(5):317, 2017. DOI: 10.1140/epjc/s10052-017-4852-3. arXiv: 1611.09661 [hep-ex].

[124] ATLAS Collaboration. Trigger Menu in 2016. ATL-DAQ-PUB-2017-001. 2017. URL: https://cds.cern.ch/record/2242069.

[125] ATLAS Collaboration. Muon reconstruction performance in early $\sqrt{s} = 13$ TeV data. ATL-PHYS-PUB-2015-037. 2015. URL: https://cds.cern.ch/record/2047831.

[126] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet User Manual. *Eur. Phys. J. C*, 72:1896, 2012. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[127] ATLAS Collaboration. Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 96:072002, 2017. DOI: 10.1103/PhysRevD.96.072002. arXiv: 1703.09665 [hep-ex].

[128] ATLAS Collaboration. Selection of jets produced in 13 TeV proton–proton collisions with the ATLAS detector. ATLAS-CONF-2015-029. 2015. URL: https://cds.cern.ch/record/2037702.

[129] Johannes Erdmann, Stefan Guindon, Kevin Kroeninger, Boris Lemmer, Olaf Nackenhorst, Arnulf Quadt, and Philipp Stolte. A likelihood-based reconstruction algorithm for top-quark pairs and the KLFitter framework. *Nucl. Instrum. Meth. A*, 748:18–25, 2014. DOI: 10.1016/j.nima.2014.02.029. arXiv: 1312.5595 [hep-ex].

[130] ATLAS Collaboration. Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C*, 76:653, 2016. DOI: 10.1140/epjc/s10052-016-4466-1. arXiv: 1608.03953 [hep-ex].

[131] ATLAS Collaboration. Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton–proton collision data. *JINST*, 2018. arXiv: 1812.03848 [hep-ex].

[132] ATLAS Collaboration. Deep generative models for fast shower simulation in ATLAS. ATL-SOFT-PUB-2018-001. 2018. URL: https://cds.cern.ch/record/2630433.

[133] ATLAS Collaboration. Monte Carlo Calibration and Combination of In-situ Measurements of Jet Energy Scale, Jet Energy Resolution and Jet Mass in ATLAS. ATLAS-CONF-2015-037. 2015. URL: https://cds.cern.ch/record/2044941.

[134] Jon Butterworth et al. PDF4LHC recommendations for LHC Run II. *J. Phys. G*, 43:023001, 2016. DOI: 10.1088/0954-3899/43/2/023001. arXiv: 1510.03865 [hep-ph].

[135] Andy Buckley, James Ferrando, Stephen Lloyd, Karl Nordström, Ben Page, Martin Rüfenacht, Marek Schönherr, and Graeme Watt. LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J.*, C75:132, 2015. DOI: 10.1140/epjc/s10052-015-3318-8. arXiv: 1412.7420 [hep-ph].

[136] LHC Higgs Cross Section Working Group. Handbook of LHC Higgs Cross Sections: 3. Higgs Properties, 2013. DOI: 10.5170/CERN-2013-004. arXiv: 1307.1347 [hep-ph].

[137] ATLAS Collaboration. Multi-Boson Simulation for 13 TeV ATLAS Analyses. ATL-PHYS-PUB-2017-005. 2017. URL: https://cds.cern.ch/record/2261933.

[138] M. Baak, G. J. Besjes, D. Côte, A. Koutsman, J. Lorenz, and D. Short. HistFitter software framework for statistical data analysis. *Eur. Phys. J.*, C75:153, 2015. DOI: 10.1140/epjc/s10052-015-3327-7. arXiv: 1410.1280 [hep-ex].

[139] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J.*, C71:1554, 2011. DOI: 10.1140/epjc/s10052-011-1554-0,10.1140/epjc/s10052-013-2501-z. arXiv: 1007.1727 [physics.data-an]. [Erratum: Eur. Phys. J.C73,2501(2013)].

[140] A. Wald. Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large. *Transactions of the American Mathematical Society*, 54(3):426–482, November 1943.

[141] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *Annals Math. Statist.*, 9(1):60–62, 1938. DOI: 10.1214/aoms/1177732360.

[142] Alexander L. Read. Presentation of search results: the CL_s technique. *J. Phys. G*, 28:2693–2704, 2002. DOI: 10.1088/0954-3899/28/10/313.

[143] Zhi Zheng. Identification of very-low transverse momentum muons with the ATLAS experiment. Technical report ATL-MUON-PROC-2018-018, CERN, Geneva, November 2018. URL: https://cds.cern.ch/record/2649299.

[144] CMS Collaboration. Search for an exotic decay of the Higgs boson to a pair of light pseudoscalars in the final state with two b quarks and two τ leptons in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$. *Phys. Lett. B*, 785:462, 2018. DOI: 10.1016/j.physletb.2018.08.057. arXiv: 1805.10191 [hep-ex].

[145] ATLAS Collaboration. Search for the Higgs boson produced in association with a vector boson and decaying into two spin-zero particles in the $H \to aa \to 4b$ channel in pp collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector. *JHEP*, 10:031, 2018. DOI: 10.1007/JHEP10(2018)031. arXiv: 1806.07355 [hep-ex].

[146] ATLAS Collaboration. The ATLAS Tau Trigger in Run 2. ATLAS-CONF-2017-061. 2017. URL: https://cds.cern.ch/record/2274201.

A Muon Isolation Maps

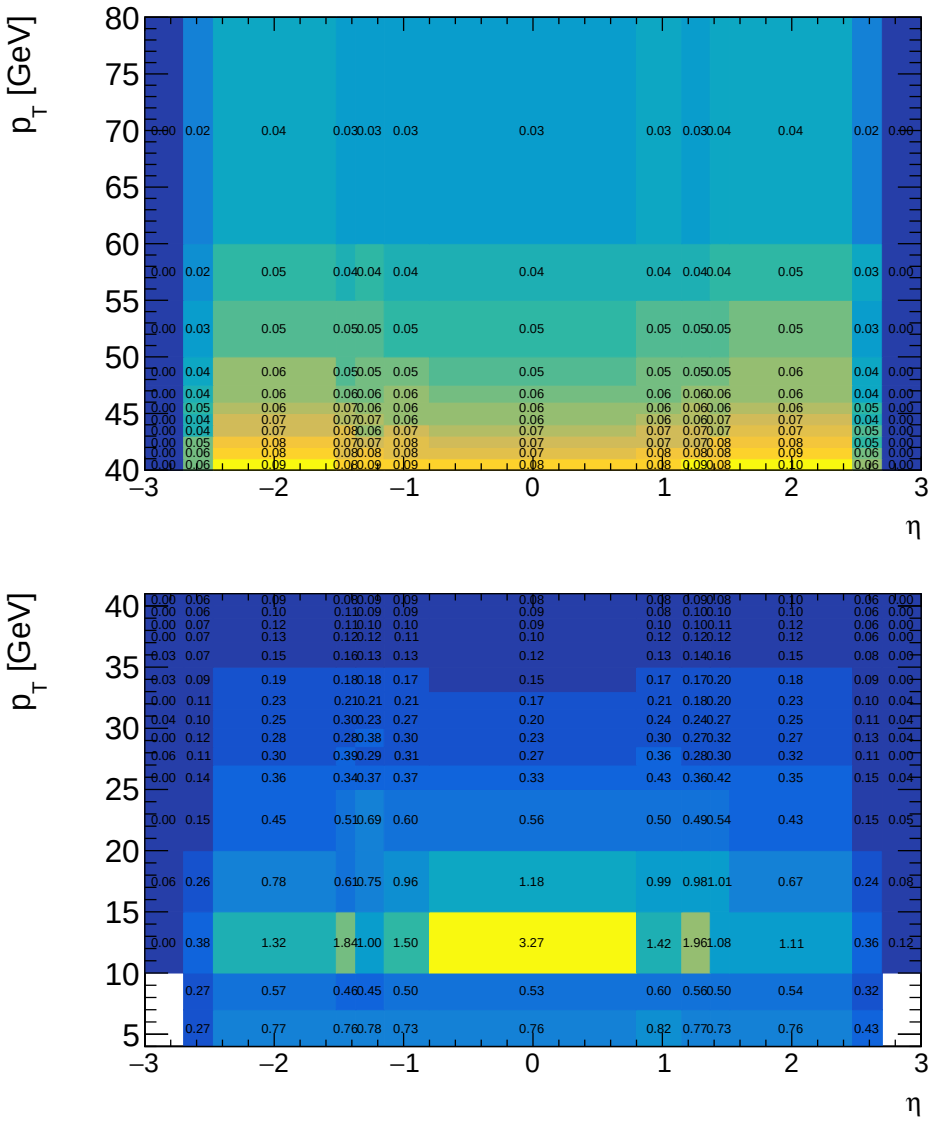

Figure 80: Muon track isolation maps of the *ptvarcone30*/p_{T} variable as a function of p_{T} and η for the 99% efficiency working point corresponding to the "Tight" isolation selections. The maps are separated into displays of $5 < p_{\mathrm{T}} < 40$ GeV and $40 < p_{\mathrm{T}} < 80$ GeV to illustrate the granularity of the binning.

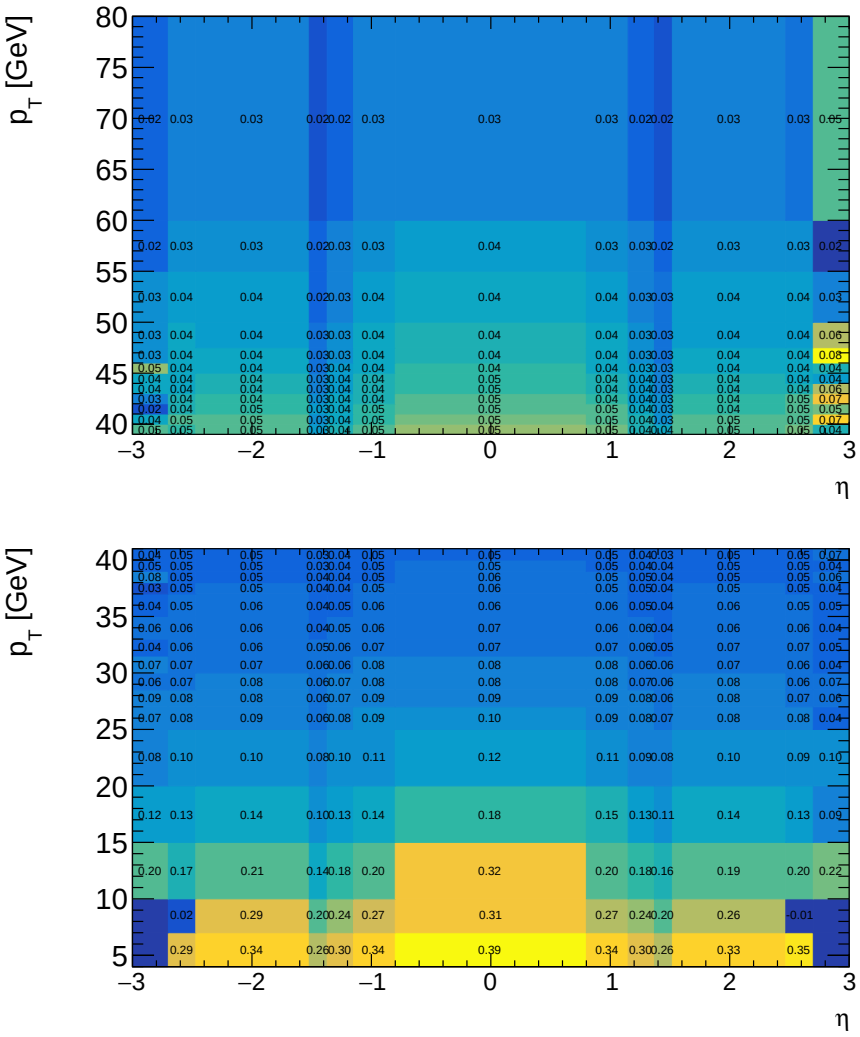

Figure 81: Muon calorimeter isolation maps of the *ettopocone20*/p_{T} as a function of p_{T} and η for the 95% efficiency working point corresponding to the "Tight" isolation selections. The maps are separated into displays of $5 < p_{\mathrm{T}} < 40$ GeV and $40 < p_{\mathrm{T}} < 80$ GeV to illustrate the granularity of the binning.

B Kinematic Likelihood Fits

The likelihood used in the analysis described in Section 4.6.2 has several terms that allow for the $m_{bb} \approx m_{\mu\mu}$ constraint to be applied. Each of the b-jets is subject to a transfer function W to modify the b-jet energy to meet the mass constraint. A transfer function encodes the conditional probability that a given reconstructed value x_{reco} corresponds to a true value x_{truth}. The reconstructed values in the case of the $bb\mu\mu$ analysis are the energy of the b-jets, which are adjusted in an attempt to meet the $m_{bb} \approx m_{\mu\mu}$ constraints.

The transfer functions are normalized such that

$$\int dx_{\text{reco}} W(x_{\text{truth}}|x_{\text{reco}}) = 1 \tag{106}$$

which ensures that there exists a response for every truth energy, outside of any losses due to acceptance or efficiency biases. The W for b-jets are parametrized as follows

$$W(\hat{E}_b, E_b) = \frac{1}{\sqrt{2\pi}(p_2 + p_3 p_5)} \left(e^{-\frac{(\Delta E - p_1)^2}{2p_2^2}} + p_3 e^{-\frac{(\Delta E - p_4)^2}{2p_5^2}} \right), \quad \Delta E = \frac{\hat{E}_b - E_b}{\hat{E}_b}, \tag{107}$$

where the parameters p_i are functions of the truth energy $\hat{E}$,

$$\begin{aligned} p_1 &= a_1 + b_1/\hat{E}, \\ p_2 &= a_2 + b_2/\sqrt{\hat{E}}, \\ p_3 &= a_3 + b_3/\hat{E}, \\ p_4 &= a_4 + b_4/\sqrt{\hat{E}}, \\ p_5 &= a_5 + b_5\hat{E} \end{aligned} \tag{108}$$

with parameters a_i and b_i obtained through fits to simulated events. The exact values of a_i and b_i used by the KLFitter package can be found in Table 38. The double Gaussian nature of the transfer function accomplishes two different types of energy corrections. The first Gaussian term uses parameters that mimic the jet resolution functions and aim to allow the jet energy to float within the experimental resolution for a given energy range. The second Gaussian term aims to correct for "lost-energy" in the jet, either from energy mis-measurement or from the energy imparted to neutrinos in the semi-leptonic decays of B and C hadrons. These functional values were derived by using $t\bar{t}$ MC simulation events to measure the response, $\Delta E/E_{\text{true}}$

Parameter	0.0 - 0.8	0.8 - 1.37	1.37 - 1.52	1.52 - 2.50
a_1	-0.0020963	-0.0126548	-0.00184705	0.0109739
b_1	3.253	5.34751	6.98626	6.02177
a_2	0.054609	0.0549931	0.0561314	0.0424557
b_2	0.461499	0.549482	0.687561	0.812251
a_3	0.179215	-0.140537	-0.296518	-0.121939
b_3	52.1964	102.086	174.964	217.198
a_4	0.288286	0.386519	0.460882	0.421104
b_4	-1.73631	-2.89377	-4.09145	-4.26916
a_5	0.234404	0.222311	0.213336	0.213189
b_5	-2.54762e-05	7.51057e-05	0.000103354	1.20014e-05

Table 38: The parameters used to form the KLFit transfer functions for b-jet in each $|\eta|$ bin used in this analysis. The last η bin is dropped, due to being outside of the ID tracking volume.

between the reconstructed jet and truth jet scales. An example fit for the $0.0 < |\eta| < 0.8$

The transfer functions for b-jets are binned as a function of $|\eta|$, with the bins [0-0.8,0.8-1.37,1.37-1.52,1.52 - 2.50]. These regions are motivated by the detector volume, with the far edge being the edge of acceptance of the ID tracker and the third bin parametrizing the barrel end-cap transition region in the calorimeter that cannot be instrumented with high granularity due services necessary for powering and cooling the ID tracker and the LAr barrel calorimeter.

An alternative to the above transfer functions is to use a Gaussian constraint with $\sigma_E = \sqrt{E}$ of the jet in the fit. In this version of the fit, the allowed energy range is set to be 10σ from the central value. While this energy range is quite large, the phase space is biased as large deviations in the jet energy transfer cause a large, negative value in the final fit result.

The Breit-Wigner function used in this analysis is defined as

$$\mathrm{BW}(m_{bb}, m_{\mu\mu}) = \frac{k}{(m_{bb}^2 - m_{\mu\mu}^2)^2 + m_{\mu\mu}^2\Gamma^2} \tag{109}$$

where k defines the normalization factor

$$k = \frac{2\sqrt{2}m_{\mu\mu}\Gamma\sqrt{m_{\mu\mu}^2(m_{\mu\mu}^2+\Gamma^2)}}{\pi\sqrt{m_{\mu\mu}^2+\sqrt{m_{\mu\mu}^2(m_{\mu\mu}^2+\Gamma^2)}}} \tag{110}$$

and Γ is the width of the Breit-Wigner function. The value of $\Gamma = 0.5$ GeV was chosen to be near the experimental resolution, as the a decay is assumed to be a narrow resonance and has negligible width compared to the experimental resolution.

C Signal Cutflows for the H→bb$\mu\mu$ Analysis

Selection Criterion	MC Events	Event Yield	Eff./Step [%]	Tot. Eff. [%]
$N_\mu = 2$ & $N_{jets} > 0$	40373	1152.721	1.000	1.000
Pass Trigger	28561	803.407	0.707	0.707
$p_T^{\mu_1} > 27$, $p_T^{\mu_2} > 7$	28047	787.334	0.982	0.695
$16 < M_{\mu\mu} < 64$	27799	780.636	0.991	0.689
OS Muons	27728	778.955	0.997	0.687
$N_{b-jets} = 2$	2350	70.834	0.085	0.058
MET < 60	2280	68.567	0.970	0.056
$\mid M^{KL}_{bb\mu\mu}$ - $M_h \mid \leq 15$	1654	48.276	0.725	0.041
Log(L) > -8	1380	39.515	0.834	0.034

Table 39: Analysis cutflow for $m_a = 20$ GeV, with event yields weighted to Br(H→aa)=100% and the BR(a→bb)BR(a→ $\mu\mu$) = 1.6×10^{-3} taken from the Type-III 2HDM+S model described in Ref. [33].

Selection Criterion	MC Events	Event Yield	Eff./Step [%]	Tot. Eff. [%]
$N_\mu = 2$ & $N_{jets} > 0$	46124.000	1312.638	1.000	1.000
Pass Trigger	31142.000	868.552	0.675	0.675
$p_T^{\mu_1} > 27$, $p_T^{\mu_2} > 7$	30427.000	846.407	0.977	0.660
$16 < M_{\mu\mu} < 64$	30412.000	845.947	1.000	0.659
OS Muons	30313.000	843.268	0.997	0.657
$N_{b-jets} = 2$	2396.000	71.147	0.079	0.052
MET < 60	2315.000	68.898	0.966	0.050
$\mid M^{KL}_{bb\mu\mu}$ - $M_h \mid \leq 15$	1848.000	52.161	0.798	0.040
$Ln(L^{max}) > -8$	1723.000	48.399	0.932	0.037

Table 40: Analysis cutflow for $m_a = 40$ GeV, with event yields weighted to Br(H→aa)=100% and the BR(a→bb)BR(a→ $\mu\mu$) = 1.6×10^{-3} taken from the Type-III 2HDM+S model described in Ref. [33].

Selection Criterion	MC Events	Event Yield	Eff./Step [%]	Tot. Eff. [%]
$N_\mu = 2$ & $N_{jets} > 0$	46472.000	1321.171	1.000	1.000
Pass Trigger	31737.000	883.172	0.683	0.683
$p_T^{\mu_1} > 27$, $p_T^{\mu_2} > 7$	31053.000	861.860	0.978	0.668
$16 < M_{\mu\mu} < 64$	31047.000	861.727	1.000	0.668
OS Muons	30968.000	859.610	0.997	0.666
$N_{b-jets} = 2$	2569.000	75.705	0.083	0.055
MET < 60	2499.000	73.687	0.973	0.054
$\mid M_{bb\mu\mu}^{KL} - M_h \mid \leq 15$	2033.000	57.795	0.814	0.044
Log(L) > -8	1899.000	53.203	0.934	0.041

Table 41: Analysis cutflow for $m_a = 50$ GeV, with event yields weighted to Br(H→aa)=100% and the BR(a→bb)BR(a→ $\mu\mu$) = 1.6×10^{-3} taken from the Type-III 2HDM+S model described in Ref. [33].

Selection Criterion	MC Events	Event Yield	Eff./Step [%]	Tot. Eff. [%]
$N_\mu = 2$ & $N_{jets} > 0$	43524	1240.568	1.000	1.000
Pass Trigger	30100	839.767	0.692	0.692
$p_T^{\mu_1} > 27$, $p_T^{\mu_2} > 7$	29484	819.587	0.980	0.677
$16 < M_{\mu\mu} < 64$	29403	817.485	0.997	0.676
OS Muons	29317	814.997	0.997	0.674
$N_{b-jets} = 2$	3390	99.184	0.116	0.078
MET < 60	3307	96.421	0.976	0.076
$\mid M_{bb\mu\mu}^{KL} - M_h \mid \leq 15$	2902	82.845	0.878	0.067
Log(L) > -8	2744	77.817	0.946	0.063

Table 42: Analysis cutflow for $m_a = 60$ GeV, with event yields weighted to Br(H→aa)=100% and the BR(a→bb)BR(a→ $\mu\mu$) = 1.6×10^{-3} taken from the Type-III 2HDM+S model described in Ref. [33].

D H→bb$\mu\mu$ Analysis Signal Regions Uncertainties

Uncertainty of channel	SR20_a20	SR20_DDZ	SR20_ttbar
Total background expectation	1.13	3.89	0.96
Total background systematic	±1.63 [144.64%]	±0.93 [23.83%]	±0.32 [33.28%]
mu_SIG	±1.65 [147.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_btag_B0	±0.21 [19.0%]	±0.00 [0.00%]	±0.03 [3.2%]
alpha_JER	±0.09 [8.0%]	±0.00 [0.00%]	±0.14 [14.2%]
alpha_sig_RenFact	±0.07 [6.0%]	±0.00 [0.00%]	±0.00 [0.00%]
gamma_stat_SR20_cuts_bin_0	±0.05 [4.3%]	±0.17 [4.3%]	±0.04 [4.3%]
alpha_sig_ISR	±0.05 [4.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_JES_SR_NP1	±0.04 [4.0%]	±0.00 [0.00%]	±0.11 [11.7%]
alpha_sig_ggFXsec	±0.04 [3.8%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_VH	±0.04 [3.5%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_PDFalphaS	±0.03 [3.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.03 [2.5%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_ISO_SYS	±0.01 [1.2%]	±0.00 [0.00%]	±0.00 [0.06%]
alpha_MU_MS	±0.01 [0.49%]	±0.00 [0.00%]	±0.05 [5.4%]
alpha_JES_SR_NP2	±0.01 [0.47%]	±0.00 [0.00%]	±0.03 [3.0%]
alpha_JES_SR_NP3	±0.00 [0.44%]	±0.00 [0.00%]	±0.01 [1.1%]
alpha_MU_ID	±0.00 [0.11%]	±0.00 [0.00%]	±0.09 [9.7%]
alpha_JES_SR_NP4	±0.00 [0.11%]	±0.00 [0.00%]	±0.03 [3.5%]
mu_Top	±0.00 [0.00%]	±0.00 [0.00%]	±0.09 [9.0%]
alpha_DDZ_subtraction	±0.00 [0.00%]	±0.12 [3.0%]	±0.00 [0.00%]
alpha_DDZ_shape	±0.00 [0.00%]	±0.54 [14.0%]	±0.00 [0.00%]
alpha_top_PDF	±0.00 [0.00%]	±0.00 [0.00%]	±0.01 [1.0%]
alpha_top_Rad	±0.00 [0.00%]	±0.00 [0.00%]	±0.05 [5.0%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±0.00 [0.00%]	±0.17 [18.0%]
alpha_top_pythHerw	±0.00 [0.00%]	±0.00 [0.00%]	±0.15 [16.0%]
mu_DDZ	±0.00 [0.00%]	±0.73 [18.8%]	±0.00 [0.00%]

Table 43: Breakdown of the systematic uncertainties on background estimates in the $m_a = 20$ GeV signal region. The diboson, single-top, and tt+V uncertainties are excluded due to negligible signal region yields. All uncertainties with less than 1% effect on any background yield have been truncated. The uncertainties may be correlated and the quadratic summation of individual uncertainties may not equate to the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

Uncertainty of channel	SR30_a30	SR30_DDZ	SR30_ttbar	SR30_st
Total background expectation	0.00	4.33	2.97	0.19
Total background systematic	±0.47 [2410429.20%]	±1.00 [23.15%]	±0.81 [27.41%]	±0.19 [96.62%]
mu_SIG	±0.47 [2410429.2%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_btag_B0	±0.00 [17.3%]	±0.00 [0.00%]	±0.04 [1.4%]	±0.00 [0.00%]
alpha_sig_RenFact	±0.00 [6.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_JER	±0.00 [5.3%]	±0.00 [0.00%]	±0.07 [2.5%]	±0.09 [47.6%]
alpha_JES_SR_NP1	±0.00 [5.3%]	±0.00 [0.00%]	±0.02 [0.67%]	±0.04 [20.0%]
alpha_sig_ISR	±0.00 [4.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
gamma_stat_SR30_bin_0	±0.00 [3.9%]	±0.17 [3.9%]	±0.12 [3.9%]	±0.01 [3.9%]
alpha_sig_ggFXsec	±0.00 [3.6%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_VH	±0.00 [3.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_PDFalphaS	±0.00 [3.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.00 [2.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_ISO_SYS	±0.00 [1.4%]	±0.00 [0.00%]	±0.00 [0.01%]	±0.00 [0.00%]
alpha_MU_ID	±0.00 [0.57%]	±0.00 [0.00%]	±0.20 [6.8%]	±0.05 [26.7%]
alpha_MU_MS	±0.00 [0.52%]	±0.00 [0.00%]	±0.07 [2.3%]	±0.05 [27.1%]
alpha_MET_Scale	±0.00 [0.47%]	±0.00 [0.00%]	±0.08 [2.7%]	±0.05 [25.5%]
alpha_MET_ResPerp	±0.00 [0.36%]	±0.00 [0.00%]	±0.04 [1.2%]	±0.10 [53.0%]
alpha_JES_SR_NP2	±0.00 [0.26%]	±0.00 [0.00%]	±0.02 [0.76%]	±0.05 [24.1%]
alpha_JES_SR_NP3	±0.00 [0.10%]	±0.00 [0.00%]	±0.02 [0.67%]	±0.06 [29.0%]
alpha_JES_SR_NP4	±0.00 [0.05%]	±0.00 [0.00%]	±0.02 [0.69%]	±0.05 [27.1%]
alpha_MU_SagRho	±0.00 [0.02%]	±0.00 [0.00%]	±0.06 [1.9%]	±0.00 [0.00%]
mu_Top	±0.00 [0.00%]	±0.00 [0.00%]	±0.27 [9.0%]	±0.00 [0.00%]
alpha_DDZ_subtraction	±0.00 [0.00%]	±0.13 [2.9%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_DDZ_shape	±0.00 [0.00%]	±0.60 [13.9%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_PDF	±0.00 [0.00%]	±0.00 [0.00%]	±0.11 [3.6%]	±0.00 [0.00%]
alpha_top_Rad	±0.00 [0.00%]	±0.00 [0.00%]	±0.15 [5.0%]	±0.00 [0.00%]
alpha_st_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.01 [5.0%]
alpha_DDZ_reweigting	±0.00 [0.00%]	±0.03 [0.68%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±0.00 [0.00%]	±0.53 [18.0%]	±0.00 [0.00%]
alpha_top_pythHerw	±0.00 [0.00%]	±0.00 [0.00%]	±0.47 [16.0%]	±0.00 [0.00%]
mu_DDZ	±0.00 [0.00%]	±0.83 [19.2%]	±0.00 [0.00%]	±0.00 [0.00%]

Table 44: Breakdown of the systematic uncertainties on background estimates in the $m_a = 30$ GeV signal region. The diboson and tt+V uncertainties are not shown due to negligible signal region yields. All uncertainties with less than 1% effect on any signal or background yield have been truncated. The uncertainties may be correlated and the quadratic summation of individual uncertainties may not equate to the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

Uncertainty of channel	SR40_a40	SR40_DDZ	SR40_ttbar	SR40_db
Total background expectation	2.33	7.10	6.55	0.02
Total background systematic	±3.04 [130.68%]	±1.74 [24.50%]	±1.79 [27.28%]	±0.04 [199.90%]
mu_SIG	±3.12 [134.2%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_btag_B0	±0.39 [16.9%]	±0.00 [0.00%]	±0.20 [3.1%]	±0.00 [0.00%]
alpha_JER	±0.20 [8.8%]	±0.00 [0.00%]	±0.38 [5.8%]	±0.04 [191.7%]
alpha_sig_RenFact	±0.14 [6.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_JES_SR_NP1	±0.12 [5.2%]	±0.00 [0.00%]	±0.35 [5.3%]	±0.01 [47.2%]
gamma_stat_SR40_bin_0	±0.11 [4.6%]	±0.32 [4.6%]	±0.30 [4.6%]	±0.00 [4.6%]
alpha_sig_ISR	±0.09 [4.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_ggFXsec	±0.08 [3.6%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_VH	±0.08 [3.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_PDFalphaS	±0.07 [3.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.06 [2.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_ISO_SYS	±0.03 [1.3%]	±0.00 [0.00%]	±0.00 [0.08%]	±0.00 [0.00%]
alpha_MU_MS	±0.03 [1.2%]	±0.00 [0.00%]	±0.15 [2.3%]	±0.00 [6.1%]
alpha_JES_SR_NP2	±0.02 [0.88%]	±0.00 [0.00%]	±0.06 [0.88%]	±0.00 [4.8%]
alpha_btag_L3	±0.01 [0.46%]	±0.00 [0.00%]	±0.05 [0.77%]	±0.00 [11.9%]
alpha_MET_Scale	±0.01 [0.38%]	±0.00 [0.00%]	±0.10 [1.5%]	±0.00 [1.2%]
alpha_MET_ResPara	±0.01 [0.36%]	±0.00 [0.00%]	±0.06 [0.86%]	±0.00 [19.9%]
alpha_MET_ResPerp	±0.01 [0.35%]	±0.00 [0.00%]	±0.12 [1.9%]	±0.00 [15.6%]
alpha_JES_SR_NP4	±0.01 [0.25%]	±0.00 [0.00%]	±0.02 [0.37%]	±0.00 [5.6%]
alpha_JES_SR_NP3	±0.00 [0.12%]	±0.00 [0.00%]	±0.03 [0.52%]	±0.00 [6.3%]
alpha_MU_SagRes	±0.00 [0.09%]	±0.00 [0.00%]	±0.12 [1.9%]	±0.00 [0.00%]
mu_Top	±0.00 [0.00%]	±0.00 [0.00%]	±0.59 [9.0%]	±0.00 [0.00%]
alpha_DDZ_subtraction	±0.00 [0.00%]	±0.28 [3.9%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_DDZ_shape	±0.00 [0.00%]	±0.99 [14.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_Rad	±0.00 [0.00%]	±0.00 [0.00%]	±0.33 [5.0%]	±0.00 [0.00%]
mu_DDZ	±0.00 [0.00%]	±1.33 [18.8%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_db_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [10.0%]
alpha_DDZ_reweigting	±0.00 [0.00%]	±0.10 [1.4%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±0.00 [0.00%]	±1.18 [18.0%]	±0.00 [0.00%]
alpha_top_pythHerw	±0.00 [0.00%]	±0.00 [0.00%]	±1.05 [16.0%]	±0.00 [0.00%]

Table 45: Breakdown of the systematic uncertainties on background estimates in the $m_a = 40$ GeV signal region. The single-top and tt+V uncertainties are not shown due to negligible signal region yields. All uncertainties with less than 1% effect on any signal or background yield have been truncated. The uncertainties may be correlated and the quadratic summation of individual uncertainties may not equate to the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

Uncertainty	SR60_a60	SR60_DDZ	SR60_ttbar	SR60_db	SR60_st
Tot. background yield	0.00	15.13	14.37	0.33	0.10
Tot. background systematic	±1.55 [2752551.04%]	±3.38 [22.32%]	±3.42 [23.82%]	±0.10 [31.42%]	±0.19 [196.34%]
mu_SIG	±1.55 [2752551.4%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_btag_B0	±0.00 [18.5%]	±0.00 [0.00%]	±0.33 [2.3%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_JER	±0.00 [8.7%]	±0.00 [0.00%]	±0.41 [2.8%]	±0.00 [0.00%]	±0.10 [102.1%]
alpha_sig_RenFact	±0.00 [6.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_JES_SR_NP1	±0.00 [5.7%]	±0.00 [0.00%]	±0.06 [0.39%]	±0.00 [0.00%]	±0.02 [15.7%]
alpha_sig_ISR	±0.00 [4.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_ggFXsec	±0.00 [3.6%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_VH	±0.00 [3.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_PDFalphaS	±0.00 [3.0%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_sig_HpTrew	±0.00 [2.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
gamma_stat_SR60_cuts_bin_0	±0.00 [2.3%]	±0.35 [2.3%]	±0.33 [2.3%]	±0.01 [2.3%]	±0.00 [2.3%]
alpha_JES_SR_NP2	±0.00 [1.1%]	±0.00 [0.00%]	±0.06 [0.40%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_EFF_SYS	±0.00 [0.97%]	±0.00 [0.00%]	±0.00 [0.03%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_MS	±0.00 [0.65%]	±0.00 [0.00%]	±0.02 [0.15%]	±0.02 [5.0%]	±0.06 [63.3%]
alpha_MU_ID	±0.00 [0.52%]	±0.00 [0.00%]	±0.11 [0.74%]	±0.00 [0.00%]	±0.06 [62.5%]
alpha_MET_Scale	±0.00 [0.36%]	±0.00 [0.00%]	±0.07 [0.51%]	±0.00 [1.2%]	±0.06 [61.2%]
alpha_MET_ResPerp	±0.00 [0.34%]	±0.00 [0.00%]	±0.08 [0.55%]	±0.06 [18.4%]	±0.00 [0.00%]
alpha_MET_ResPara	±0.00 [0.27%]	±0.00 [0.00%]	±0.13 [0.91%]	±0.08 [22.7%]	±0.00 [0.00%]
alpha_JES_SR_NP3	±0.00 [0.13%]	±0.00 [0.00%]	±0.04 [0.30%]	±0.00 [0.00%]	±0.06 [64.7%]
alpha_MU_Scale	±0.00 [0.11%]	±0.00 [0.00%]	±0.25 [1.7%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_MU_SagRho	±0.00 [0.03%]	±0.00 [0.00%]	±0.04 [0.28%]	±0.00 [0.00%]	±0.11 [110.5%]
mu_Top	±0.00 [0.00%]	±0.00 [0.00%]	±1.29 [9.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_DDZ_shape	±0.00 [0.00%]	±2.04 [13.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_PDF	±0.00 [0.00%]	±0.00 [0.00%]	±0.29 [2.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_Rad	±0.00 [0.00%]	±0.00 [0.00%]	±0.72 [5.0%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_db_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.03 [10.0%]	±0.00 [0.00%]
alpha_st_Xsec	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [5.0%]
alpha_DDZ_reweigting	±0.00 [0.00%]	±1.02 [6.8%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_PowMCnlo	±0.00 [0.00%]	±0.00 [0.00%]	±2.46 [17.1%]	±0.00 [0.00%]	±0.00 [0.00%]
alpha_top_pythHerw	±0.00 [0.00%]	±0.00 [0.00%]	±2.21 [15.4%]	±0.00 [0.00%]	±0.00 [0.00%]
mu_DDZ	±0.00 [0.00%]	±2.80 [18.5%]	±0.00 [0.00%]	±0.00 [0.00%]	±0.00 [0.00%]

Table 46: Breakdown of the systematic uncertainties on background estimates in the $m_a = 60$ GeV signal region. The tt+V uncertainties are not shown due to negligible signal region yields. All uncertainties with less than 1% effect on any background yield have been truncated. The uncertainties may be correlated and the quadratic summation of individual uncertainties may not equate to the total background uncertainty. The percentages describe the size of the uncertainty relative to the total expectation of a particular background.

E List of Monte Carlo Samples

Process	ATLAS Dataset Names and Identifiers
$gg \rightarrow \mathrm{H} \rightarrow \mathrm{bb}\mu\mu$	mc15_13TeV.341677.PowhegPy8EG_AZNLOCTEQ6L1_ggH_H125_a20a20_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341678.PowhegPy8EG_AZNLOCTEQ6L1_ggH_H125_a30a30_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341679.PowhegPy8EG_AZNLOCTEQ6L1_ggH_H125_a40a40_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341680.PowhegPy8EG_AZNLOCTEQ6L1_ggH_H125_a50a50_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341681.PowhegPy8EG_AZNLOCTEQ6L1_ggH_H125_a60a60_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
$qq \rightarrow \mathrm{H} \rightarrow \mathrm{bb}\mu\mu$	mc15_13TeV.341682.PowhegPy8EG_AZNLOCTEQ6L1_VBF_H125_a20a20_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341683.PowhegPy8EG_AZNLOCTEQ6L1_VBF_H125_a30a30_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341684.PowhegPy8EG_AZNLOCTEQ6L1_VBF_H125_a40a40_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341685.PowhegPy8EG_AZNLOCTEQ6L1_VBF_H125_a50a50_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042
	mc15_13TeV.341686.PowhegPy8EG_AZNLOCTEQ6L1_VBF_H125_a60a60_bbmumu.merge.DAOD_HIGG3D1.e4118_s2608_s2183_r7772_r7676_p3042

Table 47: ATLAS Dataset IDs for the exotic Higgs $H \rightarrow aa \rightarrow bb\mu\mu$ samples used in this analysis.

Process	ATLAS Dataset Names and Identifiers
W + Jets	mc15_13TeV.364156.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364157.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364158.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364159.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364160.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364161.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364162.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364163.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364164.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364165.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364166.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364167.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364168.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364169.Sherpa_221_NNPDF30NNLO_Wmunu_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364170.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364171.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364172.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364173.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364174.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364175.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364176.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364177.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364178.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364179.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364180.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364181.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364182.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364183.Sherpa_221_NNPDF30NNLO_Wenu_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364184.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364185.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364186.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364187.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364188.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364189.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364190.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364191.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364192.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364193.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364194.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364195.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364196.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
	mc15_13TeV.364197.Sherpa_221_NNPDF30NNLO_Wtaunu_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5340_s2726_r7772_r7676_p3042
Wt	mc15_13TeV.410015.PowhegPythiaEvtGen_P2012_Wt_dilepton_top.merge.DAOD_HIGG3D1.e3753_s2608_s2183_r7725_r7676_p3042
	mc15_13TeV.410016.PowhegPythiaEvtGen_P2012_Wt_dilepton_antitop.merge.DAOD_HIGG3D1.e3753_s2608_s2183_r7725_r7676_p3042
Zt	mc15_13TeV.410050.MadGraphPythiaEvtGen_P2012_tZ_4fl_tchan_noAllHad.merge.DAOD_HIGG3D1.e4279_s2608_s2183_r7725_r7676_p3042

Table 48: ATLAS Dataset IDs for the W+jets and single-top background samples used in this analysis.

Process	ATLAS Dataset Names and Identifiers
$t\bar{t}$	mc15_13TeV.410000.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_nonallhad.merge.DAOD_HIGG3D1.e3698_s2608_s2183_r7725_r7676_p3042
	mc15_13TeV.410009.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_dil.merge.DAOD_HIGG3D1.e4511_s2608_s2183_r7725_r7676_p3042
	mc15_13TeV.410501.PowhegPythia8EvtGen_A14_ttbar_hdamp258p75_nonallhad.merge.DAOD_HIGG3D1.e5458_s2726_r7772_r7676_p3042
	mc15_13TeV.410503.PowhegPythia8EvtGen_A14_ttbar_hdamp258p75_dil.merge.DAOD_HIGG3D1.e5475_s2726_r7772_r7676_p3042
$t\bar{t}$ Systematic	mc15_13TeV.410001.PowhegPythiaEvtGen_P2012radHi_ttbar_hdamp345_down_nonallhad.merge.DAOD_HIGG3D1.e3783_s2608_r7725_r7676_p3042
	mc15_13TeV.410002.PowhegPythiaEvtGen_P2012radLo_ttbar_hdamp172_up_nonallhad.merge.DAOD_HIGG3D1.e3783_s2608_r7725_r7676_p3042
	mc15_13TeV.410003.aMcAtNloHerwigppEvtGen_ttbar_nonallhad.merge.DAOD_HIGG3D1.e4441_s2726_r7772_r7676_p3042
	mc15_13TeV.410004.PowhegHerwigppEvtGen_UEEE5_ttbar_hdamp172p5_nonallhad.merge.DAOD_HIGG3D1.e3836_s2726_r7772_r7676_p3042
tt+X	mc15_13TeV.410155.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_ttW.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
	mc15_13TeV.410156.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_ttZnunu.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
	mc15_13TeV.410157.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_ttZqq.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
	mc15_13TeV.410218.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_ttee.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
	mc15_13TeV.410219.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_ttmumu.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
	mc15_13TeV.410220.aMcAtNloPythia8EvtGen_MEN30NLO_A14N23LO_tttautau.merge.DAOD_HIGG3D1.e5070_s2726_r7772_r7676_p3042
VV	mc15_13TeV.363355.Sherpa_221_NNPDF30NNLO_ZqqZvv.merge.DAOD_HIGG3D1.e5525_s2726_r7772_r7676_p3042
	mc15_13TeV.363356.Sherpa_221_NNPDF30NNLO_ZqqZll.merge.DAOD_HIGG3D1.e5525_s2726_r7772_r7676_p3042
	mc15_13TeV.363357.Sherpa_221_NNPDF30NNLO_WqqZvv.merge.DAOD_HIGG3D1.e5525_s2726_r7772_r7676_p3042
	mc15_13TeV.363358.Sherpa_221_NNPDF30NNLO_WqqZll.merge.DAOD_HIGG3D1.e5525_s2726_r7772_r7676_p3042
	mc15_13TeV.363359.Sherpa_221_NNPDF30NNLO_WpqqWmlv.merge.DAOD_HIGG3D1.e5583_s2726_r7772_r7676_p3042
	mc15_13TeV.363360.Sherpa_221_NNPDF30NNLO_WplvWmqq.merge.DAOD_HIGG3D1.e5805_s2726_r7772_r7676_p3042
	mc15_13TeV.363489.Sherpa_221_NNPDF30NNLO_WlvZqq.merge.DAOD_HIGG3D1.e5525_s2726_r7772_r7676_p3042
	mc15_13TeV.363490.Sherpa_221_NNPDF30NNLO_llll.merge.DAOD_HIGG3D1.e5332_s2726_r7772_r7676_p3042
	mc15_13TeV.363491.Sherpa_221_NNPDF30NNLO_lllv.merge.DAOD_HIGG3D1.e5332_s2726_r7772_r7676_p3042
	mc15_13TeV.363492.Sherpa_221_NNPDF30NNLO_llvv.merge.DAOD_HIGG3D1.e5332_s2726_r7772_r7676_p3042
	mc15_13TeV.363493.Sherpa_221_NNPDF30NNLO_lvvv.merge.DAOD_HIGG3D1.e5332_s2726_r7772_r7676_p3042
	mc15_13TeV.363494.Sherpa_221_NNPDF30NNLO_vvvv.merge.DAOD_HIGG3D1.e5332_s2726_r7772_r7676_p3042
Drell-Yan $10 < m_{\ell\ell} < 40$ GeV	mc15_13TeV.364198.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV0_70_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364199.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364200.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV70_280_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364201.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV70_280_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364202.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV280_E_CMS_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364203.Sherpa_221_NN30NNLO_Zmm_M1110_40_MAXHTPTV280_E_CMS_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364204.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV0_70_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364205.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364206.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV70_280_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364207.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV70_280_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364208.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV280_E_CMS_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364209.Sherpa_221_NN30NNLO_Zee_M1110_40_MAXHTPTV280_E_CMS_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364210.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV0_70_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364211.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364212.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV70_280_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364213.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV70_280_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364214.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV280_E_CMS_BVeto.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042
	mc15_13TeV.364215.Sherpa_221_NN30NNLO_Ztt_M1110_40_MAXHTPTV280_E_CMS_BFilter.merge.DAOD_HIGG3D1.e5421_s2726_r7772_r7676_p3042

Table 49: ATLAS Dataset IDs for the $t\bar{t}$, diboson, and Drell–Yan + jets with $m_{\ell\ell} < 40$ GeV background samples used in this analysis.

Process	ATLAS Dataset Names and Identifiers
Drell-Yan $m_{\ell\ell} > 40$ GeV	mc15_13TeV.364100.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364101.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364102.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364103.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364104.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364105.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364106.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364107.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364108.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364109.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364110.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364111.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364112.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364113.Sherpa_221_NNPDF30NNLO_Zmumu_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5271_s2726_r7772_r7676_p3042
	mc15_13TeV.364114.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364115.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364116.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364117.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364118.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364119.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364120.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364121.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364122.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364123.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364124.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364125.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364126.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364127.Sherpa_221_NNPDF30NNLO_Zee_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5299_s2726_r7772_r7676_p3042
	mc15_13TeV.364128.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV0_70_CVetoBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364129.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV0_70_CFilterBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364130.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV0_70_BFilter.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364131.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV70_140_CVetoBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364132.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV70_140_CFilterBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364133.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV70_140_BFilter.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364134.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV140_280_CVetoBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364135.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV140_280_CFilterBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364136.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV140_280_BFilter.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364137.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV280_500_CVetoBVeto.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364138.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV280_500_CFilterBVeto.merge.DAOD_HIGG3D1.e5313_s2726_r7772_r7676_p3042
	mc15_13TeV.364139.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV280_500_BFilter.merge.DAOD_HIGG3D1.e5313_s2726_r7772_r7676_p3042
	mc15_13TeV.364140.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV500_1000.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042
	mc15_13TeV.364141.Sherpa_221_NNPDF30NNLO_Ztautau_MAXHTPTV1000_E_CMS.merge.DAOD_HIGG3D1.e5307_s2726_r7772_r7676_p3042

Table 50: ATLAS Dataset IDs for the Drell–Yan+jets background samples that require $m_{\ell\ell} > 40$ GeV used in this analysis

F Triggering on Exotic Higgs Decays

Based on the 2HDM+S models described in Section 1.3.2, the dominant decay modes are the all hadronic modes (4j, 4c, 4τ, 4b, 2b2τ). Due to the overwhelming QCD background at the LHC, these modes are difficult to trigger on in the gluon-fusion production mode. To target these modes, LHC analyses have used leptonic decays of the tau [144] or have targeted the VH production mode (4b [145]). Here, a trigger that targets these decays via the VBF production mode is explored.

The VBF production mode is characterized by a large separation between the two quark jets, leading to large values of $\Delta\eta(j,j)$ and m_{jj}. These two variables for the $H \to aa \to 4b$ mode are shown in Figure 82. These two variables can be used in the ATLAS L1Topo trigger, and the m_{jj} discriminant was added to run the following trigger algorithm.

The algorithm begins with two input lists of jets, one for all J20 ROIs inside $|\eta| < 4.9$ and all J30 ROIs restricted to $|\eta| < 3.2$, which are gathered for each event. This η restriction, known as the "Non-Forward-Forward" (NFF) setup, on the J30 ROIs was implemented in order to reduce the L1 trigger input rate, as empirically there are many QCD initiated pileup events that give rise to two jets in the FCAL. In order to reduce the combinatorics and make a sufficiently fast trigger, only the 6 highest p_{T} jets are kept in each category. Overlaps in these two jet lists are not removed. All possible m_{jj} comparisons are computed until any pair of jets from the two lists contains $m_{jj} >$ X GeV. If this condition is reached, the event is accepted by the L1 trigger and passed on to the HLT. Otherwise, the event is rejected. The trigger efficiency curves as a function of offline m_{jj}, derived on MC simulation of events with a 125 GeV Higgs produced via vector boson fusion are shown in Figure 83.

During initial studies, two thresholds were compared, X= $\{400, 500\}$ GeV for the L1 trigger threshold in order to target $m_{jj} \approx 1$ TeV. The L1 rate corresponding to these two thresholds are given in Table 51.

The L1-MJJ-500-NFF threshold was chosen for running conditions in 2018 due to final rate concerns. After collecting these events at L1, a series of HLT triggers were additionally run on these events to reduce the data collection rate. All HLT chains implemented in 2018 required the presence of at least 2 HLT jets in each event. The jets were further required to have $p_{\mathrm{T}} >$ 50-70 GeV, depending on the exact decays being targeted. These thresholds were chosen by considering VBF jet p_{T} spectra seen in Figure 84, but these jets additionally help to reject the QCD background and lower the HLT output rate. At least two jets in each event are also required to pass an invariant dijet mass threshold of 900 - 1000 GeV, where only HLT jets with $p_{\mathrm{T}} > 50$ GeV used to calculate the dijet mass. These cuts at the HLT level help to

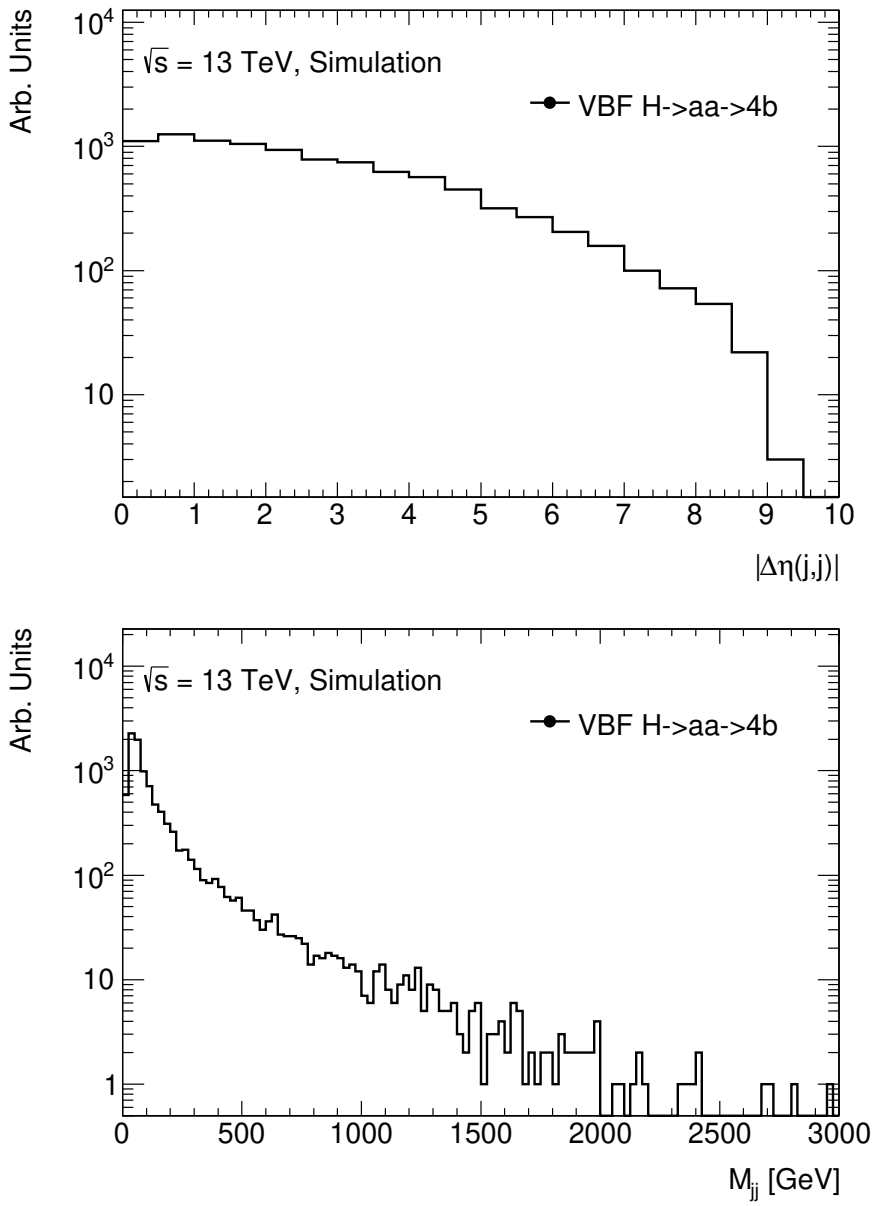

Figure 82: $\Delta\eta(j,j)$ and m_{jj} for the VBF production mode of the Higgs using a signal sample of $H \to aa \to 4b$.

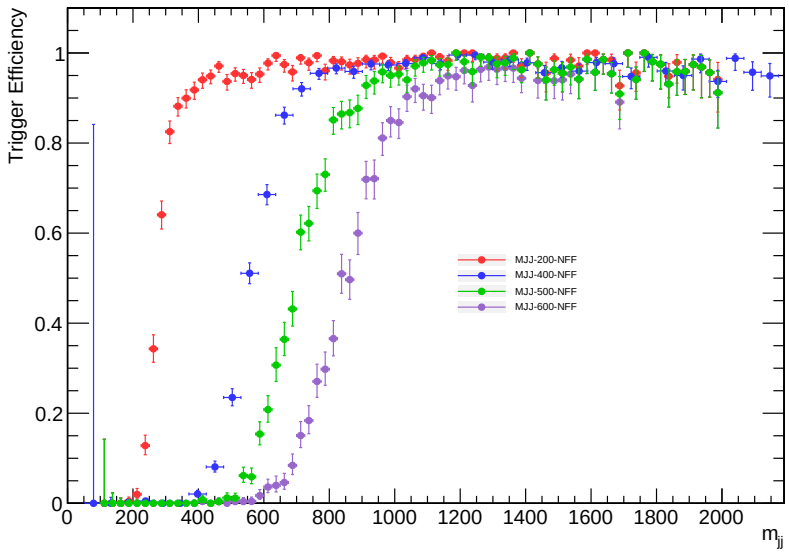

Figure 83: The trigger efficiency curves as a function of offline m_{jj}, derived on MC simulation of events with a 125 GeV Higgs produced via vector boson fusion for L1 thresholds of $m_{jj} = 200, 400, 500, 600$ GeV. The L1 efficiencies plateau at values of approximately twice the L1 threshold.

L1 Trigger	Rate [Hz]
L1-MJJ-400	4200.6 ± 43.2
L1-MJJ-400-NFF	3760.4 ± 39.6
L1-MJJ-500-NFF	2307.2 ± 19.9
L1_MU20	16151.6 ± 55.0
L1_EM22VHI	31551.7 ± 74.5

Table 51: The predicted L1 trigger rates for the MJJ-400 and MJJ-500 thresholds. As MJJ-500-NFF was initially predicted only as a backup for the MJJ-400 triggers, the non-NFF trigger was not considered for rate estimation. The L1_MU20 and L1_EM22VHI triggers correspond to the single lepton triggers in ATLAS and are placed here for comparison.

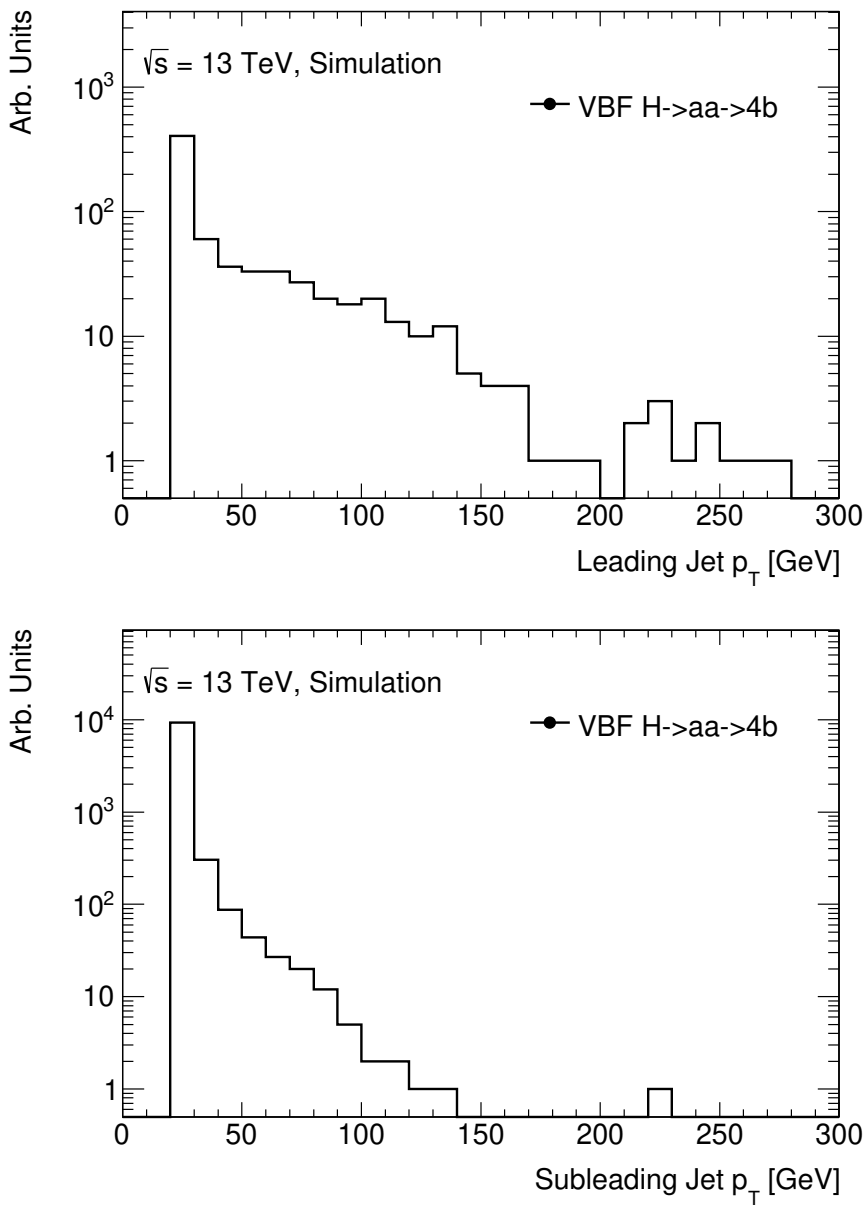

Figure 84: The leading and subleading VBF jet p_{T} for the Higgs VBF production mode. The lower p_{T} cut-off corresponds to the 20 GeV reconstruction threshold used by ATLAS. Although the selections of $p_{\mathrm{T}} > 50 - 70$ GeV seems rather tight, the acceptance of these cuts is 10% for a $m_{jj} > 500$ GeV selection. Further, these jet p_{T} selections significantly reduce the contributions from the QCD multi-jet background.

remove events which are not on the trigger plateau. Several strategies are taken to target different final states:

- **Muon Decays**: Two types of triggers have been implemented for muon decays, with one targeting the soft muons from tau decays and another targeting low mass signatures like $a \to \mu\mu$. The tau+muon triggers see some sensitivity gains by lowering the muon threshold from 14 GeV to 10 GeV, gained by tightening the isolation. The dimuon triggers lower the p_{T} of the muons from 14 GeV in the dimuon trigger threshold to 6 GeV with the VBF selections.

- **Electron Decays**: Two types of triggers have been implemented for electron decays, with one targeting the soft electrons from tau decays and another targeting low mass signatures implemented as a dielectron trigger. The VBF+electron trigger sees larger sensitivity gains compared to the inclusive tau+electron trigger, as the threshold on the electron is lowered from 17 GeV to 10 GeV and the tau leg of the HLT chain is also removed. The dielectron triggers also lower the p_{T} threshold from 17 GeV in the inclusive dielectron trigger to 5 GeV with the VBF selections.

- **Tau Decays**: For the 2b2τ final state, a hadronic tau trigger was implemented targeting two hadronic taus with $p_{\mathrm{T}}^{\tau} > 25, 20$. These taus are identified with the multivariate techniques [146] and a new approach to identify taus using tracks and calorimeter deposits as inputs in a recurrent neural network (RNN). Additionally, to target leptonic tau decays, the single and dilepton triggers above are also used due to much lower lepton p_{T} available compared to the inclusive single lepton triggers (10 compared to 26 GeV).

- **B-jet Decays**: For the 4b final state, the p_{T} of the 4 b-jets is rather low as can be seen in Figure 85 at the truth level. For comparison, the lowest unprescaled trigger targeting only b-jets requires at least two jets with $p_{\mathrm{T}} > 35$ GeV and using the 60% MV2c10 working point. With VBF events, these thresholds cannot change drastically, and the b-tagging working point can be raise to the 70% efficiency. A goal for Run 3 of the LHC is to run a 1 b-jet + VBF trigger at these low thresholds.

- **Photon Decays**: In order to cover cases with axion-like particles, a single photon trigger is implemented. The single photon trigger threshold at 35 GeV provides a large gain for low m_X particles, as these particles typically have two photon decays merging into a single photon.

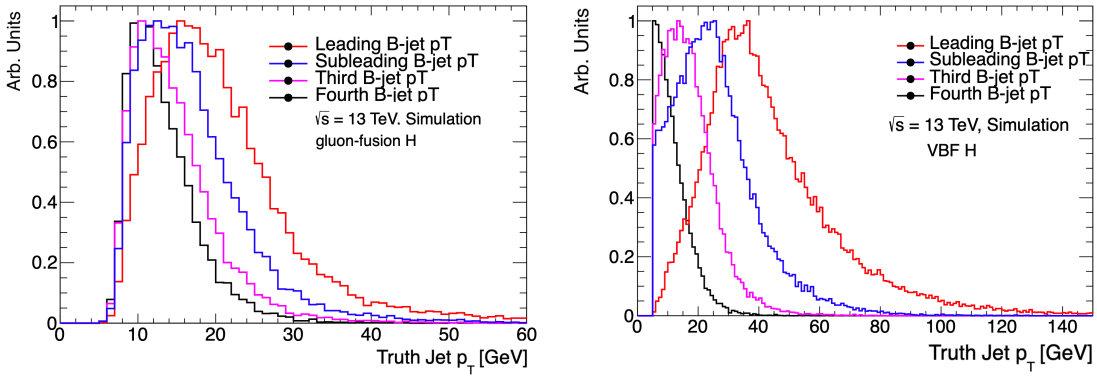

Figure 85: The transverse momentum of the 4 b-jets coming from the exotic Higgs decay to four b-jets, determined at the truth particle level in the gluon-fusion (left) and VBF (right) production modes of the Higgs boson. Most of the events will fail the ATLAS trigger thresholds of two jets with p_{T} >35 and the third and fourth jets can actually fail the ATLAS calorimeter jet reconstruction threshold of 20 GeV.

- **(Semi-)Invisible Decays**: For the cases where the Higgs boson decays invisibly or where a new supersymmetric particle decays invisibly in association with a compressed Z or W boson that decays inclusively. As the $E_{\mathrm{T}}^{\mathrm{miss}}$ threshold is 150 GeV, these triggers can provide new analysis channels with lower $E_{\mathrm{T}}^{\mathrm{miss}} > 80$ GeV and higher m_{jj}.
- **Inclusive VBF**: Rather than add additional particle content, the inclusive VBF trigger tightens the selections on the dijet system. These requirements include additional cuts on $\Delta\phi$(jet, jet) < 2.0, $\Delta\eta$(jet,jet) > 4.0, $m_{jj} > 1.1$ TeV, where only jets with $p_{\mathrm{T}} > 70$ GeV are used to calculate the invariant mass. These selections reduce the QCD background significantly and have a total efficiency of about 1% for exotic Higgs decays in the VBF topology.

The logical trigger names for the implemented triggers can be found in Table 52.

The inclusive VBF trigger, being the least dependent on physics objects, is used to determine the total trigger efficiency of the VBF Higgs signal. To make this determination, the various parts of the trigger must be characterized and the trigger turn-ons derived. The m_{jj} turn-on curve can be found in Figure 86 for 2018 ATLAS data, where the trigger efficiency is determined with respect to the lowest unprescaled single muon trigger. The nominal jet p_{T} plateau values were found to be $\{90, 80\}$ GeV for the leading/central jet and subleading/forward jet respectively. The offline

Physics Channel	Logical Name in HLT
VBF inclusive	j70_j50_0eta490_invm1100j70_dphi20_deta40
H→ invisible	j70_j50_0eta490_invm1000j50_dphi24_xe90_pufit_xe50
H→ $\tau\tau, bb\tau\tau$ leptonic	2e5_lhmedium_nod0_j70_j50_0eta490_invm900j50 2mu6_2j50_0eta490_invm900j50 e10_lhmedium_nod0_ivarloose_j70_j50_0eta490_invm900j50 mu10_ivarmedium_j70_j50_0eta490_invm900j50
H→ $\tau\tau, bb\tau\tau$ hadronic	tau25_mediumRNN_tracktwoMVA_tau20_mediumRNN _tracktwoMVA_j70_j50_0eta490_invm900j50
H→ $bb, 4b$	j70_0eta490_j50_0eta490_2j35_0eta490_invm1000j50 _AND_2j35_bmv2c1070_split
VBF Higgsino	e5_lhvloose_j70_j50_0eta490_invm1000j50_xe50_pufit e5_lhvloose_nod0_j70_j50_0eta490_invm1000j50_xe50_pufit mu4_j70_j50_0eta490_invm1000j50_xe50_pufit
VBF + γ, H→ $aa \to \gamma\gamma$jj SM H→ $bb + \gamma$	g35_medium_j70_j50_0eta490_invm900j50

Table 52: A list of VBF, L1 MJJ based triggers. The logical names contain ATLAS nomenclature for the electron and photon likelihoods (lhvloose_nod0, lhmedium_nod0, medium), lepton isolation (ivarloose, ivarmedium), the tau identification criteria (mediumRNN, tracktwoMVA), and the b-tagging WP (bmv2c1070). Split in the b-tagged chain refers to how the b-tagging was run online. The HLT and L1MJJ-500-NFF names have been suppressed for space saving.

analysis threshold for the m_{jj} for the trigger efficiency is found to be $\mathcal{O}(100$ GeV) higher than the HLT threshold.

By applying the fully efficient selections to a VBF Higgs sample, the total selection efficiency is found to be 1%. For comparison, the single lepton thresholds target the ZH production mode of the Higgs with about 70% efficiency. A breakdown of the total number of Higgs events in 36 fb^{-1} is shown in Table 53. From these estimates, the VBF inclusive trigger can target a similar number of events to the ZH production mode.

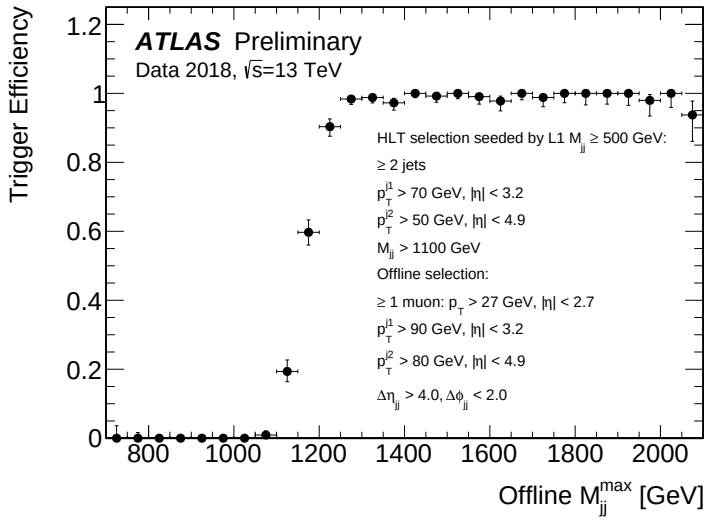

Figure 86: The trigger efficiency for the high-level trigger HLT_j70_j50_0eta490_invm1100j70_dphi20_deta40_L1MJJ-500-NFF as a function of the offline maximum dijet mass, $M_{jj}^{\max}$. The efficiency is measured using events selected using a single muon trigger with a threshold of 27 GeV. The events are additionally required to have at least two R=0.4 anti-k_t calorimeter jets, where one jet has $p_{\mathrm{T}} > 90$ GeV and $|\eta| < 3.1$, while the other selected jet must have $p_{\mathrm{T}} > 80$ GeV. The angular requirements $\Delta\phi$(jet,jet) < 2.0 and $\Delta\eta$(jet,jet) > 4.0 are also applied.

Production Mode	σ_H [pb]	Trigger, Efficiency	Events $\mathcal{L}$ =36 fb^{-1}
VBF	3.776	Inclusive VBF, 1%	1364
Z($\rightarrow \ell\ell$)H	0.06	Lepton triggers, 70%	1505

Table 53: An estimation of the number of events collected in 36 fb^{-1} for the VBF and ZH production modes. Here, the ZH production mode includes both the decays to muons and electrons. The total trigger efficiencies were determined in MC simulation samples where the exotic Higgs decay process, $H \rightarrow aa \rightarrow bbbb$, is simulated. A similar number of Higgs events are found to be triggered on between the two production modes.