

UNIVERSITÀ
DEGLI STUDI
DI PADOVA




Dipartimento
di Fisica
e Astronomia
Galileo Galilei

UNIVERSITÀ DEGLI STUDI DI PADOVA
DIPARTIMENTO DI FISICA E ASTRONOMIA "GALILEO GALILEI"

Search for rare boson decays with the CMS detector at LHC and the CMS Level-1 trigger Data Scouting

Rocco Ardino



Supervisor:	Dr. Alberto Zucchetta
Co-Supervisor:	Prof. Marco Zanetti
Co-Supervisor:	Dr. Emilio Meschi
PhD Coordinator:	Prof. Giulio Monaco

A thesis submitted to the University of Padua for the
degree of Doctor of Philosophy in Physics Research (XXXVII Cycle)

December 2024



*To Alice, Pier, Mum and Dad,
Without you, I would not be the person I am today.
Without you, I would be half the physicist I am now.
Without you, none of this would have been possible.
Thank you for walking beside me on this path,
and pulling me back up when I could no longer do it alone.*

Abstract

This thesis presents the search for the rare decays of the Z and Higgs bosons to a photon and a charmonium state J/ψ or $\psi(2S)$, which subsequently decays to a pair of muons. The channel under study is a probe of the Higgs coupling to the charm quark, thanks to a predicted direct decay process going through a charm quark loop, and it is characterized by a low amount of standard model backgrounds. The similar Z decay is included in the study as the sensitivity needed for an observation is expected to be reached first, offering a benchmark for the validation of the theoretical prediction of the branching fractions. New and unexpected physics might enter in the loop processes and affect the observed decay rates, making the study of rare Higgs boson decays to a meson and a photon a particularly promising ground for anomaly searches.

The dataset for the study corresponds to an integrated luminosity of 123 fb^{-1} of proton-proton collisions at a center-of-mass energy $\sqrt{s} = 13 \text{ TeV}$ collected during 2016-2018 with the CMS detector at the LHC. No evidence for branching fractions of these rare decay channels larger than predicted in the standard model is observed. Upper limits at 95% confidence level are set: $B(H \rightarrow J/\psi \gamma) < 2.6 \times 10^{-4}$, $B(H \rightarrow \psi(2S) \gamma) < 9.9 \times 10^{-4}$, $B(Z \rightarrow J/\psi \gamma) < 0.6 \times 10^{-6}$, and $B(Z \rightarrow \psi(2S) \gamma) < 1.3 \times 10^{-6}$. The ratio of the Higgs boson coupling modifiers κ_c/κ_γ is constrained to be in the interval $(-157, +199)$ at 95% confidence level. Assuming $\kappa_\gamma = 1$, this interval becomes $(-166, +208)$.

The analysis is statistically limited and it would profit from higher trigger efficiencies. This limitation is more generally affecting the class of Higgs boson decays to a meson and a photon and to a pair of mesons, particularly when the mesons involved are the ρ and the ϕ . This thesis presents a novel data taking and real-time analysis approach for the CMS Phase-2 upgrade at High-Luminosity LHC that allow to perform the analysis using the Level-1 trigger (L1T) reconstruction and to bypass the event rejection imposed by the L1T selection. The system implementing this idea, known as Level-1 trigger Data Scouting (L1DS), will collect and perform an online analysis at the full event collision rate on the reconstructed Phase-2 L1T objects, which will have efficiency and resolution often comparable to the offline one. The L1DS system will enable the search for signatures that feature a too-large irreducible background and therefore currently rejected by the tight requirement on the L1T output rate budget, and signatures for which the background rejection requires algorithms that do not fit within L1T latency and hardware resource constraints. Additionally, it will provide vast amount of data for fast detector diagnostic. A demonstrator of the L1DS system, assembled and commissioned in 2023, is presented. Its architecture and results after its successful deployment in production starting from 2024 are also discussed. The demonstrator will allow the test of different approaches to acquire and process online data from the L1T system, thus shaping the development of the Phase-2 L1DS on real collision data. Projections of the L1DS performance on the rare Higgs boson decays to a meson and a photon or to a pair of mesons are carried out on simulated data. The results highlight how the scouting approach is suitable for the search of these elusive decays, which might open the path to new measurements of light quark flavor couplings.

Statement of Originality

I declare that the thesis submitted for examination for the PhD degree at the University of Padua is my own work, with contributions from other members of the CMS Collaboration. Where figures or results are part of others' work, proper acknowledgment and referencing have been provided. Chapter 3: I have contributed to all the steps of the $Z, H \rightarrow J/\psi\gamma, \psi(2S)\gamma$ analysis, as main author and responsible of the analysis review within the CMS Collaboration, together with my supervisor Dr. Alberto Zucchetta. The results of the analysis are publicly available in Ref. [1] and have been submitted to Physics Letters B.

Chapter 5: I significantly contributed to the successful commissioning of the CMS Level-1 trigger Data Scouting demonstrator for Run-3. In particular, I have designed the firmware needed for the readout, concentration and pre-processing of the Level-1 trigger data on the electronic boards for the L1DS system, which are equipped with Field Programmable Gate Array (FPGA) devices. I extended the existing scouting data acquisition software to capture the output Ethernet links from the aforementioned boards, perform further processing and prepare the output fragments for the online processing. I have developed the monitoring and control application for these boards, allowing the control of the demonstrator from the CMS Run Control system. I contributed to the validation of the online processing infrastructure and to the definition of frameworks to perform online and offline analysis of the collected data. I contributed to the machine learning applications for the system, in particular I designed a multivariate algorithm that fits the trigger information sent by the muon detectors to reconstruct muon tracks and I supervised CERN summer students working on the algorithm implementation.

Chapter 6: I originally conceived the idea of searching for Higgs rare decays to a pair of ρ or ϕ mesons. I conducted the performance and sensitivity studies for the decays $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ using physics objects reconstructed by the Phase-2 L1T. My contribution includes the simulation of the signals, the development of the analysis strategy and its deployment, the projections for the analysis sensitivity, the integration of the analyses in the Phase-2 L1DS online processing demonstrator and timing benchmarks.

Acknowledgments

"Part of the journey is the end."

— TONY STARK, *Avengers: Endgame*

Early 2000s, a curious child was reading a colorful science book. The pages showed a simple zoom into the constituents of matter, with images extending all the way to atoms and quarks. After that, there was a mysterious question mark. Nobody knew, and still does not know, what hides behind it. However, that was enough to trigger an unstoppable chain reaction, which still continues to this day. A purpose filled the eyes of that child, pushing him toward a long and complex path. But there was one small yet important detail he had not realized: he was not alone. A long path ends here, it is part of the journey after all. However, the journey that began on that day is still unfolding, and I cannot continue without thanking the people who have become part of me since the beginning¹.

It all started with my mum, Tina, who bought me that book and ensured that I never lacked any fundamental thing along the way, especially the love a child could need. You have always made me feel capable of achieving everything, and that belief eventually became part of who I am. My dad, Mimmo, did the same, teaching me that nothing is guaranteed in life without the hard work, which is in the blood of the *Ardino* family. With no exception, including my brother Pier, whose grumpy character has always hidden a kind soul. Special thanks go to Gabriella and Peppe: a part of me is still wandering across the dunes of the Tunisian deserts, where one of my deepest dreams came true. You have always made me feel like your son and you welcomed me in your family. Words cannot fully express how grateful I am to both you for that. To my grandmothers, who have always waited joyfully for every of my returns during these years abroad, wishing nothing but the best for me. Thank you all for being the pillars of my life and for giving me everything I could have every dreamed of, without me even asking.

During the years in my hometown, Altamura, I met special people who eventually became my dearest friends. Domenico, describing what we have lived (or survived) together is difficult without getting emotional: bike rides randomly meeting fake intellectuals, airsoft games, watermelons carried by hand on a scooter, stylish celebrations and "kidnappings" at *Masseria Lago Cupo*. This is just a small list, but you have always been there, from the moment we realized how much we had in common. Alessandro, I am lucky to still see you without "marble documents" after the endless list of visits to the hospital. You are always ready to bring a smile and show me new points of view in the world history and dynamics. Anna Rita, with every single word you have always proven how rarely you see such a kind and sensitive person like you. Your character has always been an inspiration to me. Giacomo, your wisdom is unparalleled, but you are so much more than that. Over the years, you have opened up and shown me a very deep and eclectic personality, always offering something new to learn and experience. Luciano, you are the kind

¹In particular, after two theses without acknowledgments...

of person I would stay awake until 4.00 AM with, just to improve a Trackmania record by 0.01 sec (all *monster'ed*, of course). This says everything about what I think about you. Martina, your determination has always impressed me since the end of the high school, when we began to form the group we are today. The path you are following is admirable, and I am sure it is just the start of something even bigger. Thanks to all of you, coming back to Altamura has always had a special meaning for me, knowing that I will find you there.

The years of my bachelor's and master's in Padova were fundamental in shaping me. My character was at the same time filled by curiosity and forged by the tough bunch of lectures and exams, however it was easier with the amazing people I met there at the DFA. A huge thank you goes to the members of the so-called Joüœürdœen group. Together we faced many exams, partied (and drank) for them, spent countless hours over lunches and dinners. A particular thank you goes to Davide. I still cannot believe how lucky I have been to meet a friend like you. We have been through everything together, both hard and amazing moments, but you have never stopped trying to put a smile on those around you. With bad jokes and extremely deep thoughts, your amazing personality always shines through. Very few people ever reach the true essence of success, but I am sure you will be one of them. A big thank you also goes to Gentile, also known as Francesco, the playable character you still have to unlock. You were probability the first person I met when I arrived in Padova. With you and Davide, we formed the laboratory dream team, which was about more than just university. It was an excuse to spend time with extraordinary friends. Thanks to Vezza and MAAAAAX, the flatmates and friends I am always hyped to meet again: don't forget to restart the Minecraft server, I have another idea of farm to build.

The (almost) final thanks go had the chance to meet and get to know better during the three years of PhD between Padova and CERN. Starting from my supervisors, Alberto, Emilio and Marco: you have guided me through this challenging journey with your indispensable advice and suggestions. To the reviewers of this document, Cristina and Alex: without you, this work would not have the same value it has today. To the entire EP-CMD team, full of brilliant people both professionally and personally: you were always ready to help whenever I struggled with something completely new and unknown. To the *BoostLab* team, especially Gabriele, Matteo and Nicolò, the Black Ops dream team. I will never forget that *in gambling there are no losers, just quitters*. To Dr. James, whose grammar correction prescale set to 0 helped me a lot to let people understand me better. Thanks for all the moments we shared since day one, and for every "boooongiornooooo" bringing joy in the office. To my first flatmate, Fabian, a true German and a very precious friend always there for me during these frenetic three years: I am still wondering *why I can't drink pizza*. Perhaps one day, I will understand the reason.

Alice, you probably thought I had forgotten about you. However, you know really well that *the best comes for last*. You have been the closest person to me over the past six years and you have seen me struggle for every single achievement I got so far. Even though I have stressed you out countless of times, you have always been my first supporter, ready to shake the world for me. You have always believed in me, even when I did not believe in myself, teaching me that *there is no defeat in the heart of whom does not give up*. During these three years of PhD, we have grown together with new and often overwhelming experiences (they were actually A LOT). I am so proud of you and of the team we have become. My biggest thanks go to you, Alice. This thesis would not have been the same without you. But more than that, *I* would not be the same without the unique person you are on my side.

Contents

Abstract

Statement of Originality

Acknowledgments

Contents

Introduction	1
1 Higgs boson physics and rare decays	3
1.1 The standard model of particle physics	3
1.1.1 Fermions of the standard model and bound states	4
1.1.2 Theory symmetry group and quantum numbers	5
1.1.3 Standard model Lagrangian	5
1.2 Phenomenology of the Higgs boson	7
1.2.1 Higgs boson production at the LHC	7
1.2.2 Higgs boson decay channels and main properties measurement	8
1.2.3 Couplings to gauge bosons and fermions	10
1.3 Rare Higgs boson decays	12
1.3.1 Rare decays to a meson and a photon	12
1.3.2 Rare decays to a pair of mesons	13
2 The Compact Muon Solenoid experiment at the LHC	15
2.1 The Large Hadron Collider at CERN	15
2.1.1 LHC filling scheme	16
2.1.2 Instantaneous and integrated luminosity	17
2.1.3 Pileup collisions	18
2.2 The Compact Muon Solenoid experiment	19
2.2.1 Silicon tracker	21
2.2.2 Hadron and electromagnetic calorimeters	21
2.2.3 Muon system	22
2.2.4 Trigger and Data acquisition systems	24
2.3 Event reconstruction with the CMS detector	25
2.3.1 Tracks and vertices	25
2.3.2 Muons	26
2.3.3 Electrons, photons, jets and taus	26

Contents

3	Search for rare Z and Higgs boson decays to a charmed meson and a photon	29
3.1	Analysis strategy and datasets	29
3.1.1	Datasets for the search	30
3.1.2	Simulated signal and background samples	30
3.2	Physics objects and event preselection	31
3.2.1	Primary vertex	32
3.2.2	Trigger selection	32
3.2.3	Muon identification and isolation	33
3.2.4	$\psi(\text{nS})$ candidate reconstruction	34
3.2.5	Photon identification	34
3.2.6	Jets	35
3.2.7	Preselection summary	35
3.3	Event selection	37
3.3.1	Signal region selections	37
3.3.2	Angular variables based categorization	38
3.3.3	Higgs production mode categorization	40
3.3.4	Control region for $Z \rightarrow \ell\ell\gamma$ background modeling	40
3.3.5	Event categorization summary	41
3.4	Event modeling	42
3.4.1	Signal and resonant background modeling	42
3.4.2	QCD background modeling	44
3.5	Systematic uncertainties and corrections	49
3.5.1	QCD background model discrete uncertainty	49
3.5.2	Uncertainties affecting the predicted signal yields	50
3.5.3	Uncertainties affecting the shape of the signal model	53
3.5.4	Summary	55
3.6	Results	55
3.6.1	Upper limits at 95% CL on branching fraction	56
3.6.2	Measurement of the Z decay to $\ell\ell\gamma$	57
3.6.3	Summary and outlook of the analysis	58
4	Enhancing the CMS physics reach with Real-Time analysis at the Level-1 trigger	61
4.1	Present, future and limitations of the Level-1 trigger system	62
4.2	Performance of the CMS Phase-1 Level-1 trigger	64
4.2.1	Calorimeter Trigger	65
4.2.2	Global Muon Trigger	68
4.2.3	Global Trigger	69
4.3	The CMS Phase-2 Level-1 trigger upgrade	70
4.3.1	Calorimeter Trigger	70
4.3.2	Muon Trigger	71
4.3.3	Track Trigger	71
4.3.4	Correlator Trigger and Particle Flow	72
4.3.5	Global Trigger	74
4.4	Phase-2 Level-1 trigger Data Scouting system	75
4.4.1	Baseline system and physics potential	75

4.4.2	System architecture and requirements	77
5	The Level-1 trigger Data Scouting demonstrator	81
5.1	Demonstrator overview	81
5.2	Readout of the trigger links	83
5.2.1	Input data and formats	83
5.2.2	Trigger data pre-processing on FPGA	86
5.2.3	Scouting boards control and monitoring	91
5.2.4	Data acquisition software	94
5.3	Online processing at the LHC orbit rate	96
5.3.1	Event definition and Orbit Collection data format	96
5.3.2	Event aggregation and processing	97
5.3.3	Scouting selector modules	99
5.4	Demonstrator validation from collected data	101
5.4.1	Bunch crossing occupancies and object multiplicities	102
5.4.2	Energy and momentum distributions	102
5.4.3	BMTF stubs and KBMTF emulation	106
5.4.4	Low dimuon invariant mass resonances	108
5.4.5	Z decay to a pair of leptons	109
5.5	Machine learning inference applications for L1DS	111
5.5.1	L1 muons recalibration	111
5.5.2	Deep-learning based barrel muon track reconstruction	113
5.6	Run-3 demonstrator summary and perspectives	115
6	Study of rare Higgs boson decays with Phase-2 Level-1 trigger Data Scouting	119
6.1	Study premises and conditions	119
6.2	Generator level study	121
6.3	L1 scouting analysis	124
6.3.1	Event selections	124
6.3.2	Signal and background modelling	130
6.3.3	Projected constraints on branching fraction	130
6.4	Online processing timing benchmark	132
	Conclusions and outlook	133
	List of Figures	135
	Bibliography	139
	Acronyms	151

Introduction

The discovery of the Higgs boson in 2012 by the CMS and ATLAS experiments marked a significant milestone in understanding the Standard Model (SM) of particle physics. This breakthrough has led to a new phase of research focused on measuring the Higgs boson interactions with other SM particles. While couplings to gauge bosons and third-generation fermions have been observed and align with SM predictions, the couplings to first and second-generation fermions remain experimentally not established. Further data, particularly from the current LHC Run-3 and the upcoming High-Luminosity LHC (HL-LHC), are crucial for refining these measurements and potentially uncovering new physics.

A promising avenue of research involves studying the rare decays of the Higgs boson, which could reveal new physics through loop processes and anomalous couplings. Decays to a charmed meson such as the $\psi(nS)$ and a photon could probe the Yukawa coupling with the charm quark, with an expected branching fraction of the order of 10^{-6} . With a similar rate, decays to a ϕ meson and a photon constitute a handle to probe the Yukawa coupling with the strange quark, while decays to a ρ meson and a photon are sensitive to the coupling to the first-generation up and down quarks, with a branching fraction of the order of 10^{-5} . The same type of decay involving the Y meson, with a lower branching fraction of the order of 10^{-9} , is also of interest, as it provides an additional probe to Beyond the SM (BSM) physics. Unexpected physics might enter in the more complex class of decays of the Higgs boson to a pair of mesons, such as $H \rightarrow 2\rho, 2\phi, \phi J/\psi, 2J/\psi, 2Y$. Notably, decays to pairs of ρ and ϕ mesons, as well as $\phi J/\psi$, have never been studied before. Their branching fractions are predicted to be in the range $10^{-8} - 10^{-10}$. The predictions carry large theoretical uncertainties, providing a further challenging and promising ground for searching for possible anomalies.

Higgs boson decays involving $\rho, \phi, \psi(nS)$ and Y mesons are extremely interesting from the experimental point of view as they produce a potentially clean final state after the subsequent meson decay to pair of charged tracks, namely $\rho \rightarrow \pi\pi$, $\phi \rightarrow KK$ and $\psi(nS), Y \rightarrow \mu\mu$. They are also challenging to study as their detection requires sophisticated trigger algorithms to efficiently filter SM backgrounds. The implementation of these algorithms must fit within the available trigger rate budget and computational resources. For the case of a hardware-based trigger, the limitation stems from the limited available logic resources and the hard limit on the latency, and the available detector information to reconstruct the physics objects might only include a subset of the entire detector. For example, the L1 trigger (L1T) of the current CMS experiment cannot use the particle tracks from its tracker system. Therefore, identifying meson decays to a pair of tracks as $\rho \rightarrow \pi\pi$ and $\phi \rightarrow KK$ is not possible without dramatic signal efficiency loss.

For the HL-LHC, the CMS experiment is undergoing a significant detector upgrade in order to cope with the new extreme luminosity and pileup conditions. The upgraded detector, known as CMS Phase-2, will feature a new L1T system, which will have access to an unprecedented level of information. Advanced reconstruction algorithms will be deployed directly on the L1T Field Programmable Gate Array (FPGA) based processors, approaching the offline reconstruction

Introduction

resolution. New trigger strategies can be designed thanks to the availability at L1T of charged particle tracks and high-granularity calorimeter primitives. The searches for rare decays of the Higgs boson to a meson and a photon, in particular of $H \rightarrow \phi\gamma, \rho\gamma$, will greatly benefit from this new system. The L1T reconstruction of the meson decay needs to select the best pair of tracks among all the combinations arising from over a hundred tracks per collision event at HL-LHC. Thus, the hardware resources needed for the implementation of meson identification are not negligible, and loss in trigger efficiency due to tighter selections could be avoided by moving the analysis before the L1T decision, where collision data is unfiltered yet.

The Level-1 trigger Data Scouting (L1DS) project is a novel data acquisition and processing system, aimed at collecting and storing the trigger objects reconstructed by the L1T processors at the full LHC collision rate. The scouting approach enables the search for rare Higgs boson decay signatures without the limitations imposed by the standard trigger approach, while profiting of the quasi-offline resolution of the trigger objects. In general, a wider physics program can be probed in the L1DS system, complementing the standard program with data selected by the trigger system. Signatures as $W \rightarrow D_s\gamma, \pi\gamma, 3\pi$, Heavy Stable Charged Particles (HSCP) spanning multiple collision events, dilepton and dijet resonances at low invariant mass can be expected to be accessible to the L1DS. In this thesis, the development and results of a L1DS demonstrator, successfully commissioned during 2024, are presented. The current system description is completed by projections on Higgs boson rare decays to a meson and a photon and to a pair of mesons using the Phase-2 L1DS, which relies on the CMS Phase-2 detector and L1T, highlighting the promising capabilities of the L1DS system in this type of searches. The thesis also presents the analysis of the Higgs and Z decays to a J/ψ or $\psi(2S)$ meson and a photon using the LHC Run-2 dataset, collected from the standard trigger chain. The conducted search has not shown significant discrepancies with the SM prediction and has set constraints on the branching fraction of the processes studied and on the Yukawa coupling to the charm quark. The results highlight the fact that current searches are statistically limited, indicating that the L1DS approach may be valuable to push limits even closer to the SM predictions.

The thesis is structured as follows. Chapter 1 introduces the theoretical framework of the SM, used as reference for the following physics studies of this work. Particular focus is devoted to the phenomenology of the Higgs boson, presenting its properties, main production modes and decay channels. The physics of rare decays to a meson and a photon and to a pair of mesons is introduced. Chapter 2 describes the experiments under consideration. A brief introduction about the LHC accelerator and technical information relevant for this work is given. The CMS experiment subdetectors and event reconstruction techniques are described to help better understand the physics objects used in the analysis of the Run-2 data. Chapter 3 presents the analysis of $H, Z \rightarrow \psi(nS)\gamma$ decays, with in-depth description of the analysis strategy, event selection, modelling, treatment of systematic uncertainties and branching fraction constraint results. Chapter 4 highlights the challenges behind Higgs rare decay searches and how the L1DS approach could address their statistically limited nature. The performance of the CMS L1T is described for both its Run-3 version and the one foreseen for CMS Phase-2. Chapter 5 describes the L1DS demonstrator for Run-3, the technical choices behind its commissioning and a set of related results and applications. Chapter 6 completes this work with an estimate on the projected reach for rare Higgs boson decay searches using solely the L1 trigger information, in view of their possible deployment in the final L1DS system for HL-LHC.

Higgs boson physics and rare decays

The discovery of the Higgs boson in 2012 [2–4], achieved by the CMS and ATLAS experiments, stands as a significant milestone in the understanding of the electroweak sector within the Standard Model (SM) of elementary particles. This achievement was followed by a series of studies aimed at determining its properties, such as mass, decay width and spin. The measurement of its couplings to other SM particles has shown a remarkable agreement with respect to the SM prediction, leading to the measurement of couplings to gauge bosons and third generation fermions. Since the start of LHC Run-2, the CMS and ATLAS experiments have started to approach the sensitivity needed to study the couplings to the second generation of fermions. The first evidence of the coupling to muons has been obtained by the CMS experiment [5]. The focus is now directed towards the measurement of the couplings to the second generation charm and strange quarks. Any discrepancy between the observed results and the theoretical expectation could indicate new physics interactions or undiscovered particles. These hypothetical entities may potentially exert their influence through quantum loop processes, thereby impacting the branching fractions of rare decay events involving first and second-generation fermions. Consequently, the study of Higgs boson rare decays serves as a critical avenue for probing physics beyond the SM (BSM) and it represents a pivotal goal for the current and future physics programs at LHC and next-generation colliders [6].

This chapter is organized into three main sections. The first presents a brief introduction of the SM theory, its particle content and the fundamental interactions. The second section familiarizes the reader with the Higgs boson phenomenology, summarizing key results on the cross section of the main production modes and the branching fractions of the most common decay modes. The third section focuses on the rare decay modes relevant to this work, specifically those involving light or heavy mesons plus a photon and meson pairs, with a primary focus on charm and strange quark flavors.

1.1 The standard model of particle physics

The SM is the theory describing the electroweak and strong forces in the universe and classifying all known elementary particles [7–10]. It is formulated as a quantum field theory and its particle constituents are summarized in Fig. 1.1. The matter states are described as spin- $1/2$ fermions and the mediators of their interactions are represented as spin-1 vector bosons. The electromagnetic force is mediated by the massless photon (γ), the weak force by the massive Z and W^\pm bosons, with

masses around 91.2 GeV and 80.4 GeV respectively², the strong nuclear force by eight different color states of the massless gluon (g). In addition, the Higgs boson, which is a scalar with spin 0, interacts with both the fermions and bosons. The gravitational force, whose mediator is postulated to be the graviton, is not included in the SM. However, gravity is weak at subatomic scale with respect to the other forces and it can be safely neglected in this framework.

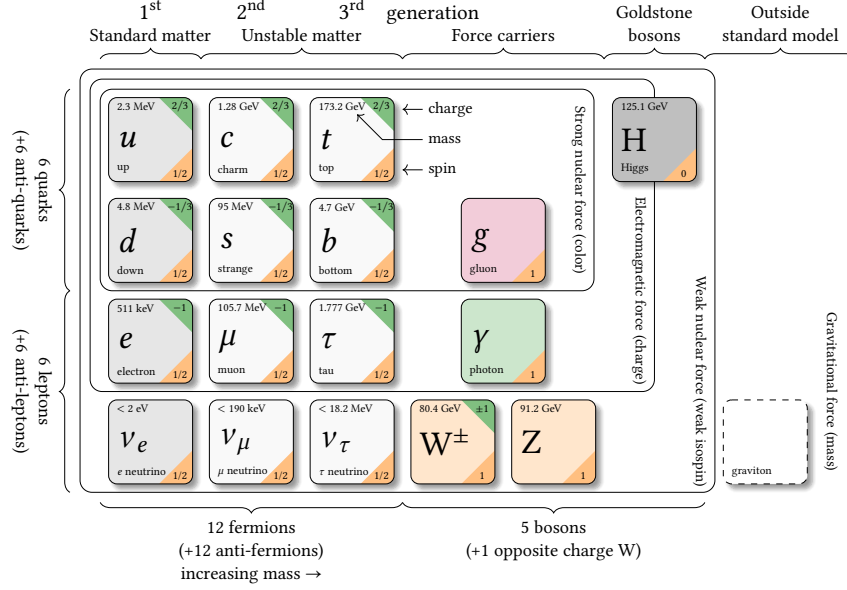


Figure 1.1 • Summary of the standard model of particle physics, with all the known elementary particles. Mass, spin, charge are reported for each particle in the relative box. Adapted from Ref. [11].

1.1.1 Fermions of the standard model and bound states

Matter states are grouped into three generations of fermions and two classes of particles: leptons (ℓ) and quarks (q).

Leptons The three generations of leptons have each a charge of -1 and include the electron (e), the muon (μ) and the tau (τ) leptons, together with their antiparticles with opposite charge. Their masses are $m_e = 0.511$ MeV, $m_\mu = 105.7$ MeV, $m_\tau = 1.777$ GeV. Every lepton flavor has a corresponding neutrino (ν_e, ν_μ, ν_τ), which is electromagnetically neutral and interacts with other particles only through the weak force. The mass of each neutrino is upper bounded at 120 eV.

Quarks Each generation of quarks is represented by a doublet of particles with charges of $+2/3$ and $-1/3$, respectively. The first generation consists of the up (u) and down (d) quarks, which are the constituents of the atomic nuclei and have masses of 2.3 MeV and 4.8 MeV, respectively. The second generation includes the charm (c) and strange (s) quarks, with masses of 1.27 GeV and 95 MeV, respectively. The bottom (b) and top (t) quarks form the third and most massive generation. The bottom quark has a mass of around 4.65 GeV, while the top is the heaviest particle in the SM with its mass of 172.52 GeV [12].

²For the mass values, natural units with $c = 1$ are assumed.

Bound quark states Quarks are observed in bound states called hadrons, which can be either triplets (barions) or doublets (mesons) of quarks. The top quark is an exception due to its short lifetime, which prevents the formation of bound states. Among the various hadrons, few heavy and light mesons are specifically mentioned or considered in this work, with their properties summarized in Ref. [13]. The ρ^0 meson is a quantum mixture of $u\bar{u}$ and $d\bar{d}$ states, with a mass of around 770 MeV. The ϕ meson, with a mass of 1.019 GeV, originates from a bound state of strange quark pairs. The J/ψ is a bound state of a pair of charm quarks with mass 3.097 GeV, while the $\psi(2S)$ is its first excited state, with a mass of 3.686 GeV. These two mesons are collectively referred to as $\psi(nS)$. The $Y(nS)$ mesons, where $n = 1, 2, 3$, are three bound states of a pair of bottom quarks, with masses of 9.460 GeV, 10.023 GeV and 10.355 GeV, respectively. In this work, the $Y(nS)$ and $\psi(nS)$ mesons are considered as heavy states, while the ρ and ϕ mesons are considered as light ones.

1.1.2 Theory symmetry group and quantum numbers

The SM theory can be described by a Lagrangian that is symmetric under transformations of the gauge group [7]:

$$\mathcal{G}_{\text{SM}} = \text{SU}(3)_C \otimes \text{SU}(2)_L \otimes \text{U}(1)_Y \quad . \quad (1.1)$$

In Eq. 1.1, $\text{SU}(3)_C$ represents the color symmetry group underlying the Quantum Chromodynamics (QCD) theory, which describes the strong interactions and is assumed to be an exact symmetry [7]. The $\text{SU}(2)_L \otimes \text{U}(1)_Y$ group corresponds to the symmetry group of the electroweak theory, which describes the interaction between matter fields through the exchange of massless vector bosons. However, this symmetry is spontaneously broken into:

$$\mathcal{G}_{\text{SM}}^{\text{SSB}} = \text{SU}(3)_C \otimes \text{U}(1)_{\text{em}} \quad , \quad (1.2)$$

where $\text{U}(1)_{\text{em}}$ is the symmetry group associated with the electromagnetic charge conservation. Through the process of Spontaneous Symmetry Breaking (SSB) process, it is possible to explain why gluons and photons remain massless after SSB, while the other gauge bosons and the Higgs boson acquire mass [10, 14].

The quantum numbers of the SM matter fields under the \mathcal{G}_{SM} symmetry group are listed in Table 1.1. Due to the chiral structure of the weak interaction, the SM matter fields have both left and right-handed components, which are assigned to different representations of the symmetry group [7]. Neutrinos are the only exception to this scheme, as they are observed only in their left-handed component or in their right-handed one for their antiparticles.

1.1.3 Standard model Lagrangian

The SM theory is constructed by specifying the particle content of the theory, its gauge group \mathcal{G} and the possibility of spontaneous symmetry breaking of \mathcal{G} into a smaller group. Combining these elements introduced in the previous subsections, it is possible to present the mathematical form of the SM Lagrangian [15] in a short-handed form, highlighting the different sectors based on the type of interaction:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{kin}}^{\text{g.b.}} + \mathcal{L}_{\text{kin}}^{\text{f}} + \mathcal{L}_{\text{kin}}^{\text{h}} + \mathcal{L}_{\text{pot}}^{\text{h}} + \mathcal{L}_Y \quad . \quad (1.3)$$

Generation	1 st	2 nd	3 rd	U(1) _{em}	SU(3) _C	SU(2) _L	U(1) _Y
Leptons	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L$	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L$	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$	0	-1/2	2	1
	e_R	μ_R	τ_R	-1	-1	1	1
Quarks	$\begin{pmatrix} u \\ d \end{pmatrix}_L$	$\begin{pmatrix} c \\ s \end{pmatrix}_L$	$\begin{pmatrix} t \\ b \end{pmatrix}_L$	+2/3	+1/6	2	3
	u_R	c_R	t_R	+2/3	+2/3	1	1
	d_R	s_R	b_R	-1/3	-1/3	1	3

Table 1.1 • The quantum numbers of SM fermions under the $\mathcal{G}_{\text{SM}} = \text{SU}(3)_C \otimes \text{SU}(2)_L \otimes \text{U}(1)_Y$ symmetry group, for the left and right-handed components. The right-handed chirality is omitted for neutrinos, given that they are observed only in their left-handed component, or right-handed one for their antiparticles.

Gauge bosons kinetic term $\mathcal{L}_{\text{kin}}^{\text{g.b.}}$ contains all the kinetic terms for the gauge bosons, including the non-Abelian interactions between gauge bosons:

$$\mathcal{L}_{\text{kin}}^{\text{g.b.}} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} \quad , \quad (1.4)$$

where $F_{\mu\nu}$ is the field strength tensor of the gauge boson field.

Fermions kinetic term $\mathcal{L}_{\text{kin}}^{\text{f}}$ contains the fermion kinetic terms as well as the interactions between gauge bosons and fermions. It is expressed as:

$$\mathcal{L}_{\text{kin}}^{\text{f}} = i\bar{\psi}\gamma^\mu D_\mu\psi + \text{h.c.}^3 \quad , \quad (1.5)$$

where ψ represents the matter fields and $\bar{\psi}$ their adjoint, while $D_\mu = \partial_\mu - iqA_\mu$ denotes their covariant derivative.

Higgs boson kinetic term $\mathcal{L}_{\text{kin}}^{\text{h}}$ is the kinetic term for the Higgs field, which gives masses to the gauge bosons Z and W through the SSB mechanism, once a vacuum state $\langle\varphi\rangle$ of the Higgs potential is chosen. This term includes the interaction between the Higgs field and the gauge bosons and it is expressed as:

$$\mathcal{L}_{\text{kin}}^{\text{h}} = (D_\mu\varphi)^\dagger(D^\mu\varphi) \quad , \quad (1.6)$$

where φ represents the Higgs field. Specifically, the φ field is a complex SU(2)_L doublet with a charged and a neutral component:

$$\varphi = \begin{pmatrix} \varphi^+ \\ \varphi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \varphi_1 + i\varphi_2 \\ \varphi_3 + i\varphi_4 \end{pmatrix} \quad , \quad (1.7)$$

where $\{\varphi_i\}_{i=1,\dots,4}$ are real scalar fields used to rewrite the complex components φ^+ and φ^0 .

³The term “h.c.” in the Lagrangian indicates the hermitian conjugate of the preceding terms.

Higgs boson potential $\mathcal{L}_{\text{pot}}^h$ represents the Higgs potential, which gives mass to the Higgs field by self-interaction:

$$\begin{aligned}\mathcal{L}_{\text{pot}}^h &= -V(\varphi^\dagger \varphi) \\ &= -\mu^2 \varphi^\dagger \varphi + \lambda |\varphi^\dagger \varphi|^2.\end{aligned}\quad (1.8)$$

The Higgs potential V , with the choice of $\mu^2 > 0$ and $\lambda > 0$, has an arbitrary non-trivial minimum, whose selection determines the SSB mechanism, as explained in Section 1.1.2.

Yukawa sector \mathcal{L}_Y represents the Yukawa sector, which is of central importance to this work as it describes the interaction between fermions and the Higgs field, thereby generating all the fermion masses. In a compact form, it is expressed as:

$$\mathcal{L}_Y = \bar{\psi} y \varphi \psi + \text{h.c.}, \quad (1.9)$$

where the y terms are the Yukawa couplings between the Higgs boson and the matter fields, and they are proportional to the fermion masses. Consequently, the scattering amplitude of a decay process like $H \rightarrow f\bar{f}$, namely the decay of a Higgs boson to a pair of fermions, will be proportional to the fermion mass m_f at the Leading-Order (LO), and the cross section will be proportional to m_f^2 .

1.2 Phenomenology of the Higgs boson

Given its central role in this work, this section describes the physics of the Higgs boson at the Large Hadron Collider (LHC), focusing on its core properties. A brief introduction on the LHC accelerator complex can be found in Section 2.1. From a phenomenological perspective, the Higgs boson can be produced in experiments through various processes. Their production cross section depends on the experimental conditions, namely the type of colliding particles and the collision energy in the center-of-mass, denoted as \sqrt{s} . The SM predicted values for the production cross section allow for the calculation of the expected rate of produced bosons based on the amount of data acquired. This calculation is essential for every analysis targeting a specific Higgs boson decay channel, as it yields the expected amount of signal events for statistical treatment.

1.2.1 Higgs boson production at the LHC

Each production mode is characterized by a distinct final state topology, which requires a tailored search strategy. Under the conditions of LHC Run-2 with $\sqrt{s} = 13$ TeV, the different modes are listed below in descending order of production cross section for a Higgs boson mass $m_H = 125.00$ GeV [16]:

- **gluon-gluon fusion (ggF):** this is the dominant production mode at the LHC. The Higgs boson is produced from the interaction of two gluons through a quark loop, primarily involving top or bottom quarks. The Feynman diagram is shown in Fig. 1.2a and the cross section is:

$$\sigma(\text{ggF}; \sqrt{s} = 13 \text{ TeV}) = 48.58_{-6.7\%}^{+4.6\%} (\text{th.}) \pm 3.9\% (\text{PDF} \oplus \alpha_s) \text{ pb}, \quad (1.10)$$

where the uncertainties are due to theoretical calculations, Parton Distribution Functions (PDF) and the α_s measurement. This value is calculated at N³LO QCD.

- **Vector Boson Fusion (VBF)**: in this mode, the Higgs boson is produced from the interaction of two vector bosons emitted by the initial state partons. The Feynman diagram is shown in Fig. 1.2b and the production cross section at NNLO QCD is:

$$\sigma(\text{VBF}; \sqrt{s} = 13 \text{ TeV}) = 3.782^{+0.43\%}_{-0.33\%} (\text{scale}) \pm 2.1\% (\text{PDF} \oplus \alpha_s) \text{ pb} \quad . \quad (1.11)$$

The final state is characterized by a pair of forward and angularly separated jets from the hadronization of the two additional partons. The dijet system will also have a high invariant mass, allowing for more efficient discrimination from the multijet background.

- **Higgs-strahlung (VH)**: in this process, the Higgs boson is produced in association with a Z (ZH) or W (WH) vector boson, as shown in the Feynman diagram in Fig. 1.2c. The cross section predictions for these modes at NNLO QCD are:

$$\sigma(\text{WH}; \sqrt{s} = 13 \text{ TeV}) = 1.373^{+0.5\%}_{-0.7\%} (\text{scale}) \pm 1.9\% (\text{PDF} \oplus \alpha_s) \text{ pb} \quad , \quad (1.12)$$

$$\sigma(\text{ZH}; \sqrt{s} = 13 \text{ TeV}) = 0.8837^{+3.8\%}_{-3.1\%} (\text{scale}) \pm 1.6\% (\text{PDF} + \alpha_s) \text{ pb} \quad . \quad (1.13)$$

The additional vector bosons are useful for triggering this process, particularly targeting their subsequent leptonic decay.

- **Associated production with a pair of top or bottom quarks (ttH, bbH)**: in this process, similar to VBF production, the Higgs boson is created from the interaction of an initial state top or bottom quark-antiquark pair, as shown in Fig. 1.2d. The ttH and bbH production modes have a similar cross section values:

$$\sigma(\text{ttH}; \sqrt{s} = 13 \text{ TeV}) = 0.5071^{+5.8\%}_{-9.2\%} (\text{scale}) \pm 3.6\% (\text{PDF} \oplus \alpha_s) \text{ pb} \quad , \quad (1.14)$$

$$\sigma(\text{bbH}; \sqrt{s} = 13 \text{ TeV}) = 0.4880^{+20.2\%}_{-23.9\%} (\text{scale} \oplus \text{PDF} \oplus \alpha_s) \text{ pb} \quad . \quad (1.15)$$

with Eq. 1.14 computed at NLO QCD+EW and Eq. 1.15 at NNLO QCD. The final state includes at least two jets from the hadronization of bottom quarks, which are useful for tagging this particular production mode.

The cross sections of the production modes for different the center-of-mass energy \sqrt{s} and assuming a different Higgs boson mass are shown in Fig. 1.3.

1.2.2 Higgs boson decay channels and main properties measurement

The study and observation of the decay widths and branching fractions of the various Higgs boson decay modes are crucial for the understanding of its properties. The most frequent modes are presented in Table 1.2, along with the associated sources of systematic uncertainty [13].

The Higgs boson primarily decays into a pair of bottom quarks and of W bosons, with branching fractions of approximately 58.2% and 21.5%, respectively. However, even though some decay channels might have smaller branching fractions, their topologies may be particularly clean from SM backgrounds, making them more suitable for the experimental observation. A clear example is the decays $H \rightarrow ZZ^* \rightarrow 4\ell$, known as the “golden channel”, and $H \rightarrow \gamma\gamma$. These decays are

1.2. Phenomenology of the Higgs boson

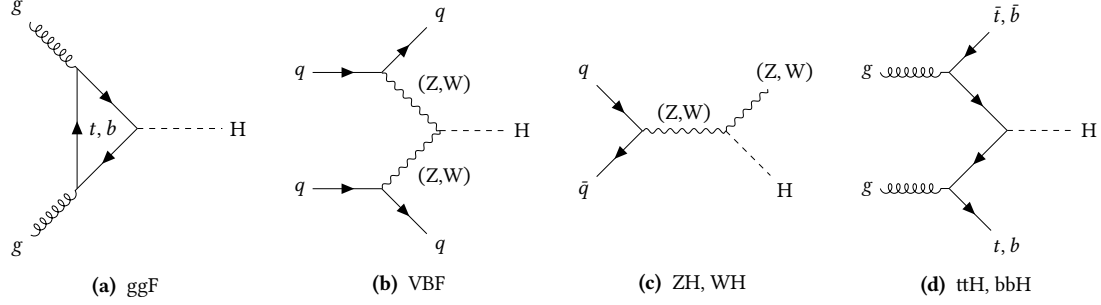


Figure 1.2 • Feynman diagrams of the main production modes of the Higgs boson at the LHC: gluon-gluon Fusion (1.2a), Vector Boson Fusion (1.2b), Higgs-strahlung (1.2c) and ttH/bbH associated productions (1.2d).




Figure 1.3 • Higgs boson production cross section in proton-proton collisions at the LHC for $m_H = 125$ GeV as a function of the center-of-mass energy \sqrt{s} (left) and for a center-of-mass energy $\sqrt{s} = 13$ TeV as a function of the Higgs boson mass (right). Taken from [16].

Decay channel	\mathcal{B}	Relative uncertainty
$H \rightarrow b\bar{b}$	5.82×10^{-1}	+1.2% -1.3%
$H \rightarrow WW$	2.14×10^{-1}	$\pm 1.5\%$
$H \rightarrow gg$	8.19×10^{-2}	$\pm 5.1\%$
$H \rightarrow \tau\tau$	6.27×10^{-2}	$\pm 1.6\%$
$H \rightarrow c\bar{c}$	2.89×10^{-2}	+5.5% -2.0%
$H \rightarrow ZZ^*$	2.62×10^{-2}	$\pm 1.5\%$
$H \rightarrow \gamma\gamma$	2.27×10^{-3}	$\pm 2.1\%$
$H \rightarrow Z\gamma$	1.53×10^{-3}	$\pm 5.8\%$
$H \rightarrow \mu\mu$	2.18×10^{-4}	$\pm 1.7\%$

Table 1.2 • Main decay modes of the Higgs boson with branching fraction and relative uncertainty [13].

among the most interesting processes for studying the main properties of the Higgs boson. All their final state particles can be precisely measured and m_H can be reconstructed with excellent resolution, usually in the order of 1-2%. In particular, the Higgs boson mass has been measured with the highest accuracy by the CMS Collaboration using the $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channel and combining the results from LHC Run-1 and Run-2 [17], yielding:

$$m_H = 125.08 \pm 0.10 \text{ (stat.)} \pm 0.07 \text{ (syst.) GeV} \quad . \quad (1.16)$$

The Higgs boson mass peak and the combined result are shown in Fig. 1.4. Using the same decay channel, the CMS and ATLAS Collaborations have indirectly measured the total decay width of the Higgs boson at a 95% Confidence Level (CL) from the ratio of on-shell to off-shell production [17, 18]. The CMS result reports:

$$\Gamma_H = 2.9^{+2.3}_{-1.7} \text{ MeV} \quad , \quad (1.17)$$

with an upper limit of 7.9 MeV.




Figure 1.4 • Observed data and the expected pre-fit distributions of the four-lepton invariant mass in the inclusive final state (left) and a summary of the observed H boson mass measurements using the four-lepton final state (right). Taken from Ref. [17].

1.2.3 Couplings to gauge bosons and fermions

The Higgs boson couplings to other known particles are reflected in the rates of Higgs boson production and its branching fractions into the various final states. Given a certain production ($pp \rightarrow H$) and decay ($H \rightarrow ff$) mode, the starting point for studying these couplings is to compute the expected number of Higgs bosons produced and decaying through this particular channel [19]:

$$N_{\text{ev}}(pp \rightarrow H \rightarrow ff) \approx \sigma(pp \rightarrow H) \cdot \mathcal{B}(H \rightarrow ff) = \sigma_i \cdot \mathcal{B}_f \approx \frac{\sigma_i \Gamma_f}{\Gamma_{\text{tot}}} \quad , \quad (1.18)$$

where Γ_f and Γ_{tot} are the partial and total decay widths, respectively. The first approximation is possible due to the small decay width of the Higgs boson, which allows the production and decay mechanisms to be factorized. It is a common practice to simplify the notation by introducing

signal strength modifiers to parametrize any potential discrepancies from the SM prediction:

$$\mu_i^f = \frac{\sigma_i \mathcal{B}_f}{(\sigma_i \mathcal{B}_f)^{\text{SM}}} \quad (1.19)$$

The effects of production and decay can be disentangled in Eq. 1.19 using the κ -framework notation [16, 20], where a coupling modifier κ is introduced as:

$$\kappa^2 = \begin{cases} \frac{\sigma_i}{\sigma_i^{\text{SM}}} & \text{for a Higgs boson production,} \\ \frac{\Gamma_i}{\Gamma_i^{\text{SM}}} & \text{for a Higgs boson decay.} \end{cases} \quad (1.20)$$

The CMS and ATLAS Collaborations have measured the couplings to the SM gauge bosons and third-generation fermions with excellent precision [21, 22]. As summarized in Fig. 1.5, a remarkable agreement with the SM expectation has been observed.

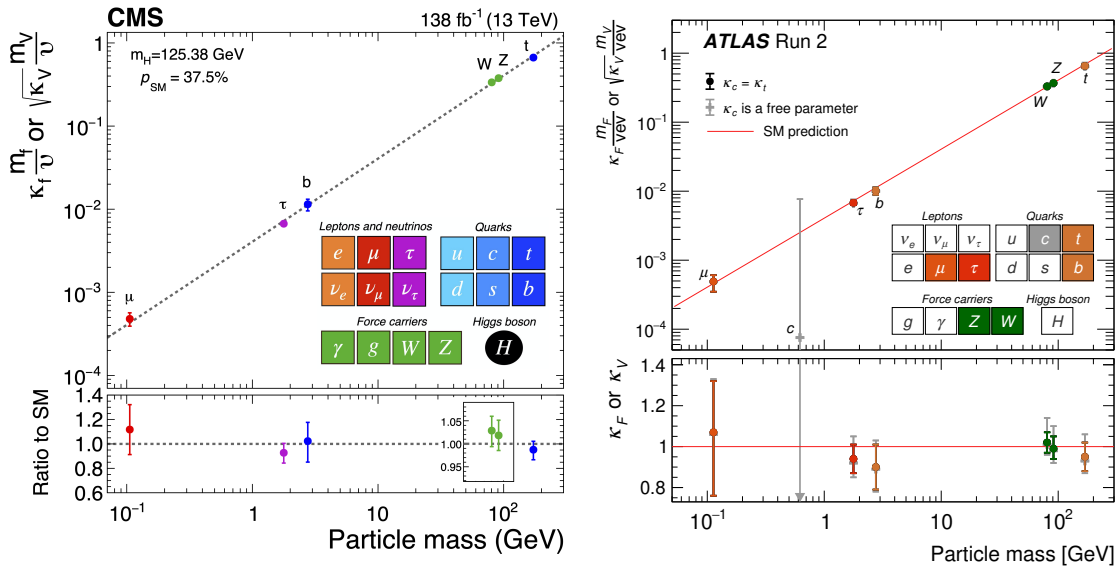


Figure 1.5 • Higgs boson coupling strength modifiers and their uncertainties as measured by CMS (left) and ATLAS (right) Collaborations [21, 22]. Loop-induced processes are assumed to have the SM structure.

One of the goals of the current physics program at the LHC is to measure the couplings to the second generation of fermions. While experiments are starting to approach the measurement of κ_μ , the couplings to the second generation of quarks, κ_c and κ_s , remain unobserved and require more data for their measurement. A constraint at 95% CL of $1.1 < |\kappa_c| < 5.5$ has been set on the coupling modifier for the charm quark by searching for direct $H \rightarrow c\bar{c}$ decays [23, 24]. For the κ_s , a constraint of $|\kappa_s| < 1700$ has been set by investigating the associated Higgs production mode with a W boson [25]. In the same search, constraints on the couplings to the first-generation quarks were set at $|\kappa_u| < 16000$ and $|\kappa_d| < 17000$.

1.3 Rare Higgs boson decays

The Higgs boson decay modes that have already been observed cover more than 90% of its total width. However, there are still rare and yet unobserved decay channels where new and unexpected physics might manifest itself through processes containing loops of SM particles. Studying these rare decays at hadron colliders is experimentally challenging, since it requires a large amount of data and advanced discrimination techniques to identify a potential signal amidst backgrounds arising from QCD processes. Among the possible channels, this work focuses on the decays to a heavy or light meson and a photon or to a pair of mesons, specifically $H \rightarrow \rho\gamma, \phi\gamma, \psi(nS)\gamma, Y(nS)\gamma$ and $H \rightarrow 2\rho, 2\phi, \phi J/\psi, 2J/\psi, Y(nS)Y(mS)$.

1.3.1 Rare decays to a meson and a photon

The SM predicts Higgs boson decays to a meson and a photon via direct and indirect diagrams [26]. In the direct case, the decay occurs through a loop of quarks of the same meson flavor. In the indirect case, the decay proceeds through a $\gamma\gamma^*$ state, where the γ^* converts into a quark-antiquark pair, producing the target meson. Among the possible $H \rightarrow Q\gamma$ channels, Q can be a heavy meson like the $\psi(nS)$ or the $Y(nS)$, or a light meson as the ρ or ϕ . The Feynman diagrams for both the direct and indirect processes for each channel are shown in Fig. 1.6. For the experimental search, the subsequent decays $\psi(nS), Y(nS) \rightarrow \mu\mu$, $\rho \rightarrow \pi\pi$, $\phi \rightarrow KK$ are considered. Thus, the decay branching fraction must be multiplied by a factor corresponding to the $Q \rightarrow ff$ branching fraction, as listed in Table 1.3.

Process	\mathcal{B}_{obs}
$\rho \rightarrow \pi\pi$	$\sim 100\%$
$\phi \rightarrow KK$	$(4.89 \pm 0.05) \times 10^{-1}$
$J/\psi \rightarrow \mu\mu$	$(5.961 \pm 0.033) \times 10^{-2}$
$\psi(2S) \rightarrow \mu\mu$	$(8.0 \pm 0.6) \times 10^{-3}$
$Y(1S) \rightarrow \mu\mu$	$(2.48 \pm 0.05) \times 10^{-2}$
$Y(2S) \rightarrow \mu\mu$	$(1.93 \pm 0.17) \times 10^{-2}$
$Y(3S) \rightarrow \mu\mu$	$(2.18 \pm 0.21) \times 10^{-2}$

Table 1.3 • Branching ratios of the $\rho \rightarrow \pi\pi$, $\phi \rightarrow KK$ and $J/\psi, \psi(2S), Y(nS) \rightarrow \mu\mu$ decays [13].




Figure 1.6 • Leading-Order Feynman diagrams for Higgs rare decays to $Q\gamma$, where $Q = \rho, \phi, \psi(nS), Y(nS)$. The direct process is shown in 1.6a and the indirect processes are shown in 1.6b, 1.6c and 1.6d. For the experimental search, the subsequent decays $\rho \rightarrow \pi\pi$, $\phi \rightarrow KK$ and $\psi(nS), Y(nS) \rightarrow \mu\mu$ are considered.

The calculation of the branching fraction for these decay processes relies on the approximation:

$$\frac{\mathcal{B}(\mathrm{H} \rightarrow Q\gamma)}{\mathcal{B}(\mathrm{H} \rightarrow \gamma\gamma)} \approx \frac{\Gamma(\mathrm{H} \rightarrow Q\gamma)}{\Gamma(\mathrm{H} \rightarrow \gamma\gamma)} , \quad (1.21)$$

decoupling the calculation from the measurement of the Higgs total decay width. For the calculation of $\Gamma(\mathrm{H} \rightarrow Q\gamma)$, a QCD factorization approach is used. For the cases $Q = \psi(\mathrm{nS}), \mathrm{Y}(\mathrm{nS})$, the decay width of these processes can be expressed as the sum of the decay amplitudes for the direct $\mathcal{A}_{\mathrm{dir}}$ and indirect $\mathcal{A}_{\mathrm{ind}}$ contributions:

$$\Gamma(\mathrm{H} \rightarrow Q\gamma) = |\mathcal{A}_{\mathrm{dir}} + \mathcal{A}_{\mathrm{ind}}|^2 , \quad (1.22)$$

where $\mathcal{A}_{\mathrm{dir}}$ and $\mathcal{A}_{\mathrm{ind}}$ are numerically computed using a combination of Non-Relativistic QCD (NRQCD) and Long-Distance Matrix Elements (LDME) approaches [6, 26–29]. The calculation for $\mathrm{H} \rightarrow \rho\gamma, \phi\gamma$ is done using a combination of Soft-Collinear Effective Theory (SCET) and Light-Cone Distribution Amplitude (LCDA) frameworks [6, 30], approximating:

$$\Gamma(\mathrm{H} \rightarrow Q\gamma) = \frac{\alpha f_Q^2}{8m_{\mathrm{H}}} \left(|F_1^Q|^2 + |F_2^Q|^2 \right) , \quad (1.23)$$

where α is the fine structure constant, f_Q is a decay constant that provides information on the QCD interaction between the valence quark and antiquark of the involved meson, and F_i^Q are form factors describing the internal structure of the light meson. The numerical values of the branching fractions for the different decay processes, as well as the best upper limits at 95% CL set to date, are listed in Table 1.4.

Process	$\mathcal{B}_{\mathrm{pred}}$	Upper limit (95% CL)
$\mathrm{H} \rightarrow \rho\gamma$	$1.68_{-0.08}^{+0.08} \times 10^{-5}$	3.74×10^{-4} [31]
$\mathrm{H} \rightarrow \phi\gamma$	$2.31_{-0.11}^{+0.11} \times 10^{-6}$	2.97×10^{-4} [31]
$\mathrm{H} \rightarrow \mathrm{J}/\psi\gamma$	$3.01_{-0.15}^{+0.15} \times 10^{-6}$	2.0×10^{-4} [32]
$\mathrm{H} \rightarrow \psi(2\mathrm{S})\gamma$	$1.03_{-0.06}^{+0.06} \times 10^{-6}$	9.9×10^{-4} [1]
$\mathrm{H} \rightarrow \mathrm{Y}(1\mathrm{S})\gamma$	$5.22_{-1.70}^{+2.02} \times 10^{-9}$	2.5×10^{-4} [32]
$\mathrm{H} \rightarrow \mathrm{Y}(2\mathrm{S})\gamma$	$1.42_{-0.57}^{+0.72} \times 10^{-9}$	4.4×10^{-4} [32]
$\mathrm{H} \rightarrow \mathrm{Y}(3\mathrm{S})\gamma$	$9.1_{-3.8}^{+4.8} \times 10^{-9}$	3.0×10^{-4} [32]

Table 1.4 • Predicted branching fractions for Higgs boson rare decays to a heavy or light meson and a photon, along with the corresponding upper limit at 95% CL.

1.3.2 Rare decays to a pair of mesons

Rare Higgs boson decays to a pair of mesons can occur through a direct and an indirect processes, similar to decays to a meson and a photon, as shown in the diagrams in Fig. 1.7. The available theoretical predictions are dominated by large uncertainties [33–37] and the experimental analysis is more challenging due to the more varied and complex final states.

In the case of a decay to a pair of heavy mesons like $\psi(\mathrm{nS})$ and $\mathrm{Y}(\mathrm{nS})$, the selected final state contains four leptons, particularly four muons (4μ) or two muons and two electrons ($2\mu 2e$). If the decay involves light mesons (ρ, ϕ), the final state consists solely of charged hadrons (π^\pm ,




Figure 1.7 • Leading-Order Feynman diagrams for Higgs rare decays to a pair of mesons. An example of direct process is shown in 1.7a, while a mixed direct and indirect process is shown in 1.7b. Examples of indirect processes are shown in 1.7c and 1.7d.

K^\pm), and the best meson candidates must be identified among all the track combinations. The current trigger technology is not able to select these events. As a matter of fact, no results or theoretical predictions are currently available for the $H \rightarrow 2\rho, 2\phi$ decays. Only a branching fraction prediction is available for the mixed decay $H \rightarrow \phi J/\psi$, for which a potential analysis strategy could target a $2K2\ell$ final state. A list of channels with available upper limit results and predicted branching fractions is provided in Table 1.5.

Process	$\mathcal{B}_{\text{pred}}$	Upper limit (95% CL)
$H \rightarrow \phi J/\psi$	1.0×10^{-9}	—
$H \rightarrow 2J/\psi$	1.5×10^{-10}	3.8×10^{-4} [38]
$H \rightarrow J/\psi \psi(2S)$	$\mathcal{O}(5) \times 10^{-11}$	2.1×10^{-3} [38]
$H \rightarrow 2\psi(2S)$	$(5.1 \pm 2.0) \times 10^{-11}$	3.0×10^{-3} [38]
$H \rightarrow J/\psi Y(1S)$	1.6×10^{-11}	—
$H \rightarrow 2Y(1S)$	2.3×10^{-9}	1.7×10^{-3} [38]
$H \rightarrow 2Y(2S)$	$(1.0 \pm 0.2) \times 10^{-10}$	—
$H \rightarrow 2Y(3S)$	$(5.7 \pm 1.2) \times 10^{-11}$	—
$H \rightarrow Y(mS)Y(nS)$	—	3.5×10^{-4} [38]

Table 1.5 • Predicted branching fractions for Higgs boson rare decays to a pair of mesons, along with best upper limit at 95% CL for every process. Only the processes where an upper limit result or a branching fraction prediction is available are listed.

The incentive to study these rare decays lies in the possibility of BSM physics phenomena affecting very rare partial decay widths. In particular, the channels involving ϕ and $\psi(nS)$ mesons are sensitive to the Yukawa couplings to charm and strange quarks, which have not yet been measured and may therefore exhibit anomalies relative to the SM expectation. Although the branching fractions of these channels are currently beyond experimental sensitivity, unexpected physics could enhance the production rate of these processes, potentially providing hints of extensions to the SM.

2

The Compact Muon Solenoid experiment at the LHC

Probing rare decays of the Higgs boson is an extremely complex challenge that requires a large amount of high-energy collision events. This demand is met by the Large Hadron Collider (LHC) accelerator, which produces hadron collisions at unprecedented energies and luminosity. Four main experiments are housed along the LHC ring, where beams of proton or heavy ion bunches collide. Among them, the CMS experiment has proven particularly successful in this task, leading to precision SM measurements and, together with the ATLAS experiment, the Higgs boson discovery [2–4].

After a general introduction to the LHC in Section 2.1, this chapter focuses on the CMS detector, given its central role in the accomplishment of this work, and on the techniques to reconstruct higher-level physics objects from the low-level raw data produced by its components. The design of the detector has been improved since the start of collider operations in 2008, evolving from the legacy “Phase-0” version used during LHC Run-1 (2011–2012) to its upgraded “Phase-1” design during Run-2 (2015–2018) and Run-3 (2022–2026). Its final design, known as CMS “Phase-2”, will begin operations for the High-Luminosity LHC (HL-LHC) campaign starting in 2030.

2.1 The Large Hadron Collider at CERN

The LHC at the European Organization for Nuclear Research (CERN) is the largest accelerator built to date. It is designed to produce proton-proton and heavy-ion collisions at unprecedented center-of-mass energies, with a maximum value of $\sqrt{s} = 14$ TeV, and high collision rates [39]. The average instantaneous luminosity for Run-3 is around $2 \times 10^{34} \text{ cm}^2\text{s}^{-1}$, and it is expected to reach a maximum of $7.5 \times 10^{34} \text{ cm}^2\text{s}^{-1}$ for the upcoming HL-LHC campaign [40]. The accelerator complex was built with the goal of deepening the understanding of the SM through the discovery of the Higgs boson and precision measurements, as well as pushing the boundaries of physics beyond the SM.

The LHC consists of a two-ring superconducting proton-proton collider housed in a 27 km tunnel, already constructed for the older Large Electron-Positron (LEP) collider. The two high-energy counter-rotating proton beams travel close to the speed of light c . This relativistic regime is progressively reached with the help of other CERN facilities, following the chain LINAC4(160 MeV) \rightarrow PSB(2 GeV) \rightarrow PS(26 GeV) \rightarrow SPS(450 GeV) prior to injection into the LHC. Thousands

of superconducting magnets are used to steer and focus the beams around the accelerator. Dipole magnets bend the beams, while quadrupole magnets act as magnetic lenses to focus them. Before reaching the collision points, where the main experiments are located, the beams are further squeezed to increase the probability of producing collision events of interest. A schematic overview of the accelerator complex, the beam path, and the main experiments along the ring is shown in Fig. 2.1, with technical details available in Refs. [39, 40].

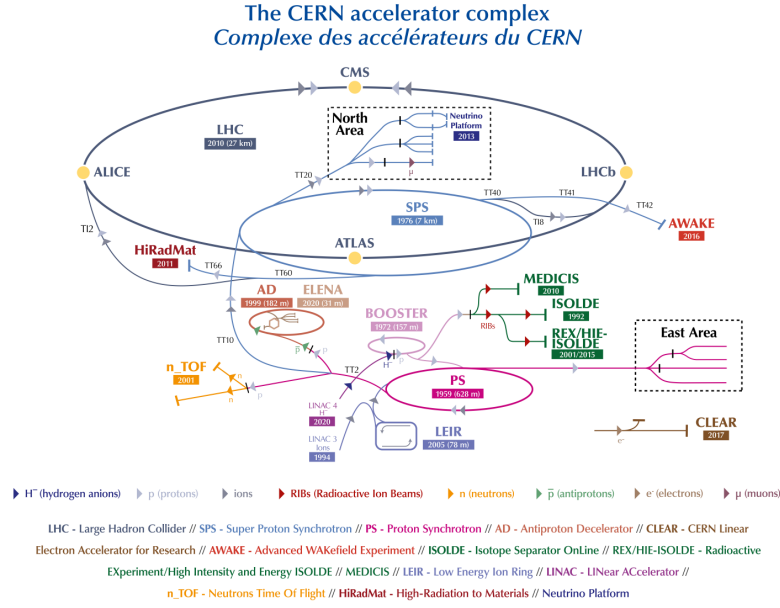


Figure 2.1 • Schematic view of CERN facilities. The LHC is the last ring (dark blue line) in a complex chain of particle accelerators, alongside smaller experiments designed for specific physics purposes. Taken from Ref. [41].

2.1.1 LHC filling scheme

The two circulating beams in the LHC consist of bunches, each containing approximately 10^{11} protons under nominal conditions, with a transverse distribution that follows a two-dimensional gaussian profile. The bunches are injected into the accelerator with a temporal spacing of 25 ns, completing one revolution around the LHC tunnel in 89.1 μs . This revolution is referred to as an LHC orbit, which can contain up to 3564 bunches. The bunch crossing (BX) frequency between the two beams is $f_{\text{BX}} = 40$ MHz. However, the bunch structure of the beam is not continuous and a maximum of around 2400 bunches is filled during nominal operations of Run-3. The bunches are consecutively organized into “trains”, separated by gaps, though some may be placed individually between two trains as “isolated bunches”. The final few hundred bunches in the orbit are part of the “orbit gap”, where no collisions are expected. The gap is used for synchronization, calibration data acquisition, and to reset the front-end electronics of the experiment where the bunches collide. A schematic representation of the structures in a typical LHC filling scheme is shown in Fig. 2.2.




Figure 2.2 • Overview of LHC filling scheme structure, showing trains of bunches, isolated bunches, gaps between consecutive trains, and the orbit gap at the end of the orbit.

2.1.2 Instantaneous and integrated luminosity

For any given physics process, the event rate is proportional to both the production cross section of the process and the luminosity of the machine. Generally, the production cross section for events of interest varies depending on the center-of-mass energy. The luminosity of the machine depends solely on the beam parameters and its instantaneous value for a gaussian beam distribution can be expressed as:

$$\mathcal{L}_{\text{inst}} = \frac{N_b^2 n_b f_{\text{rev}} \gamma_r}{4\pi \epsilon_n \beta^*} \cdot F \quad , \quad (2.1)$$

where N_b is the number of protons per bunch, n_b is the number of colliding bunches per beam, f_{rev} is the orbit revolution frequency (11.2 kHz), γ_r is the relativistic Lorentz factor and ϵ_n represents the normalized emittance. β^* is the beta function at the interaction point, namely the distance from the interaction point to the point where the beam width is twice as wide as at the interaction point, and a typical value during normal operations is 0.30 cm. F is the geometric luminosity reduction factor due to the crossing angle at the interaction point.

The instantaneous luminosity is measured by the experiments along the LHC ring using detectors positioned at small angles relative to the interaction points. It is not constant during a data taking run as both beam intensity and emittance degrade over time of operation. The total amount of data delivered to an experiment is measured by integrating the instantaneous luminosity during a data taking run:

$$\mathcal{L}_{\text{int}} = \int_{t_{\text{start}}}^{t_{\text{end}}} \mathcal{L}_{\text{inst}}(t) dt \quad . \quad (2.2)$$

The CMS experiment recorded an integrated luminosity of approximately 138 fb^{-1} during Run-2, and by 2024, combining data from Run-2 and Run-3, the total has surpassed 300 fb^{-1} , as shown in Fig. 2.3.

Based on the average instantaneous luminosity recorded by the CMS experiment and the Higgs boson production cross section reported in Section 1.2.1, the LHC has produced approximately 1 Higgs boson per second during LHC Run-2. Under Run-3 conditions, this estimate is slightly higher and equal to around 1.2 bosons per second. With the HL-LHC upgrade, this rate will increase to around three Higgs bosons per second, assuming an average instantaneous luminosity of $5.0 \times 10^{34} \text{ cm}^2 \text{ s}^{-1}$ and the inclusive Higgs production cross section at $\sqrt{s} = 14 \text{ TeV}$ [16]. The total number of rare Higgs boson decays $H \rightarrow Q\gamma$ during Run-2, along with the projected estimate for the Run-3 using 300 fb^{-1} of integrated luminosity, and for one year of HL-LHC operations,




Figure 2.3 • Cumulative measured luminosity by the CMS experiment during LHC Run-2 and Run-3. The plot is taken from the public results of the CMS Luminosity Physics Object Group [42], where similar plots are available for each year of data-taking.

corresponding to an acquired amount of 400 fb^{-1} , are presented in Table 2.1. These estimates illustrate the LHC role as “Higgs factory”, enabling searches of the rare Higgs decays under study in this work, which would not be possible to explore in other machines such as electron-positron colliders [6].

Process	Approximate number of expected events		
	Run-2 (138 fb^{-1})	Run-3 ($\sim 300 \text{ fb}^{-1}$)	HL-LHC (1y, $\sim 400 \text{ fb}^{-1}$)
$H \rightarrow \rho\gamma$	1.3×10^2	3.0×10^2	4.2×10^2
$H \rightarrow \phi\gamma$	1.8×10^1	4.1×10^1	5.8×10^1
$H \rightarrow J/\psi\gamma$	2.3×10^1	5.4×10^1	7.5×10^1
$H \rightarrow \psi(2S)\gamma$	7.9	1.8×10^1	2.6×10^1
$H \rightarrow Y(1S)\gamma$	4.0×10^{-2}	9.3×10^{-2}	1.3×10^{-1}
$H \rightarrow Y(2S)\gamma$	1.1×10^{-2}	2.5×10^{-2}	3.6×10^{-2}
$H \rightarrow Y(3S)\gamma$	7.0×10^{-2}	1.6×10^{-1}	2.3×10^{-1}

Table 2.1 • Approximate number of expected $H \rightarrow Q\gamma$ events during LHC Run-2, Run-3 and one year of HL-LHC. The estimates do not include detector geometric acceptance factors.

2.1.3 Pileup collisions

While beam squeezing near the collision point increases the likelihood of interactions, it also introduces a significant experimental challenge: an additional number of proton-proton collisions per bunch crossing, known as pileup (PU). When two proton bunches collide, multiple collisions overlap in the detector, surrounding the primary collision of interest. This overlap complicates the reconstruction of the hard scattering vertex and degrades the overall performance of physics object reconstruction. As a result, pileup subtraction algorithms are applied to mitigate these effects, when necessary.

2.2. The Compact Muon Solenoid experiment

In 2018, the CMS experiment recorded an average pileup of $\langle\mu\rangle = 32$. An average of $\langle\mu\rangle = 57$ was observed in 2024. During the operation of HL-LHC, an average peak pileup of $\langle\mu\rangle = 200$ is expected. The pileup distribution for different years of data acquisition is shown in Fig. 2.4.




Figure 2.4 • Pileup distribution observed by the CMS experiment across different years of data taking. Taken from Ref. [42].

2.2 The Compact Muon Solenoid experiment

The CMS experiment is a general-purpose detector built around a huge solenoid magnet, shaped like the form of a cylindrical coil of superconducting cable that generates a 3.8 Tesla field in the inner region of the solenoid. The field is confined by a steel “yoke” that forms the bulk of the detector, which weighs about 14000 tons and was built in 15 sections reassembled on-site. The detector gets its name from its compact dimensions (15 meters high and 21 meters long), from its design specifically made to detect muons with high accuracy and efficiency, and from its powerful solenoid magnet.

The detector is shaped like a cylindrical onion, with several concentric layers of components, as seen in Fig. 2.5 from a 3D view and in Fig. 2.6 from an octant of its transverse section. These components are designed to record a collision event by determining the properties of the particles produced in that particular collision. This task is done by:

- **Bending the trajectory** of charged particles as they move outwards from the collision point through the solenoidal magnet. This bending is required to identify the charge and measure the momentum of a particle. Positively and negatively charged particles curve in opposite directions in the same magnetic field, and high momentum particles have a larger bending radius compared to low-momentum ones.
- **Identifying tracks**, which correspond to the trajectories of charged particles. Precise tracking is essential for good reconstruction accuracy and resolution. This task is performed by the silicon pixel detector, silicon strip tracker, and muon detectors.

Chapter 2. The Compact Muon Solenoid experiment at the LHC




Figure 2.5 • 3D model of the CMS detector with the 2017 geometry in SketchUp, a 3D modelling program. Taken from Ref. [43].

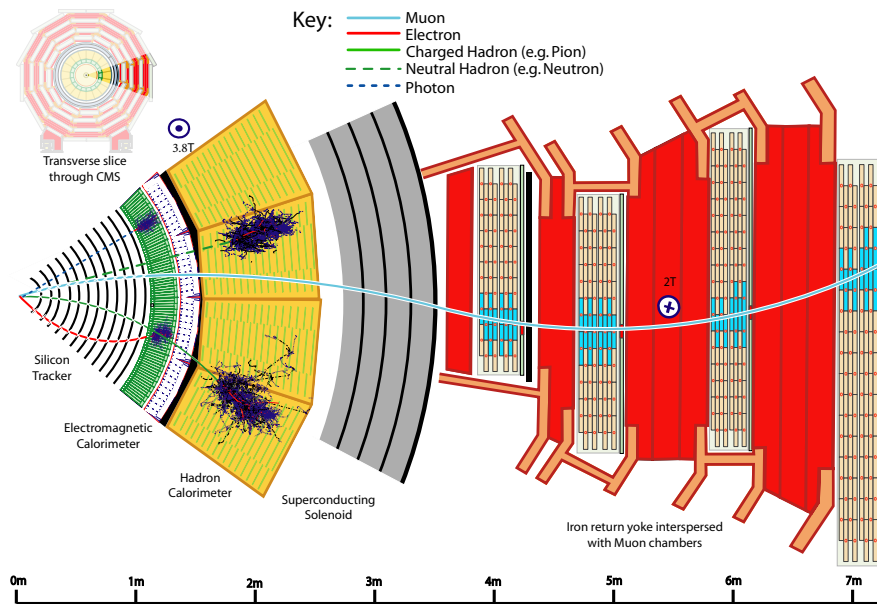


Figure 2.6 • Slice of the CMS detector showing sample particles it can detect. Taken from Ref. [44].

- **Measuring the energy** of product particles, which is done by two types of calorimeter detectors. The inner one is the Electromagnetic Calorimeter (ECAL), which measures the energy of electrons and photons by stopping them completely. Hadrons and jets, which are showers of particles resulting from parton hadronization, pass through the ECAL and are stopped in the Hadron Calorimeter (HCAL), which surrounds the ECAL.
- **Detecting muons**, which penetrate both calorimeters without being stopped. Their momentum can be measured inside the superconducting coil by the tracking devices, if the muon is not displaced from the interaction point, and outside of it by the muon chambers.

By combining the information from all these components, the detector can measure and identify the energy and momentum of most particles produced in the collisions, except for neutrinos, which are detected indirectly by measuring the missing energy in the transverse plane (MET). The following discussion describes the main detector components and their reconstruction performance.

2.2.1 Silicon tracker

The CMS Silicon Tracker allows to track the paths of particles emerging from a collision with extreme accuracy and in three-dimensions. This reconstruction poses a complex challenge, as the number of particles passing through the detector near the interaction point is huge. In addition to this reconstruction task, the detector must meet strict technical specifications to withstand the harsh radiation over long acquisition times.

The Silicon Tracker is composed of different substructures, namely a pixel detector and a silicon microstrip detector. Since the pixel detector was upgraded between the 2016 and 2017 campaigns, it is necessary to distinguish the descriptions of the pixel detector between the two versions.

The first pixel detector version, used for 2016 data acquisition, measured charged particles within the pseudorapidity range $|\eta| < 2.5$. It consisted of 1440 silicon pixel modules, arranged in 3 barrel layers at radii of 44, 73, and 102 mm, and two endcap disks on each end at distances of 345 and 465 mm from the interaction point. The strips are arranged in four inner barrel layers (TIB), two inner disks (TID), six concentric outer barrel layers (TOB), and two endcaps of nine layers (TEC), which close off the tracker. The upgraded tracker, used for 2017 and 2018 data acquisition, measures charged particles within the pseudorapidity range $|\eta| < 3.0$ and is shown in Fig. 2.7. In this case, the new pixel detector consists of 1856 silicon pixel modules, arranged in four barrel layers at radii of 29, 68, 109, and 160 mm, and three disks on each end at distances of 291, 396, and 516 mm from the center of the detector [45].

2.2.2 Hadron and electromagnetic calorimeters

The Electromagnetic Calorimeter (ECAL) detects photons and electrons and measures their energies with high accuracy. These particles are particularly interesting when they are very energetic because they could indicate the presence of a Higgs boson or other interesting SM particles in a collision event.

The ECAL is divided into two regions: the barrel (EB), providing coverage in the region $|\eta| < 1.479$ and shown in Fig. 2.8a, and the endcaps (EE), covering the region $1.479 < |\eta| < 3.000$ and shown in Fig. 2.8b. About $61200 \times 2 \times 2 \times 23 \text{ cm}^3$ lead tungstate (PbWO_4) crystals cover the

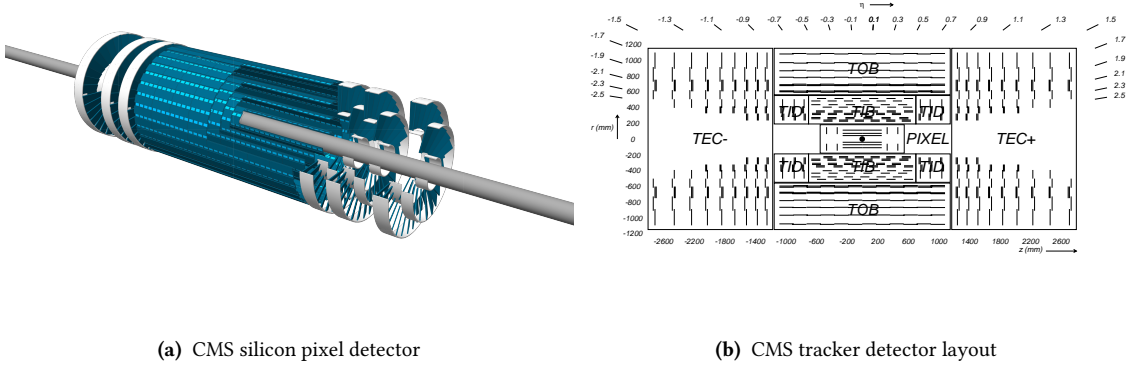


Figure 2.7 • 3D model of the CMS silicon tracker detector (2.7a), taken from [43], and general layout (2.7b), taken from [46]. In particular, in 2.7a the silicon pixel detector with the geometry of Phase-1 is visualized.

barrel region, while $14648 \times 3 \times 22 \text{ cm}^3$ crystals of the same material are placed in the endcaps. Visible light is produced when a photon or electron traverses these dense components. The light is then collected and converted into an electric signal through an avalanche photo-diode. The signal is finally converted into a measure of energy through a calibration procedure.

The Hadronic Calorimeter (HCAL) is responsible for measuring the energy of hadrons and it is shown in Fig. 2.8c. This component is important since it provides an indirect measurement of the presence of non-interacting neutral particles, such as neutrinos, and can provide information on the production of new particles that decay into products that may not leave traces in the detector. The HCAL must contain all particles produced by the collisions to detect those invisible particles, so it must be “hermetic”. When a particle is observed in a direction of the detector but is not balanced by any other particle, resulting in an imbalance of transverse momentum and energy, it can be deduced that invisible particles have been produced in the collision.

The HCAL detector is composed of five distinct components: the barrel (HB), covering the region $|\eta| < 1.3$, two endcaps (HE), covering $1.3 < |\eta| < 3.0$, the outer barrel (HO), complementing the HB detector in the pseudorapidity region $|\eta| < 1.4$, and the Hadron Forward calorimeter (HF), completing the pseudorapidity coverage in $3.0 < |\eta| < 5.0$. To prevent familiar particles from escaping undetected, the HCAL is composed of staggered layers alternating between “absorber” and fluorescent “scintillator” materials. This configuration forms the basis of a sampling calorimeter. When a hadron passes through the active material, it produces a rapid light pulse, which is collected by an optical fiber and sent to a photodetector, where its electronic signal is then amplified. The total amount of light generated by the passage of a particle can be converted into a measure of energy through a specific calibration.

2.2.3 Muon system

Muons can penetrate several meters of iron without interacting, and they are not stopped by any of the CMS calorimeters. The muons are measured in the pseudorapidity range $|\eta| < 2.4$ using detection planes that employ different technologies: drift tubes, cathode strip chambers, resistive plate chambers and, starting from Run-3, gas electron multipliers. These components are placed

2.2. The Compact Muon Solenoid experiment




Figure 2.8 • 3D models of the CMS calorimeters, taken from [43]. In particular, in 2.8a and 2.8b the ECAL barrel and endcap are shown, respectively, while the HCAL is shown in 2.8c.

externally to the other detectors of CMS since muons are the only particles likely to produce a signal in this part of the structure. This specialized system for muon detection is displayed in Fig. 2.9, with a quarter of its transverse section along the beam direction shown in Fig. 2.9b.

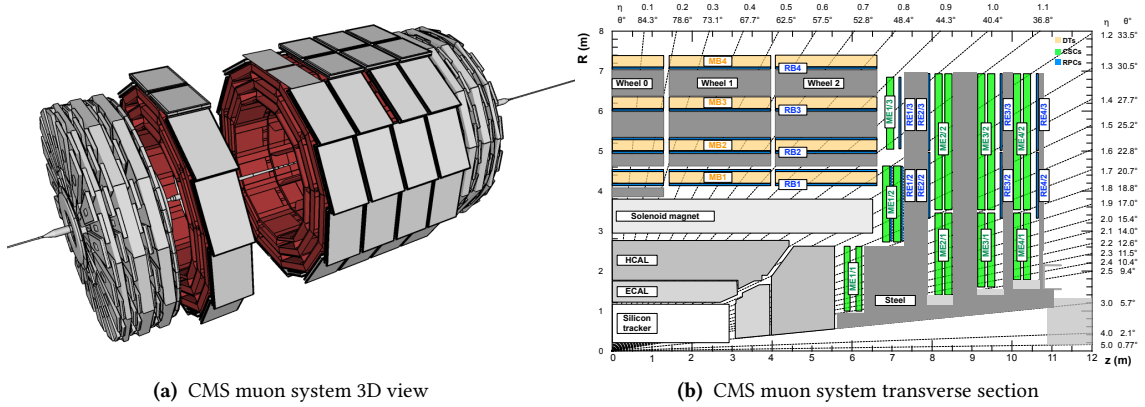


Figure 2.9 • In 2.9a, a 3D view of the CMS muon system, taken from [43]. In 2.9b, a quadrant of the section of the CMS detector along the beam direction, showing the configuration of DTs (orange), CSCs (green) and RPCs (blue) in the detector design.

The Drift Tubes (DTs) system measures muon positions in the barrel region of the detector. Each cell is 4.2 cm wide, filled with a mixture of Argon and CO_2 , and contains a stretched wire. When a muon or any charged particle passes through the gas volume, it creates electron-ion pairs from the gas atoms. The electrons (ions) follow the electric field, ending up at the positively (negatively) charged wire. By knowing the drift time inside the gas, two-dimensional spatial information is obtained and combined with signals from other tubes. For this reason, the DTs are grouped in $2 \times 2.5 \text{ m}^2$ chambers, each consisting of 12 aluminum layers arranged in three groups of four, with each containing up to 60 cells. The middle group measures the coordinate along the direction parallel to the beam, and the two outer groups measure the perpendicular coordinates.

Cathode Strip Chambers (CSC) are arrays of positively charged anode wires crossed with negatively charged copper cathode strips within a gas volume. This component is used in the endcap disks where the magnetic field is uneven and particle rates are high. When a muon hits the chamber, it ionizes the gas, creating an electron avalanche towards the anode wire. The positively

charged ions, in turn, migrate perpendicularly, providing two-dimensional spatial information. Each CSC module contains six layers, enabling accurate identification of muons and matching their tracks to those in the tracker.

Resistive Plate Chambers (RPC) consist of two parallel plates, a positively charged anode and a negatively charged cathode, made of a very high-resistivity plastic material and separated by a gas volume. When a muon passes through the chamber, it ionizes the gas, creating an electron avalanche. The electron signal is collected by external metallic strips after a small but precise time delay. The pattern of hit strips provides a quick measure of the muon momentum, which is then used by the trigger system to make immediate decisions about whether the event is worth keeping.

Since the start of Run-3, the muon system has been upgraded with the addition of Gas Electron Multiplier (GEM) detectors in the pseudorapidity region $1.6 < |\eta| < 2.2$ [47]. The GEMs are trapezoidal-shaped chambers filled with a gaseous mixture of Argon and CO_2 , segmented into eight η sections with 384 readout strips, each with an angular width of $230 \mu\text{rad}$, radially placed with a $463 \mu\text{rad}$ pitch. A total of 72 GEM chambers are placed in each endcap, improving the spatial and timing resolution of muon detection.

2.2.4 Trigger and Data acquisition systems

Given the average number of additional collisions per bunch crossing, the average event size is approximately 1 MB for LHC Run-2 and 2 MB for Run-3. When LHC is operating at its peak, the event rate can approach 30 MHz, resulting in a raw data rate of $\mathcal{O}(10)$ TB/s. It is not feasible to store all these event records under these conditions with the current storage and detector technology. Given a raw data bandwidth limit of $\mathcal{O}(2)$ GB/s, an event rate reduction factor of $\mathcal{O}(10^5)$ has to be applied. This task is accomplished by CMS through a two-level trigger system, selecting only the most interesting events for further analysis

Level-1 Trigger The first level (L1T) consists of Field Programmable Gate Array (FPGA) based processors and uses information from the calorimeters and muon detectors to select events at a rate of around $\mathcal{O}(100)$ kHz within a fixed latency of about $4 \mu\text{s}$ [48–50]. The limited rate budget is driven by the sub-detectors readout, which cannot operate at a higher rate. The L1 trigger is further divided into three major subsystems: the muon trigger, the calorimeter trigger and the global trigger. These systems and their performance for CMS Phase-1 and Phase-2 will be described more in depth in Chapter 4.

High-Level Trigger The second level, known as the High-Level Trigger (HLT), consists of a farm of processors running a version of the full event reconstruction software optimized for fast processing, and it reduces the event rate to around 1 kHz before data storage [48–51]. Specifically, the rate reduction begins with an acceptance signal from the L1 trigger. Then, the raw data from the detectors are transferred to an event builder system. The reconstructed event is transferred to a filter stage where the HLT algorithms are applied. The output is divided into several online streams depending on the trigger algorithm and they are sent to the storage system.

2.3 Event reconstruction with the CMS detector

The raw data saved on permanent storage undergoes a step of particle identification, handled by the Particle Flow (PF) event reconstruction algorithm [52–54], which correlates the information from all the subdetectors. The reconstruction workflow starts with the unpacking and decoding of low-level detector data, such as the calorimetric towers and hits in the silicon and muon detectors. Pattern recognition, track finding and track fitting algorithms are used to obtain higher-level objects for physics analysis, whose reconstruction is described more in detail in the next subsections. An example of a particle identification task performed by Particle Flow is shown in Fig. 2.10.




Figure 2.10 • Examples of Particle Flow algorithm application, taken from Ref. [52]. A jet reconstruction in a simulated “dijet” event is shown, with the particles clustered in the two PF jets displayed with thicker lines, while ECAL and HCAL deposits are in orange and purple, respectively.

2.3.1 Tracks and vertices

Charged particle tracks are reconstructed using a Combinatorial Track Finder (CTF) based on a Kalman Filter (KF) method [55, 56], applied in three stages:

- Seed generation with a few hits compatible with a charged-particle trajectory.
- The initial estimate of the trajectory parameters is propagated layer by layer through the detector, looking for compatible hits, using the KF technique.
- Final fitting to determine the track parameters and removal of incompatible hits collected during the pattern recognition step.

The reconstructed tracks are clustered into sets likely originating from the same interaction points to identify primary (PV) and secondary vertices (SV). The constituents of each cluster are then refitted again to improve the resolution on the vertex position. The resulting vertices are ranked based on the quadratic sum of the transverse momenta of the daughter tracks, with the primary

vertex identified as the one with the highest score for the quadratic sum. The challenge of this approach is the increasing number of candidate vertices with increasing pileup, which degrades the vertex resolution.

With the pixel detector up to 2016, for non-isolated particles with $1 < p_T < 10$ GeV and $|\eta| < 1.4$, the track resolutions were typically 1.5% in p_T and 25–90 (45–150) μm in the transverse (longitudinal) impact parameter [57]. With the new pixel detector from 2016, for non-isolated particles of $1 < p_T < 10$ GeV and $|\eta| < 3.0$, the track resolutions are typically 1.5% in p_T and 20–75 μm in the transverse impact parameter [58].

2.3.2 Muons

The CMS detector is designed to achieve excellent muon reconstruction performance. The employed algorithm distinguishes between independent reconstruction in the silicon tracker, where the PF approach is used to produce “tracker tracks”, and in the muon chambers, where an independent approach produces “standalone tracks”. The two types of tracks are then combined to obtain three different categories of reconstructed muons:

- **Standalone muons**, built using muon segments obtained only from hits in the muon system.
- **Tracker muons**, which are tracks with $p_T > 0.5$ GeV extrapolated from the tracker to the muon system and matched with at least one muon segment from a standalone muon.
- **Global muons**, built from a standalone track matched to a tracker track, extrapolated to a common surface, and refitted with the hits from the matching tracker and standalone tracks.

The efficiency to reconstruct and identify muons is greater than 96%. Matching muons to tracks measured in the silicon tracker results in a relative transverse momentum resolution of 1% in the barrel and 3% in the endcaps for muons with p_T up to 100 GeV. The p_T resolution in the barrel is better than 7% for muons with p_T up to 1 TeV [59].

2.3.3 Electrons, photons, jets and taus

Electrons and photons These two types of particle deposit almost all of their energy in the ECAL, unlike hadrons, which are expected to deposit most of their energy in the HCAL. Unlike electrons, photons do not leave a track in the tracker system since they are neutral. Both manifest as showers of multiple electrons and photons rather than as single particles, as they interact with the material in front of the ECAL during propagation. Thus, electrons radiate bremsstrahlung photons, and photons convert into electron-positron pairs, producing energy clusters that are broader in φ and narrower in η .

The electrons and photon reconstruction algorithm begins with the clustering of crystals where energy significantly higher than the electronic noise was deposited. A seed cluster is then identified as the one containing the most energy deposited in a specific region. A group of clusters around the seed is combined into a supercluster to include products of photon conversions and electron bremsstrahlung. Electrons are expected to have a compatible tracker track, so trajectory seeds in the pixel detector, compatible with the supercluster position and an electron trajectory,

are used to seed a Gaussian Sum Filter (GSF) tracking algorithm [60]. All tracks reconstructed in the event compatible with an electron hypothesis are added as seeds of the GSF tracking. Next, the ECAL clusters, superclusters, GSF seed tracks, including those likely to come from photon conversions and their associated clusters, are given to the PF reconstruction. This algorithm returns blocks of particles and identifies a candidate as an electron if it satisfies a set of loose selection requirements with an associated GSF track, or as a photon if there is no associated GSF track. To differentiate the e/γ from hadrons, tighter selections are applied to determine if they are accepted as electrons or isolated photons.

Thanks to this combined approach, the energy measurement of electrons and photons achieves remarkable resolution. For electrons from $Z \rightarrow ee$ decays with $p_T \approx 45$ GeV it ranges between 1.6 and 5%, generally better in the barrel region than in the endcaps [61, 62]. An energy resolution of about 1% is achieved for unconverted or late-converting photons in the barrel region and in the tens of GeV energy range. The remaining barrel photons are reconstructed with a resolution of about 1.3% up to $|\eta| = 1$, changing to about 2.5% at $|\eta| = 1.4$. In the endcaps, the resolution is estimated to be about 2.5% for unconverted or late-converting photons and between 3% and 4% for the other photons [63].

Jets For each event, hadronic jets are clustered from reconstructed charged hadrons using the infrared and collinear safe anti- k_T algorithm [64, 65] with a cluster parameter of $R = 0.4$. Jet momentum is determined as the vector sum of all particle momenta in the jet and is found from simulation to be, on average, within 5% to 10% of the true momentum over the entire p_T spectrum and detector acceptance. The additional proton-proton collisions from pileup can contribute extra tracks and calorimetric energy depositions to the jet momentum. To mitigate this effect, charged particles identified as originating from pileup vertices are discarded, and an offset correction is applied to account for remaining contributions. Jet energy corrections are derived from simulations to bring the measured response of jets in line on average with that of “particle-level” jets⁴. In situ measurements of the momentum balance in dijet, $\gamma + \text{jet}$, $Z + \text{jet}$, and multijet events are used to account for any residual differences in the jet energy scale (JES) between data and simulation [66].

The jet energy resolution (JER) typically amounts to 15–20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV [66]. Additional selection criteria are applied to each jet to remove those potentially dominated by anomalous contributions from various subdetector components or reconstruction failures.

Hadronic taus Hadronically decaying τ particles, referred to as τ_h , are reconstructed from jets using the “hadrons-plus-strips” algorithm [67]. This algorithm combines 1 or 3 tracks with matching energy deposits in the calorimeters, covering the main τ fully leptonic and hadronic decay modes. To distinguish genuine τ_h decays from jets originating from the hadronization of quarks or gluons, as well as from electrons or muons, the DEEPTAU algorithm is used [68]. Its misidentification rate depends on the p_T and quark flavor of the jet. In simulated events from W

⁴By “particle-level” definition it is implied a collection of stable particles from full matrix element plus parton shower generators, without any simulation of the interaction of these particles with the detector components or any additional proton-proton interactions.

Chapter 2. The Compact Muon Solenoid experiment at the LHC

boson production in association with jets, the misidentification rate is estimated to be 0.43% for a genuine τ_h identification efficiency of 70%. The misidentification rate for electrons (muons) is 2.60% (0.03%) for a genuine τ_h identification efficiency of 80% (>99%).

3

Search for rare Z and Higgs boson decays to a charmed meson and a photon

This chapter describes the search for decays of the Higgs and Z bosons to a J/ψ or $\psi(2S)$ meson (collectively referred to as $\psi(nS)$) and a photon, one of the rare $H \rightarrow Q\gamma$ decays introduced in Section 1.3.1. The significance of this channel lies in its potential to probe the coupling of the Higgs boson to the charm quark, as the direct decay is expected to occur through a charm quark loop. The channel provides a better signal-to-background ratio than the Higgs boson decay to a pair of c quarks hadronizing into jets [23, 24], thanks to the very low SM backgrounds and the excellent muon momentum and photon energy resolution of the CMS detector. It is also very rare, thus requiring a large amount of data. The yet outstanding observation of the $Z \rightarrow \psi(nS)\gamma$ process, which is predicted to occur at a rate $\mathcal{B}(Z \rightarrow J/\psi\gamma) = 8.96_{-1.38}^{+1.51} \times 10^{-8}$ and $\mathcal{B}(Z \rightarrow \psi(2S)\gamma) = 4.83_{-0.91}^{+1.02} \times 10^{-8}$ [27], would represent an intermediate goal that can be achieved with a reasonable amount of data, due to the Z production cross section being 10^3 times higher than that of the Higgs boson. This observation would provide a crucial benchmark for the theoretical prediction framework, so this additional channel is included in the search.

The search for the decays $Z \rightarrow \psi(nS)\gamma$ and $H \rightarrow \psi(nS)\gamma$ was previously performed by both the ATLAS and CMS Collaborations using data collected in 2016 [69, 70] and the ATLAS Collaboration recently updated its search using the complete Run-2 dataset [32]. The present search is performed with data from proton-proton collisions at a center-of-mass energy of 13 TeV, corresponding to an integrated luminosity of 123 fb^{-1} , collected by the CMS detector during the LHC Run-2. Since more collision events and more efficient trigger algorithms are required to achieve the sensitivity needed for a potential observation, the aim of this work is to set an upper limit on the cross section of these processes. The experimental observation of these rare decays of the Z and Higgs bosons might become a reality with the Run-3 and HL-LHC campaigns, supported by the possible implementation of dedicated trigger algorithms or the application of an innovative system for the real-time analysis of all collision events, as described in the next chapters.

3.1 Analysis strategy and datasets

The analysis targets the subset of $Z, H \rightarrow \psi(nS)\gamma$ decays where the meson subsequently decays to a pair of muons, resulting in a final state of $\mu\mu\gamma$. Thus, the expected signal is characterized by two narrow resonances visible in the $m_{\mu\mu}$ and $m_{\mu\mu\gamma}$ invariant mass spectra. The former will

peak around the $\psi(\text{nS})$ meson mass, the latter around the decaying H or Z boson mass. These resonances can be fully reconstructed thanks to the excellent resolution of the CMS detector in measuring photon energy and muon momentum, as highlighted in Section 2.3.

The two narrow signal peaks are expected to appear over different SM background contributions:

- A QCD continuum background, where the photon is emitted from charged quarks, and the muons either come from heavy flavor semileptonic decays or from punch-through or decays in flight of charged hadrons. This background does not exhibit a peak in either $m_{\mu\mu}$ or in $m_{\mu\mu\gamma}$ invariant mass spectra.
- A QCD continuum background where a real $\psi(\text{nS})$ meson decay is present, exhibiting the corresponding peak in $m_{\mu\mu}$, and a photon comes from the initial or final state.
- Electroweak backgrounds that produce a pair of muons and a photon, consisting of Z and H bosons that decay through quark loops without the intermediate resonant $c\bar{c}$ state, referred to as “Dalitz decays”, and Z decays to a pair of muons with a Final State Radiation (FSR) photon. These backgrounds resonate in $m_{\mu\mu\gamma}$ but not in $m_{\mu\mu}$.
- A non-resonant electroweak background where the Z boson decays to a pair of muons and the photon is emitted as Initial State Radiation (ISR).

While the third and fourth contributions can be modeled using Monte Carlo samples as for the signals, a sufficiently large and reliable sample is not available for the first two QCD backgrounds, making it necessary to estimate them using a data-driven approach. Additionally, cascade decays $\psi(2S) \rightarrow J/\psi + X \rightarrow \mu\mu + X$, although they provide a non-negligible number of events compared to the direct decays $\psi(2S) \rightarrow \mu\mu$, are not considered in this work. As a matter of fact, the momentum carried away by the additional particles (X) causes a shift towards smaller $m_{\mu\mu\gamma}$ values, not contributing to the peak of the expected signal. Moreover, the additional X hadron would reduce the selection efficiency as it would reduce the isolation of the muon pair from hadronic activity, which is a key signal-versus-background discriminating factor in the search.

3.1.1 Datasets for the search

The sample analyzed has been collected by the CMS experiment with a “single muon plus photon” ($\mu + \gamma$) trigger algorithm that selects events with at least one muon with $p_T > 17$ GeV and a photon with $p_T > 30$ GeV. This trigger was not active at the beginning of 2017, resulting in the loss of 14 fb^{-1} of data, leading to a total integrated luminosity for this analysis of 123 fb^{-1} . As a cross check, an alternative dataset is considered for the search requiring the presence of two muons with an invariant mass value in the region of the J/ψ or $\psi(2S)$ and a combined $p_T > 25$ GeV.

3.1.2 Simulated signal and background samples

The $Z/H \rightarrow \psi(\text{nS})\gamma$ signal Monte Carlo (MC) simulated events are produced at LO accuracy in QCD with the `MADGRAPH5_AMC@NLO v2.6.5` matrix element generator [71], directly including the spin correlations [72] using a model based on Refs. [27, 30] instead of reweighting the simulated event as in previous analyses [32, 69, 70]. Whereas the $Z \rightarrow \psi(\text{nS})\gamma$ production is simulated through quark-antiquark annihilation ($q\bar{q} \rightarrow Z$), the simulation of the H boson is performed

separately for each production mode, namely through gluon-gluon fusion (ggF), vector-boson fusion (VBF), vector boson associated production (WH and ZH), and bottom and top quark associated productions (bbH and ttH). For each boson and production mode, two different samples are generated, imposing the Z or H boson to decay to $J/\psi\gamma$ or $\psi(2S)\gamma$, where $\psi(nS)$ is forced to decay to a pair of muons. The SM electroweak backgrounds are generated at NLO in QCD using the same matrix element generator.

The Z boson cross section is normalized to the NNLO, and the Higgs boson cross section is calculated at N³LO [16] with FEWZ v3.1 [73]. The branching fractions used for the signals are those listed in Tables. 1.3-1.4. For the SM electroweak backgrounds, the generator branching fractions $\mathcal{B}(H \rightarrow \mu\mu\gamma) = 3.92 \times 10^{-5}$ and $\mathcal{B}(Z \rightarrow \mu\mu\gamma) = 1.32 \times 10^{-5}$ are used as reference [71].

In all simulated samples, parton showering and hadronization processes are simulated by interfacing the event generators with PYTHIA 8.230 with the CP5 tune [74]. The NNPDF3.1 NNLO [75, 76] PDFs are used to model the momentum distribution of the partons inside the colliding protons. Pileup interactions are included, with a distribution that is chosen to match that observed in the data. The response of the CMS detector to the generated particles is modeled using the GEANT4 toolkit [77].

Simulated events are reconstructed with the same algorithms as those used for data. To correctly handle the normalization for the processes modeled from MC samples, an event weight is introduced to account for the simulated process production rate and the needed corrections due to imperfect modeling relative to the observed data. Denoting w_i as the weight for the i^{th} event in the samples, its mathematical expression is:

$$w_i = \begin{cases} 1 & \text{if event} \in \text{observed data} \\ \epsilon_{\text{lumi}} \cdot \frac{\sigma_{\text{prod}} \mathcal{B} \mathcal{L}_{\text{int}}}{N_{\text{ev}}} \cdot \epsilon_{\text{corr}} & \text{if event} \in \text{Monte Carlo} \end{cases}, \quad (3.1)$$

where:

- $\sigma_{\text{prod}} \mathcal{B} \mathcal{L}_{\text{int}}$ is the product of the production cross section of the heavy boson, the branching fraction of the decay under study and the integrated luminosity of the analysis sample. Hence, it is equal to the expected number of events for the simulated process over the period considered, typically a year of data taking.
- ϵ_{corr} is a correction factor accounting for the imperfect modeling in the simulation of the pileup distribution, objects identification and reconstruction, and trigger efficiency. The factors composing this term are discussed in Section 3.5.

3.2 Physics objects and event preselection

The physics objects used for the analysis are introduced along with the definition of the relevant features. The chosen working points and their implicit selections are also defined, serving as basic preselection criteria that the observed and simulated events must meet to be considered for further selection.

3.2.1 Primary vertex

Due to the additional interactions per bunch crossing, several primary vertices are reconstructed in an event. The candidate PV for the events is chosen as the one with the highest sum of the p_T^2 of the associated clustered particles, identified leptons, and missing transverse momentum. Additionally, it must have:

- At least one associated track.
- A number of degrees of freedom $n_{\text{dof}} > 4$, where n_{dof} is the sum of the track weights [78,79].
- A distance from the center of the nominal interaction point along the beam pipe $|z_{\text{vtx}}| < 24$ cm.
- A distance in the transverse plane from the center of the nominal interaction point along the beam pipe $d_0 < 2$ cm.

If the PV does not satisfy these criteria, the event is rejected.

3.2.2 Trigger selection

The High-Level Trigger algorithm selected for the analysis is a single muon plus photon trigger ($\mu + \gamma$), requiring one muon and one photon with transverse momentum $p_T^\mu > 17$ GeV and $p_T^\gamma > 30$ GeV respectively. As alternative, a dimuon ($\mu\mu$) trigger requiring a J/ψ or $\psi(2S)$ charmonium candidate with transverse momentum $p_T^{\mu\mu} > 25$ GeV has also been considered, providing a useful cross-check for the Z decay search. This alternate trigger yields a performance on the final upper limit results that is slightly worse with respect to the main $\mu + \gamma$ trigger.

The efficiency of the $\mu + \gamma$ trigger is computed for both data and simulation using a topology with two different flavor leptons ($e\mu$) that are prompt and isolated from hadronic activity, in addition to a photon. The processes that mostly contribute are $t\bar{t} + \gamma$ and $Z + \gamma$, with the electron and muon originating from the leptonic W decays in the former and from $Z \rightarrow \tau\tau e\mu + X$ decays in the latter. The dataset chosen for this measurement must be orthogonal to the sample used for the analysis to avoid the introduction of any bias in the efficiency calculation. Thus, it is selected from single electron triggers and filtered using the following criteria:

- The events satisfy a single, isolated electron trigger.
- At least one electron passes a medium cut identification criteria (ID) with $p_T > 28$ GeV.
- The probe muon is required to satisfy a medium ID and to be isolated from hadronic activity, as described in Section 3.2.3.
- The probe photon is required to pass the multivariate ID at 80% efficiency working point, introduced in Section 3.2.5, and must be more energetic than the electron to avoid a possible bias in the trigger efficiency.
- The electron, the muon and the photon must be angularly separated with $\Delta R > 0.5$.

The denominator of the efficiency consists of all events passing the above selections. The numerator definition is the same, with the addition of passing the $\mu + \gamma$ trigger considered for the analysis. The efficiencies are evaluated as a function of the p_T and η of both the muon and photon. The correction factors deriving from differences in the computed trigger efficiency between data and simulation are presented in Section 3.5.2.

3.2.3 Muon identification and isolation

Events are required to have at least two muons. The leading and subleading muons must satisfy $p_T^{\mu_1} > 18$ GeV and $p_T^{\mu_2} > 5$ GeV, respectively. Their four-momenta are calibrated by applying the Rochester corrections [80], which account for residual differences between data and simulation in the muon momentum scale and resolution.

During the CMS event reconstruction, identification criteria are applied to muons, and only those that meet a predefined reconstruction quality criteria are retained. A complete list and definition of the IDs can be found in [52, 59]. For this analysis, an event must have a pair of opposite-sign “medium prompt ID” muons. The identification criterion requires each muon to satisfy the following conditions:

- The muon is identified as a PF muon.
- It is either a global or a tracker muon, whose definitions are explained in Section 2.3.2.
- The fraction of valid tracker hits is greater than 80%.
- If the muon is reconstructed as a tracker muon, its segment compatibility must exceed 0.451. For a global muon, the segment compatibility must exceed 0.303, and additional criteria apply: the global fit must have $\chi^2/n_{\text{dof}} < 3$, the position match between the tracker and standalone muons must have $\chi^2 < 12$, and the maximum χ^2 from a kink-finding algorithm must be below 20.
- It satisfies additional cuts on the longitudinal and transverse impact parameters of $d_z < 100 \mu\text{m}$ and $d_{xy} < 20 \mu\text{m}$.

A requirement on the isolation of muons is set to reject muons from semileptonic decays of hadrons within jets or any jet that punches through the calorimeters, mimicking a muon signature. The relative isolation variable of a muon is determined by the PF algorithm by constructing a cone of size ΔR around the muon direction [52]. The isolation variable is computed by summing the transverse energy (as measured in the calorimeters) and the transverse momentum (as measured in the silicon tracker) of all objects within this cone, excluding other muons, and then dividing by the transverse momentum of the muon. In addition to this calculation, a correction of the energy deposit of neutral particles (hadrons and photons) from pileup vertices is applied. The final formula is:

$$\text{PFRelIso}(\mu) \equiv \mathcal{I}^\mu = \frac{1}{p_T^\mu} \left(\sum_i p_T^{h^\pm, i} + \max \left[0, \sum_i p_T^{h^0, i} + \sum_i p_T^{\gamma, i} - \Delta\beta \sum_{i, \text{pileup}} p_T^{h^\pm, i} \right] \right), \quad (3.2)$$

with:

- $\sum_i p_T^{h^\pm, i}$ the scalar p_T sum of charged hadrons h^\pm originating from the primary event vertex within a cone of size ΔR around the muon.
- $\sum_i p_T^{h^0, i}$ the scalar p_T sum of neutral hadrons h^0 within the same cone.
- $\sum p_T^\gamma$ the scalar p_T sum of the photons inside the cone.
- $\Delta\beta \sum_{i, \text{pileup}} p_T^{h^\pm, i}$ the pileup correction, where $\Delta\beta = \frac{1}{2}$, approximately representing the ratio of neutral particle to charged hadron production in inelastic proton-proton collisions, as estimated from simulation.

For this work, a cone of size $\Delta R = 0.4$ is used. Since the signal events are expected to have a low $m_{\mu\mu}$ invariant mass and high p_T^μ , it is common for the isolation cone of the second muon to overlap with that of the first muon. In this case, the contribution of the nearby muon is not accounted for, as the isolation algorithm does not take into account other PF candidates classified as muons. The working point chosen for this analysis is a tight isolation requirement, corresponding to an upper threshold of 0.15 on the relative isolation variable in order to consider the muon as isolated.

3.2.4 $\psi(\text{nS})$ candidate reconstruction

An event may have more than one possible pair of muons, for example in the case of VH Higgs boson production and subsequent leptonic decay of the Z or W boson. The three criteria tested to reconstruct the $\psi(\text{nS})$ candidate select the opposite-sign pair of muons that either has the highest transverse momentum (p_T), the smallest invariant mass (m_{inv}), or the closest angular separation (ΔR). These association criteria are tested on different signal samples with varying lepton populations and compositions. Table 3.1 shows the fraction of events where the correct pair of muons is selected, using the generator-level information as the ground truth.

H production	ggF	ZH	WH
p_T	0.975	0.903	0.871
m_{inv}	0.971	0.945	0.938
ΔR	0.976	0.946	0.939

Table 3.1 • Fraction of events with the correct selection of the muons used to reconstruct the J/ψ candidate under each association criterion, depending on the signal production mode.

The difference among the criteria are particularly important when other prompt muons are present, which often have a larger momentum than those emerging from the $\psi(\text{nS})$ decay. For instance, in WH events the ΔR criterion selects the correct muon pair in 94% of the events, compared to 87% for the less effective p_T criterion. The selection scheme with the highest efficiency and greatest robustness against the presence of other muons in the event is the ΔR criterion, which has been adopted throughout the analysis.

3.2.5 Photon identification

The most energetic photon in an event must satisfy $p_T^\gamma > 32$ GeV and not fall in the ECAL transition region of $1.479 < |\eta^\gamma| < 1.653$. It must be reconstructed with sufficiently high quality and resolution. This condition is enforced through a Multivariate Classifier (MVA) identification

criterion with an 80% working point, meaning a photon selection efficiency of 80% is chosen [61]. The MVA includes isolation variables that account for the hadronic activity in a cone around the photon. The MVA classifier employs a Boosted Decision Tree (BDT) implemented in the TMVA framework [81], providing an excellent separation between signal (prompt photons) and background from misidentified jets. More details on its implementation and training are provided in [61]. Additionally, the reconstructed photon must pass a “Pixel Seed Veto” (PSV) to improve rejection of non-genuine photons from a conversion vertex. This veto ensures cleaner events for the final selection by checking for the absence of a matching track in the pixel detector, indicating no associated electron track [61].

3.2.6 Jets

As described in Section 2.3, the jets are clustered using the FASTJET package with the anti- k_T clustering algorithm and a clustering parameter $R = 0.4$, whose products are called AK4 jets. Contributions from pileup are mitigated with the Charged Hadron Subtraction (CHS) algorithm [82]. In this analysis, AK4 jets are considered if their p_T is larger than 30 GeV, they lie within the tracker acceptance ($|\eta| < 2.4$), pass a tight jet ID [83], and are at least $\Delta R > 0.4$ away from other isolated photons and muons.

VBF jets To enhance the identification of the Higgs VBF production mode, AK4 jets in the forward region of the detector and with $p_T > 50$ GeV are also considered in addition to the aforementioned central jets. These jets should not geometrically overlap with the muons or the photons used in the reconstruction of the H or Z candidate. If an event has at least a pair of these jets with invariant mass m_{jj} greater than 350 GeV, it is tagged as a candidate VBF event.

b -tagged jets The algorithm used to tag jets from b -quarks throughout this analysis is the DEEPCSV algorithm [84, 85], an extension of the Combined Secondary Vertex (CSV) algorithm [86] that utilizes the tracks and secondary vertices associated with the jets as inputs to a neural network (NN). The NN classifies jets assigning a probability for the jet to belong to one of the following categories: exactly one b hadron, two or more b hadrons, exactly one c hadron with no b hadrons, two or more c hadrons with no b hadrons, or none of the aforementioned categories.

Three working points are defined for b -tagging based on the level of mis-tagging [87]. The medium working point has been selected for this analysis, as it provides a balance between b -tagging efficiency (around 70% for genuine b -jets) and low mis-tagging of light jets (around 1%) [88].

3.2.7 Preselection summary

The preselection efficiencies on observed and simulated datasets are reported in Table 3.2. The distributions of a few key variables after preselection are shown in Fig. 3.1.

Chapter 3. Search for rare Z and Higgs boson decays to a charmed meson and a photon

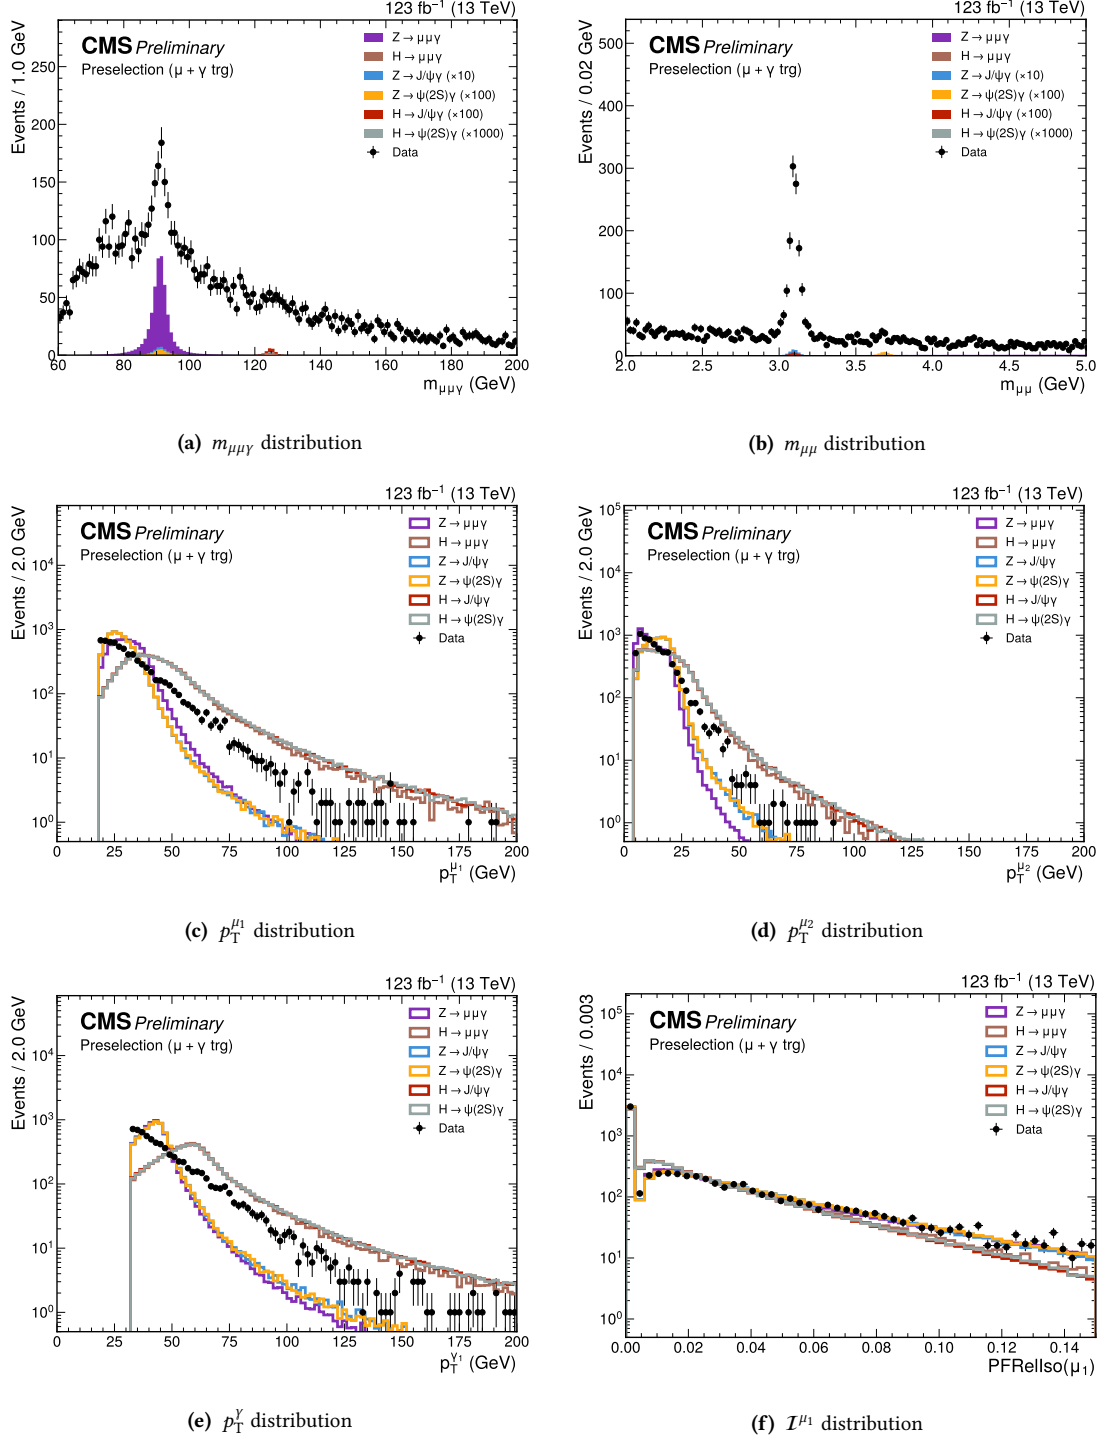


Figure 3.1 • Distributions $m_{\mu\mu\gamma}$, $m_{\mu\mu}$, $p_T^{\mu_1}$, $p_T^{\mu_2}$, p_T^{γ} and I^{μ_1} after preselection. For signal and background simulations, the distributions of $p_T^{\mu_1}$, $p_T^{\mu_2}$, p_T^{γ} and I^{μ_1} are normalized to match the number of observed events for comparison.

Sample	Before preselection	After preselection	Efficiency (%)
Data	3.84×10^8	7.84×10^3	2.0×10^{-3}
$Z \rightarrow J/\psi \gamma$	44.0	5.09	12
$H \rightarrow J/\psi \gamma$	1.36	2.37×10^{-1}	17
$Z \rightarrow \psi(2S) \gamma$	2.04	3.78×10^{-1}	17
$H \rightarrow \psi(2S) \gamma$	6.28×10^{-2}	1.06×10^{-2}	17
$Z \rightarrow \mu\mu \gamma$	1.30×10^4	5.10×10^2	3.9
$H \rightarrow \mu\mu \gamma$	2.67×10^2	6.25	2.3

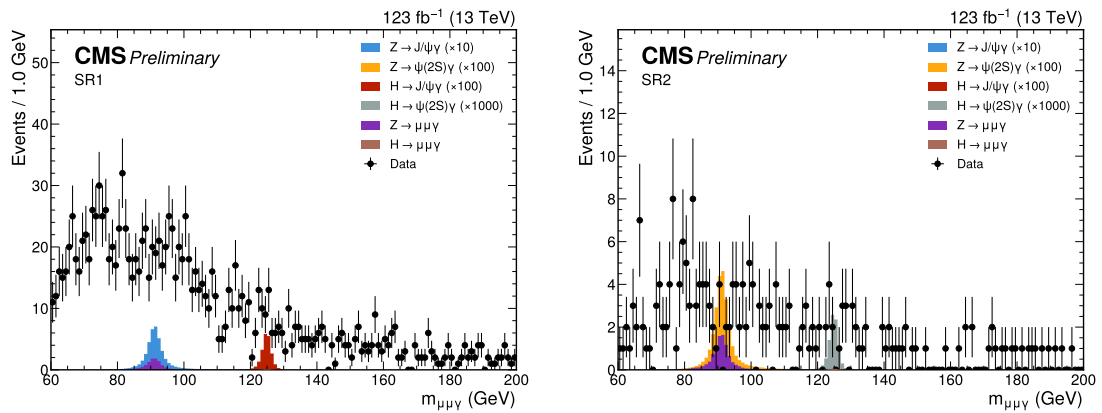
Table 3.2 • Number of observed and simulated events before and after the application of the preselection.

3.3 Event selection

This section introduces the event selection criteria for the analysis, based on the analysis objects defined in the previous section. The selection criteria are organized into two levels, each with a specific purpose and granularity: the “Signal Region” (SR) selections, which include the preselection and additional criteria based on higher-level features, as well as event categorization; the “Control Region” (CR) selections, a special set of filters defined for the $Z \rightarrow \mu\mu \gamma$ background control. The defined levels and criteria for categorization are described below.

3.3.1 Signal region selections

The second set of selection criteria is defined on top of the preselection and adds higher-level variables to maximize the signal-to-background ratio. Events are divided into two signal regions based on the value of $m_{\mu\mu}$: events with $3.0 < m_{\mu\mu} < 3.2$ GeV are considered to be J/ψ candidates and belong to the a signal region called SR1; events with $3.60 < m_{\mu\mu} < 3.75$ GeV are assumed to originate from a $\psi(2S)$ state and belong to a separate signal region SR2. The invariant mass distribution of the final state particles $m_{\mu\mu\gamma}$ for the two SRs is shown in Fig. 3.2. Depending on the signal process, the defined categories are slotted into subcategories to exploit their physical features and maximize the sensitivity, as presented in this section.

**Figure 3.2** • Distribution of $m_{\mu\mu\gamma}$ after the application of the SR1 (left) and SR2 (right) criteria.

3.3.2 Angular variables based categorization

Three main angular variables, illustrated in Fig. 3.3, characterize the topology of the process under study: the cosine of the angle between the momentum directions of the $\psi(\text{nS})$ meson and of the beam in the rest frame of the Higgs or Z boson ($\cos \theta^*$); the cosine of the angle between the momentum directions of the positively charged muon and the $\psi(\text{nS})$ in the rest frame of the $\psi(\text{nS})$ ($\cos \theta_1$); the angle between the decay plane of the two muons and the plane of the $\psi(\text{nS})$ and the photon (φ_1). Their distributions for the observed data and simulated samples after the preselection are shown in Fig. 3.4.




Figure 3.3 • Topology of $Z, H \rightarrow \psi(\text{nS})\gamma \rightarrow \mu\mu\gamma$ decay, with the main angular variables $\theta^*, \theta_1, \varphi_1$ highlighted.

The distributions of $\cos \theta^*$ and $\cos \theta_1$ are particularly important as they depend on the polarization of both the candidate heavy boson and the candidate $\psi(\text{nS})$ meson. The absence of polarization is indicated by a flat distribution, while transverse or longitudinal polarization is evident from a convex or concave distribution, respectively. Mathematically, denoting $f(\cos \theta)$ as the distribution under study, its analytic expression is:

$$f(\cos \theta) = \begin{cases} \frac{1}{2} & \text{No polarization} \\ \frac{3}{4} (1 + \cos^2 \theta) & \text{Transverse polarization} \\ \frac{3}{2} (1 - \cos^2 \theta) & \text{Longitudinal polarization} \end{cases} . \quad (3.3)$$

The joint information of the $\cos \theta^*$, $\cos \theta_1$ and φ_1 variables can be used to implement an angular discriminator between signal and backgrounds, which can be exploited for event categorization. This strategy is also physically motivated since these variables are not correlated with the invariant mass distributions $m_{\mu\mu}$ and $m_{\mu\mu\gamma}$, as shown in Fig. 3.4d. After preselection, the angular distributions for the observed and simulated data are modeled using polynomial functions f_s and f_b for the signal and background, respectively. Then, an angular Likelihood Discriminator (LD) variable is constructed as a likelihood ratio between the signal and the backgrounds:

$$\begin{cases} L_s = f_s(\cos \theta^*) \cdot f_s(\cos \theta_1) \cdot f_s(\varphi_1) \\ L_b = f_b(\cos \theta^*) \cdot f_b(\cos \theta_1) \cdot f_b(\varphi_1) \end{cases} \Rightarrow L = \frac{L_s}{L_s + L_b} . \quad (3.4)$$

3.3. Event selection

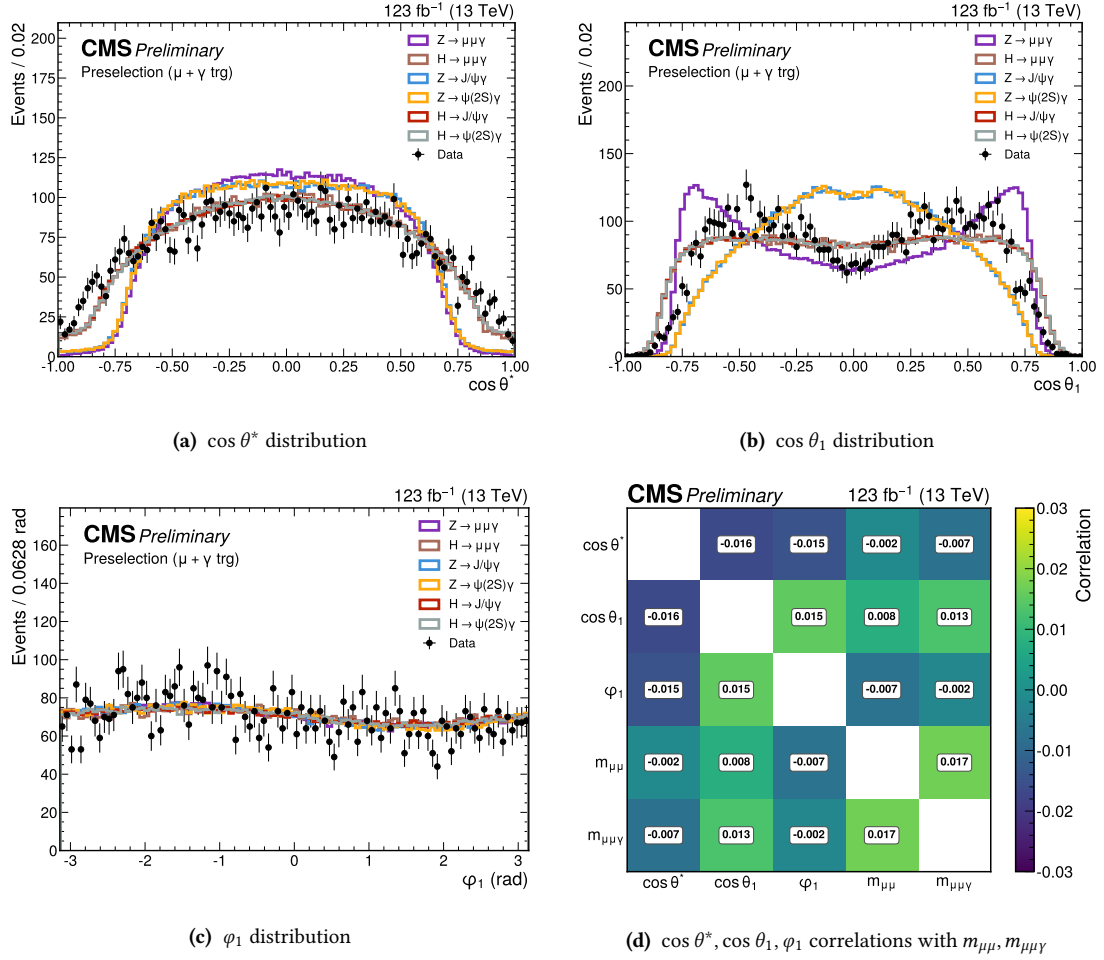


Figure 3.4 • Distributions of the angular variables $\cos \theta^*$, $\cos \theta_1$, ϕ_1 after preselection. Their correlation with $m_{\mu\mu\gamma}$, $m_{\mu\mu}$ is shown in 3.4d. Note that the angular distributions are sculpted at their edges by the thresholds in p_T on the muons and the photon.

A variable is defined for each decaying boson, discriminating the corresponding signal from the main QCD background. The distributions of these variables after the preselection are shown in Fig. 3.5. From the distributions, it is evident that the LD variable, owing to the intrinsic spin correlations of the particles considered, provides a significant signal-background separation for the Z searches. Thus, the variable $LD_{\text{bkg vs Z sig}}$ is used in the $Z \rightarrow J/\psi\gamma$ signal search to define two independent categories: a High Purity (Z-HP) category, for events that have an LD score greater than 0.5, belonging to the signal-enriched part of the discriminator distribution; a Low Purity (Z-LP) category, for the events with an LD score lower than 0.5. The limited amount of data samples available in the $Z \rightarrow \psi(2S)\gamma$ search after the SR2 selection prevents the application of this discriminator to its respective signal region.

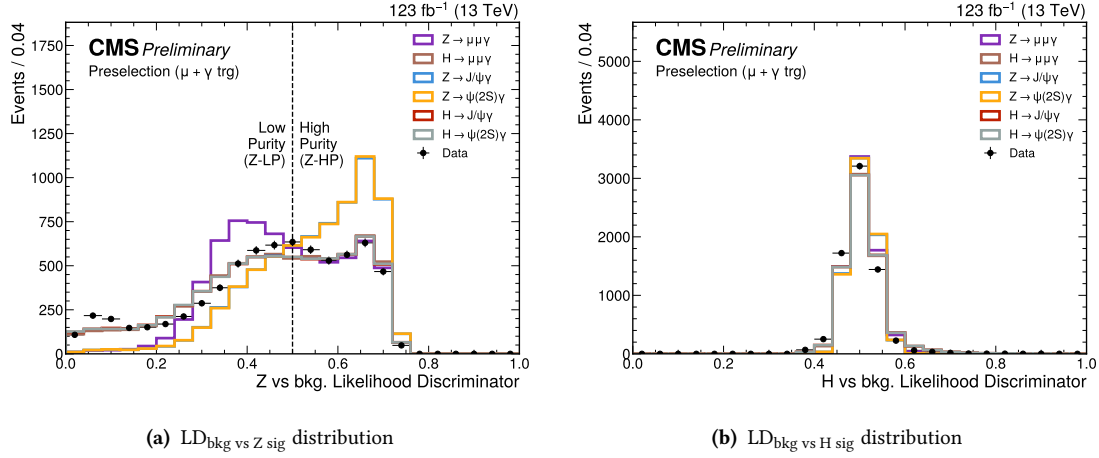


Figure 3.5 • Distributions after preselection of $LD_{bkg} \text{ vs } Z \text{ sig}$ (3.5a) and $LD_{bkg} \text{ vs } H \text{ sig}$ (3.5b) variables.

3.3.3 Higgs production mode categorization

In the case of a Higgs signal, for which the LD variable does not present a significant signal-background discrimination power, the sensitivity can be improved by further categorizing events based on the characteristics of the different Higgs boson production modes, introduced in Section 1.2.1. A VBF category contains events with at least a pair of VBF jets with $m_{jj} > 350$ GeV, as described in Section 3.2.6. A Heavy-Flavor (HF) category accepts events with at least one b -tag jet, targeting Higgs boson production in association with bottom and top quarks. Events that do not fall into the previous categories constitute the inclusive category (ggF), which mainly targets gluon-gluon fusion production. This category is further divided in high (ggF-HP) and low purity (ggF-LP) depending on the value of the angular variable $|\cos \theta^*|$: high purity for $|\cos \theta^*| > 0.5$ and low purity for $|\cos \theta^*| < 0.5$.

A category targeting the WH and ZH Higgs boson production with additional leptons was explored. However, too few events would fall into this category to allow reliable background estimation, so it was merged into the inclusive category. The same categorization based on the production mode is not applicable to the Z boson decay search, as it is predominantly produced through $q\bar{q}$ annihilation. Moreover, the limited amount of events passing SR2 selections for the $H \rightarrow \psi(2S)\gamma$ process, does not allow to apply the same event categorization, as it would result in categories with too few events for the background estimation fit.

3.3.4 Control region for $Z \rightarrow \ell\ell\gamma$ background modeling

With the SR selections, most of the Monte Carlo events for $Z, H \rightarrow \mu\mu\gamma$ are discarded by the $m_{\mu\mu}$ window cut, leading to a loss of precision in the modeling of the $m_{\mu\mu\gamma}$ distribution. Since the resonant background is expected to peak at the same mass values of the Higgs and Z bosons, it is essential to control its normalization with sufficient precision. Therefore, a control region (CR) is defined as the set of events that do not belong to either of the two signal regions but satisfy $2 < m_{\mu\mu} < 8$ GeV. A plot for the $m_{\mu\mu\gamma}$ invariant mass spectrum obtained with the CR selections is shown in Fig. 3.6. It can be observed that, after these selections, the dataset presents a resonance peaked at the Z boson mass, due to the Z FSR background.




Figure 3.6 • Distribution of $m_{\mu\mu\gamma}$ after the application of the CR selections, where the resonance peaked at the Z boson mass value is due to the Z FSR background.

3.3.5 Event categorization summary

The event selection and categorization introduced in the previous subsections are summarized in Table 3.3. The number of events falling in each category is shown in Fig. 3.7.

	SR1						SR2	CR
	Z-HP	Z-LP	ggF-HP	ggF-LP	VBF	HF		
$p_T^{\mu_{1,2}}$			> 18 GeV (leading μ_1), > 5 GeV (subleading μ_2)					
p_T^γ			> 32 GeV					
$ \eta^{\mu_{1,2}} $			< 2.4					
$ \eta^\gamma $			< 2.5 , excluding ECAL transition region ($1.479 < \eta^\gamma < 1.653$)					
$m_{\mu\mu}$ (GeV)			$[3.0, 3.2]$				$[3.6, 3.75]$	$[2.0, 3.0]$, $[3.2, 3.6]$, $[3.75, 8.0]$
$\cos \theta^*$	—	—	> 0.5	< 0.5	—	—	—	—
LD	> 0.5	< 0.5	—	—	—	—	—	—
b -jets	—	—	0	0	0	≥ 1	—	—
VBF candidates ($m_{jj} > 350$ GeV)	—	—	0	0	≥ 1	0	—	—

Table 3.3 • Summary of the selections of the seven signal regions and the control region.




Figure 3.7 • Summary of categorization and amount of observed events per category, highlighted for each search considered in this work.

3.4 Event modeling

This section describes the modeling of the signal and background components, as introduced in Section 3.1. The RooFit framework [89] is employed to perform unbinned Maximum Likelihood (ML) fits. In the minimizer algorithm, every data point is considered with a weight factor, as described in Eq. 3.1.

3.4.1 Signal and resonant background modeling

The signal and resonant background contributions are estimated from the Monte Carlo samples after applying the corresponding SR and CR selections, respectively. This operation is accomplished by fitting a parametric function to the $m_{\mu\mu\gamma}$ distribution. A Double-Sided Crystal Ball (DSCB) function [90] is chosen, with an additional Gaussian added to improve the fit. The functional form of the DSCB is characterized by a Gaussian core (2 parameters) and two exponential tails (4 parameters), and it reads:

$$\text{DSCB}(x|\mu, \sigma, \alpha_1, \alpha_2, n_1, n_2) = \begin{cases} A_1 \cdot \left(B_1 - \frac{x - \mu}{\sigma}\right)^{-n_1} & \frac{x - \mu}{\sigma} \leq -\alpha_1 \\ \exp\left[-\frac{(x - \mu)^2}{2\sigma^2}\right] & -\alpha_1 < \frac{x - \mu}{\sigma} < \alpha_2 \\ A_2 \cdot \left(B_2 + \frac{x - \mu}{\sigma}\right)^{-n_2} & \frac{x - \mu}{\sigma} \geq \alpha_2 \end{cases}, \quad (3.5)$$

with:

$$A_i = \left(\frac{n_i}{|\alpha_i|}\right)^{n_i} \cdot \exp\left[-\frac{\alpha_i^2}{2}\right], \quad i = 1, 2, \quad (3.6)$$

$$B_i = \frac{n_i}{|\alpha_i|} - |\alpha_i|$$

3.4. Event modeling

The results of the fit for the signal samples filtered with the SR selections are shown in Figs. 3.8a-3.8b for two categories, provided as examples. The results of the fit for the resonant background samples, using the CR selections, are shown in Figs. 3.8c-3.8d.

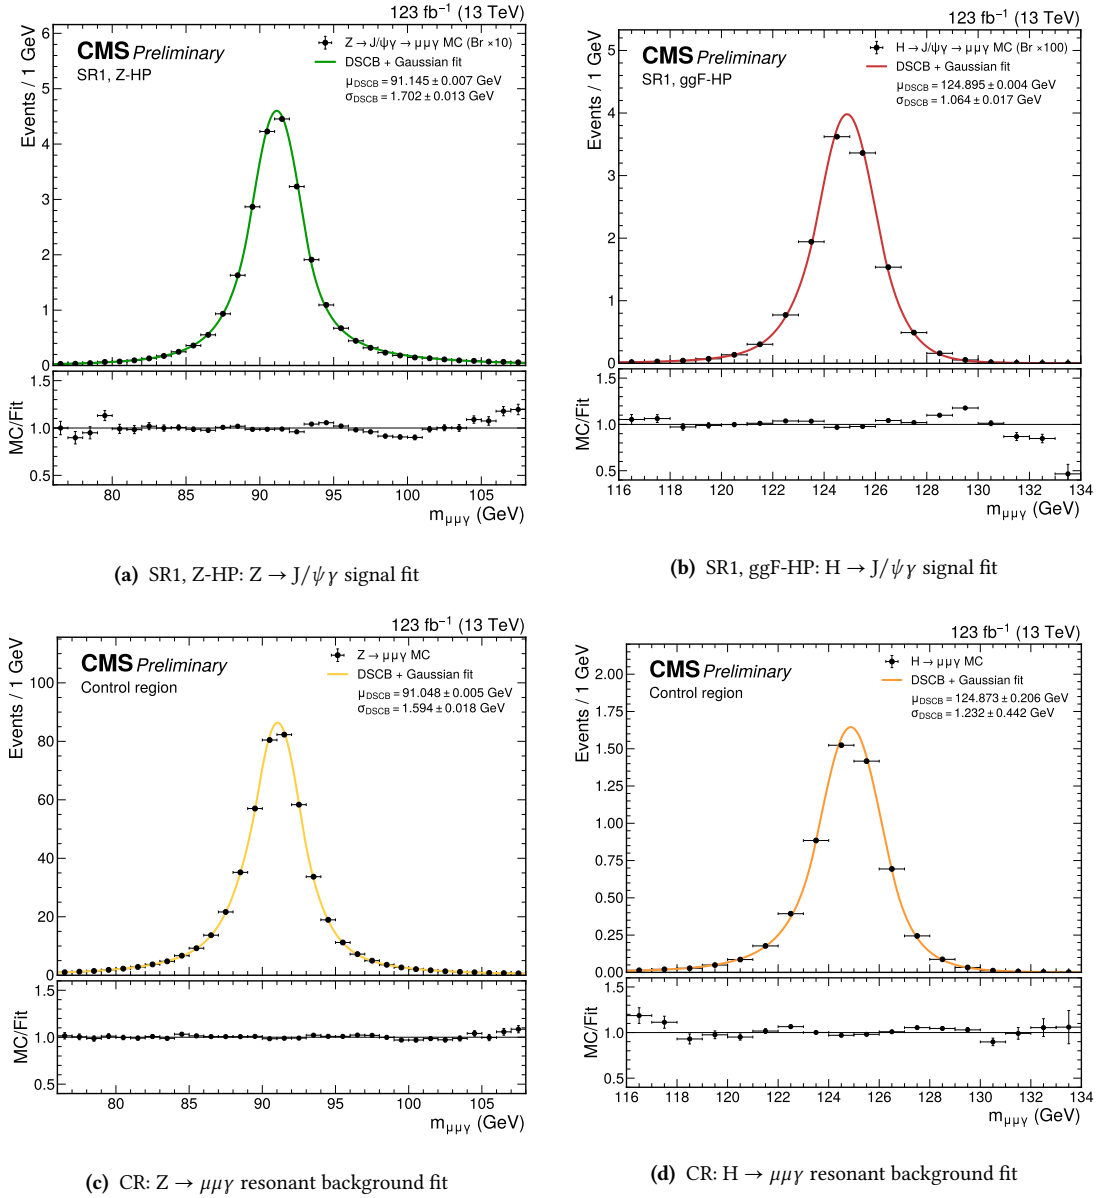


Figure 3.8 • Examples of results of signal and resonant background MC samples $m_{\mu\mu\gamma}$ distribution fit. The model employed is the DSCB, defined in Eq. 3.5, plus an additional Gaussian, while the datasets are obtained using the corresponding SR selections for 3.8a-3.8b, and the CR selections for 3.8c-3.8d.

3.4.2 QCD background modeling

The main continuous background from QCD processes is modeled directly from the data⁵ in the absence of a reliable MC sample simulating this contribution. The upper limits obtained from the statistical analysis will be highly dependent on the background model employed. For this reason, multiple choices of function families are considered in the statistical analysis and, for each family, models with different number of parameters are included. The aim of this strategy is not only to find the best-fitting model, but also to minimize potential biases introduced by the choice of the function in the analysis. The functional forms considered in this work are listed below, omitting the normalization parameter:

- **Power-law functions** $f_{\text{power}}^{(n)}$, with a number of shape parameters n ranging from 1 to 4. The parameterization of the power-law is designed to minimize the correlations between the parameters:

$$f_{\text{power}}^{(1)}(x) = x^{-a_1} \quad , \quad (3.7)$$

$$f_{\text{power}}^{(2)}(x) = x^{-a_1} \cdot (1-x)^{-a_2} \quad , \quad (3.8)$$

$$f_{\text{power}}^{(3)}(x) = x^{-a_1-a_3 \log x} \cdot (1-x)^{-a_2} \quad , \quad (3.9)$$

$$f_{\text{power}}^{(4)}(x) = x^{-a_1-a_3 \log x - a_4 \log^2 x} \cdot (1-x)^{-a_2} \quad . \quad (3.10)$$

The x variable represents the quantity $m_{\mu\mu\gamma}/\sqrt{s}$.

- **Exponential functions** $f_{\text{exp}}^{(n)}$, with a number of shape parameters n ranging from 1 to 4:

$$f_{\text{exp}}^{(1)}(x) = \exp[b_1 x] \quad , \quad (3.11)$$

$$f_{\text{exp}}^{(2)}(x) = \exp[b_1 x^{b_2}] \quad , \quad (3.12)$$

$$f_{\text{exp}}^{(3)}(x) = \exp[b_1 x^{b_2} + b_3(1-x)] \quad , \quad (3.13)$$

$$f_{\text{exp}}^{(4)}(x) = \exp[b_1 x^{b_2} + b_3(1-x)^{b_4}] \quad . \quad (3.14)$$

The x variable has the same definition as for the power-law functions.

- **Bernstein polynomials** $f_{\text{bern}}^{(n)}$, with a number of shape parameters n ranging from 3 to 5. Their general form is expressed as a linear combination:

$$f_{\text{bern}}^{(n)} = \sum_{i=0}^n d_i B_{i,n}(x) \quad , \quad (3.15)$$

where the set $\{B_{i,n}(x)\}$ is the $n+1$ Bernstein basis polynomials of degree n , whose elements are computed as:

$$B_{i,n}(x) = \binom{n}{i} \cdot x^i (1-x)^{n-i} \quad . \quad (3.16)$$

⁵Before the unblinding of the analysis, the data inside the invariant mass windows centered on $m_{Z,H}$, namely [86, 96] GeV and [120, 130] GeV, have not been considered in ML fit.

The choice of the optimal number of shape parameters for each functional family is resolved by applying an F -test. Given two models with $p_{(1)}$ and $p_{(2)}$ shape parameters, where $p_{(2)} > p_{(1)}$, this test answers the question of whether the model (2) does not provide a significantly better fit to the data than model (1), which is the null hypothesis. Given a binned dataset with a number of bins n_{bins} , the F -test statistics is computed as:

$$F = \frac{\left(\frac{\text{RSS}_{(1)} - \text{RSS}_{(2)}}{p_{(2)} - p_{(1)}} \right)}{\left(\frac{\text{RSS}_{(2)}}{n_{\text{bins}} - p_{(2)}} \right)}, \quad (3.17)$$

where $\text{RSS}_{(i)}$ is the Residual Sum of Squares of the model (i). The null hypothesis is rejected if the F statistics calculated from the data is lower than a critical value $\alpha = 0.05$ of the F -distribution with $(p_{(2)} - p_{(1)}, n_{\text{bins}} - p_{(2)})$ degrees of freedom, thus the model (2) provides a significantly better description of the data. In the case considered here, the test is applied iteratively to each pair of models with $p_{(1)}$ and $p_{(2)} = p_{(1)} + 1$ shape parameters. The F -test optimal results for each family of functions and for all the categories are reported in Table 3.4.

Function	SR1					SR2		CR
	Z-HP	Z-LP	ggF-HP	ggF-LP	VBF	HF		
Power-law	3	3	3	4	2	2	2	3
Exponential	2	4	2	2	2	2	2	3
Bernstein	5	2	5	5	4	4	4	4

Table 3.4 • Optimal number of parameters from F -test for all families of functions used for the QCD background modeling.

The PDF obtained from the fit and normalized to the number of observed events is summed to the fitted resonant background contributions, with their normalization fixed to the SM expectation in the case of the SRs. For the CR, only the normalization of $H \rightarrow \mu\mu\gamma$ is fixed to the SM expectation, while the $Z \rightarrow \mu\mu\gamma$ normalization is left floating and obtained from the fit, as the resonance is observable from the data. The resulting model represents the entire background contribution, shown in Figs. 3.9-3.10 for all the SR categories and in Fig. 3.11 for the CR selections.

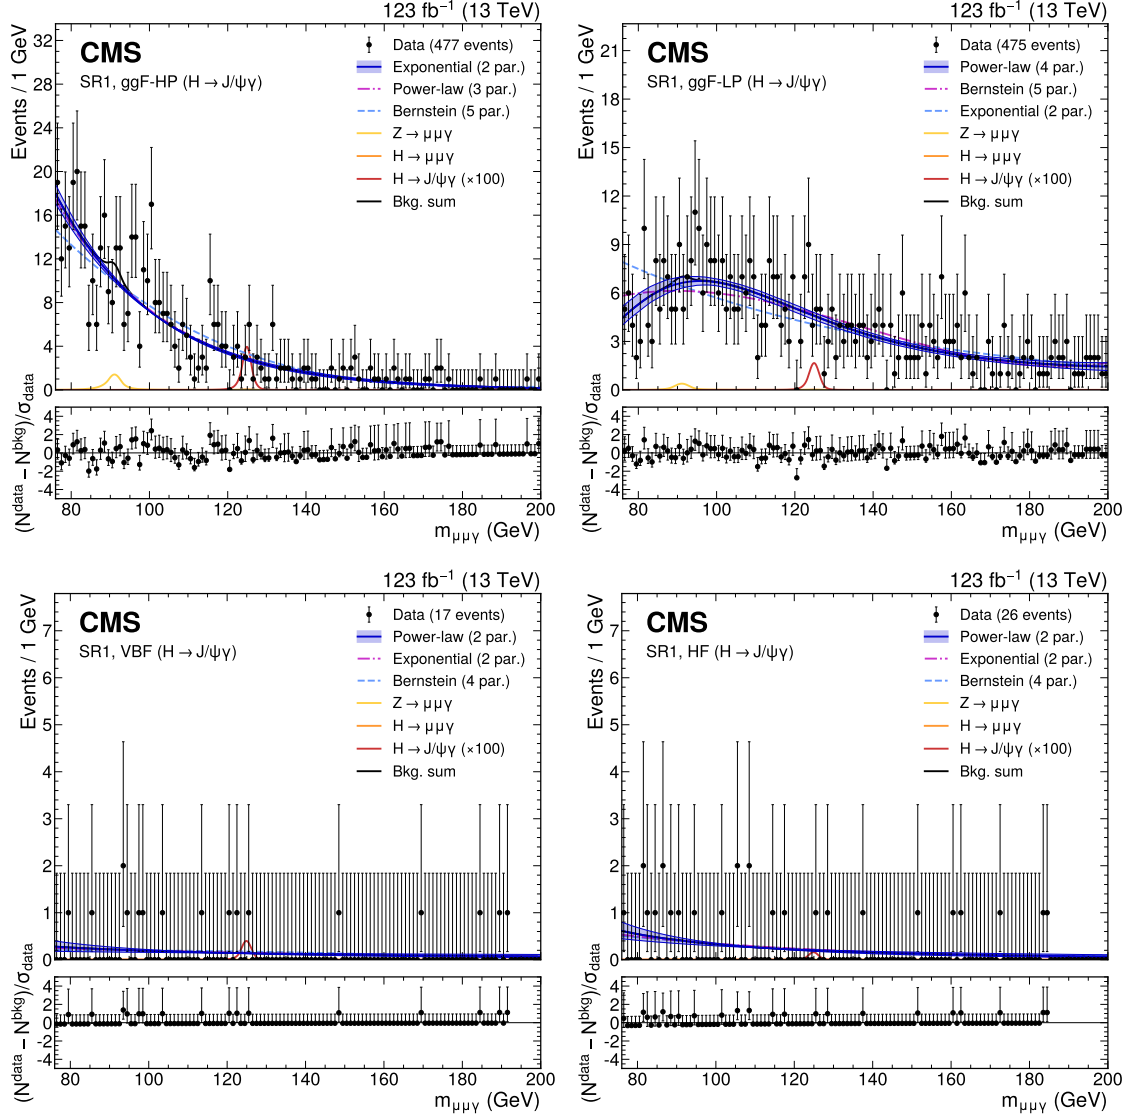


Figure 3.9 • The upper panels show the background-only fit using the power-law, exponential, and Bernstein polynomial functions with the optimal number of parameters determined with the F -test, and ranked according to the fit χ^2 . The resonant background contributions are added with the normalization fixed to the SM expectation. The vertical error bars on the data points show the statistical uncertainty. The bottom panels show the pulls for the fit result for the model with lowest χ^2 . Upper left: $H \rightarrow J/\psi\gamma$ ggF-HP category. Upper right: $H \rightarrow J/\psi\gamma$ ggF-LP category. Lower left: $H \rightarrow J/\psi\gamma$ VBF category. Lower right: $H \rightarrow J/\psi\gamma$ HF category.

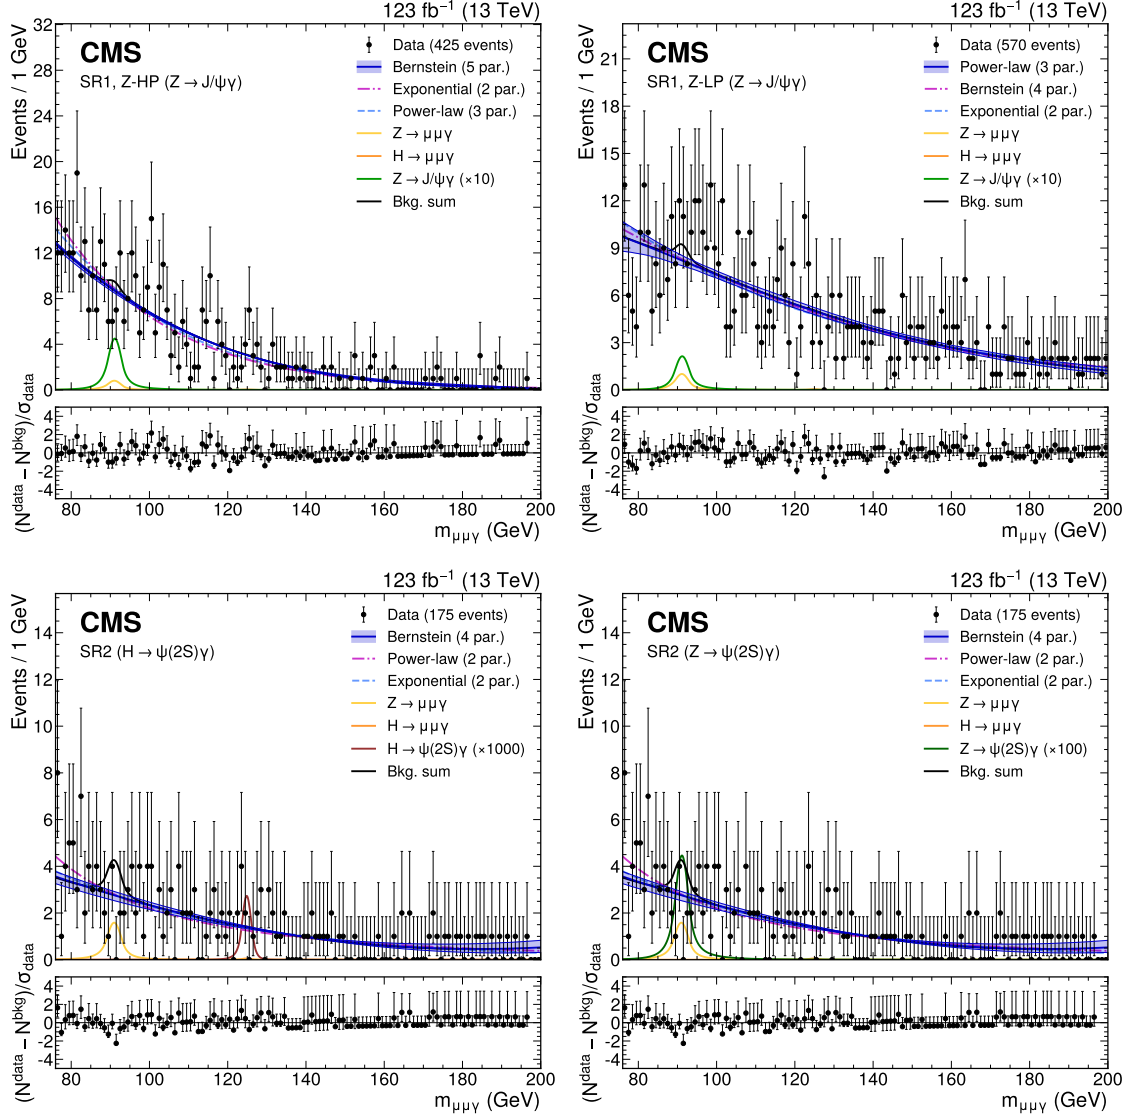


Figure 3.10 • The upper panels show the background-only fit using the power-law, exponential, and Bernstein polynomial functions with the optimal number of parameters determined with the F -test, and ranked according to the fit χ^2 . The resonant background contributions are added with the normalization fixed to the SM expectation. The vertical error bars on the data points show the statistical uncertainty. The bottom panels show the pulls for the fit result for the model with lowest χ^2 . Upper left: $Z \rightarrow J/\psi\gamma$ Z-HP category. Upper right: $Z \rightarrow J/\psi\gamma$ Z-LP category. Lower left: $H \rightarrow \psi(2S)\gamma$ category. Lower right: $Z \rightarrow \psi(2S)\gamma$ category.




Figure 3.11 • The upper panel shows the results of the main background fit in the CR, using the power-law, exponential, and Bernstein polynomial functions with the optimal number of parameters determined with the F -test, and ranked according to the fit χ^2 . The vertical error bars on the data points show the statistical uncertainty. The bottom panel shows the pulls for the fit result for the model with lowest χ^2 .

3.5 Systematic uncertainties and corrections

Systematic uncertainties may affect the results and are due to incomplete knowledge of:

- The true model of the invariant mass distribution $m_{\mu\mu\gamma}$ for the QCD multijet background.
- The detector simulation.
- The theoretical prediction of the signal production mechanism.

For the first point, Section 3.5.1 introduces the “discrete profiling” method to take into account the uncertainty in the QCD background model choice. The systematic uncertainties of the last two points are evaluated by varying contributing sources within their corresponding uncertainties and propagating every uncertainty to the signal yield or shape using the signal MC sample. In particular, as described in Section 3.1.2 and in Eq. 3.1, the simulated event weights require correction factors, known as Scale Factors (SF), to account for the imperfect modeling of the simulation with respect to the observed data. The next subsections describe these corrections, which are in general selected according to the standard Run-2 recommendations provided by the involved Physics Object Groups (POG).

3.5.1 QCD background model discrete uncertainty

The true model of the QCD background is not known and it constitutes one of the main sources of systematic uncertainty. In fact, the choice of the function used to model the QCD background can lead to non-negligible variations in the upper limit results. To address this source of uncertainty in a simple and elegant way, the “discrete profiling” (or “common envelope”) method is chosen [91].

This method was designed as a way to estimate the systematic uncertainty associated with choosing a particular analytic function to fit the background $m_{\gamma\gamma}$ distribution for the $H \rightarrow \gamma\gamma$ analysis [92]. A set of candidate function families is considered, including the ones introduced in Section 3.4.2 (exponentials, power-laws and Bernstein polynomials). These functions are fitted to the $m_{\mu\mu\gamma}$ distribution in the mass range defined by the event selections. For each family of functions, an F -test is performed to determine proper number of parameters to be used in the fit, while requiring an acceptable goodness of fit to the data.

When fitting these functions to the background $m_{\mu\mu\gamma}$ distribution, the value of twice the negative logarithm of the likelihood (2NLL) is minimized. A penalty term is added to 2NLL to account for the number of floating parameters in each candidate function. When measuring a given parameter of interest, the discrete profiling method determines the envelope of the lowest values of 2NLL profiled as a function of that parameter. In this way, the choice of the background function is treated as a discrete nuisance parameter in the fit, accounting for the uncertainty associated with the arbitrary choice of function. The method is already implemented in the CMS statistical analysis tool, Combine [93].

A signal injection test is performed to assess the robustness of this method. After defining a signal strength modifier $\mu_{\text{true}} = \sigma/\sigma_{\text{SM}}$ for the signal to inject, a signal plus background toy is generated, using for the background component generation a function selected among the possible families after the F -test. The generated toy dataset is then fitted using the discrete profile method, with the signal normalization free to float in the fit. The result of the fit is the measured signal strength μ_{fit} , which should ideally correspond to the injected fraction μ_{true} . The uncertainty

σ_{fit} of the measured signal strength modifier is determined as well to compute a pull value, defined as:

$$\text{pull} = \frac{\mu_{\text{fit}} - \mu_{\text{true}}}{\sigma_{\text{fit}}} . \quad (3.18)$$

The procedure is repeated 100 times and applied for each possible combination of signal process under study and function family used for the generation of the background component of the toy dataset. The procedure is further repeated for multiple μ_{true} points.

A conservative criterion for determining whether the fit is unbiased is that the distribution of the pull values from all the toys, for a given combination of signal process, generator family and μ_{true} point, follows a Gaussian distribution with an absolute mean value less than a generic threshold of 0.20 and variance around 1. This criterion ensures that any possible bias in the fitted signal strength is at least 20% smaller than its statistical uncertainty and can therefore be considered as acceptable [94]. The results of this test are presented in Fig. 3.12 and do not show in general any significant bias introduced in the analysis.




Figure 3.12 • Bias studies with discrete profiling method: pulls distribution depending on toy and on signal injected. The orange line is the Gaussian mean of the pulls distribution for a certain signal fraction injected, while the red lines are the $\pm 1\sigma$ bands.

3.5.2 Uncertainties affecting the predicted signal yields

Luminosity The uncertainty on the integrated luminosity is the recommended value of 1.2% in 2016 [95], 2.3% in 2017 [96], and 2.5% in 2018 [97], when considering the uncertainties as correlated across years. Accounting for the relative contribution of the different years, the total uncertainty for the full Run-2 is 1.6%, which is propagated to the normalization of the signal and $H \rightarrow \mu\mu\gamma$ background.

High-Level-Trigger The SFs accounting for differences in the $\mu + \gamma$ trigger efficiency between simulation and data are computed as the ratio of efficiencies for data and simulation, following the methodology described in Section 3.2.2 for the efficiency calculation. The SFs as a function of the p_T and η of the muon and photon are shown in Fig. 3.13. Since no dependence on η is observed, the SFs are derived as a function of p_T . The SFs and the corresponding uncertainties, integrated over the p_T and η phase space considered in the analysis, are summarized in Table 3.5.

The uncertainties on the signal and $H, Z \rightarrow \mu\mu\gamma$ background yields are estimated by varying the trigger efficiency SFs by $\pm 1\sigma$ of their uncertainties. The largest variation in the yields is reported as the systematic uncertainty.

Year	Scale Factor
2016	0.982 ± 0.059
2017	0.919 ± 0.061
2018	0.972 ± 0.045
Run-2	0.978 ± 0.035

Table 3.5 • Scale factors for the $\mu + \gamma$ trigger, separately by year and integrated over the full Run-2. The uncertainty is statistical only.

Identification and Isolation of muons The SF deriving from the application of the muon identification and isolation criteria are provided by the Muon POG and are derived using a tag-and-probe approach on $Z \rightarrow \mu\mu$ events [98]. The uncertainties on the signal and $H, Z \rightarrow \mu\mu\gamma$ background yields are derived by running the analysis and varying the muon identification and isolation SFs by $\pm 1\sigma$ of their uncertainties. The maximum variation in the yields is quoted as the systematic uncertainty.

Identification of photons The SF values for the photon identification with a MVA technique and for the Pixel Seed Veto are provided by the e/γ POG [99]. As for the muon identification and isolation, the uncertainties on the yields are derived by running the analysis and varying the photon MVA identification SFs by $\pm 1\sigma$ of their uncertainty. The maximum variation in the yields is quoted as the systematics.

Pileup reweighting The Monte Carlo samples include the simulation of additional interactions per bunch crossing for a more realistic comparison. However, there are differences between simulated and observed pileup that must be accounted for in the analysis. In the observed data, the pileup can be estimated based on the luminosity and proton-proton inelastic cross section. Since the instantaneous luminosity may vary during data taking, the pileup distribution in the simulation needs to be aligned with the data. This alignment is achieved through a set of SFs correcting the simulated event weight and included in Eq. 3.1.

The SF values and their uncertainties are derived from histograms of the pileup distributions for data and simulation [100], taking into account a proton-proton inelastic cross section of 69.2 mb with an uncertainty of 4.6%. The cross-section is varied by $\pm 4.6\%$, and the analysis is then run with the varied weights. The maximum difference in the yields of signal and background simulations compared to the nominal yield is quoted as the systematic uncertainty.

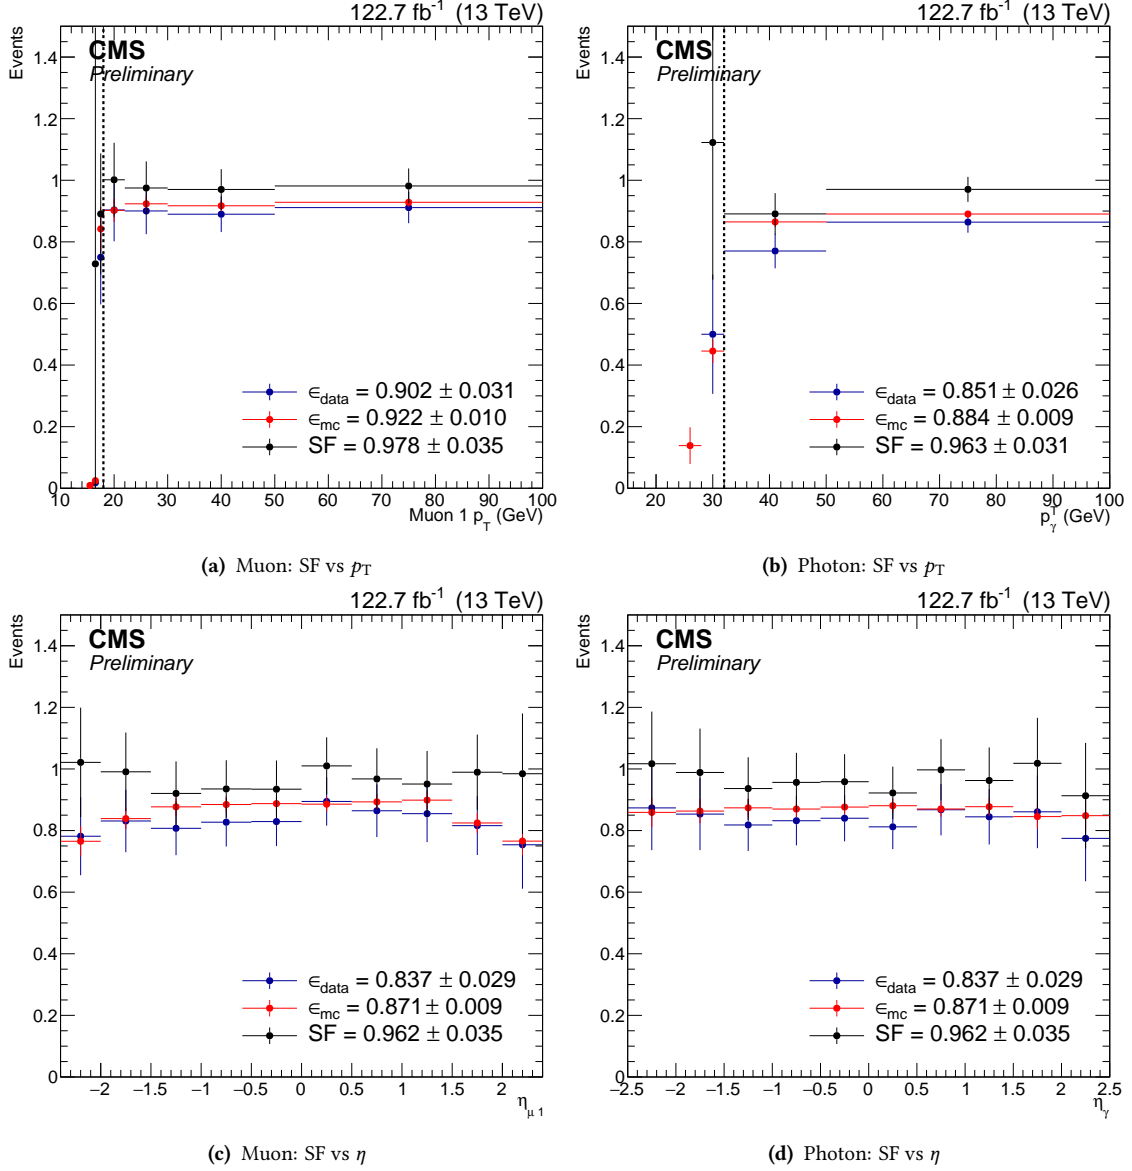


Figure 3.13 • Efficiency of the $\mu + \gamma$ trigger as a function of p_T and η of the photon and muon legs.

3.5.3 Uncertainties affecting the shape of the signal model

Choice of the signal model Given that the choice of the function used to fit the signal is arbitrary, a systematic uncertainty is introduced to account for potential biases related to this choice. The uncertainty is estimated by fitting the signal with an alternative function, such as a “Voigtian” PDF, and comparing the difference in the fitted signal strength with respect to the case of a DSCB plus a Gaussian. This uncertainty is generally small, typically below 1%. For this reason, a flat 1% uncertainty is considered in the main fit in Section 3.6.

Muon momentum scale and resolution The Rochester corrections, mentioned in Section 3.2.3, come with their own set of uncertainties, which arise from the limited amount of available data samples. These uncertainties are applied to the analysis and yield a very small difference in the resolution of the reconstructed H or Z peak. Other effects, such as thresholds effects, are negligible. The impact of these uncertainties on the signal model shape is reported in Table 3.6.

Photon energy scale and resolution The photon uncertainties mainly influence the shape of the $m_{\mu\mu\gamma}$ distribution and affect the mean and width of the signal model. The source of uncertainties are:

- **Photon energy scale and resolution:** corrections are applied to the photon energy scale and to the energy resolution in simulation. The uncertainties related to these corrections are computed using signal events. The magnitude of the resulting uncertainty is below 1%.
- **Non-linearity of the photon energy scale:** any remaining differences in the linearity of the photon energy scale between data and simulation are covered by this uncertainty, which is estimated using Lorentz-boosted events. An uncertainty of 0.1% on the photon energy scale is assigned in this analysis. The procedure for estimating this uncertainty is detailed in [101].
- **Non-uniformity of light collection:** within a given ECAL crystal, there is an uncertainty associated with the modeling of the light collection as a function of the emission depth. This uncertainty is estimated by comparing simulation with the longitudinal shower profile estimates, and the procedure is detailed in [101].
- **Modeling of material in front of the ECAL:** the behavior of electromagnetic showers is affected by the amount of material present in front of the ECAL. This behavior may not be well modeled in simulation, and thus special samples with variations in amount of upstream material are used to compute the impact on the photon energy scale. For most central photons, the magnitude of this uncertainty ranges from 0.02-0.05%, increasing to approximately 0.24% for photons in the endcap [63].

The impact of these uncertainties on the signal model shape is reported in Table 3.6.

Jet energy scale and resolution The JES and JER corrections are measured using the p_T balance of jets with Z bosons and photons in $Z \rightarrow ee, \mu\mu$, and $\gamma + \text{jets}$ events, as well as the p_T balance between jets in dijet and multijet events [83]. At the momentum scale of the Higgs boson, like in the case of this analysis, the uncertainty in the jet energy scale amounts to a few

Source	Mean	Sigma
Muon corrections	$\pm 0.05\%$	$\pm 0.8\%$
Photon energy scale	$\pm 0.04\%$	$\pm 3.4\%$
Photon energy resolution	—	$\pm 1.8\%$

Table 3.6 • Impact of the shape uncertainties on the signal mean and width.

percent and depends on p_T and η . The impact of jet energy scale uncertainties on event yields is evaluated by varying the jet energy corrections within their uncertainties and propagating the effect to the variable on which the final cut is applied. The uncertainties among different years are conservatively considered as correlated.

Only the VBF category is significantly affected by jet reconstruction and measurement, as it is the only one requiring more than one jet that could fall in the forward region outside the tracker geometrical acceptance. The event yield variations due to the JES and JER systematic uncertainties amount to 2.9% and 1.0%, respectively.

The heavy flavor category also requires one jet, but no requirement on the p_T is applied besides the basic $p_T > 30$ GeV threshold, making the effect of jet energy scale and resolution negligible. Other categories do not include selections on jets at all, and thus they are independent on JES and JER uncertainties.

Heavy flavor tagging Uncertainties in the jet b -tagging arise from differences in efficiency between data and MC, which are measured centrally by the b -tagging and vertexing POG [86]. These uncertainties include the statistical component of the estimate of the fraction of heavy and light flavor jets in both data and simulation.

Events are reweighted using the official SFs for Run-2 [87]. The systematic uncertainty on the b -tagging efficiency is derived from the uncertainty in the SF measurement. These uncertainties are evaluated simultaneously for all years.

Only the HF category is affected by the b -tagging efficiency uncertainty, resulting in a $\pm 5.4\%$ variation in the normalization of the Higgs boson signal. The variations in other categories are negligible because they do not make explicit selections based on b -tagging, and the fraction of signal events containing genuine b quarks (ttH, bbH) is very small compared to the inclusive cross section.

Trigger prefire An additional uncertainty is introduced to account for a gradual shift in the timing of the inputs of the ECAL L1 trigger in the region at $|\eta| > 2.0$, which caused a specific trigger inefficiency during 2016 and 2017 data taking. This inefficiency can affect photons and, to a greater extent, jets.

Adopting a conservative approach, the variation resulting from the application of the prefire weights is taken as the full uncertainty. The resulting uncertainty is largest for the categories targeting VBF production, with a maximum impact on the yield of 2.7% and 0.6% for the H and Z searches, respectively.

Source	Type	Backgrounds		Signals	
		QCD	$Z \rightarrow \mu\mu\gamma$	$Z \rightarrow \psi(nS)\gamma$	$H \rightarrow \psi(nS)\gamma$
Integrated luminosity	norm.	—	—	1.6%	1.6%
Fact. and renorm. scales	norm.	—	—	3.5%	5.6%
PDF	norm.	—	—	1.7%	3.2%
Pileup modeling	norm.	—	—	0.9%	0.8%
Trigger efficiency	norm.	—	—	3.5%	3.5%
Trigger timing shift	norm.	—	—	0.6%	2.7%
Muon identification/isolation	norm.	—	—	2.0%	1.4%
Photon identification	norm.	—	—	1.5%	1.5%
Electron veto	norm.	—	—	1.0%	1.0%
Jet energy scale (VBF category)	norm.	—	—	—	2.9%
Jet energy resolution (VBF category)	norm.	—	—	—	1.0%
Jet b tagging (HF category)	norm.	—	—	—	5.4%
QCD background model	shape	discrete	—	—	—
Resonant bkg. model	norm.	—	1%	—	—
Signal model	norm.	—	—	1%	1%
Size of MC samples	$m_{\mu\mu\gamma}$ mean	—	0.06%	0.06%	0.06%
Size of MC samples	$m_{\mu\mu\gamma}$ width	—	0.1%	0.1%	0.1%
μ calibration	$m_{\mu\mu\gamma}$ mean	—	0.05%	0.05%	0.05%
μ calibration	$m_{\mu\mu\gamma}$ width	—	0.8%	0.8%	0.8%
γ energy	$m_{\mu\mu\gamma}$ mean	—	0.04%	0.04%	0.04%
γ energy	$m_{\mu\mu\gamma}$ width	—	3.4%	3.4%	3.4%
γ resolution	$m_{\mu\mu\gamma}$ width	—	1.8%	1.8%	1.8%

Table 3.7 • Sources and types of systematic uncertainties. Normalization (norm.) uncertainties yield a variation in the number of events equal to the reported value. Other uncertainties also vary the shape of the mean and the width of the resonant background and signal $m_{\mu\mu\gamma}$ distribution by the indicated amount. Uncertainties in the same line are treated as correlated.

3.5.4 Summary

A summary of the systematic uncertainties is provided in Table 3.7. The theoretical prediction and trigger efficiency uncertainties present the most significant impact on the signal yield. The former may be reduced through refined calculations of the branching fraction predictions for the signal processes under study and from more precise H production cross section predictions. The estimation of the latter is limited by the low amount of events available for its computation, as discussed in Section 3.5.2. The shape systematic uncertainties contribute less to the overall uncertainty, and their impact on the final result is considerably lower compared to other sources and to the discrete systematic uncertainty originating from the QCD background model choice.

3.6 Results

As no significant excess is observed with respect to the SM expectation, exclusion limits are set by performing a background-only fit and a combined signal-plus-background fit to the $m_{\mu\mu\gamma}$ mass distributions, independently for the Z and Higgs boson signals, and the J/ψ and $\psi(2S)$ hypotheses. The Higgs and Z boson signals are fitted separately to the H and Z SRs, respectively. Each fit also includes the CR, which constrains the normalization of the Z boson Dalitz backgrounds. In contrast, the H boson Dalitz backgrounds are fixed to the SM background expectation, as the number of selected events is not sufficient to constrain this process.

In the fit, based on a profile likelihood, the parameters and the normalization of the background in each category are left free to float, while for the $H \rightarrow \mu\mu\gamma$ background the normalization is fixed to the MC expectation and for the $Z \rightarrow \mu\mu\gamma$ the normalization is constrained by the CR. The systematic uncertainties in the signals are treated as nuisance parameters, with gaussian constraints for the muon and photon energy scale and resolution, and log-normal constraints for the normalization uncertainties. These uncertainties are profiled in the statistical interpretation [102].

3.6.1 Upper limits at 95% CL on branching fraction

Upper limits at 95% CL are set using the CL_s criterion [103, 104] in the asymptotic approximation [105], and cross-checked with those obtained using the generation of pseudo-data [102]. The limits derived with the two methods are found to be in agreement within 5%.

The results are shown in Fig. 3.14 and summarized in Table 3.14, where the branching fraction and signal strength values for each channel are reported. The upper limits on the branching fractions are compared to the previous results from the CMS and ATLAS Collaborations [70, 106] in Table 3.9. No statistically significant excess is observed. A deficit in the data consistent with a downward fluctuation with a significance of 2.2 standard deviations is seen in the $Z \rightarrow \psi(2S)\gamma$ search channel.

Process	$\mu_{\text{obs}}(\mu_{\text{exp}})$	$\sigma \times \mathcal{B}_{\text{obs}}(\sigma \times \mathcal{B}_{\text{exp}})$ (pb)	$\mathcal{B}_{\text{obs}}(\mathcal{B}_{\text{exp}})$
$Z \rightarrow J/\psi\gamma$	$7.2 \left(8.6^{+4.1}_{-2.7}\right)$	$3.8 \left(4.4^{+1.9}_{-1.3}\right) \times 10^{-2}$	$0.6 \left(0.7^{+0.3}_{-0.2}\right) \times 10^{-6}$
$Z \rightarrow \psi(2S)\gamma$	$29 \left(68^{+36}_{-22}\right)$	$8 \left(19^{+8}_{-6}\right) \times 10^{-2}$	$1.3 \left(3.1^{+1.4}_{-0.9}\right) \times 10^{-6}$
$H \rightarrow J/\psi\gamma$	$88 \left(62^{+30}_{-19}\right)$	$1.4 \left(1.0^{+0.5}_{-0.3}\right) \times 10^{-2}$	$2.6 \left(1.8^{+0.9}_{-0.6}\right) \times 10^{-4}$
$H \rightarrow \psi(2S)\gamma$	$970 \left(781^{+417}_{-259}\right)$	$5.5 \left(4.4^{+2.3}_{-1.5}\right) \times 10^{-2}$	$9.9 \left(8.0^{+4.2}_{-2.6}\right) \times 10^{-4}$

Table 3.8 • Observed (expected) upper limits at 95% CL on the normalized values with respect to the SM expectation, denoted as the signal strength parameter μ , the product of the cross section σ and the branching fraction \mathcal{B} of the $(H, Z) \rightarrow \psi(nS)\gamma$ decays, and the branching fraction assuming a SM Z and H boson cross section.

Process	This work (123 fb ⁻¹) $\mathcal{B}_{\text{obs}}(\mathcal{B}_{\text{exp}})$	CMS (36 fb ⁻¹) [70] $\mathcal{B}_{\text{obs}}(\mathcal{B}_{\text{exp}})$	ATLAS (139 fb ⁻¹) [106] $\mathcal{B}_{\text{obs}}(\mathcal{B}_{\text{exp}})$
$Z \rightarrow J/\psi\gamma$	$0.6 \left(0.7^{+0.3}_{-0.2}\right) \times 10^{-6}$	$1.5 \left(1.7^{+0.7}_{-0.5}\right) \times 10^{-6}$	$1.2 \left(0.7^{+0.3}_{-0.2}\right) \times 10^{-6}$
$Z \rightarrow \psi(2S)\gamma$	$1.3 \left(3.1^{+1.4}_{-0.9}\right) \times 10^{-6}$	—	$2.4 \left(3.0^{+1.3}_{-0.8}\right) \times 10^{-6}$
$H \rightarrow J/\psi\gamma$	$2.6 \left(1.8^{+0.9}_{-0.6}\right) \times 10^{-4}$	$7.6 \left(5.2^{+2.4}_{-1.6}\right) \times 10^{-4}$	$2.0 \left(1.8^{+0.8}_{-0.5}\right) \times 10^{-4}$
$H \rightarrow \psi(2S)\gamma$	$9.9 \left(8.0^{+4.2}_{-2.6}\right) \times 10^{-4}$	—	$10.5 \left(8.1^{+3.6}_{-2.3}\right) \times 10^{-4}$

Table 3.9 • Comparison of observed (expected) upper limits at 95% CL on the normalized values with respect to the SM with the previous ones by CMS and ATLAS Collaborations [70, 106].

The upper limits on the $H \rightarrow J/\psi\gamma$ signal strength are interpreted within the κ -framework [107] to constrain the charm quark Yukawa coupling. The ratios of the signal strengths for the $H \rightarrow J/\psi\gamma$ and $H \rightarrow \gamma\gamma$ decays and their SM expectations are approximated as the ratio of their respective partial decay widths, normalized to their SM expectations Γ^{SM} . Writing the partial $H \rightarrow J/\psi\gamma$ decay width in terms of both direct and indirect amplitudes \mathcal{A}_{dir} and \mathcal{A}_{ind} from the diagrams




Figure 3.14 • Observed and expected (with ± 1 , ± 2 standard deviation bands) exclusion limits on the branching fraction of the $(H, Z) \rightarrow \psi(nS)\gamma$ decays.

reported in Fig. 1.6, the ratio κ_c/κ_γ between the coupling modifiers for the charm quark Yukawa coupling and the effective coupling of the Higgs boson to photons can be written as:

$$\begin{aligned} \frac{\mu(H \rightarrow J/\psi\gamma)}{\mu(H \rightarrow \gamma\gamma)} &\approx \frac{\Gamma(H \rightarrow J/\psi\gamma)/\Gamma^{\text{SM}}(H \rightarrow J/\psi\gamma)}{\Gamma(H \rightarrow \gamma\gamma)/\Gamma^{\text{SM}}(H \rightarrow \gamma\gamma)} = \\ &= \frac{|\mathcal{A}_{\text{ind}}\kappa_\gamma + \mathcal{A}_{\text{dir}}\kappa_c|^2}{\kappa_\gamma^2} \cdot \frac{1}{\Gamma^{\text{SM}}(H \rightarrow J/\psi\gamma)} = \frac{|\mathcal{A}_{\text{ind}} + \mathcal{A}_{\text{dir}}\kappa_c/\kappa_\gamma|^2}{\Gamma^{\text{SM}}(H \rightarrow J/\psi\gamma)}. \end{aligned} \quad (3.19)$$

In this interpretation, the signal strength for $H \rightarrow \gamma\gamma$ comes from Ref. [108], while the values for the direct and indirect amplitudes for $H \rightarrow J/\psi\gamma$ are extracted from Ref. [29]. This interpretation follows the methodology used by the ATLAS Collaboration [106], which is based on Refs. [29, 109]. The observed constraint on the κ_c/κ_γ ratio at 95% CL corresponds to the interval $(-157, +199)$, while the expected interval from pseudo-experiments is $(-121, +161)$, with the likelihood scan over κ_c/κ_γ shown in Fig. 3.15a.

An alternative interpretation is considered. Under the pure SM assumption, κ_γ is equal to 1, and consequently $\mu(H \rightarrow \gamma\gamma) = 1$. Thus, the equation Eq. 3.19 simplifies to:

$$\mu(H \rightarrow J/\psi\gamma) \approx \frac{|\mathcal{A}_{\text{ind}} + \mathcal{A}_{\text{dir}}\kappa_c|^2}{\Gamma^{\text{SM}}(H \rightarrow J/\psi\gamma)}. \quad (3.20)$$

In this interpretation, the constraint is applied to κ_c directly. The observed and expected intervals at 95% CL are $(-166, +208)$ and $(-128, +169)$, respectively. The likelihood scan over the κ_c is shown in Fig. 3.15b.

3.6.2 Measurement of the Z decay to $\ell\ell\gamma$

A dedicated discussion has to be reserved for the control region. Here, it is possible to perform a direct fit to the resonant $Z \rightarrow \mu\mu\gamma$ background, excluding other regions and signals. The CR selected events allow to perform a SM measurement by verifying if the number of events from

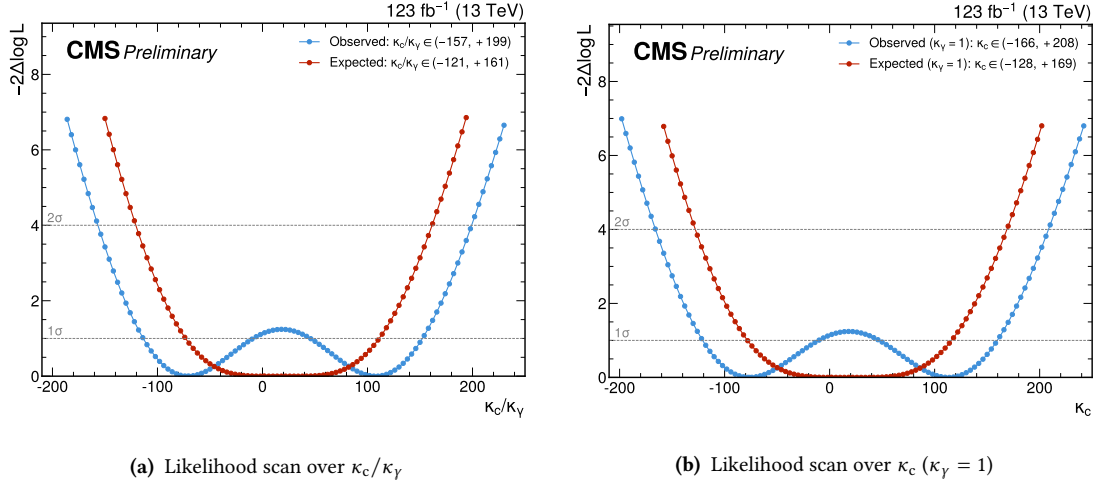


Figure 3.15 • Likelihood scan over κ_c/κ_γ (3.15a) and over κ_c with $\kappa_\gamma = 1$ (3.15b) with expected and observed constraints on the respective parameter of interest at 95% CL. The bump structure at the center of the scan performed on the observed dataset is due to fluctuations in the data within the $m_{\mu\mu\gamma}$ signal region. The scan over pseudo-datasets is not centered at 0 due to the contribution of the indirect amplitude to the process.

the $Z \rightarrow \mu\mu\gamma$ resonance is consistent with the SM expectation. The same background modeling approach used in the main searches is applied also here. The uncertainties are the same as those described in Section 3.5. The ratio between the modeled and expected number of events is:

$$\mu_{Z \rightarrow \mu\mu\gamma} = \frac{\sigma}{\sigma_{\text{SM}}} = 1.18 \pm 0.12 \quad . \quad (3.21)$$

The measured cross section of the $Z \rightarrow \mu\mu\gamma$ background is statistically compatible with the SM expectation.

3.6.3 Summary and outlook of the analysis

With respect to the previous CMS search [70], the sensitivity of this analysis shows an improvement that surpasses the expectations from a simple increase in the integrated luminosity. This enhancement is due to a more sophisticated event selection, including the use of angular variables in the decays, the categorization of the Higgs boson depending on the production mode, and the adoption of a more advanced statistical treatment, represented by the discrete profiling method described in Section 3.5.1. When compared with the 2016 CMS search updated with the larger data set used in this work, this analysis achieves an improvement by a factor of 2 in the expected limit in the case of the $H \rightarrow J/\psi\gamma$ decay, and of 30% in the case of the $Z \rightarrow J/\psi\gamma$ decay.


The same method used to compute upper limits is applied to project upper limits under the hypothesis of combining the Run-2 and Run-3 available datasets, totaling approximately $123+300 \text{ fb}^{-1}$ by the end of Run-3 campaign. The same signal and background models are used for this purpose, with normalization adjusted to match the yield from the total integrated luminosity. In these projections, the H and Z production cross sections are fixed to their predicted values for $\sqrt{s} = 13 \text{ TeV}$ for Run-2, and $\sqrt{s} = 13.6 \text{ TeV}$ for Run-3. Systematic uncertainties are conservat-

ively left unchanged with respect to the Run-2 analysis. The observed data are replaced with a pseudo-experiment generated from the total background model, with event counts reflecting the considered integrated luminosity and maintaining the same event selection. Under these assumptions, the expected upper limits at 95% CL on the branching fractions for the processes under study are shown in Table 3.10.

Process	Analysis results (123 fb ⁻¹)		Analysis projection (123+300 fb ⁻¹)	
	μ_{exp}	\mathcal{B}_{exp}	μ_{exp}	\mathcal{B}_{exp}
$Z \rightarrow J/\psi \gamma$	$8.6^{+4.1}_{-2.7}$	$0.7^{+0.3}_{-0.2} \times 10^{-6}$	$4.5^{+2.1}_{-1.3}$	$0.4^{+0.2}_{-0.1} \times 10^{-6}$
$Z \rightarrow \psi(2S) \gamma$	68^{+36}_{-22}	$3.1^{+1.4}_{-0.9} \times 10^{-6}$	34^{+17}_{-11}	$1.6^{+0.7}_{-0.5} \times 10^{-6}$
$H \rightarrow J/\psi \gamma$	62^{+30}_{-19}	$1.8^{+0.9}_{-0.6} \times 10^{-4}$	30^{+14}_{-9}	$0.9^{+0.4}_{-0.3} \times 10^{-4}$
$H \rightarrow \psi(2S) \gamma$	781^{+417}_{-259}	$8.0^{+4.2}_{-2.6} \times 10^{-4}$	356^{+173}_{-111}	$3.6^{+1.7}_{-1.1} \times 10^{-4}$

Table 3.10 • Comparison between observed (expected) upper limits at 95% CL on the signal strength μ and the branching fraction \mathcal{B} of the (H, Z) $\rightarrow \psi(nS)\gamma$ decays from this work using 123 fb⁻¹ of collision data, and their expected projection using a pseudo-dataset, mimicking the Run-2 and Run-3 datasets combination and corresponding to approximately 123+300 fb⁻¹ of integrated luminosity.

The same methods as in Eqs. 3.19-3.20 are used for the interpretation of the upper limit projections to set constraints on κ_c/κ_γ and κ_c . The expected constraint at 95% CL on κ_c/κ_γ from the $H \rightarrow J/\psi \gamma$ projected upper limit is found to be the interval $(-79, +120)$, as shown in the likelihood scan in Fig. 3.16a. Under the $\kappa_\gamma = 1$ hypothesis, the projected constraint on κ_c is set to the interval $(-84, +125)$, as shown in the likelihood scan in Fig. 3.16b. Both cases are compared with the scans for the expected constraints using the Run-2 dataset.



(a) 123+300 fb⁻¹ projection of likelihood scan over κ_c/κ_γ (b) 123+300 fb⁻¹ projection of likelihood scan over κ_c ($\kappa_\gamma = 1$)

Figure 3.16 • Projection to 123+300 fb⁻¹ of collision data of likelihood scan over κ_c/κ_γ (3.16a) and over κ_c with $\kappa_\gamma = 1$ (3.16b) with expected constraints on the respective parameter of interest at 95% CL. The scans of the projection study are compared to the scans for the expected constraints of this work, obtained using the Run-2 dataset and shown in Fig. 3.15.

The projection study highlights how the upper limit results are expected to improve proportionally to the square root of the integrated luminosity. It is important to note that the same systematic uncertainties discussed in Section 3.5 have been used and a more sophisticated treatment may

improve the results. Additionally, improvements in the signal efficiency are anticipated due to a refined event selection tailored to the features available in the Run-3 dataset. Therefore, a future iteration of this analysis is expected to yield better results than this projection, potentially approaching the sensitivity needed to observe the $Z \rightarrow J/\psi \gamma$.

4

Enhancing the CMS physics reach with Real-Time analysis at the Level-1 trigger

Searching for rare Higgs boson decays is a challenging and often statistically limited task, particularly when the final state involves products from meson decays. Triggering on these specific signatures is complex, because the final state might resemble a common pp collision. While a dedicated trigger can be easier to implement for the case of $H \rightarrow J/\psi \gamma$, thanks to the presence of muons in the final state, more complex strategies are needed for decays like $H \rightarrow \phi \gamma$, where the absence of muons makes triggering more difficult. The subsequent decay of a meson to a pair of charged hadrons produces a two-track plus photon final state to select within a high number of tracks in the event. Moreover, setting a higher trigger threshold on the photon transverse energy to reduce the trigger rate might result in a significant loss of signal efficiency. For searches where the number of expected signal events is relatively low and thus statistically limited, maintaining the signal efficiency as high as possible results in a significant improvement in the analysis sensitivity and contributes to a potential observation of the rare decay under study.

A more complex challenge arises when trigger handles like a photon are absent in the signature of interest. An example is a final state composed only of generic charged tracks, such as kaons or pions. In this context, the SM predicts rare Higgs decays to a pair of mesons, which the CMS experiment has already explored for J/ψ or Y pairs [38]. No public results exist to date for decays to a pair of ρ or ϕ mesons, due to the difficulty of developing an efficient trigger for these signatures, or to $\phi J/\psi$ [6], where in principle a trigger strategy could exploit the decay $J/\psi \rightarrow \mu\mu$.

With the CMS Phase-2 upgrade for HL-LHC, tracker tracks and PF candidates will be available at the Level-1 trigger, opening new possibilities for these searches and improving the sensitivity to the Higgs decays treated in this work. Triggering on these signatures within the latency and hardware resource constraints presents another challenge, as a higher average pileup of 200 will increase the amount of collision products to analyze. Thus, keeping a high signal efficiency while controlling the trigger algorithm rates is a challenging task.

This chapter builds on the analysis described in the previous, generalizing it to a wider physics landscape and introducing the novel paradigm of “Data Scouting” at the first level of the CMS trigger chain to extend the current physics boundaries. Given the central role in this work of the CMS L1T, a detailed description of the system will be presented with an overview of the current Phase-1 L1T in Section 4.2 and a summary of the capabilities of its future Phase-2 upgraded

version in Section 4.3. The Level-1 trigger Data Scouting (L1DS), enabling real-time analysis using the subset of information available at L1, will be introduced in Section 4.4 to mitigate the limitations discussed above.

4.1 Present, future and limitations of the Level-1 trigger system

The current CMS Phase-1 L1T trigger is designed to select interesting events among all the collisions and reduce the event rate from the 40 MHz bunch crossing frequency to around 100 kHz. The event reconstruction at L1 is performed using a subset of the detector information, producing coarse resolution physics objects, which are the input to a set of algorithms that determine whether the event will undergo finer reconstruction at the next trigger level (HLT).

The trigger selection, starting with the decision at L1, inherently introduces a bias in the collected dataset. As the Run-2 data taking campaign has excluded a big portion of the phase-space for traditional BSM models, from the start of Run-3 the CMS experiment has extended the physics program to more exotic signatures, such as displaced particles originating from the decay of Long-Lived Particles (LLP). In particular, new triggers have been implemented to extend the sensitivity to signatures with displaced muon and jets [110, 111], enhancing the sensitivity on scenarios like Dark Matter (DM) decays. These signatures are summarized in Fig. 4.1.

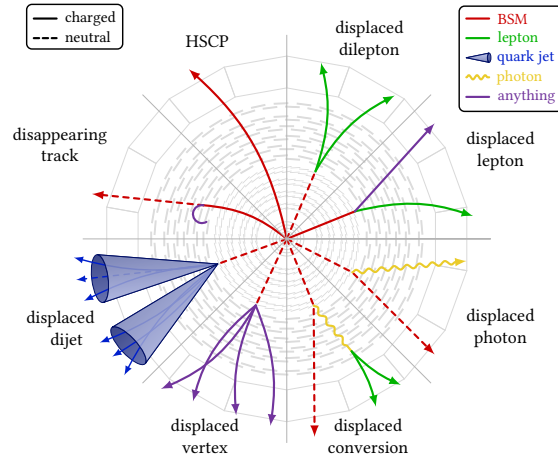


Figure 4.1 • Illustration of common BSM signatures in exotic physics searches: Heavy Stable Charged Particles (HSCP), disappearing tracks, displaced dijet or dilepton, displaced vertex, displaced photon conversion, displaced photon radiation or lepton. Charged (neutral) particles are portrayed with solid (dashed) lines. BSM particle content is shown in red, leptons in green, quark jets with blue cones, photons in yellow curly lines, everything else in purple.

For the HL-LHC upgrade, the CMS Phase-2 L1T is designed to maintain or improve the acceptance for all physics objects under the HL-LHC pileup conditions, and can profit from an increase in the maximum output rate from 100 to 750 kHz, thanks to the upgrade of the readout systems of the CMS subdetectors. A larger amount of information will be accessible by the L1 reconstruction algorithms, including higher granularity primitives and tracker tracks, thereby enhancing the physics capabilities of the trigger system and advancing towards the ambitious goal of performing a real-time analysis of the collision data.

4.1. Present, future and limitations of the Level-1 trigger system

Despite the significant extensions in the CMS physics program with the new Run-3 L1T algorithms and the further enhancements expected with the Phase-2 L1T upgrade, CMS may not be able to fully cover the entire phase-space of physics processes. Once an event is discarded at L1, it is lost forever and cannot be analyzed further. The same consideration applies for the selection at HLT, where the discarded events are not available for offline reconstruction and analysis. Several approaches have been developed at HLT to mitigate or solve this limitation, as schematized in Fig. 4.2:

- **Data Parking:** storing a large amount of raw detector data collected by algorithms with low trigger thresholds to be processed later when sufficient computational power is available to handle such data. This approach takes advantage of the large storage throughput available at the DAQ, allowing the collection of data even though the offline reconstruction cannot keep up with the rate. The offline reconstruction applied to this raw data is the same applied to the standard triggered data, thus the event content is the same.
- **Data Scouting:** moving the physics analysis at the trigger level, by relaxing the HLT selection and running the HLT online reconstruction on more events. This approach brings significant advantage whenever the information and the resolution of the objects reconstructed at HLT is sufficient. No offline reconstruction is applied and only a tiny event record, compared to the full detector information, is stored for further analysis.

The Data Scouting approach was first introduced at the HLT level since the start of Run-1 and further developed during Run-2 and Run-3. HLT scouting has notably extended exclusion limits and constraints on dijet and dilepton searches [112], and led to the observation of SM physics processes not visible with the standard triggered data [113].

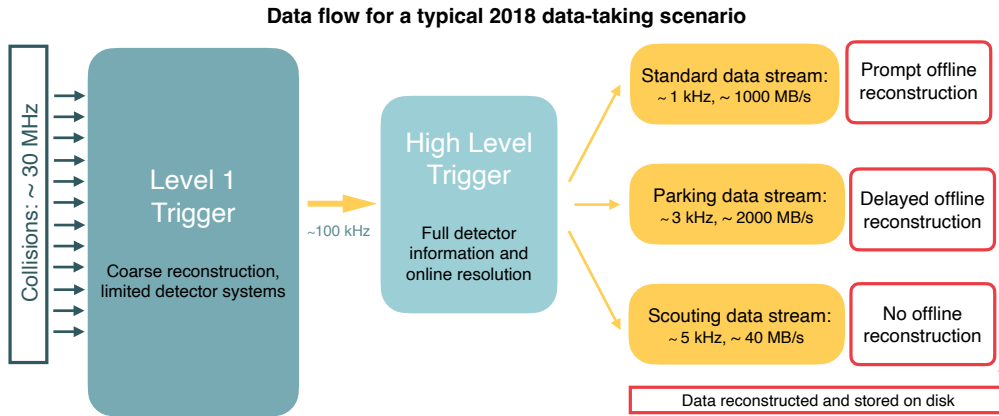


Figure 4.2 • A schematic view of the typical Run-2 data flow during 2018 showing the data acquisition strategy with scouting and parking data streams, along with the standard data stream. A value of instantaneous luminosity of $1.2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ over a typical 2018 fill, corresponding to an average pileup of 38, is considered. Taken from [112].

The scouting approach has been introduced and commissioned at L1 from the end of Run-2 and during Run-3, as described in Chapter 5. The Run-3 demonstrator of the L1DS collects the muon and calorimetric objects reconstructed by the Phase-1 Global Muon and Calorimeter

Trigger subsystems [114], respectively, and it is sketched in Fig. 4.3. The L1DS system will be fully integrated into the CMS Phase-2 L1T project [115], and it will collect and perform an online analysis at the full event collision rate on the reconstructed Phase-2 L1T objects, which will have efficiency and resolution often comparable to the offline one. The system will enrich the standard physics program as presented in the Phase-2 L1T technical design report. Additionally, it can explore statistically limited signatures to increase their selection efficiency, as for instance the Higgs exclusive rare decays involving light and heavy mesons presented in this work. A feasibility study performed using Phase-2 L1T objects from the simulation of these decays will be presented in Chapter 6.




Figure 4.3 • A schematic view of the Level-1 trigger Data Scouting flow inside the CMS data acquisition and L1 trigger chains.

4.2 Performance of the CMS Phase-1 Level-1 trigger

The Phase-1 L1T is designed to cope with the LHC Run-2 and Run-3 luminosity and pileup conditions, replacing the legacy system used during Run-1 [50]. The system is engineered to produce an accept decision within approximately $4 \mu\text{s}$ and to reduce the event accept rate to approximately 100 kHz. As it is the first stage in the data acquisition chain, the L1T implementation is designed to be flexible, scalable, and tunable to the varying LHC running conditions and continuously evolving physics program.

The L1T processors are FPGA boards in a Micro Telecommunications Computing Architecture (μTCA) form factor. These processors receive inputs known as trigger primitives (TPs), produced in situ by the front end electronics of the detector. In the Phase-1 CMS detector, only a subset of the detector information generates TPs, specifically from the muon and calorimetric systems. Thus, the L1 event reconstruction is fast and coarse, producing limited resolution objects for the purpose of selecting interesting collisions. During the reconstruction, the input TPs are processed by dedicated trigger algorithms, which reconstruct higher-level objects to provide L1 candidate muons, jets, e/γ , τ_h and energy sums. These L1 objects are then input into the global trigger system, where several hundred selection algorithms are applied to target generic physics

4.2. Performance of the CMS Phase-1 Level-1 trigger

signatures, producing a set of output bits. The final logic “OR” of these bits constitutes the output of the μ GT system and provides the event accept signal for the subsequent full acquisition of the detector data, which is then fed to the HLT for further filtering prior to storage. The trigger chain, encompassing all local, regional and global systems, is schematized in Fig. 4.4. The main subsystems are described in the following subsections.

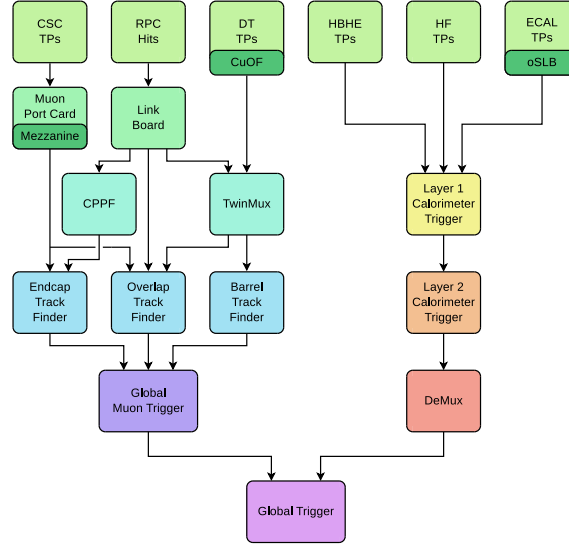


Figure 4.4 • Diagram of the CMS Phase-1 Level-1 trigger system. The muon trigger subsystem reconstructs L1 muon candidates from the TPs provided by the CSC, RPC and DT detectors. The energy deposits in the ECAL and HCAL are combined in the calorimeter trigger subsystem to reconstruct L1 jets, e/γ , τ_h and energy sums. The muon and calorimeter trigger output candidates are then combined in the global trigger, where the event accept or reject decision is made.

4.2.1 Calorimeter Trigger

The Calorimeter Trigger (CT) employs a Time-Multiplexed Trigger (TMT) architecture [116, 117], schematized in Fig. 4.5. In this design, a single processing board receives data from the entire calorimeter, taking more time to fully process the event and improving the overall reconstruction performance under high pileup conditions. Consecutive events are distributed to different processing boards, such that the latency constraints can be achieved.

The CT architecture consists of two layers. The Layer-1 receives data on 18 Calorimeter Trigger Processor (CTP7) boards, consisting of TPs from the calorimeters with a size of $\eta \times \phi \approx 0.087 \times 0.087$. Each CTP7 covers a sector of $\Delta\phi = 20^\circ$ and produces a trigger primitive known as Trigger Tower (TT), which encodes the sum of ECAL and HCAL energy deposits. The output of Layer-1 for the entire calorimeter region and for a single event is sent to the second layer, where the reconstruction algorithms are applied. The Layer-2 consists of 9 MP7 processing FPGA boards plus a spare redundant one. The output of the reconstruction consists of up to 12 jets, 12 e/γ , 12 τ_h and global quantities such as total and missing transverse energy sums. A demultiplexing stage runs on a single MP7 board before sending the reconstructed objects to the Global Trigger.

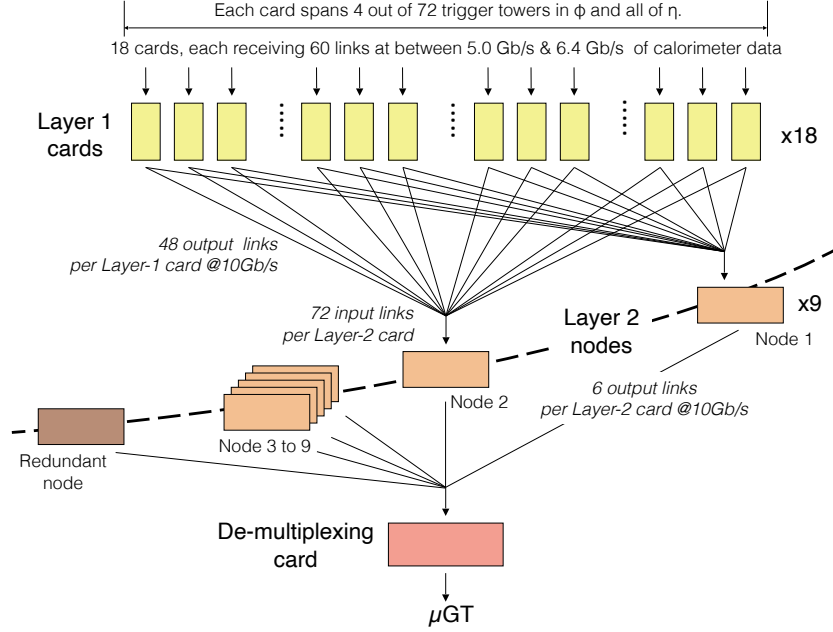


Figure 4.5 • Scheme of the time-multiplexed architecture of the CMS Phase-1 calorimeter trigger [48]. The 18 CTP7 cards of Layer-1 collect the information of the calorimeters. Their output is sent to 10 MP7 cards of the Layer-2, which reconstructs L1 jets, e/γ , τ_h and energy sums. These objects are then sent to the μ GT after a demultiplexing stage running on a single MP7 card.

L1 jets, e/γ and τ_h reconstruction L1 jets are reconstructed using a clustering technique within a 9×9 window (AK4) around jet seeds, which are TTs with the highest local energy deposits [48]. Pileup contributions were corrected up to 2023 using the “chunky doughnut” algorithm, which estimates the pileup from a roughly toroidal region around the jet candidate. From 2024, it was replaced by the “phi-ring” subtraction, which uses a ring of TTs at the same η coordinate for improved pileup mitigation. The pileup-subtracted L1 jet candidates are then calibrated based on their energy and η , with the best 12 candidates being forwarded to the μ GT. Moreover, new algorithms introduced in Run-3 allow for the identification of delayed and displaced jets, enhancing searches for long-lived particles (LLPs) [111]. The performance and efficiency of the L1 reconstruction during 2022 is shown in Fig. 4.6a [118].

Photons and electrons are reconstructed as e/γ objects, as they are indistinguishable at this stage in the Phase-1 L1T. The reconstruction begins with identifying seeds from TTs with significant local energy deposits. A dynamic clustering technique is then applied, which includes neighboring towers satisfying an energy threshold energy, forming the e/γ object. A shape veto is incorporated to discriminate against jets, exploiting the narrower profile of electromagnetic showers compared to the broader shapes associated with QCD background events. To mitigate pileup effects, an isolation criterion is applied based on the energy deposits in the ECAL and HCAL within a specified region. This criterion leverages the lower hadronic activity typically surrounding e/γ candidates. The performance of the e/γ reconstruction and identification algorithm observed in early Run-3 data is summarized in Fig. 4.6b.

4.2. Performance of the CMS Phase-1 Level-1 trigger

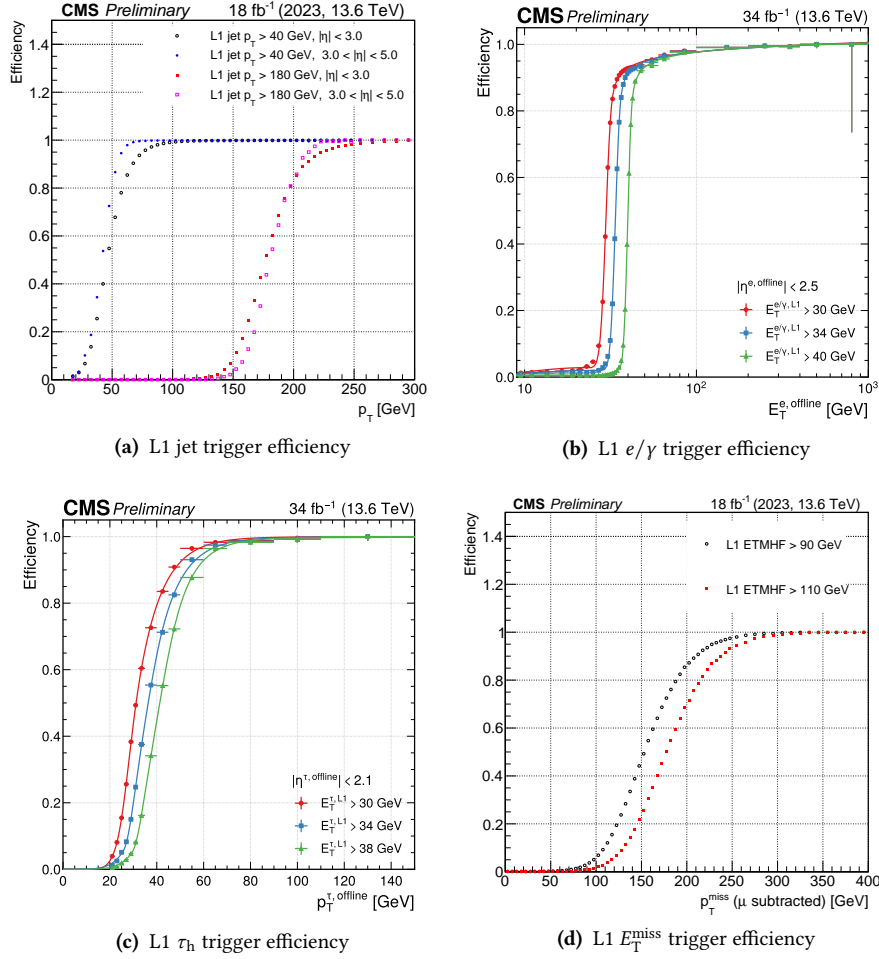


Figure 4.6 • Trigger efficiencies for different thresholds used at L1 for reconstruction of L1 jets (4.6a), e/γ (4.6b) and τ_h (4.6c), taken from Refs. [111, 119, 120]. E_T^{miss} reconstruction efficiency for typical L1 thresholds as a function of offline E_T^{miss} after excluding the contribution of muons (4.6d), taken from Ref. [118].

The L1T algorithm for hadronically decaying τ leptons (τ_h) is designed to select τ_h decays, which typically involve one, two, or three charged or neutral pions [48, 119]. Due to the magnetic field of the CMS detector, these pions may spread across different TTs, requiring a dynamic clustering technique adapted from the e/γ trigger to accurately capture the energy deposits. The raw energy of the τ_h candidate is then calibrated based on factors like the raw energy and η of the candidate. An isolation algorithm is applied to reject background contributions, taking advantage of the fact that τ_h decays tend to create more localized energy deposits compared to QCD jets, which have higher surrounding activity. The reconstruction performance and trigger efficiencies for usual L1 thresholds, measured with tag-and-probe techniques on $Z \rightarrow \tau_\mu \tau_h$, is shown in Fig. 4.6c.

L1 energy sums reconstruction After the reconstruction of the L1 jets, e/γ and τ_h candidates, their information is combined to calculate energy sums of interest. Two energy sum types have particular physics relevance: the missing transverse energy (E_T^{miss} , MET) and the total transverse

energy sum (H_T , HT) of the L1 jets. The MET is determined by calculating the magnitude of the vector sum of E_T across all the TTs and can provide crucial information on invisible particles as neutrinos. A pileup subtraction and calibration algorithm is also applied, with corrections evaluated on an event-by-event basis by analyzing the activity in the central region of the calorimeter barrel. The efficiency of the reconstruction of these energy sums during Run-3 is shown in Fig. 4.6d.

4.2.2 Global Muon Trigger

The Phase-1 Global Muon Trigger (μ GMT) enhances the Run-1 version, where the muon candidates were reconstructed independently for the DTs, RPCs and CSCs. The upgraded system applies a reconstruction on the combined information of the muon detectors in the Barrel Muon Track Finder (BMTF), Endcap Muon Track Finder (EMTF), and Overlap Muon Track Finder (OMTF) subsystems. Each of these regional subsystems can produce up to 36 regional muon candidates, which are then sent to μ GMT processor, composed of a single MP7 board. The μ GMT sorts these candidates, ranking them by p_T and quality, removes duplicates, and forwards the top eight highest-ranked muons to the Global Trigger.

The μ GMT also applies extrapolation of the muon φ and η coordinates back to the collision point, as they are initially reconstructed at the second muon station. The extrapolation parameters are stored in Look-Up-Tables (LUTs) on the FPGA processors and are derived from MC simulation as a function of p_T , φ , η and charge of the muon candidate. The muon parameters extrapolated at the nominal center of the detector are then added to the muon data sent to the μ GT inputs to improve the algorithms that rely on the momentum at vertex of the muons. The reconstruction in the three regions is summarized and the reconstruction performance is presented.

BMTF reconstruction The reconstruction in the muon barrel covers the pseudorapidity region $|\eta| < 0.83$ and utilizes a combination of spatial information from the DTs and timing information from the RPCs. The hits in these detectors are processed by the TwinMux into higher-level objects known as super-primitives, referred as “BMTF stubs” in this work. The TwinMux sends these super-primitives to 12 MP7 processors of the BMTF. The i^{th} BMTF processor receives stubs from its corresponding wedge and the two adjacent wedges, allowing each processor to perform an independent reconstruction.

During Run-2, the tracking of muon candidates from the muon stubs was performed using a track extrapolation method that moved from the inner to the outer muon stations [121]. This method was replaced at the start of Run-3 with a Kalman Filter (KF) method, initially developed for the CMS Phase-2 L1T upgrade [122]. The KF works as a mean squared error minimizer, mathematically resembling a χ^2 fit [55]. Tracks are initiated with a stub from the outermost available station and then propagated inward, with adjustments made at each station step to account for the detector geometry and the magnetic field. Two separate transverse momentum measurements are ultimately stored: one without a collision point constraint (p_T^{unc}) and the other one with the constraint (p_T).

The implementation of the KF in the BMTF has significantly improved the trigger performance compared to the legacy BMTF algorithm. In particular, for muons displaced by more than 50 cm from the CMS center, the efficiency is enhanced by a factor of four [122, 123]. The efficiency observed in early data from Run-3 is shown in Fig. 4.7a.

4.2. Performance of the CMS Phase-1 Level-1 trigger

OMTF reconstruction The OMTF operates in the overlap region ($0.83 < |\eta| < 1.24$), where it combines data from the DTs, RPCs, and CSCs to reconstruct muon tracks. The OMTF is divided into 6 sectors in φ on both the positive and negative side of η , each managed by a Modular Track Finder (MTF7) processor. The reconstruction algorithm is based on pattern recognition, using a set of predefined patterns loaded onto the MTF7 boards, and functions similarly to a naive Bayes classifier [124]. Muon tracks are seeded by reference hits in the inner layers of the detector. The algorithm then matches these seeds with hits in other layers using the predefined patterns, which encode expected track propagation and hit distributions. When multiple matches are found, a statistical estimator helps resolve ambiguities, ensuring that the best candidate is selected. The efficiency of the OMTF reconstruction observed in early data from Run-3 is shown in Fig. 4.7a.

EMTF reconstruction The EMTF covers the endcap region ($1.24 < |\eta| < 2.4$) using data from the CSCs, RPCs, and GEM detectors, the latter introduced in Run-3. Similar to the OMTF, the EMTF divides the region into 6 sectors per endcap, each managed by an MTF7 board. The reconstruction algorithm relies on pattern recognition to match hits across stations, particularly focusing on stations 1 and 2 where the magnetic field is stronger, allowing for more precise p_T assignment. A BDT regression model [125] manages the cases where low p_T muons experience multiple scattering and energy loss, and high p_T muons initiate electromagnetic showers. In Run-3, the EMTF introduced a neural network, implemented directly in the FPGA logic, to improve p_T and transverse impact parameter estimation for displaced muon tracks. Additionally, a new hadronic shower flag has been added to identify LLPs that might produce showers in the endcap muon systems [114]. The EMTF reconstruction efficiency observed during the early phase of Run-3 is shown in Fig. 4.7c.

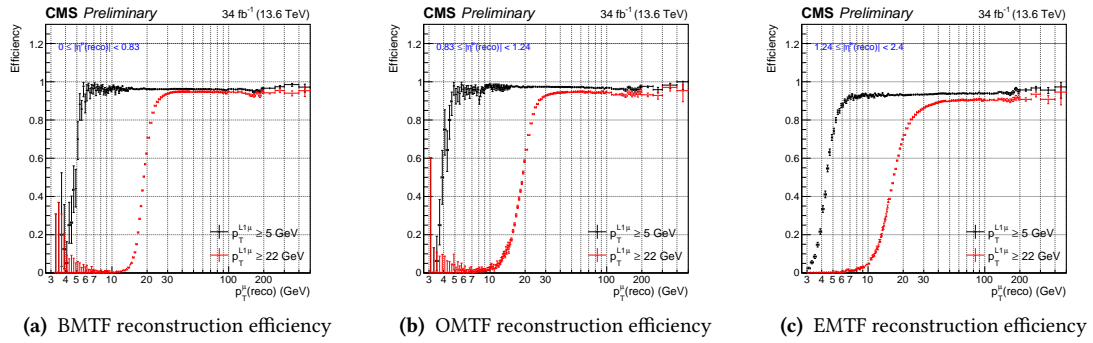


Figure 4.7 • CMS Phase-1 L1 muon reconstruction efficiencies [126], separately for BMTF (4.7a), OMTF (4.7b) and EMTF (4.7c). The data used for the study is from the start of Run-3 and only highest quality muons are considered.

4.2.3 Global Trigger

The μ GT combines the highest ranked 8 muons from the μ GMT, 12 jets, 12 e/γ , 12 τ_h and global quantities such as energy sums from the CT, and makes a trigger decision based on a set of algorithms distributed across six independent MP7 boards. The decisions are OR'ed across the boards to form the final L1 accept signal that is sent to all detector front-ends. The acceptance

rate of the L1 algorithms can be controlled through prescaling, which involves accepting only one out of every N events, where N is the prescale factor. The decision bits for the accepted event are forwarded to the HLT.

The trigger menu contains between 400 and 500 seeds, with inclusive single, double, triple objects triggers or prescaled triggers with lower thresholds for calibration or trigger efficiency measurements. Additionally, the trigger menu includes dedicated triggers designed to target specific physics signatures, such as: triggers computing two-body invariant masses, b -jet tagging triggers requiring a muon in close proximity to a jet [127], displaced muon triggers using the new unconstrained p_T reconstruction, and anomaly detection algorithms [128–131].

4.3 The CMS Phase-2 Level-1 trigger upgrade

The CMS Phase-2 L1T, schematized in Fig. 4.8, is designed to maintain or enhance the CMS acceptance for all physics objects under the HL-LHC pileup conditions of up to 200 simultaneous collisions per bunch crossing. These improvements profit from an increase in the maximum output rate from 100 to 750 kHz [115], thanks to the upgrade of the readout systems of the CMS subdetectors. The Global Calorimeter Trigger (GCT) and Global Muon Trigger (GMT) subsystems will use higher granularity TPs to produce standalone candidates. Particle tracks will be provided at 40 MHz from the tracker back-end to the GMT, to correlate them with muon TPs, and a novel Global Track Trigger (GTT), to perform vertex fitting. The Correlator Trigger (CT) will combine the information from the upstream subsystems and the particle tracks by applying a PF reconstruction directly in hardware. The Global Trigger system (GT) will make a selection decision within a total latency window of around 12 μ s, determining which collision events will be read out for further reconstruction. This decision is based on a set of trigger paths that involve cut-based selections or neural network inference on the reconstructed objects from the upstream subsystems.

The upgraded trigger utilizes modern FPGA-based processors to implement sophisticated algorithms targeting specific final states. The processing boards use the Advanced Telecommunications Computing Architecture (ATCA) standard for the electronics, while the CMS Standard Protocol (CSP) has been adopted for the serial optical link at approximately 25 Gb/s [115]. Research and development efforts are ongoing to further refine the architecture outlined in the technical design report, as well as the algorithms to deploy for L1 object reconstruction. The next subsections describe the state-of-the-art algorithms, available physics objects, and reconstruction performance evaluated on simulated events under the HL-LHC conditions with the CMS Phase-2 detector.

4.3.1 Calorimeter Trigger

The Phase-2 Calorimeter Trigger receives the TPs from the Barrel Calorimeter (BC), the Hadron Forward (HF) and High-Granularity Calorimeter (HGCAL) [115, 132]. The BC forwards the information from ECAL crystals to the Barrel Calorimeter Trigger (BCT), which clusters energy deposits in the calorimeter and merges the crystal TPs into TTs. The Global Calorimeter Trigger (GCT) receives this processed information, creates e/γ , jets, hadronically decaying τ_h , and energy

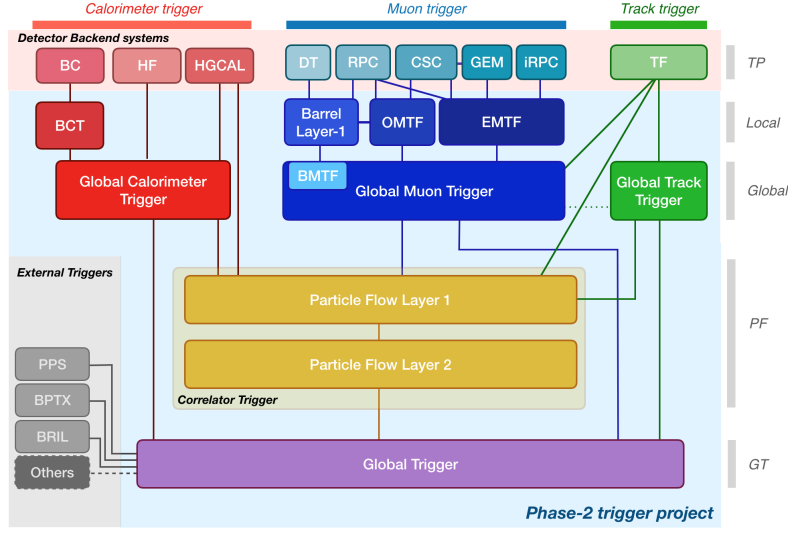


Figure 4.8 • Architecture of the CMS Phase-2 L1 trigger [115]. The various levels of processing are indicated on the right: Trigger Primitives (TP), local and global trigger reconstruction, Particle Flow trigger reconstruction (PF), and global decision (GT).

sum trigger objects, and sends them to the CT for input into the PF algorithm and to the GT. More details on the reconstruction algorithms and performance of the standalone GCT objects are reported in Ref. [115].

4.3.2 Muon Trigger

The Phase-2 Global Muon Trigger receives the TPs from the upgraded DTs, RPCs, CSC and GEMs. As in the Phase-1 trigger, these detectors cover different regions in pseudorapidity, with the same three track finders (barrel, overlap, endcap). The algorithms employed for the three muon track finders will be based on those used during Run-3. The BMTF will use the KF version developed for the Phase-2 trigger FPGA processors and profiting of higher resolution muon stubs as input. The OMTF will improve the baseline naive Bayes classifier used during Run-2 and Run-3. The Phase-2 algorithm for EMTF (EMTF++) will enhance the pattern recognition and neural network p_T assignment thanks to the addition of hits from the new improved RPCs (iRPCs) and GEMs detectors.

The muon track finders generate standalone candidates, which are collected by the GMT. The GMT receives tracker tracks from the GTT to perform muon track-tracker track matching. This algorithm significantly improves the p_T resolution of the muon candidates, reducing the mismeasurement of standalone low p_T muons as high p_T ones. Both the standalone and tracker muons are sent to the CT and to the GT. A more detailed description of the GMT reconstruction and the performance of the standalone objects is available in Ref. [115].

4.3.3 Track Trigger

The identification of charged particle trajectories at the LHC bunch crossing rate of 40 MHz is made possible thanks to the design of the Phase-2 Outer Tracker (OT) [133]. The “ p_T modules” within the OT perform hit correlations between pairs of closely-spaced silicon sensors to form

stubs and estimate their p_T . The OT stubs with a minimal p_T of around 2 GeV are sent to the backend L1 track finding system, where track reconstruction is performed through a Kalman filter approach. Thanks to the applied minimal p_T threshold, the number of stubs to process is significantly reduced and the track reconstruction can be performed at 40 MHz. The resulting tracks are forwarded to the GTT for primary vertex reconstruction and for the identification of tracker-only jets and energy sums.

Since the pixel detector information is not available at the L1 stage, it is not possible to fully identify every interaction within an event, as would be done during the offline reconstruction. The primary objective at this stage is to identify the hard scatter primary vertex and its associated tracks or PF candidates. The proposed algorithm for this purpose is the histogram-based “FastHisto” [134]. The input tracks for FastHisto are filtered to remove fake tracks and select prompt tracks. These filtered tracks are sorted based on their position along the z axis and weighted by their p_T . This process is equivalent to generating a histogram, with the primary vertex determined by the p_T -weighted average of the z position of a bin window of configurable size that has the highest scalar sum of p_T . With the vertex position determined, tracks can be assigned to the vertex based on their z_0 position. Fig. 4.9 shows an example of FastHisto application with $t\bar{t}$ simulated event with 200 pileup interactions.

This step is crucial for downstream systems to associate reconstructed objects with the highest energy interaction in a bunch crossing, thereby reducing hadronic pileup contributions. Moreover, it reduces the amount of objects to process, allowing for the application of more sophisticated algorithms.




Figure 4.9 • Example application of FastHisto algorithm on $t\bar{t}$ simulated events at 14 TeV and 200 pileup [115]. The plot shows a p_T -weighted histogram of the track positions along the z axis. The gray area represents the scalar sum of the track p_T within each bin, while the colored markers indicate the position of each track, including those from pileup, fake tracks and tracks associated with the primary vertex.

4.3.4 Correlator Trigger and Particle Flow

The trigger objects previously introduced are transmitted to the CT, where their information is combined. A first layer in the CT (CL1) takes as inputs standalone muons from the GMT, L1 tracks from the GTT, 3D clusters from HGCal, barrel ECAL clusters and HCAL TTs from the GCT. The

CL1 then performs the PF reconstruction, producing candidate muons, electrons, photons, neutral hadrons and charged hadrons. The reconstruction is regionalized in η and ϕ to enable parallel processing across multiple regions and thus meet the trigger latency requirements [135, 136]. The CL1 PF algorithm is illustrated in the sketch in Fig. 4.10 and involves a series of steps for reconstructing the aforementioned higher-level physics objects.

Calorimeter objects that are outside the tracker acceptance are promoted to neutral hadrons or photons based on their energy deposition pattern. As described in Section 4.3.2, standalone muons are processed first and matched to tracker tracks using ΔR and p_T criteria. Once matched, these muon tracks are excluded from further processing. The remaining tracks are matched to the reconstructed objects from the GCT separately for the barrel and endcap regions. The barrel region, which has finer granularity in the collection of electromagnetic energy deposits, undergoes three reconstruction steps. First, electron and photon candidates are created by matching tracker tracks to electromagnetic clusters. In the second step, the remaining electromagnetic clusters are associated with the nearest hadronic cluster. Finally, in the third step, hadronic clusters are linked to tracks to identify charged and neutral hadrons. In the endcap, the HGCAL 3D clusters are separated in pileup, pions, e/γ objects by means of 2 BDTs.

Pileup subtraction is handled by the Pileup per Particle Identification (PUPPI) algorithm, which uses the reconstructed PV from the GTT to match the previously reconstructed objects to the candidate hard scatter vertex. Objects with charged tracks not associated with the PV are removed, while neutral candidates are either removed or assigned a lower weight depending on their distance along z from the PV. The performance of the PUPPI algorithm is closely tied to the PV assignment quality.

The output of the first CT layer consists of PUPPI candidates. These physics objects are a factor of 10 less than the PF ones. Thus, the PUPPI candidates can be sent to the second layer (CL2), where a global event reconstruction is carried out using identification and isolation algorithms running in parallel.

Multiple jet reconstruction algorithms have been considered for implementation in firmware. These include a pure calorimeter-based algorithm, utilizing the information from the barrel ECAL and HCAL, HGCAL, and the forward HF system, similar to the current Phase-1 jet clustering. Another approach is based on clustering tracks using a lightweight clustering algorithm, which produces tracker-only jets. Finally, jet clustering can be performed using PF objects as input, as demonstrated using an implementation on a small platform. The Seeded Cone algorithm [137] has been recently proposed for the CL2; this algorithm clusters PUPPI candidates and has performance similar to the AK4 offline jet reconstruction [64]. The performance of the aforementioned algorithms are shown and compared in Fig. 4.11a, where the PF version shows superior trigger efficiency compared to the calorimeter-only and tracker-based methods.

The missing transverse energy sum E_T^{miss} is computed across all PUPPI candidates, enhancing the identification of neutrinos and the study of invisible decays. The hadronic transverse energy H_T is calculated from the reconstructed jets. In this case as well, the PUPPI algorithm outperforms those based purely on tracks or calorimeter data, as shown in Fig. 4.11b.

Hadronically decaying τ_h leptons can be reconstructed using a variety of algorithms to maximize the use of the detector systems and ensure robustness. The options explored include calorimeter-only inputs or tracks matched to e/γ clusters, and using PUPPI candidate inputs.

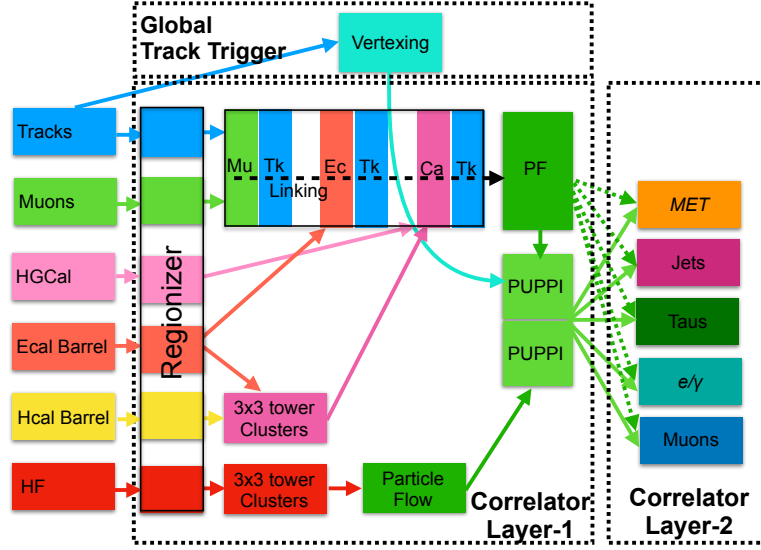


Figure 4.10 • Overview of the PF+PUPPI algorithm, taken from Ref. [115]. The time ordering of the operations is indicated roughly by motion from left to right and the solid black boxes indicate the grouping of the algorithm across different boards (CL1, CL2 and GTT). Each shaded box, regionizer, linker, and clustering denote firmware blocks needed for the algorithm. The top green box labeled “PF” indicates Particle Flow. CL2 is the algorithmic layer and constitutes a broad class of algorithms that can be run in parallel on separate boards. While PUPPI particle inputs are considered the default for CL2, there is also the possibility of using PF candidate inputs, as indicated by the green dashed line.

Several implementations of the τ_h reconstruction based on PUPPI candidates have been examined, including a neural network-based approach. A summary of the performance of these algorithms is shown in Fig. 4.11c.

4.3.5 Global Trigger

The Phase-2 GT will process a significant amount of information from the subsystems described earlier. The increase in number of object collections and variety is expected to result in an extended trigger menu of around 1000 algorithms, evaluated conventionally without time-multiplexing architecture and distributed across multiple GT boards [138]. These algorithms will involve conditions based on single or multiple objects, as well as topological correlations between objects, such as calculating the invariant mass of any pair of objects from the same collection, their transverse mass, or the transverse momentum of a hypothetical parent particle by reconstructing the kinematic of the process. Additionally, ML-based algorithms as NNs and BDT will be integrated on hardware to enhance the ability to discriminate particular signal topologies, for instance separating VBF $H \rightarrow b\bar{b}$ from the LHC multijet background. The new trigger architecture and hardware will also support the deployment of algorithms capable of analyzing data from a window of up to ± 3 BXs, enabling the study of inter-BX object correlations to trigger more efficiently on HSCP and LLP exotic topologies.

4.4. Phase-2 Level-1 trigger Data Scouting system

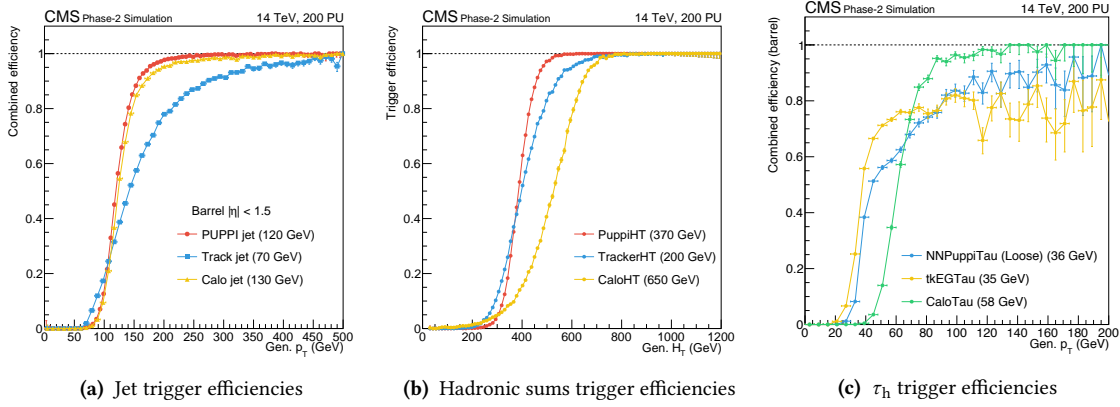


Figure 4.11 • Summary of CMS Phase-2 L1T efficiencies for jets (4.11a), hadronic transverse energy sum (4.11b) and τ_h (4.11c), compared for different algorithms [115].

The results from these algorithms will be sent to the Phase-2 Timing and Control Distribution System (TCDS-2). As with the Phase-1 system, the GT will allow dynamic configuration of prescale factors for specific algorithms, enabling only 1 in N trigger accept signals to be sent to the TCDS. It will also allow the application of a bunch mask to restrict the operation of the algorithms to a subset of bunch crossings and monitor the algorithm accept rate, both before and after accounting for dead time.

4.4 Phase-2 Level-1 trigger Data Scouting system

The output primitives of the CMS Phase-2 L1T will be collected at the full 40 MHz LHC bunch crossing rate by the L1DS system [115, 139–141]. The L1DS approach will enable the search for signatures that feature a too-large irreducible background and therefore currently rejected by the tight requirement on the L1T output rate budget, and signatures for which the background rejection requires algorithms that do not fit within L1T latency and hardware resource constraints. The system takes full advantage of the upgraded L1T, which produces objects with resolution often comparable to the current offline reconstruction. Additionally, the L1DS will provide a vast amount of data for monitoring purposes, such as detector diagnostic with real-time heatmaps of regional TPs, offering instantaneous indications of transient detector issues. The scouting approach will also benefit per-bunch luminosity measurements, simplifying the comparison with other experiments.

4.4.1 Baseline system and physics potential

The L1DS system represents a second, parallel readout chain processing the L1 intermediate data streams at the full LHC bunch crossing rate of 40 MHz. Trigger data will be extracted directly from the L1 processors using spare outputs from the boards, utilizing the same high-speed serial optical links used for the L1 interconnects and the same communication protocol. No handshake with the L1 control system is required, but the scouting system will receive Run Control commands and

will be fully integrated in the operational sequences of the experiment. If one or more downstream components cannot handle the input throughput, data will simply be lost. No attempt to modulate the input traffic, which would result in throttling the L1T, is made.

The planned architecture for the L1DS system is designed to be scalable and it is shown in Fig. 4.12. In a first baseline stage of the system (Stage-1), the decision outputs of the GT are captured by the scouting Decision System (sDS), and the output physics objects from the four global subsystems are captured by the scouting Global System (sGS). The Stage-1 architecture has relatively modest throughput requirements, providing essential trigger diagnostic capabilities for the GT and enabling new physics analyses. This architecture can later be extended including a scouting Local System (sLS) to capture local muon and regional barrel calorimeter trigger objects and the endcap calorimeter primitives. Additionally, a scouting Track System (sTS) would capture the output from the L1 track finder. The sLS and sTS would form a second stage (Stage-2) of the scouting system, which would have higher I/O and processing demands.

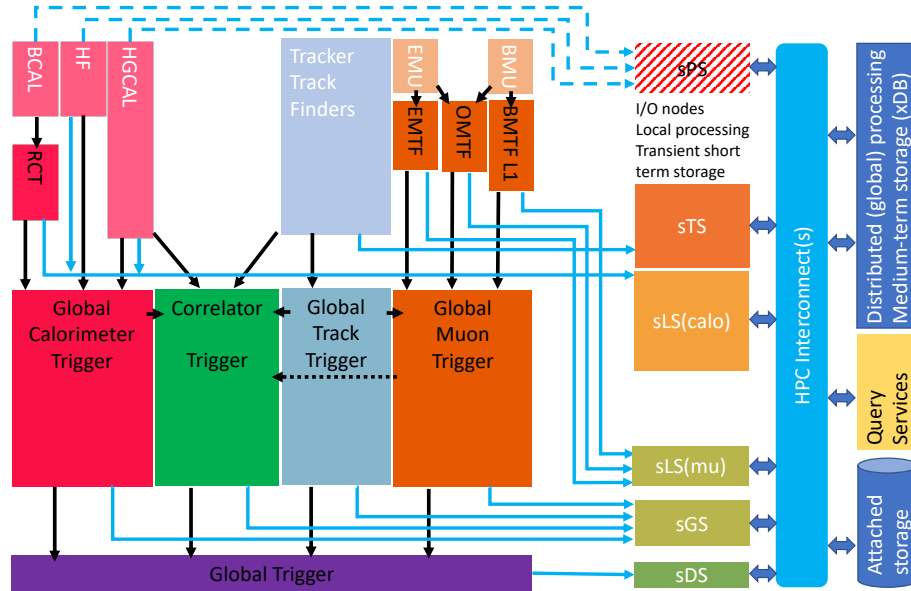


Figure 4.12 • Architecture of the CMS Phase-2 Level-1 trigger Data Scouting system [115, 139], with the Stage-1 scouting Decision System (sDS) and scouting Global System (sGS). A subsequent Stage-2 would include a scouting Local System (sLS) and Track System (sTS).

Examples of physics processes that could potentially benefit from a scouting-based analysis using the L1DS include:

- **Soft dilepton resonances** ($ee, \mu\mu, \tau\tau$), such as in DM models with a dark photon decaying to muon pairs [142], where removing p_T thresholds of the trigger selection could enhance the physics sensitivity at low invariant mass.
- **Dijet resonance searches at low invariant mass**, as this region of the phase-space is inaccessible to standard L1 because of higher jet E_T thresholds required to control the accept rate [112].
- **Multiple soft jet final states**, for example originating from hidden sectors with compressed mass spectra [143].

- **Soft unclustered energy patterns (SUEP)**, with high multiplicity unclustered hadronic final states [144].
- **Exclusive rare Higgs decay channels**, such as $H \rightarrow \rho\gamma, \phi\gamma, J/\psi\gamma$ and $H \rightarrow 2\rho, 2\phi, \phi J/\psi$. In the absence of tracks at the trigger level, the first class of these decays has been selected up to now using single isolated photon triggers. Due to the threshold on the photon energy, the trigger selection is highly inefficient for the signal, making these searches statistically limited. For the second class of decays, the L1 selection requires the reconstruction of two mesons, and the algorithm implementation in hardware might result in high FPGA logic resources utilization. Both classes of decays are extensively studied in this work in Chapter 6.
- **SM hadronic decays of the W**, such as the fully hadronic channel $W \rightarrow 3\pi$ and radiative decays $W \rightarrow D_s\gamma, \pi\gamma$, for which the same argument made in the previous point motivates the scouting approach.
- **B hadron physics involving τ particles**, such as $B_s \rightarrow \tau\tau$ with $\tau \rightarrow 3\pi + X$ [145], where high efficiency τ selection at low p_T is necessary, and can be achieved with the Stage-2 system using the L1 tracker tracks.
- **Flavor anomalies**, such as the $\tau \rightarrow 3\mu$ [146], where the acceptance of low p_T muons is key to improve the sensitivity.
- **Heavy Stable Charged Particles (HSCP)**, leaving a signature across multiple consecutive bunch crossings.

4.4.2 System architecture and requirements

The L1DS baseline architecture is illustrated in Fig. 4.13. The L1 objects will be collected via optical links from the data sources presented in the previous subsections and operating at around 25 Gb/s. The DAQ-800⁶ board, designed for the CMS Phase-2 central DAQ readout [147], has been selected to receive these L1 links. The DAQ-800 can accept up to 48 L1 input links via FireFly optical receivers [148] for a total input bandwidth of 1.2 Tbps. The nominal maximum theoretical output bandwidth is 1 Tbps via ten 100 Gb Ethernet (GbE) links using on-board Quad Small Form-factor Pluggables (QSFPs). To maintain a steady output data rate of 800 Gb/s, a moderate data reduction will be required. The data pre-processing logic will be implemented on the two Xilinx Ultrascale+ VU35P FPGAs mounted on the board and chosen for their built-in High-Bandwidth Memory (HBM). Each FPGA device contains 8 GB of HBM partitioned over two stacks, each divided into eight blocks. Each block contains two slots with 256-bit alignment, managed by the same memory controller. The HBM provides sufficient data buffering between the LHC-synchronous back-end and the commercial-off-the-shelf (COTS) switched network.

The L1DS will utilize a custom firmware implementation of the TCP/IP protocol [149, 150] on the DAQ-800 board FPGAs to output the pre-processed and zero-suppressed L1 objects. The choice of a reliable standard protocol enables a direct connection to the COTS network, which will convey data to commercial servers optimized for I/O, the Ingestion Units (IU). These will

⁶The “800” in DAQ-800 refers to the maximal average bandwidth of 800 Gb/s to the event builder network when used for the central DAQ [147].

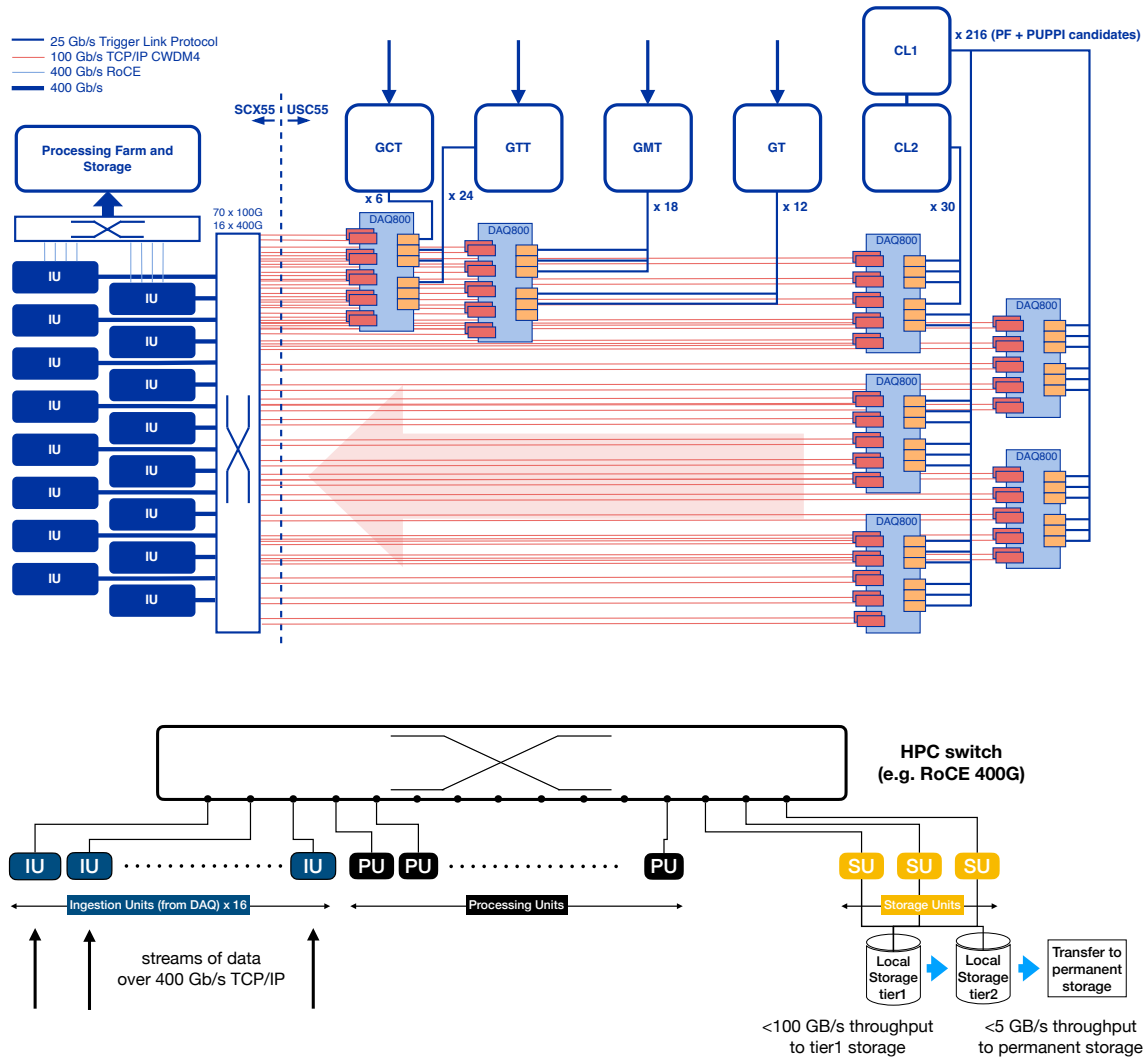


Figure 4.13 • Baseline architecture of the CMS Phase-2 Level-1 trigger Data Scouting system. Data from the trigger links will be collected and concentrated using the DAQ-800 data acquisition boards. The pre-processed L1 data from the DAQ-800 outputs are buffered and formatted in a set of Ingestion Units (IUs), connected with all the other nodes via High-Performance Computing (HPC) switched network. A set of Processing Units (PUs) performs event building and processing, running online analysis of a set of generic or exclusive signatures of interest. The output datasets are sent to permanent storage from a set of Storage Units (SUs).

4.4. Phase-2 Level-1 trigger Data Scouting system

serve as temporary data buffers and will format the data for ingestion into the event building in the Processing Units (PU). Each PU will run multiple analysis processes in parallel, focusing on tasks such as re-calibration, classification and combination of items. The use of columnar data formats and GPU or FPGA accelerators can dramatically speed-up this kind of operations, reducing the CPU workload.

The processed data from the PU analysis processes will then be aggregated prior to storage in cluster of Storage Units (SUs). The baseline system is expected to store data at a rate of up to 5 GB/s, which corresponds to 10% of the estimated central DAQ bandwidth.

Standard distributions of quantities from object combinations can be produced efficiently online, in a manner similar to the current Data Quality Monitoring (DQM) of the experiment, and predefined analyses can run in real-time, storing the full content of selected BXs, possibly with adjacent BXs included to study the correlations across consecutive collision events. Any remaining bandwidth may be filled by permanently storing a prescaled fraction of complete orbits, providing an unbiased dataset to develop new analysis techniques.

A demonstrator of a simplified version of this architecture has been developed during LHC Run-3. This demonstrator has been connected to a subset of the CMS Phase-1 L1T processors, making possible the evaluation of the demonstrator performance on real collision data. The Run-3 system, which is described in detail in the following chapter, provides the first demonstration of its applications for potential physics analyses using data from the L1T.

5

The Level-1 trigger Data Scouting demonstrator

To test and validate the concept of the L1DS, a demonstrator was assembled to capture intermediate objects reconstructed by the CMS Level-1 trigger during LHC Run-3, with the goal of collecting real collision data. The Run-3 L1DS demonstrator collects muon and calorimetric trigger objects, which constitute the inputs to the global trigger decision algorithms, along with the muon barrel superprimitives introduced in Section 4.2.2, which are the inputs to the KBMTF algorithm. Although this information represents only a small fraction of the detector data for each collision event, it is potentially sufficient to carry out analyses that are otherwise inaccessible using data collected by the standard trigger chain or HLT Data Scouting. The data collected by the demonstrator after its commissioning in 2024 is being actively studied for specific physics measurements.

After a brief introductory Section 5.1, this chapter provides a detailed description of the firmware and software components for trigger data readout and pre-processing, outlined in Section 5.2. Section 5.3 describes the software for real-time data processing, designed to enable online analyses in the L1DS system and produce datasets for further offline analysis. Section 5.4 presents the studies done for the validation of the demonstrator, with a focus on SM candles, such as $Z \rightarrow \ell\ell$. Section 5.5 discusses potential applications of machine learning techniques in the L1DS system to enhance the performance of the trigger objects data analysis. Section 5.6 summarizes the future prospects of the Run-3 demonstrator system and its potential application for physics measurements.

5.1 Demonstrator overview

The architecture of the demonstrator is illustrated in Fig. 5.1. The system receives copies of the following trigger objects via spare 10 Gb/s optical output links from the L1 trigger processors:

- Up to 8 muon candidates from the μ GMT MP7 processor, sent via 4 links.
- Up to 12 jet, 12 e/γ , 12 τ_h candidates and energy sums from the calorimeter DeMux board, sent via 7 links.

- A subset of the input super-primitives sent by the TwinMux to the 12 BMTF processors. Each BMTF board selects the best muon stubs in the central sector it covers. These stubs are ranked by highest quality and lowest bending, then duplicated and sent to the scouting system via 24 links, 2 from each BMTF processor.
- The final decision bits returned by the algorithms running on the μ GT. Each μ GT processor sends 3 links to the scouting system, totaling 18 links.

A total of 43 input links provides an input throughput of approximately 430 Gb/s to concentrate and pre-process. While additional trigger processors can be connected to the scouting system, the selected sources represent a balance between physics potential, data storage optimization and availability of links and resources to feed them to the scouting system.

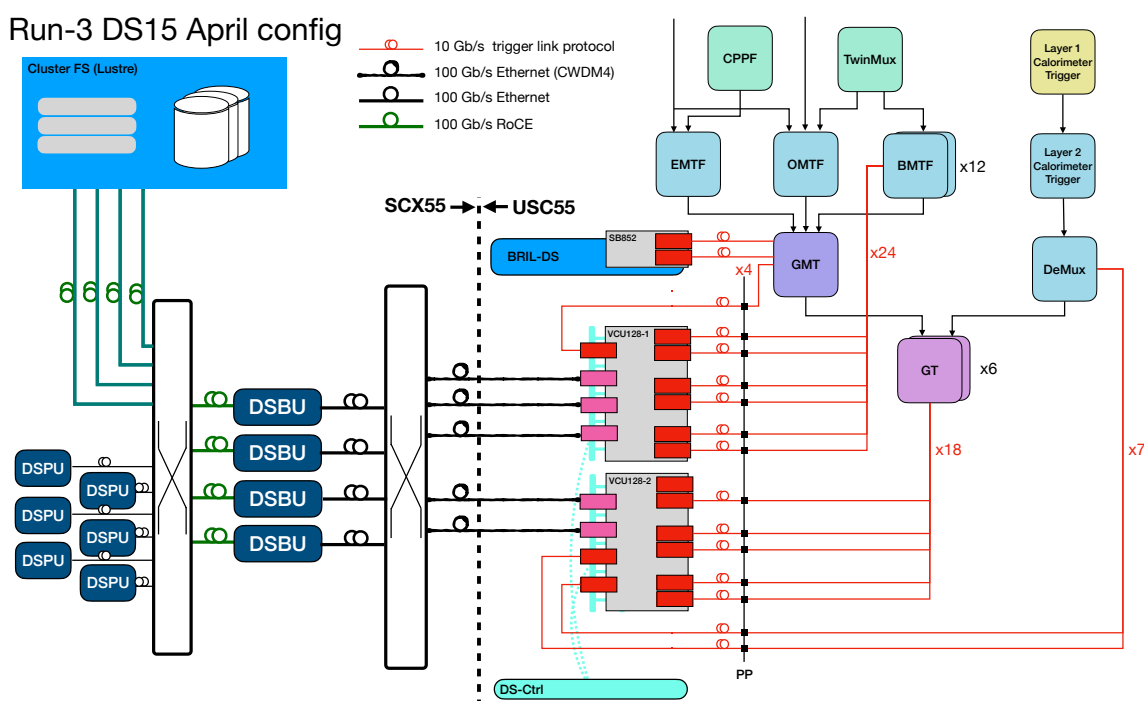


Figure 5.1 • Schematics of the architecture of the L1DS Run-3 demonstrator [114]. It captures the trigger objects from the Global Muon Trigger (μ GMT) and the Calorimeter Trigger (DeMux), the Global Trigger (μ GT) decision output and the input muon stubs to the Barrel Muon Track Finder (BMTF). After basic zero-suppression, the received data is propagated via Ethernet to compute nodes (DSBUs and DSPUs) for data buffering and online processing, and subsequently sent to long-term storage.

The readout of the trigger links is performed with a pair of commercial FPGA boards, consisting of two Xilinx VCU128 development kits. The connection is unidirectional, with no back-pressure to the trigger, making it impossible by design for the scouting system to interfere with the standard trigger operation. The scouting boards concentrate the trigger links and apply basic pre-processing on the trigger data, such as zero-suppression. Room for more sophisticated algorithms besides data reformatting and reduction is available in the FPGA logic. Recalibration and track finding algorithms using an ML-based implementation, explored during the development of the demonstrator, are discussed in Section 5.5.

The pre-processed data is buffered in the integrated HBM of the FPGA before being sent to a set of computing nodes in the surface data center via 100 GbE optical links. The first stage of computing nodes comprises the Data Scouting Buffer Units (DSBU), where a data acquisition software receives the ethernet streams sent by the boards and formats the received raw data in fragments of events. These fragments are aggregated in a second block of the demonstrator responsible for the online processing. The Data Scouting Processing Units (DSPU) run event-based processing and online selection under the CMS software framework (CMSSW). Several data streams are produced in this step and stored in a Lustre filesystem by a set of Storage Units (SUs). The data is finally sent to the Tier-0 permanent storage of the CMS experiment. Here it is repacked in an Analysis Object Data (AOD) format suitable to be distributed through the CMS Data Aggregation System (CMS-DAS) and made available to users for analysis.

The stored L1 muon and calorimetric objects are reconstructed with the methods described in Section 4.2.1 and Section 4.2.2 and can be used for physics studies after recalibration. On the other hand, the BMTF stubs are lower-level objects that require the application of pattern recognition and track fitting techniques before being usable for analysis purposes. A custom reconstruction without latency and logic resource constraints can be implemented to improve the p_T and transverse impact parameter resolution.

Scouting the μ GT decision bits is redundant, as most of the μ GT algorithms can be emulated using the collected muon and calorimetric objects. Therefore, this scouting source is not enabled for continuous operation in the production system for storage optimization reasons. However, since the decision bits can be useful for trigger diagnostics under special running conditions, the data acquisition from this source can be temporarily enabled when necessary.

5.2 Readout of the trigger links

This section provides an in-depth description of the first block of the scouting demonstrator, which captures and concentrates the data from the trigger sources at the LHC bunch crossing rate of 40 MHz. The hardware to receive the trigger links, the implemented firmware running on it and the control and data acquisition software are introduced. The main task of this block is to aggregate and align data from the L1 trigger links and organize it in event fragments. The components in this block are orchestrated by the scouting Function Manager (FM), part of the CMS Run Control. The production system can be configured, started or stopped together with the rest of CMS. Issues in these operations can be automatically handled with dedicated recovery procedures.

5.2.1 Input data and formats

μ GMT muons The muon objects provide the p_T , assuming a primary vertex constraint in the muon track fitting and without this constraint (p_T^{unc}), the charge sign assignment and its validity, the η_{st2} and φ_{st2} coordinates measured at the second muon station, and the coordinates η_{ext} and φ_{ext} extrapolated at the collision vertex. In addition, the index of the TF originating the candidate and a truncated version of the impact parameter, assigned by the subsystem from which the muon comes from, are included. A quality code encodes the information of the stations used in the reconstruction. The complete format, with variables, scales and definitions, is summarized in Table 5.1.

Parameter	Bits	Range	Step	Notes
p_T	9	0...256 GeV	0.5 GeV	p_T assignment with vertex constraint
p_T^{unc}	8	0...256 GeV	1.0 GeV	p_T assignment without vertex constraint
φ_{st2}	10	0... 2π rad	~ 0.011 rad	φ at second muon station
η_{st2}	1+8	-2.45...2.45	0.010875	Pseudorapidity at the second muon station
φ_{ext}	10	0... 2π rad	~ 0.011 rad	φ extrapolated at the vertex
η_{ext}	1+8	-2.45...2.45	0.010875	Pseudorapidity extrapolated at the vertex
Quality	4	0...15		Muon quality code
Charge sign	1	sign		Charge of the muon (0: μ^+ ; 1: μ^-)
Charge valid	1	bool		Validity of charge assignment
Index	7	0...107		Track finder index idx $\in [36, 70]$: barrel idx $\in [17, 35] \cup [71, 89]$: overlap idx $\in [0, 16] \cup [90, 107]$: endcap
HS trigger	1	bool		Hadronic Shower trigger
d_{xy}	2	0...3		Impact parameter

Table 5.1 • Features, scales and definitions of the μ GMT muon primitives transmitted from μ GMT to the scouting system.

Calorimeter jets, e/γ , τ_h The calorimeter objects contain basic kinematic features, such as E_T , η , φ . For the e/γ and τ_h an isolation variable indicates if the reconstructed object is isolated from other hadronic activity and pileup. The jets have an additional bit encoding the displacement information, as described in Section 4.2.1. The features, scales and definitions of these objects are summarized in Table 5.2.

Calorimeter energy sums The energy sums are provided in the transverse plane, with only their E_T and φ values. The most important sums collected by the demonstrator are the total transverse energy collected by the calorimeter (E_T), the total E_T in the ECAL (E_T^{em}), the total transverse energy from jets (H_T), the missing transverse energy (E_T^{miss}), and the missing transverse energy calculated only from jets (H_T^{miss}). Another important available quantity is the number of firing trigger towers (n_{TT}), which is a measure of the amount of pileup. The features, scales and definitions of the energy sums are summarized in Table 5.3

BMTF input stubs The muon barrel super-primitives are provided with their station and wheel, while the sector number is implicit in the BMTF board from which the object is sent. The stub position is given in local φ and η coordinates relative to the sector and wheel. In particular, the η coordinate is encoded as a 7-bit pattern showing which of 7 areas in the chamber along the η direction has at least 3 SL hits, as illustrated in the example sketch in Fig. 5.2. An η quality pattern is also available in the stub format, with the i^{th} bit set to “1” in case of 4 SL hits.

The relative bending of the stub is a measure of the transverse momentum of the candidate muon track. A hardware quality code is encoded in a 3-bit feature, where “111₂” is unused, “000₂”, “001₂”, “010₂” indicate an uncorrelated DT measurement, namely uncorrelated tracks between inner and outer SL, and “011₂”, “100₂”, “101₂” and “110₂” indicate a correlated DT measurement. The features, scales and definitions of the stubs are summarized in Table 5.4.

5.2. Readout of the trigger links

Object	Parameter	Bits	Range	Step	Notes
Jet	E_T	11	0...1024 GeV	0.5 GeV	
	η	7+1	-5...5	0.0435	
	φ	8	0... 2π rad	0.0435 rad	
	Displacement	1			Delayed jet flag
e/γ	E_T	9	0...256 GeV	0.5 GeV	
	η	7+1	-5...5	0.0435	
	φ	8	0... 2π rad	0.0435 rad	
	Isolation	2	0...3		"01 ₂ " or "11 ₂ " for isolated e/γ
Tau	E_T	9	0...256 GeV	0.5 GeV	
	η	7+1	-5...5	0.0435	
	φ	8	0... 2π rad	0.0435 rad	
	Isolation	2	0...3		"01 ₂ " or "11 ₂ " for isolated τ_h

Table 5.2 • Features, scales and definitions of the calorimeter objects transmitted from the DeMux board to the scouting system.

Object	Parameter	Bits	Range	Step	Notes
ET	E_T	12	0...2048 GeV	0.5 GeV	Total E_T
	E_T^{em}	12	0...2048 GeV	0.5 GeV	Total E_T in ECAL
HT	E_T	12	0...2048 GeV	0.5 GeV	Total hadronic E_T
	Tower count	13	0...8191		Sum of firing trigger towers
ET_{miss}	E_T	12	0...2048 GeV	0.5 GeV	Total missing E_T
	φ	8	0... 2π rad		φ direction of missing E_T
HT_{miss}	E_T	12	0...2048 GeV	0.5 GeV	Total hadronic missing E_T
	φ	8	0... 2π rad		φ direction of hadronic missing E_T

Table 5.3 • Scale and definitions of the calorimeter global sums transmitted from the DeMux board to the scouting system. Only a selected subset of all the sums is reported.




Figure 5.2 • Pattern matching scheme in η . Each η chamber is segmented in 7 areas. The information coming from the TwinMux is delivered as a 14-bit pattern, containing one hit-bit and one quality-bit for each of the seven adjacent chamber areas. A muon must leave a trace in at least 3 out of the 4 planes of the η SL, for a hit-bit to be set to "1". Image adapted from Ref. [122].

Parameter	Bits	Range	Notes
Valid	1	flag	Keep or Zero Suppress stub
φ	12	-2048...2047	Relative position of a segment inside a sector
φ_b	10	-512...511	Bending angle of stub
Quality	3	0...7	Encodes the number of SLs used
η hits	7	pattern	Each bit corresponds to one chamber area 0: no hit (less than 3 SLs hit) 1: hit (3 or more SLs hit)
η quality	7	pattern	Each bit corresponds to one chamber area 0: 3 SLs hit 1: 4 SLs hit
Station	2	0...3	Station of stub
Wheel	3	-2...2	Wheel of stub

Table 5.4 • Scale and definitions of the super primitive parameters transmitted from BMTF to the scouting system.

5.2.2 Trigger data pre-processing on FPGA

The trigger readout and pre-processing are handled by FPGA devices on commercial boards. The AMD Xilinx VCU128 development kit [151], equipped with a Virtex Ultrascale+ VU37P FPGA, was chosen for its similarity to one of the two units of the DAQ-800 board and it offers a realistic test of the Phase-2 design, albeit with reduced input bandwidth.

The board is equipped with 4 QSFP transceivers on its Printed Circuit Board (PCB), with 6 additional slots available through a *HT-Global* mezzanine card [152], shown in Fig. 5.3a along with the VCU128 board. A QSFP slot can be used to receive up to 4 trigger links at either 10 (Run-3) or 25 (Phase-2) Gb/s, or for a single 100 GbE output. The FPGA chip has also 8 GB of integrated HBM, providing large buffers before the output stage.

The VCU128 boards are housed in a PCIe crate from *One Stop Systems* [153], shown in Fig. 5.3b. The PCIe bus of the crate is connected to a Dell R320 server, used to monitor and configure the boards. The output QSFPs of the VCU128 boards in the experiment service cavern are connected to the surface over 200 meters of single mode optical fibers. The other end of the link is connected to a network switch, which connects all the compute nodes of the demonstrator on the surface.

The firmware for the VCU128 boards has been designed to be modular and flexible to allow receiving different and multiple trigger data sources on a single board. The design is divided into 3 main blocks:

- An input block receiving the trigger links and scalable to up to 40 input links per board.
- A processing pipeline specific for each scouting source, in charge of aligning the input streams, applying basic pre-processing such as data formatting and reduction or machine learning algorithms.
- An output stage buffering the concentrated data into the HBM, before sending it via TCP/IP over 100 GbE links using dedicated IP cores developed in the CMS DAQ group [149, 150].

The details of each step are described in the following paragraphs.




Figure 5.3 • AMD Xilinx VCU128 development kit used for the L1 trigger processors readout and HT-Global mezzanine (5.3a); PCIe crate housing the VCU128 boards in the service cavern of the experiment (5.3b).

Input stage The Phase-1 L1 trigger processors send data to the scouting system via 10 Gb/s links. The protocol works with 32-bit words at 250 MHz word frequency, with additional control bits carrying the information about special words. The BX data is a fixed record of 6×32 -bit consecutive data words per input link. To match the protocol frequency, 10 MHz of “filler” words are added and they are signaled by the control bits.

Data words are replaced by special “comma” words in a gap of few BXs near the end of the orbit, needed for the alignment of the links and to signal the start and end of a packet. Two special words are added before the start of the gap: the link identification (LID) and the Circular Redundancy Check (CRC). The former identifies the provenance of the trigger link, the latter is compared with a CRC recomputed on the received data to check the integrity of the packet. In case a mismatch is found, the packet is flagged and a mismatch counter is updated. The orbit header at the start of each packet carries the information of the number of the next LHC orbit.

The input stage of the scouting firmware is illustrated in Fig. 5.4. It extracts the data and special words from the received packets and format the data stream to simplify the subsequent steps. Each input stream operates on its own clock, so they need to be synchronized to a common clock before moving on to the next stages. This is done by writing each stream into an intermediate buffer (FIFO), which is then read for all streams using the common 250 MHz clock at which the HBM is programmed to operate. The synchronized streams are then sent to the next processing stage for aggregation and further pre-processing.

Processing stage The design of the processing pipeline is illustrated in Fig. 5.5. The first step in the chain is the alignment of the HBM clock-synchronized streams to the same collision event. This alignment is based on the fixed size of the BX record. Each data stream is written into

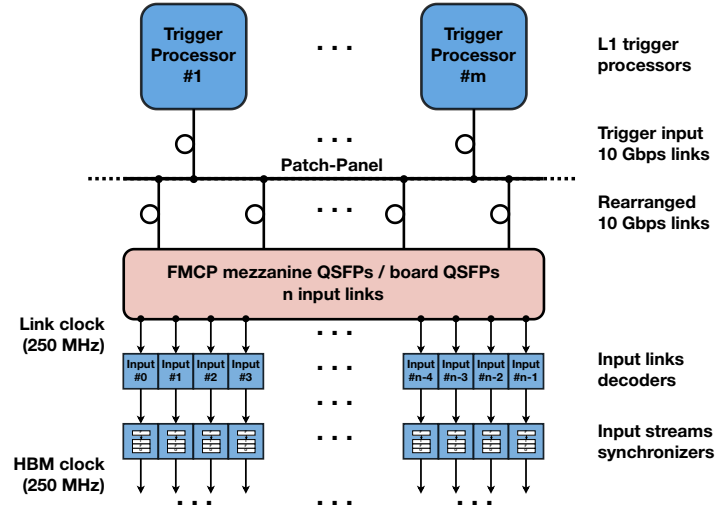


Figure 5.4 • Schematics of scouting firmware design for input stage. The received trigger links are decoded and synchronized to the common HBM clock.

a separate FIFO buffer. A Finite State Machine (FSM) controls the reading from these buffers, ensuring that at least a full BX record is available in each before reading, and assigns the BX number.

After alignment, the output streams are merged into groups, each corresponding to a TCP/IP stream in the output logic. The data groups are then processed by a zero-suppression module that works differently depending on the type of trigger object:

- For L1 muons, it checks that at least one of the 8 possible candidates has p_T greater than zero, indicating the presence of at least one muon in the BX.
- For jets, e/γ , and τ_h objects the same approach is applied, checking if any of the 12 possible objects has E_T greater than zero.
- Calorimeter sums are global quantities, so their record is kept only if there is at least one jet or e/γ in the same BX record.
- For BMTF stubs, the record of a group is kept only if there is at least one valid stub, based on the valid bit field in their data words, as described in Table 5.4.
- For the μ GT decision bits, the data reduction algorithm works differently. It uses a bitwise OR operation across the 512-bit fields from the six μ GT boards. The result determines whether the BX record is kept. A configurable 512-bit field can be loaded via PCIe at runtime to mask algorithms that are not of interest.

The grouped streams are then padded to accommodate the 256-bit alignment of the HBM buffers. A trailer with a 3564-bit field along 16×256 -bit frames is prepared to encode information about the BXs that were suppressed by the data reduction algorithm, allowing the data acquisition software to efficiently unpack the events and their BX numbers.

The final stage of the processing pipeline is designed to run algorithms on the trigger objects. Examples include neural network inference for L1 muon recalibration [154] and BMTF stub tracking, which will be discussed in Section 5.5. Deploying these applications directly on the FPGA is possible using the hls4ml package [155, 156], which translates higher-level neural network models into HLS code that can be converted to hardware description language using the Vivado suite [155, 156]. After applying the algorithm, the group record is converted to a compact format to reduce the number of needed 256-bit frames, optimizing therefore the output throughput.

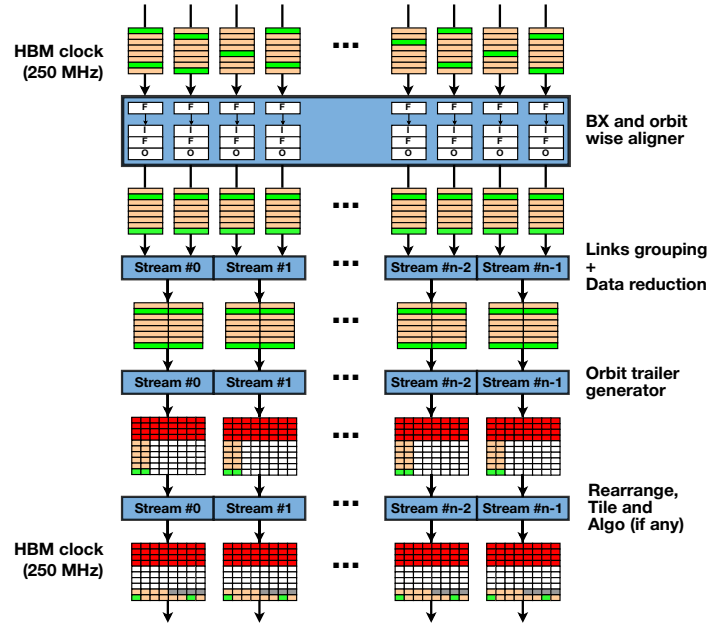


Figure 5.5 • Schematics of firmware design for FPGA processing pipeline. The collision record fragments across the trigger data streams are aligned, aggregated into groups and zero-suppressed to reduce the final output throughput.

Output stage The final stage is illustrated in Fig. 5.6a. The aggregated and pre-processed trigger data is organized into “scouting blocks”, with a header and trailer encapsulating a programmable number of orbits. The header contains information about the block size for efficient unpacking in the DAQ software, along with other useful metrics.

Before writing the scouting data blocks to the HBM, a 256-bit frame is reserved for the block header. After the payload and trailer have been written, the necessary counters to place in the header are finalized, and the HBM control logic places the frame in the reserved space. The orbits within the payload are treated as fragments of the block and are further encapsulated with a fragment header and the orbit trailer generated in the previous processing stage.

Each group stream has its own logic for preparing scouting blocks. The logic for each instance is connected to the HBM controller, which monitors the occupancy of the corresponding HBM buffer. If the occupancy in a HBM slot crosses a threshold of about 95%, equal to around 240 MB, a back-pressure signal is sent to the group stream pipeline, causing orbits to be dropped until the HBM buffer occupancy falls below approximately 80%. The information about the dropped orbits is recorded by creating an empty fragment with just a flagged header. The number of dropped

orbits is saved in a register that can be accessed via PCIe. This approach ensures efficient and reliable event aggregation in the online processing step, described in Section 5.3.2, where all the sources will unpack the same number of orbits.

At the other end of the HBM buffers, a FSM controls reading from the buffer and manages the TCP/IP logic. Multiple TCP/IP streams are multiplexed into a single 100 GbE core, which transmits the packets to a configurable destination. When an Ethernet packet is correctly received by the DAQ software, described in Section 5.2.4, an acknowledgement is sent to the 100 GbE core and the corresponding HBM buffer pointers are updated. If the acknowledgement is not received within a timeout window, the packet is retransmitted.

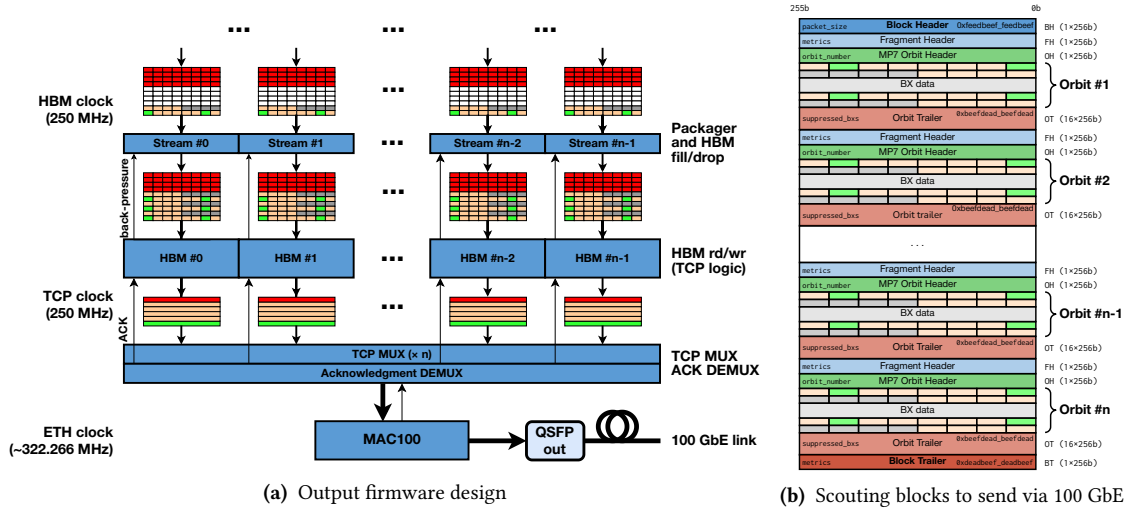


Figure 5.6 • Schematics of FPGA firmware design for scouting block packets preparation, back-pressure handling and 100 GbE output via TCP/IP (5.6a). Format of the scouting blocks sent via TCP/IP to the scouting DAQ software (5.6b).

Firmware and resource usage The amount of allocated input/output resources, HBM blocks and DAQ units is designed to be configurable, such that these numbers can be easily tuned for the readout of specific trigger sources. Two main firmware design combinations have been developed for the 2024 L1DS demonstrator.

The first design, illustrated in Fig. 5.7a, processes 4 input links from μ GMT and 24 links from BMTF, concentrating them into two separate processing pipelines. The μ GMT pipeline writes all muons that pass the zero-suppression filter to a single HBM buffer, resulting in a single TCP/IP stream sent via a DAQ unit. The BMTF pipeline aggregates the input data into 12 streams, each corresponding to a BMTF wedge. These aggregated streams are written to 12 HBM buffers, which then become 12 TCP/IP streams on the read side of the HBM. The streams are multiplexed into 2 DAQ units, with 6 streams per core.

The second design, shown in Fig. 5.7b, processes 7 links from the calorimeter DeMux board and 18 links from the μ GT. The calorimeter scouting pipeline aggregates the input data based on the type of object, writing jets, e/γ , τ_h , and energy sums to four different HBM buffers. The resulting

four TCP/IP streams are then multiplexed into a single DAQ unit, producing the highest output throughput among all the designs. Due to the high data reduction factor in the μ GT pipeline, only a single HBM buffer and DAQ unit are required.

The resource usage of these two main designs is reported in Table 5.5, with the floorplans on the VU37P chip shown in Figs. 5.7d-5.7d. Overall, resource usage is below 30%, with no significant congestion of logic resources. Both designs achieve timing closure without requiring aggressive optimization of the logic resource routing. Since neural network inference or reconstruction algorithm modules are not included in the baseline, the usage of Digital Signal Processors (DSPs) is negligible. Given that the DAQ-800 board for the Phase-2 L1DS readout is equipped with VU35P FPGA devices, the resource usage numbers are projected to the available resources of the VU35P. Although this chip has fewer resources compared to the VU37P, the resource usage is generally low enough to allow the allocation of more HBM buffers, TCP/IP stream logic, and DAQ units.

Design	Resource	VU37P			VU35P extrapolation	
		Available	Utilization	%	Available	%
μ GMT + BMTF	LUT	1303680	227126	17.4	871680	26.1
	FF	2607360	408834	15.7	1743360	23.5
	BRAM	4032	1168	29.0	2688	43.5
	URAM	960	52	5.4	640	8.1
	DSP	9024	6	0.07	5952	0.1
μ GT + DeMux	LUT	1303680	151516	11.6	871680	17.4
	FF	2607360	286248	11.0	1743360	16.4
	BRAM	4032	713	17.7	2688	26.5
	URAM	960	20	2.1	640	3.1
	DSP	9024	4	0.04	5952	0.07

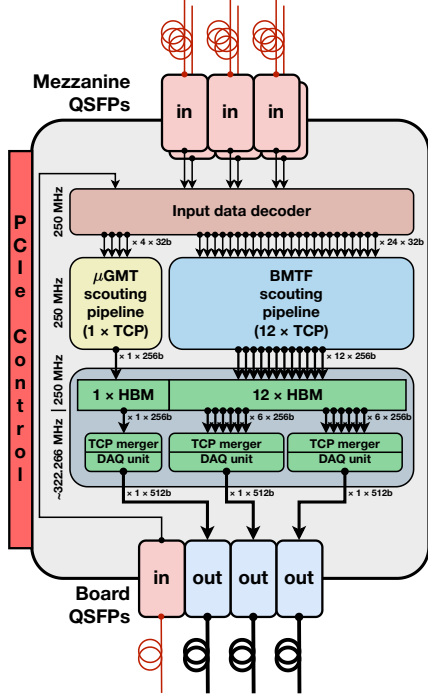
Table 5.5 • Target firmware resource utilization for a VU37P FPGA and extrapolation to a VU35P chip.

5.2.3 Scouting boards control and monitoring

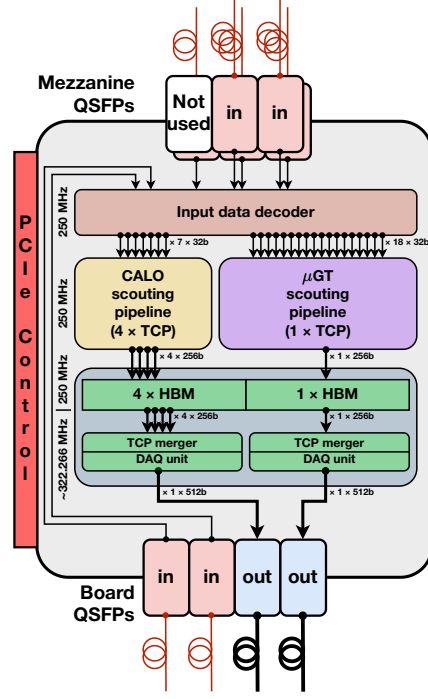
The scouting boards must be configured prior to start of the data acquisition, and their health and functionality must be continuously monitored to ensure the correct operation of the trigger readout and to detect any potential issues. This process involves accessing internal registers of the logic design, such as reading metrics or writing to registers to control specific functions. The read/write operations are managed using a DMA core in the firmware, interfaced via PCIe with the Dell R320 control server.

Configuring the scouting FPGA boards, particularly the 100 GbE cores and the TCP/IP logic, requires a complex series of register accesses, which are handled through higher-level sequences known as “actions”. The Scouting Configuration Endpoint (SCONE) was implemented to manage these operations through a user-friendly interface. The scouting FM in the CMS Run Control system interacts with this interface to set the boards in the correct state during transitions in the global run status.


The SCONE software architecture is shown in Fig. 5.8. The interface is implemented as a REST API, using the Flask framework. Its primary function is to handle board control and monitoring through web requests. When a web request is made by the scouting FM, it is queued for execution




(a) Firmware design for μ GMT+BTMF scouting



(b) Firmware design for DeMux+ μ GT scouting



(c) μ GMT+BTMF scouting firmware floorplan



(d) DeMux+ μ GT scouting firmware floorplan

Figure 5.7 • Main firmware designs for μ GMT muons and BTMF input stubs scouting (5.7a) and for calorimeter objects and μ GT decision bits scouting (5.7b). Floorplan of the design on the VU37P FPGA for “ μ GMT+BTMF” (5.7c) and “DeMux+ μ GT” (5.7d) designs.

on the targeted board. The list of available boards, actions, readable and writable registers, and firmware parameters are stored in JSON files, loaded onto the control server during FPGA firmware deployment. At startup, SCONE loads a module that implements the requested actions as a series of read/write operations. The results of these actions, along with the status of the board after execution, are placed in a response queue and returned to the scouting FM. If needed, recovery actions can be triggered to restore the system to a functional state.

There are four main actions available for board control by the scouting FM in the production system:

- **RESET:** it resets all the main modules of the boards, including the input decoder logic, the scouting processing pipelines, the HBM controller, and the output logic.
- **CONFIGURE:** it resets all counters in the firmware, disables the injection of trigger data into the scouting processing pipeline, configures the TCP/IP parameters and the destination of the output streams, and establishes the communication between the 100 GbE cores and the destination DSBU.
- **START:** it opens the connection with the DSBU and enables the injection of data into the scouting processing pipelines.
- **STOP:** it masks the processing pipelines, flushes any remaining data in the HBM, and closes the connection.

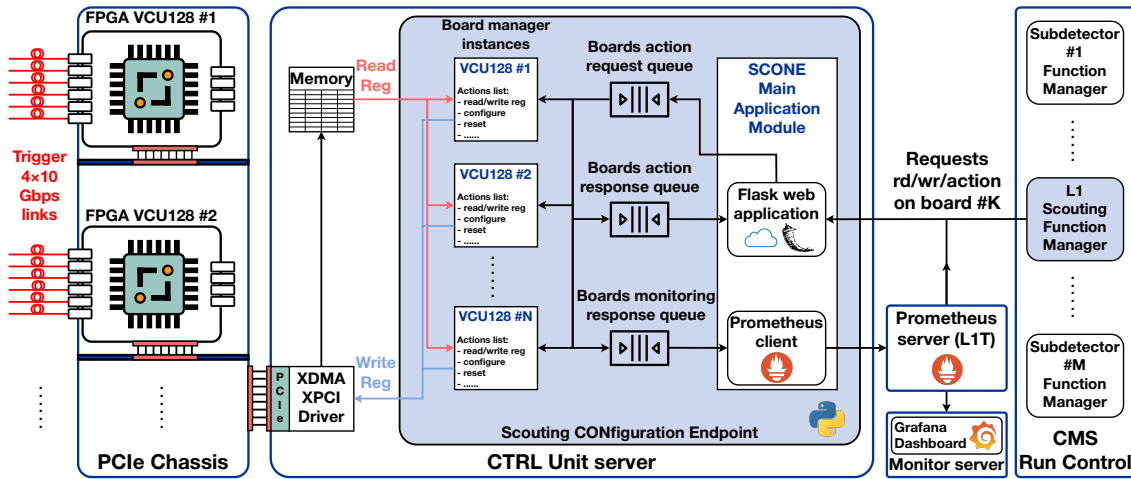


Figure 5.8 • Scouting Configuration Endpoint (SCONE) software for control and monitoring of scouting FPGA boards. The web application is based on the Flask framework. The monitoring part is based on Prometheus, which scrapes a set of exposed registers of the boards and provides the retrieved data to a Grafana dashboard.

Monitoring requests for specific board metrics are handled separately: a Prometheus server, part of the Level-1 trigger system, periodically sends monitoring requests to a client process within SCONE. This client retrieves relevant metrics, exposes them to Prometheus, and sends them to the server, which tracks the current status of the system. These metrics are displayed on a Grafana dashboard, as shown in Fig. 5.9, providing a user-friendly view of the health of the system for easier debugging.

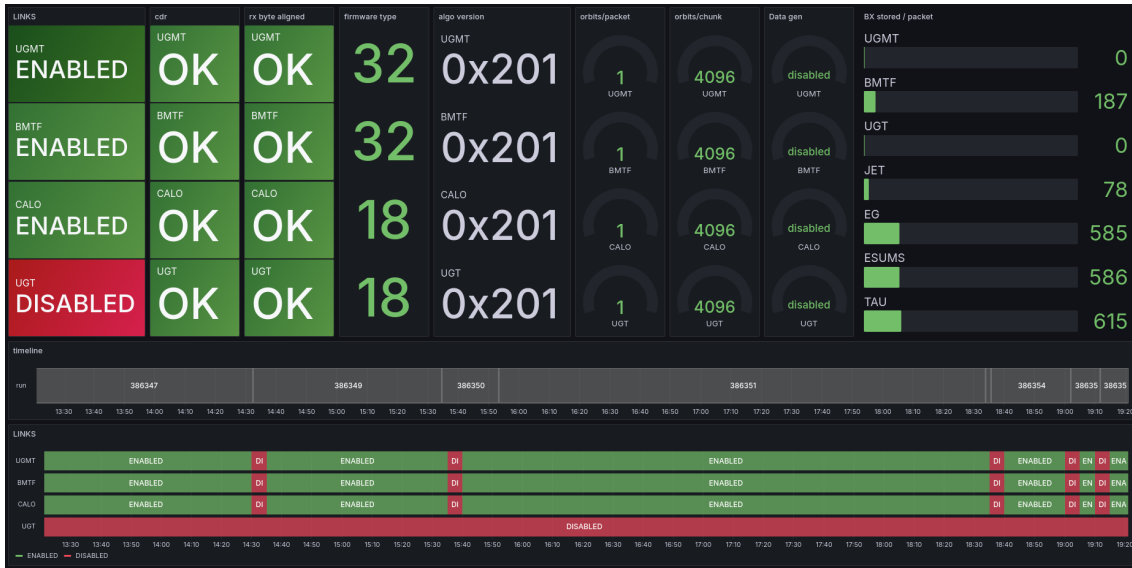


Figure 5.9 • Grafana monitoring dashboard for Run-3 Level-1 trigger Data Scouting demonstrator, reporting the current status of the L1DS readout block.

5.2.4 Data acquisition software

The output TCP/IP streams from the scouting boards are sent to the surface network switch, which redirects them to the DSBU. These are Dell R7515 servers with an AMD EPYC 7502P 32-Core CPU, optimized for data buffering using large RAM disks. A scouting data acquisition software (SCDAQ) runs on the DSBU to receive and further process the data streams. The software is built using Intel Threading Building Blocks (TBB) and follows a “pipeline” design pattern. An overview of the software is shown in Fig. 5.10.

SCDAQ creates a pipelines for every stream it receives on a single DSBU. These pipelines consist of sequential data filters. The first, “input” filter waits for the connection to be established with the board, which is managed by the scouting control application SCONE. At the start of data taking, the data is read from the TCP/IP streams, and the first expected frame is the header of the first scouting block, as shown in Fig. 5.6b. A 256-bit read is performed to decode the packet size from the received header frame, followed by a second read of size indicated, to receive the block payload and trailer. The input throughput of SCDAQ during normal LHC operations is shown in Fig. 5.11a, with the highest one observed from the streams containing the calorimeter objects.

The received packets are enqueued in memory slices and passed to a processor filter. The latter decodes the fragments in the blocks to unpack the raw data. For each fragment, the first step is to retrieve the orbit number from the orbit header and the information of the BX records that were retained from the orbit trailer. Each BX record is written to a memory slice in a compact format, with header words encoding the information of the orbit and BX number, along with the number of objects in the record. Only valid objects in the BX data are kept, optimizing the memory usage for storing the output slices.

The output slice is closed after it reaches a configurable number of orbits, which are then queued for writing to a file by the third, “output” filter. In the 2024 system, a number of 2^{12} orbits per output file was chosen, to create reasonably small data fragments for aggregation in the next

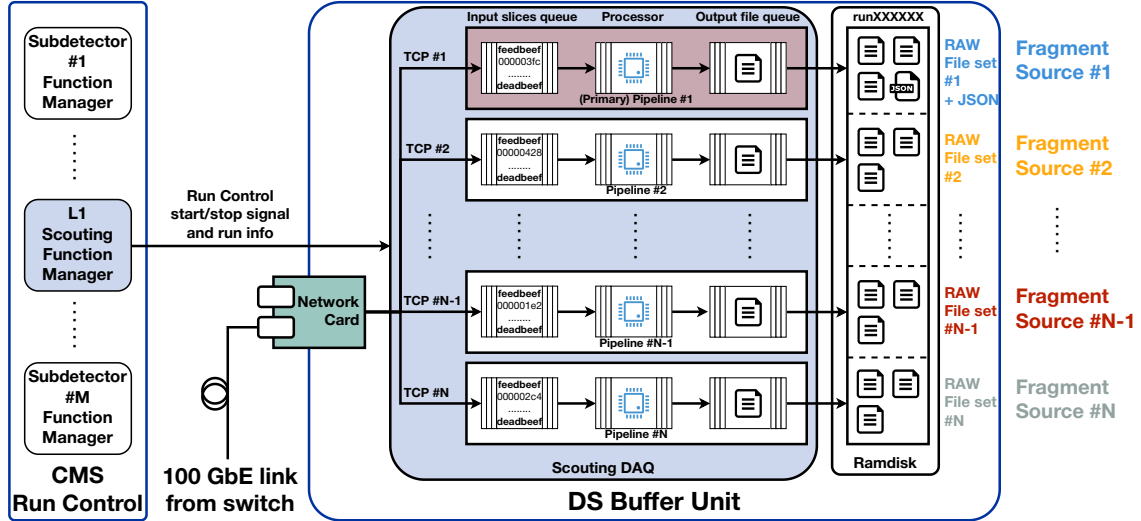


Figure 5.10 • Scouting data acquisition software (SCDAQ), based on Intel-TBB. The architecture exploits a set of pipeline pattern instances, each of them with an input receiver, stream processor and output raw data producer.

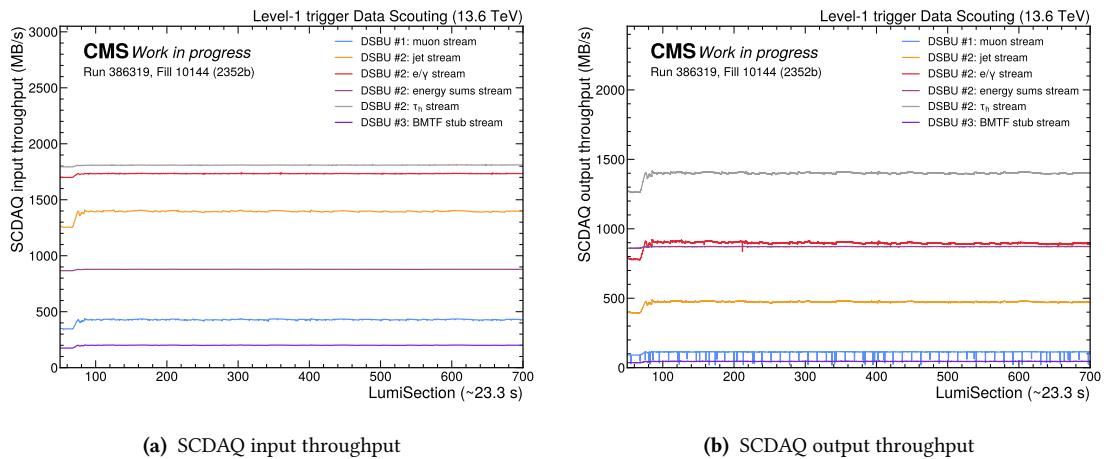


Figure 5.11 • Raw data throughput at SCDAQ input (5.11a) and output (5.11b), observed during normal LHC operations, with a fill containing 2352 colliding bunches at CMS.

step of online processing. These fragments are written to the RAM disk on the DSBUS where the SCDAQ instance is running. The output throughput of SCDAQ during normal LHC operations is shown in Fig. 5.11b, with the highest one observed from the streams containing the calorimeter objects.

The total number of events and the size of the output files are updated with each received orbit. When the output module receives the last orbit in a section of luminosity (LumiSection, LS), which contains 2^{18} orbits, the counters integrated over the LS are written to an “End-of-LS” (EoLS) metadata file. This metadata is needed by the online processing infrastructure to know how many orbits are available from a data source. Similar counters are integrated over the entire data taking run and written to an “End-of-Run” (EoR) metadata file to ensure the correct closure of the data-taking operations after the run ends. Only the first and primary pipeline writes these metadata.

5.3 Online processing at the LHC orbit rate

The second block of the scouting demonstrator receives the data fragments produced by the first block and performs event aggregation and online analysis. The software components of this block are integrated into the CMS offline software framework (CMSSW) and are described below. Also in this case the infrastructure is orchestrated by the scouting FM, which configures the processing nodes on the surface and triggers the release of the allocated resources when the data taking ends.

5.3.1 Event definition and Orbit Collection data format

A single BX record collected by the L1DS demonstrator during Run-3 has a size on the order of $\mathcal{O}(1)$ kB, which is approximately 10^3 times smaller than the average event size collected by the central DAQ. Defining the event as the BX record would result in suboptimal readout at 40 MHz, inefficient event storage, and more challenging processing for exploring correlations between consecutive BXs. To address these issues, a different event definition was chosen for L1 scouting: the event corresponds to the full set of BXs in an LHC orbit. By defining an event as an LHC orbit, rather than a single collision event, the processing is performed at a stable LHC orbit rate of 11.2 kHz.

To efficiently store the collected trigger objects under this event definition, a dedicated data structure, known as the scouting “Orbit Collection”, was developed. This structure stores the trigger objects in an orbit in a time-ordered flat vector, where the elements are dedicated data structures containing the trigger object features. The mapping of these objects to their respective BX record is encoded in an additional index vector belonging to the orbit collection. The i^{th} and $(i + 1)^{\text{th}}$ elements of this vector contain the information about the start and end positions, respectively, of the i^{th} BX record in the data vector.

A schematic view of the orbit collection and of its working principles is shown in Fig. 5.12. To simplify the extraction of BX records from an event, “getter” methods are available in the orbit collection class. These methods allow for the retrieval of a specific BX record or a single object instance at a given position.

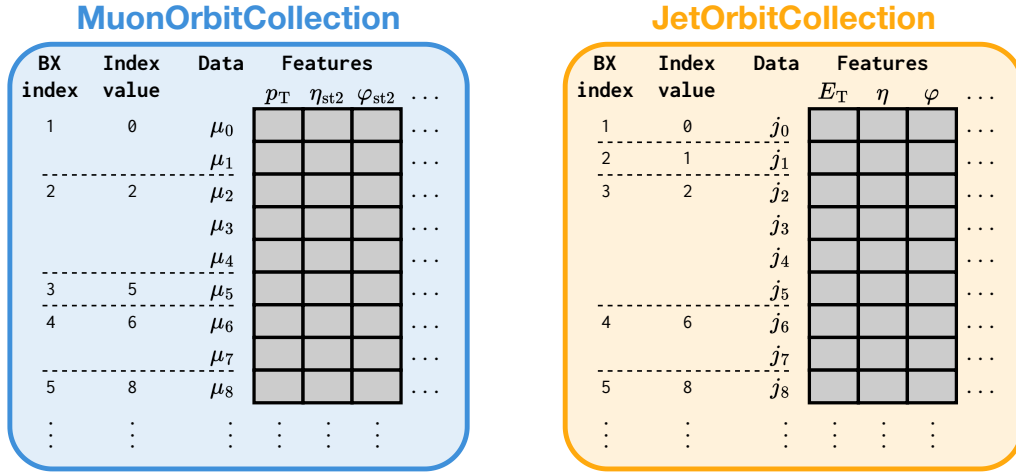


Figure 5.12 • Schematic view of Scouting Orbit Collection data format used to store the L1 trigger primitives read out by the L1DS. In particular, the collections for muons and jets are shown.

5.3.2 Event aggregation and processing

The event aggregation and processing flow is summarized in Fig. 5.13. This step is executed on the DSPU nodes, which are Dell R6525 machines with two AMD EPYC 7763 64-Core CPUs, optimized for intensive processing workloads.

A file broker software, used also in the central DAQ for the HLT, moves data fragments with the same event range from the RAM disk of the DSBUs to the DSPU nodes. This approach allows the workload for different event ranges to be distributed across multiple DSPUs, ensuring optimal parallelization of event processing. The output of this step is a chunk of data with aligned information from all sources in the same event.

Inside a DSPU, the event fragments are injected into CMSSW in dedicated raw data buffers. Raw data unpackers extract the trigger objects from all the event fragments for each event. The objects belonging to the same event are then aligned into the orbit collections.

The output of the unpacker modules is handed over to a set of analyzer modules, which perform the online analysis of the collected data. These analyzers loop through the possible BXs in an orbit, searching for specific signatures of interest for physics analyses. The output is a list of BX numbers that meet the required selections, which is then added to the event content. A filter module then loops through the event content, keeping only the BXs that were selected by at least an analyzer module.

The data chunks aggregated on the DSPUs are transferred back to a second RAM disk on the DSBUs designated as the primary unit. The chunks are then merged by a merger process into units of LSs. The data is temporarily sent to a Lustre filesystem in a “streamer” format and then copied to the Tier-0 storage for repacking into an AOD format. From there, the data is distributed to the CMS data aggregation system, where it can be accessed by users. These steps are illustrated in Fig. 5.14, where it can be seen that multiple dataset streams can be produced by the online processing infrastructure, serving different purposes.

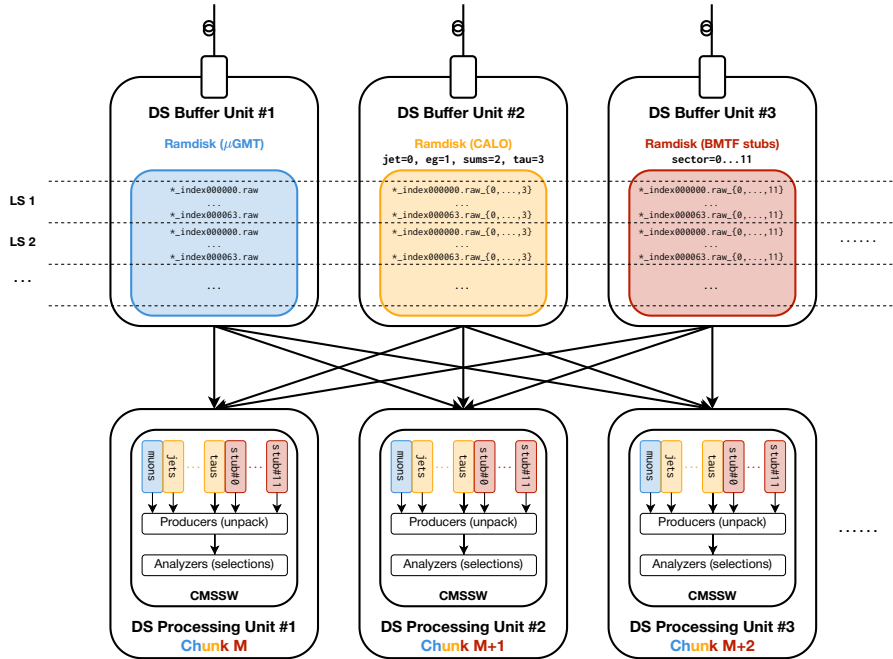


Figure 5.13 • Schematic of event aggregation and processing on the DSPU, after collecting fragments from the RAM disk of the DSBU.

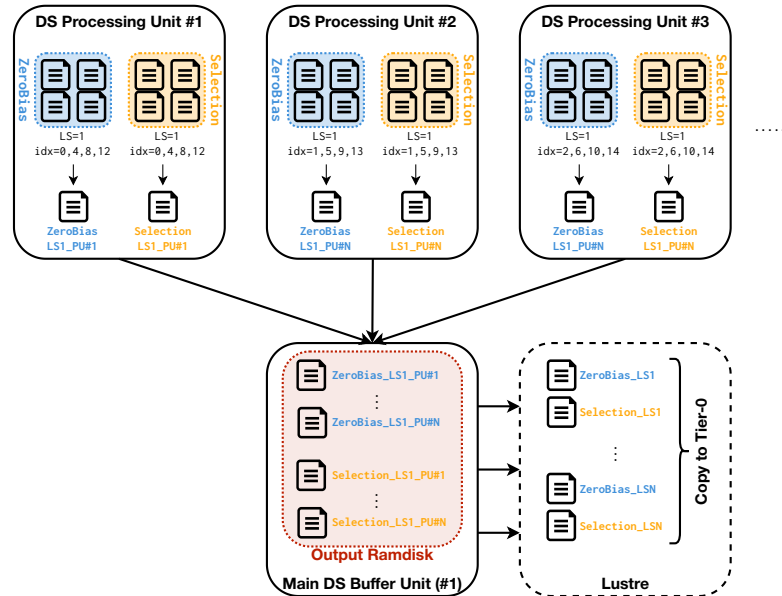


Figure 5.14 • Overview of the merging on the primary DSBU of the data chunks produced by the DSPU. Two main streams are produced, namely the scouting ZeroBias and Selection streams, and sent to a Lustre filesystem. The streams are then copied to the Tier-0 storage, where they are repacked and distributed to the data aggregation system of CMS.

5.3. Online processing at the LHC orbit rate

Two main CMSSW output data streams have been designed for the Run-3 demonstrator, producing two different output datasets for physics analysis and monitoring purposes. The first, known as “scouting ZeroBias” stream, is an unbiased dataset containing all the L1 event content sent by the scouting boards⁷. The average ZeroBias event size from all the orbit collections, after repacking, is shown in Table 5.6 for an LHC fill with about 2400 colliding bunches.

Given an average event compressed size of $\mathcal{O}(100)$ kB, the average raw ZeroBias data rate after repacking would be $\mathcal{O}(2)$ GB/s. Since the allocated storage for the Run-3 demonstrator for 2024 and 2025 operations is 3 PB of disk and 3 PB of tape storage, the limit of $\mathcal{O}(200)$ MB/s for the repacking output bandwidth has been imposed, necessitating data reduction for the dataset. This reduction for the ZeroBias dataset is accomplished by applying an event prescale factor N , which was kept around $\mathcal{O}(20)$ during 2024 operation. Thus, the dataset is not biased by any selections and can be used to study additional selector modules to add in the online analysis step, as well as for L1 diagnostic purposes.

Object	Uncompressed size (kB/ev)	Compressed size (kB/ev)	Objects per event
Muon	37	4.3	474
Jet	99	15	5265
e/γ	222	33	13000
τ_h	414	59	24965
E_T sums	226	30	2519
BMTF stubs	22	1.6	223
Total	1020	142	

Table 5.6 • Size of objects per event from scouting ZeroBias stream after data repacking and average number of objects per event.

The second dataset, known as the scouting “Selection” stream, is produced by the online analysis step. This dataset is intended for potential physics analyses and contains selected event records based on the signatures identified by the selector modules, which will be discussed in Section 5.3.3. The data reduction needed to meet the bandwidth constraints is achieved by discarding event content that does not meet the selector requirements and slimming the dataset by removing event collections with less potential for physics analysis. In particular, the L1 τ_h collections are discarded due to their lower purity compared to other physics objects, achieving a 50% throughput reduction.

5.3.3 Scouting selector modules

The selector modules developed for the 2024 operations focus on L1 muon and jet objects, which have higher physics potential given the Run-3 L1 trigger reconstruction. Each selector module searches for specific signatures based on single or multiple object conditions or correlations between different object types. Parameters are provided to configure thresholds in p_T and η of trigger objects and their multiplicity in a BX. This flexible implementation allows the module to be

⁷A ZeroBias dataset is also available from the standard CMS triggered data. It is acquired through the use of random triggers and a BX mask that considers only the BXs where a collision is expected. The prescale factor for this dataset ranges from 10^4 to 10^6 , depending on the running conditions.

reused with different conditions to produce various selections. The selector modules developed for the 2024 data-taking and their working points are described in this subsection, with the selection rate and throughput shown in Fig. 5.15.

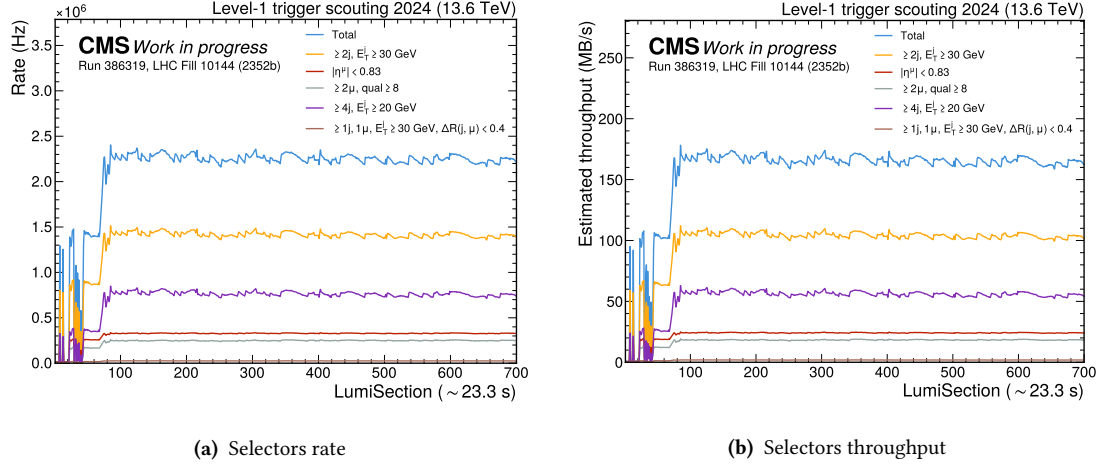


Figure 5.15 • Rate (5.15a) and throughput (5.15b) from selector modules deployed for L1DS demonstrator 2024 data taking: single muon in BMTF, dimuon with muons of quality at least 8, dijet with leading and subleading jets of $E_T \geq 30 \text{ GeV}$, high-multiplicity jet with at least four jets of $E_T \geq 20 \text{ GeV}$, muon inside jet cone of $R = 0.4$ and $E_T \geq 30 \text{ GeV}$.

Jet selector This module requires the presence of at least a configurable number of jets in a BX, meeting a specific E_T threshold and η range requirement. Two selection instances are defined from the same module:

- A “**dijet**” selection, requiring at least two jets with energy and pseudorapidity requirements. The chosen working point for 2024 is $E_T \geq 30 \text{ GeV}$ for both the jets and no η requirement.
- A “**high-multiplicity**” jet selection, requiring an arbitrary number of jets meeting E_T and η requirements. The chosen working point for 2024 is 4 jets with $E_T \geq 20 \text{ GeV}$ and no η requirement.

Muon selectors The muon selectors are implemented similarly to the jet selectors. The two selection classes deployed using the muon selector are:

- A “**single muon**” selection, which requires at least a muon with a certain hardware quality within a pseudorapidity range encoded in the track finder index quantity, and meeting a p_T threshold. For 2024, the single muon selection covers only the barrel region and is inclusive in p_T and muon quality.
- A “**double muon**” selection, where the chosen working point is inclusive in η and p_T , requiring both muons to have a quality equal or greater than 8.

Muon-tagged jet selector The final implemented selector correlates information from the jet and muon collections in the same bunch crossing to identify potential jets originating from b quarks. The module looks for a muon overlapping with a jet using a ΔR criterion with a configurable threshold. In this selection, the ΔR criterion is applied only to jets meeting E_T and η requirement and to muons with sufficiently high quality, p_T and η . The chosen working point for 2024 requires at least one muon overlapping a jet within $\Delta R < 0.4$ and the jet must have $E_T \geq 30$ GeV. A simple event display of a selected BX from this module is shown in Fig. 5.16.




Figure 5.16 • Transverse $r-\phi$ (left) and longitudinal $r-z$ (right) views of a collision event with L1 objects and collected by the Run-3 L1DS. The event contains eight jets (yellow), several e/γ (magenta), a muon within a jet cone (red curve), with its trajectory extrapolated at the vertex, and MET (red arrow). τ_h objects are omitted in the display.

5.4 Demonstrator validation from collected data

The demonstrator has been validated using data collected during 2024 from both the ZeroBias and Selection streams. The following checks are described in more detail below:

- Occupancy plots for every bunch crossing in the orbit and object multiplicity per bunch crossing.
- Distributions of kinematic features, such as p_T and E_T , which are useful for the evaluation of the trigger performance for different regional subsystems, and the effects of calibration procedures for triggering purposes.
- The emulation of trigger reconstruction algorithms on lower-level objects, such as the application of the KBMTF algorithm on the collected scouting stubs.

Further validation, and an assessment of the trigger performance, is provided by reconstructing known SM candles. In particular, low dimuon resonances such as $J/\psi \rightarrow \mu\mu$ and Z decays to a pair of leptons can be investigated, and they constitute useful benchmarks for a potential analysis targeting a new measurement.

5.4.1 Bunch crossing occupancies and object multiplicities

The L1 trigger object rate for each BX number in the LHC orbit is expected to reflect the LHC filling scheme, with peaks in the colliding bunches and depleted regions where no collisions are expected. For the μ GMT muons, the occupancy is stacked for the three regional subsystems in Fig. 5.17a, highlighting that most of the L1 muon rate comes from the EMTF subsystem. Activation effects from this region can be observed before the start and after the end of a bunch train, as shown in a zoomed-in view of the occupancy plot in Fig. 5.17b. In general, regions where no colliding bunches are expected are dominated by cosmic background, particularly visible in the BMTF contribution. A gap at the end of the orbit is visible, where multiple BXs are masked by the L1 trigger protocols, as described in Section 5.2.2. Focusing on the BMTF region, the occupancy of the BMTF stubs for different quality codes is reported in a restricted region in Fig. 5.17c, where higher-quality stubs are more likely to originate from real muons rather than background.

The L1 jets occupancy for the same BX region is reported in Fig. 5.17d as an example of calorimetric objects. The same distribution for e/γ and τ_h objects is omitted, as a similar pattern can be observed. The contributions in the HB, HE, and HF regions are highlighted separately. In this case, the rate is also dominated by the endcap regions, with visible drops and spikes at the end and start of a train, respectively. This feature is explained by the pedestal subtraction method used in the lower-level TPs, which removes out-of-time contributions from the BXs readout window [157]. Due to the lower contribution of out-of-time pileup from previous (subsequent) BXs at the start (end) of a train, a positive (negative) bias is introduced in the signal amplitude, producing this effect.

The multiplicity of objects per BX can also be used to evaluate the average size of an event, as reported in Table 5.6. The multiplicities for L1 muons, jets, e/γ , τ_h , and BMTF stubs per BX for colliding bunches are shown in Fig. 5.18. The observed averages may vary between different data-taking runs due to changing LHC and pileup conditions.

For muons, an average of 1 colliding BX out of approximately 10 contains at least one muon, consistent with the muon rate observed in the BX occupancy plots. The maximum multiplicity for stub objects in a BX is 96, as up to 8 stubs per wedge are sent to the L1DS demonstrator by the BMTF processors. The observed average for colliding BXs is around 0.1 in the run considered in this section, reflecting the small collection size in the stored data.

For calorimeter objects, the average multiplicity is significantly higher, explaining their larger collection sizes in the stored data. For jets, an average of approximately 2.2 objects per BX is observed, while about 5.4 e/γ objects are observed on average per colliding BX. For τ_h , there is a saturation towards the maximum number of 12 objects sent by the DeMux board, with nearly every collision event presenting a saturated number of τ_h objects. The multiplicity of sums is not reported as it is directly related to the presence of at least one jet or e/γ in the bunch, and therefore, sums are non-trivial in almost every colliding BX.

5.4.2 Energy and momentum distributions

The distribution of the L1 muon transverse momentum, shown in Fig. 5.19a, highlights the contributions from the three regional subsystems, with OMTF providing only a coarser p_T assignment. A peak at 255.5 GeV is visible, used in the trigger as an overflow bin for p_T .

5.4. Demonstrator validation from collected data

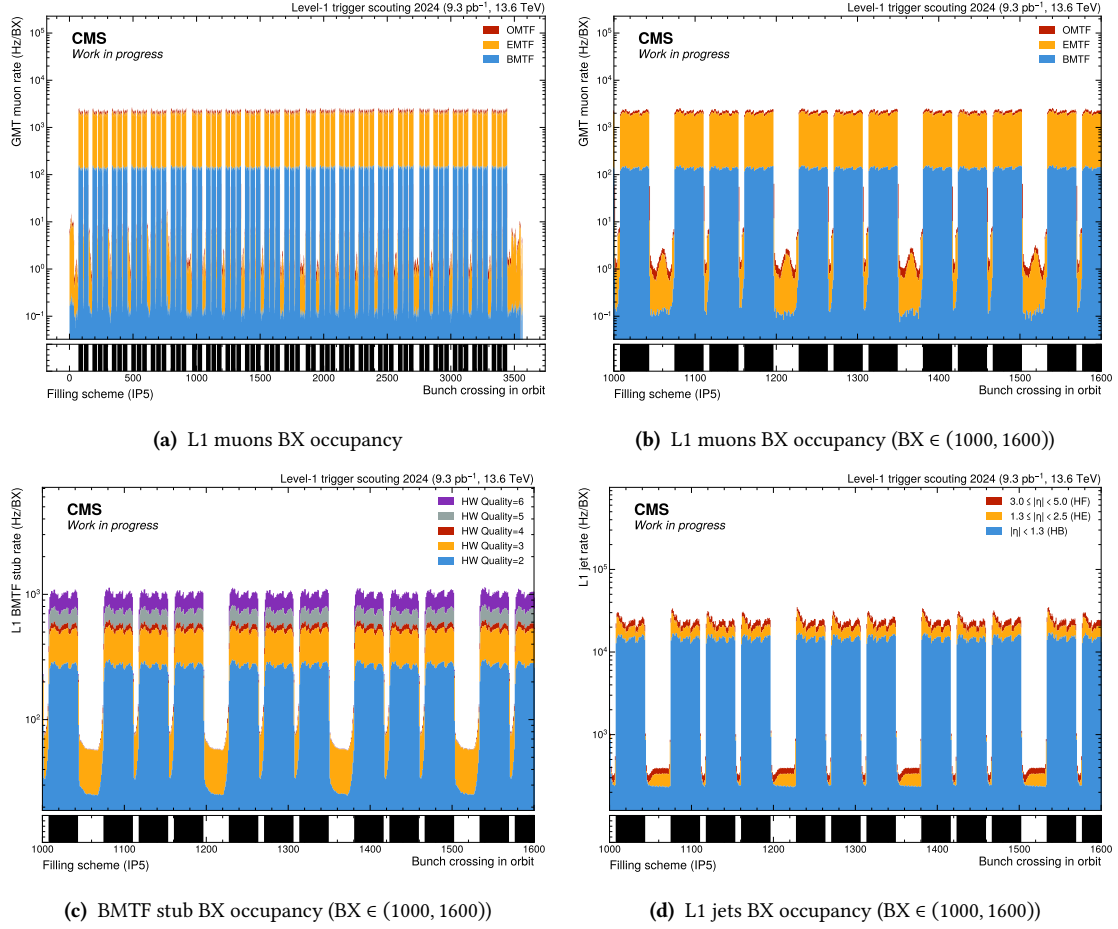


Figure 5.17 • L1 muons occupancy per bunch crossing in LHC orbit, over the full BX range (5.17a) and zoomed in the BX range of (1000, 1600) (5.17b). The contribution from BMTF, EMTF and OMTF are stacked. BMTF stubs occupancy in the BX range of (1000, 1600), stacked by quality code (5.17c). L1 jets occupancy in the BX range of (1000, 1600), stacked by regional subsystem (5.17d).

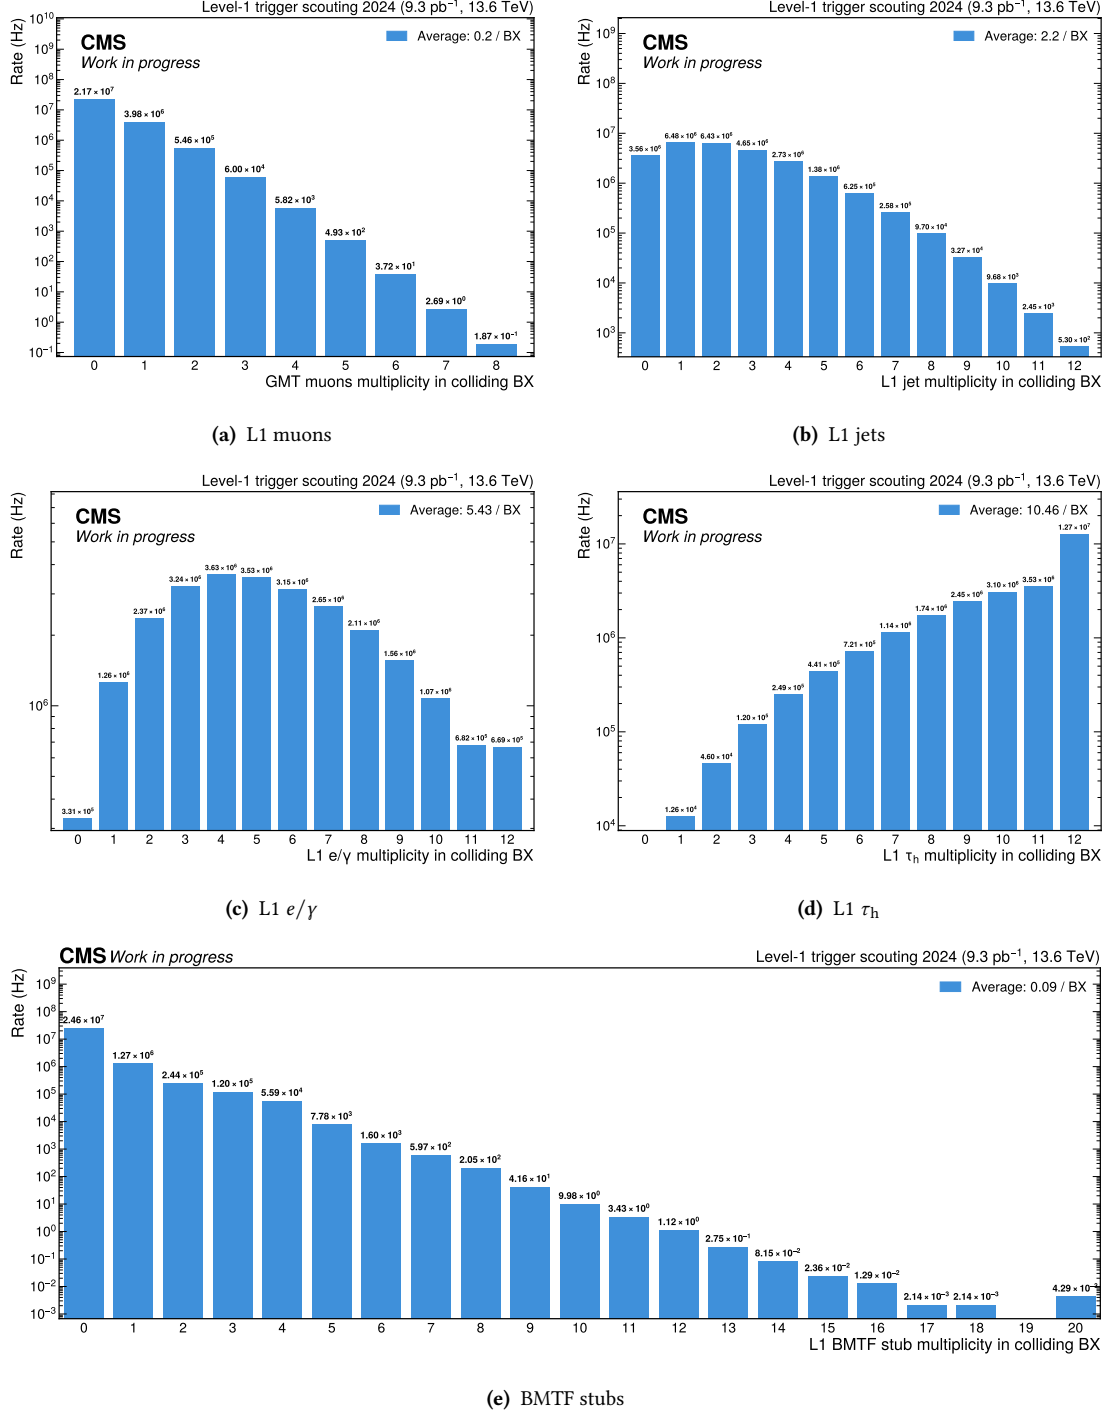


Figure 5.18 • L1 object multiplicities per BX, including only colliding bunches, for L1 muons (5.18a), jets (5.18b), e/γ (5.18c), τ_h (5.18d) and BMTF stubs (5.18e). The average number of objects is reported in the legend. Saturation at the maximum number of 12 τ_h objects is observed.

5.4. Demonstrator validation from collected data

Fig. 5.19 displays the E_T distributions of the e/γ and τ_h calorimeter objects, as well as the energy sums, where several distinct structures can be identified. For e/γ , a step at 128 GeV is observed, corresponding to the release of trigger selections during the reconstruction of the objects [48]. At the upper end of all the E_T distributions, a saturation bin collects all objects in which saturated towers were present in the cluster used for reconstruction. This saturation bin is also present in the energy sums E_T distribution, where the energy sum is set to the maximum value if at least one saturated tower is found in the BX. This behavior is driven by the inability to reliably calculate the total energy and energy balance in the collision event when saturation occurs.

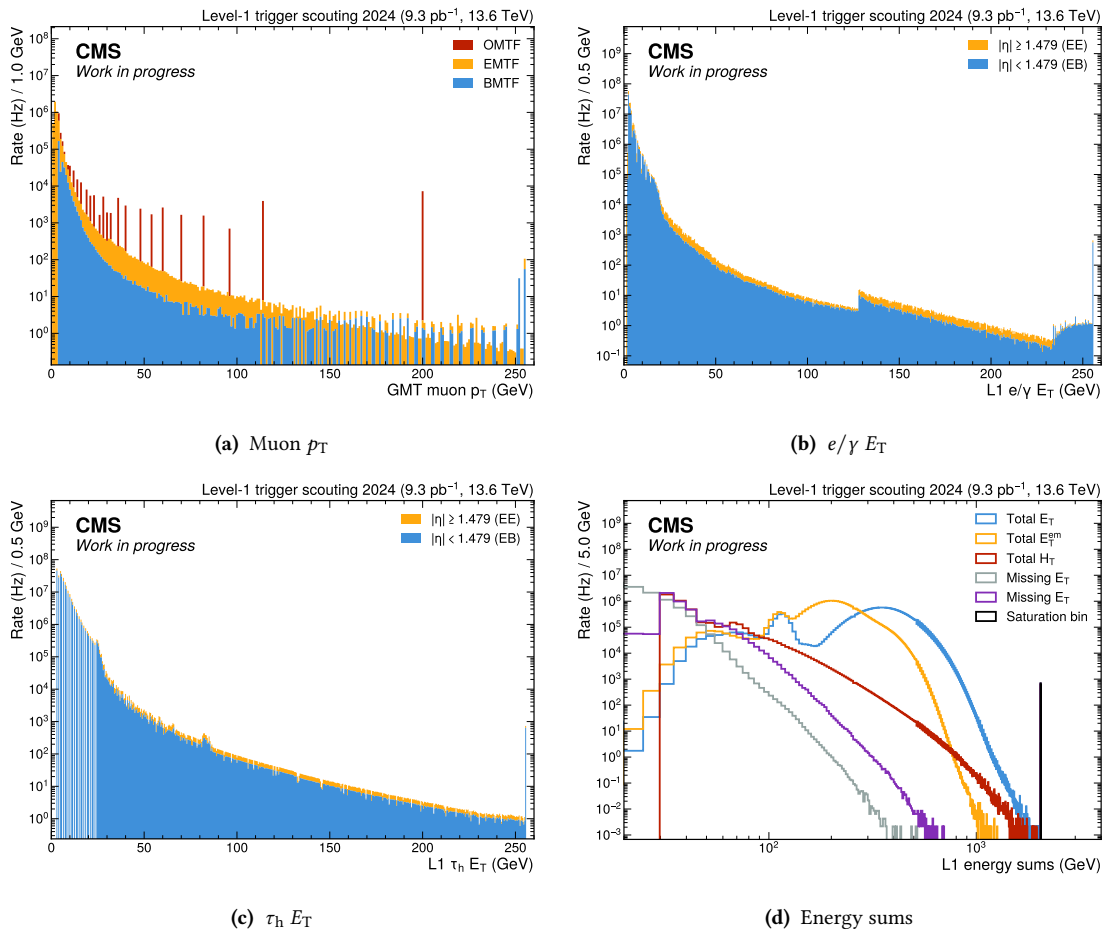


Figure 5.19 • p_T distribution of L1 muons, stacked for the three regional subsystems (5.19a). E_T distribution of L1 e/γ (5.19b) and τ_h (5.19c), stacked by barrel and endcap subsystems. E_T distribution of all the energy sums acquired by the L1DS (5.19d).

A dedicated discussion is reserved for the sum of TTs, which, as highlighted in Section 5.2.1 serves as a measure of pileup. The distribution of this quantity in each bunch crossing is shown in Fig. 5.20 for the BX range [1220, 1360]. For comparison, the pileup profile for each BX as reported by the CMS Online Monitoring System (OMS) is superimposed. The distribution of firing TTs shows a trend comparable to the pileup profile, with the exception of the spikes and drops at the start and end of each train, respectively, as described in Section 5.4.1.




Figure 5.20 • Number of TTs distribution versus BX in the BX range [1220, 1360]. The pileup information for every BX taken from the CMS OMS is superimposed for comparison, as the number of TTs is a measure of the pileup information.

For the L1 jets, a depleted region around 40-50 GeV and localized jet-enriched regions between 100 and 400 GeV in the E_T spectrum are present due to the effects of L1 calibrations, as shown in Fig. 5.21 for dijet events in the barrel. In general, these features are well modeled in the trigger simulation with MC samples. A good agreement between data and simulation is observed across the E_T spectrum, particularly at higher energies. The ratio between data and simulation normalization in the regions of the localized bumps and the saturation bin is acceptable. Reliable modeling of the L1 trigger simulation is crucial for any physics search to understand the background contributions from QCD processes and pileup, especially for signatures involving jets. These signatures are among the most promising for exploration with the L1DS demonstrator, as discussed in Section 5.6.

5.4.3 BMTF stubs and KBMTF emulation

Unlike the muon and calorimetric objects, the parameters of the BMTF stubs do not have a straightforward conversion from hardware to physical units. To investigate the spatial distribution of stubs within the detector, the stub acquisition rate in sectors, wheels, and stations is analyzed. Fig. 5.22 displays these rates as two-dimensional histograms. Fig. 5.22a shows the stub acquisition rate as a function of wheel and sector, with wheels ± 1 exhibiting the highest super-primitives acquisition rate. Fig. 5.22b focuses on the acquisition rate segmented by muon station and wheel. The innermost stations of the outer wheels (station 1, wheels ± 2) are excluded from the BMTF inputs and directed solely to the OMTF system. Consequently, the outer wheels do not contribute the highest rate of stubs to the BMTF. The acquisition rates decrease for the outermost stations, which are farther from the interaction point.

The ability to collect lower-level primitives offers an opportunity for custom reconstruction of L1 object candidates without the constraints of latency and resources. It also provides a way to emulate trigger reconstruction offline into higher-level objects, bypassing any discretization that would be applied during the online trigger reconstruction. The BMTF input stubs collected

5.4. Demonstrator validation from collected data




Figure 5.21 • Distribution of transverse energy E_T for jets collected by the L1DS demonstrator using the Selection stream [158]. Only the bunch crossing records with at least two jets with $E_T \geq 30$ GeV and $|\eta| < 1.2$ (barrel) for both the jets are stored for the purpose. The data are compared to QCD Monte Carlo samples generated for different hard-scattering p_T . Calibration artifacts are visible as bumps in the spectrum between 100 and 400 GeV. The bottom ratio panel shows the ratio between the data and simulation.

by the L1DS demonstrator serve as an example. While the first possibility will be explored in Section 5.5.2, the results from running the L1 trigger software emulator on these super-primitives are summarized here.

The KBMTF algorithm for muon reconstruction in the barrel region is typically executed on FPGA in the BMTF trigger processors. A software version of this algorithm, which emulates its hardware counterpart, is available within CMSSW, allowing access to its full resolution. While some quantities are truncated or approximated in the hardware version before being sent to the μ GMT due to bandwidth constraints, the software emulator avoids these limitations, permitting a more detailed analysis.

A study was conducted using the stubs collected by the L1DS during 2024, injecting them into the software emulator of the KBMTF to assess whether it is possible to recover the same muon candidate as the one sent to the μ GMT and collected by the L1DS. This comparative study serves as a final verification and proof of principle that it is possible to run software emulators on the collected data from the demonstrator with acceptable agreement. The methodology for matching BMTF muon candidates with μ GMT muons is based on their angular distance ΔR , with matching considered successful if $\Delta R < 0.1$. The p_T , p_T^{unc} , $\varphi_{\text{st}2}$, $\eta_{\text{st}2}$ and quality code from the matched pairs are compared, and discrepancies are evaluated as shown in Fig. 5.23a. Some mismatches are observed, particularly in $\eta_{\text{st}2}$, which can be explained by the fact that the scouting system does not capture the full set of stubs and due to differences in internal data formats used in the emulation. For further investigation, the distribution of $\Delta\eta_{\text{st}2} - \Delta\varphi_{\text{st}2}$ between BMTF and software emulation

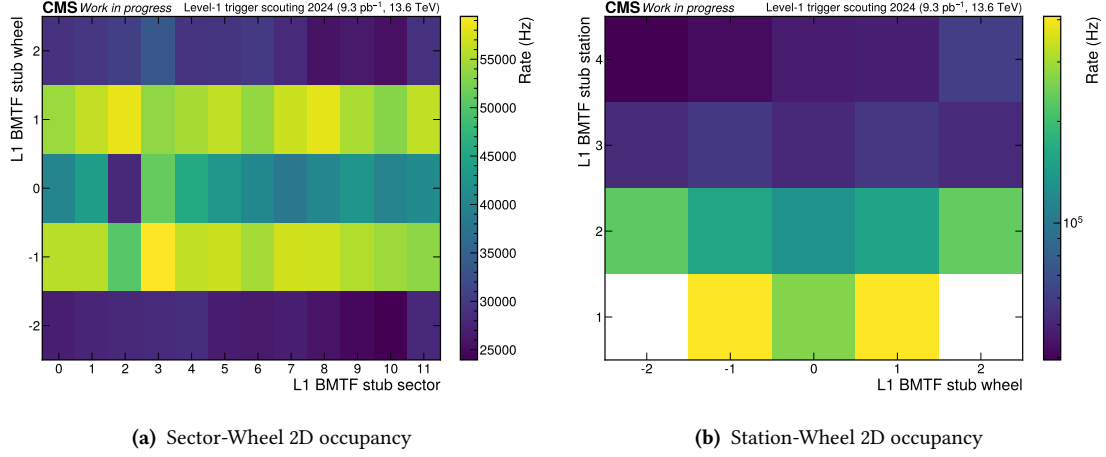


Figure 5.22 • 2D sector-wheel and station-wheel occupancies for BMTF stubs, collected with the Run-3 L1DS demonstrator during 2024.

is shown in Fig. 5.23b and the p_T distributions are compared in Fig. 5.23c, demonstrating that even in the case of mismatch, the differences are small and often shifted by just a single hardware quantity step.




Figure 5.23 • Mismatches between matched BMTF and software emulation from collected muon barrel stubs (5.23a). Distributions of $\Delta\eta_{st2} - \Delta\phi_{st2}$ (5.23b) and p_T (5.23c) between BMTF and KBMTF emulation.

5.4.4 Low dimuon invariant mass resonances

SM resonances at low $m_{\mu\mu}$ invariant mass include the decay of the J/ψ meson into a pair of muons. Observing this resonance allows for the evaluation the L1 trigger performance and resolution in the reconstruction of muons in the low momentum regime, which can be valuable for new physics searches, particularly in Dark Matter models.

To enhance sensitivity to $J/\psi \rightarrow \mu\mu$ decays, a set of event selection criteria is applied, starting from approximately 1.2 fb^{-1} of integrated luminosity collected during 2024 with the “double muon” selection stream, described in Section 5.3.3. Events with at least two muons are considered, and the muon pairs are filtered with the following requirements:

- The transverse momentum of the muons greater than 8 GeV.

5.4. Demonstrator validation from collected data

- A muon hardware quality code of at least 12.
- Hardware transverse impact parameter equal to 0.
- A dimuon candidate constructed as the pair of muons with the smallest angular distance ΔR , ensuring $\Delta R > 0.05$ to remove possible muon duplicates.
- A candidate dimuon transverse momentum greater than 16 GeV, driven by the generator-level cut $p_T^{\mu\mu} > 8$ GeV in the MC sample of $J/\psi \rightarrow \mu\mu$ used for comparison with the data.

The resulting invariant mass $m_{\mu\mu}$ is shown in Fig. 5.24 for both Opposite Sign (OS) and Same Sign (SS) pairs. The comparison with the $J/\psi \rightarrow \mu\mu$ simulation at L1 trigger level for OS pairs shows an acceptable agreement, given the L1 trigger resolution and the fact that no SFs are applied to correct the simulated event weights.




Figure 5.24 • $m_{\mu\mu}$ invariant mass distribution for dimuon final state, for opposite sign (OS, blue) and same sign (SS, orange) pairs. The opposite sign distribution is compared to simulated $J/\psi \rightarrow \mu\mu$ events normalized to the SM expectation (gray) and their subtraction is shown (red).

5.4.5 Z decay to a pair of leptons

As with the low dilepton resonances, searching for higher mass resonances in dilepton final states serves as a useful benchmark for the performance of the L1 trigger. Specifically, analyzing decays of the Z boson into a pair of muons or electrons can be used as a reference to compute calibrations for jets, such as in DY plus jets events where a jet recoils against the dilepton candidate.

A comparison with the simulation of these processes, focusing solely on L1 information, was performed using early data collected during 2024, corresponding to approximately 110 pb^{-1} of integrated luminosity. Drell-Yan (DY) to a lepton pair plus jets simulated samples were employed as the SM reference, with a production cross section $\sigma(Z \rightarrow \ell\ell) = 6244.8 \text{ pb}$ for the phase space region where $m_{\ell\ell} > 50 \text{ GeV}$ [159].

For the $\mu\mu$ final state, the event selection applies the following criteria to L1 muons to enhance the signal-to-background ratio:

- The muons are required to be from the muon barrel, so their track finder index must be in the range $[36, 71]$, or equivalently, their pseudorapidity must be $|\eta| < 0.83$.
- The momentum of the muons must be greater than 8 GeV.
- The hardware transverse impact parameter must be equal to 0.

The result of the selection for OS pairs is shown in Fig. 5.25a. For the ee final state, the event selection criteria for e/γ objects are as follows:

- The e/γ pair must be detected in the ECAL barrel, meaning they must be within the pseudorapidity range $|\eta| < 1.44$.
- An energy threshold of $E_T \geq 20$ GeV is imposed on the e/γ objects of the pair.
- Each e/γ has to be isolated from hadronic activity, so the hardware isolation code must be either 1 or 3.

The observed invariant distribution $m_{e/\gamma e/\gamma}$ is shown in Fig. 5.25b.

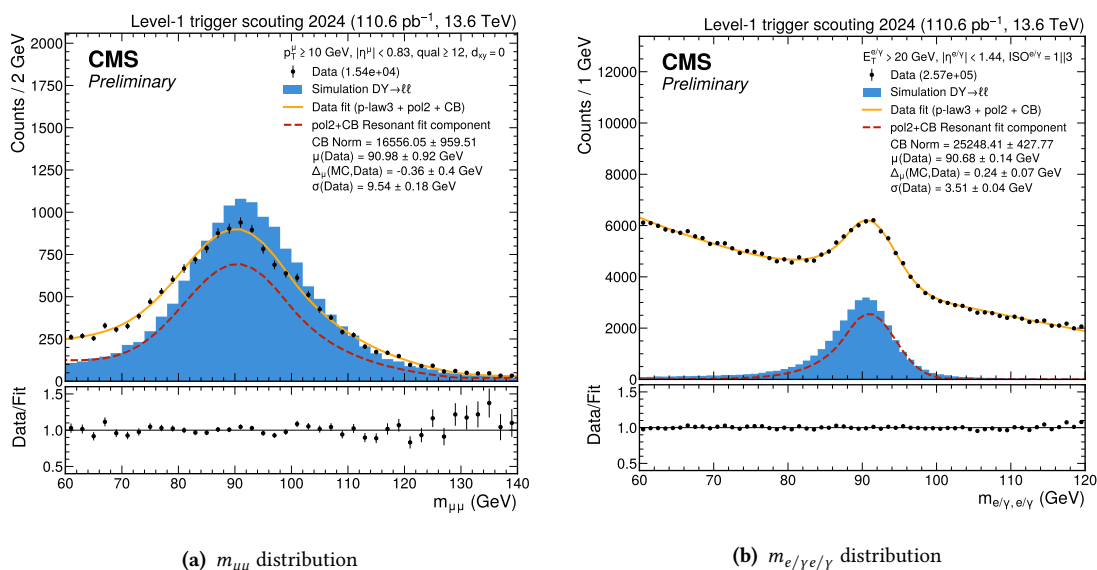


Figure 5.25 • Distribution of the invariant mass (black) from opposite sign pairs of μ GMT muons reconstructed by the BMTF (5.25a) and of pair of e/γ reconstructed by the CT (5.25b), collected by the L1DS demonstrator during 2024. Both the distributions are modelled using a resonant function as the DSCB summed to a polynomial function to model mismatched dilepton events and background contributions (orange), with the resonant component summed to a polynomial for mismatched modelling plotted separately (red).

In both studies, a fit is performed on the observed data using a DSCB resonant component for the signal, summed with a polynomial function to model mismatched events, and a power-law function for the background. The DSCB shape parameters, apart from width and mean, and the shape of the polynomial are derived from simulated events and fixed in the fit to the data.

Discrepancies in the resonance position between data and simulation are expected due to calibration artifacts in the L1 collections. The shift in the resonance mean is modeled through an additional parameter Δ_μ added to the mean of the DSCB in the fit. The discrepancy in normalization is attributed to the absence of SFs correcting the simulated event weight, such as corrections for trigger efficiency differences between data and simulation. These discrepancies are particularly pronounced in the muon final state, where the p_T has been corrected by a factor of 1.16, reflecting the fact that, on average, the p_T of L1 muons is about 20% higher to enhance trigger efficiency. A better agreement is found for the e/γ final state, where no E_T corrections have been applied, and the MC yield is closer to that observed in the data.

5.5 Machine learning inference applications for L1DS

The L1DS system also serves as a testbed for applying machine learning techniques to the acquired L1 trigger objects. These applications could be implemented directly on the FPGA readout board of the L1DS, enabling real-time inference and enhancing the physics capabilities of the system. This can be achieved by translating higher-level neural networks and boosted decision trees into hardware description language to integrate into an FPGA design, using tools such as hls4ml [155, 156] and conifer [160]. This section explores the use of neural networks for the recalibration of L1 muons, for fake/real L1 muon pair classification, and for BMTF stubs tracking.

5.5.1 L1 muons recalibration

L1 trigger objects are calibrated to achieve a certain efficiency at a momentum threshold, which makes them unsuitable for direct physics analysis. To address this, neural networks have been developed and trained to recalibrate the parameters of these objects, with the goal of deploying them on the scouting FPGA boards or using them in the offline analysis. This subsection focuses on the recalibration of L1 muons and improving the extrapolation to the collision vertex performed by the μ GMT.

The chosen network architecture consists of a “feed-forward” model with three hidden layers, each containing 32 neurons. L2 regularization [161] and Batch Normalization [162] layers are included to enhance performance. The model takes as input the p_T , η_{st2} , η_{st2} , charge, and hardware quality of the muons and returns a set of correction terms Δp_T , $\Delta \eta$, $\Delta \phi$ to be applied to the kinematic features for recalibration and extrapolation. The logarithmic hyperbolic cosine loss function has been identified as yielding the best performance.

The neural network is trained on offline collision data collected during the 2022 data-taking period using random colliding BX triggers, ensuring that the data used for the application is unbiased by any trigger selection. μ GMT muons are matched to offline reconstructed muon tracks, whose kinematic parameters are used as targets for the network. The matching between offline reconstructed and L1 muons requires their angular distance at the second muon station to be $\Delta R < 0.1$. Since the offline reconstructed muons are provided at the vertex, their tracks are propagated to the second muon station using a track propagator.

The correction terms predicted by the network on a test dataset, not used for training, are applied to the corresponding L1 muons, and the corrected kinematic quantities are compared to the matched offline candidates, as shown in Fig. 5.26. The difference between L1 and offline quantities is superimposed for comparison.

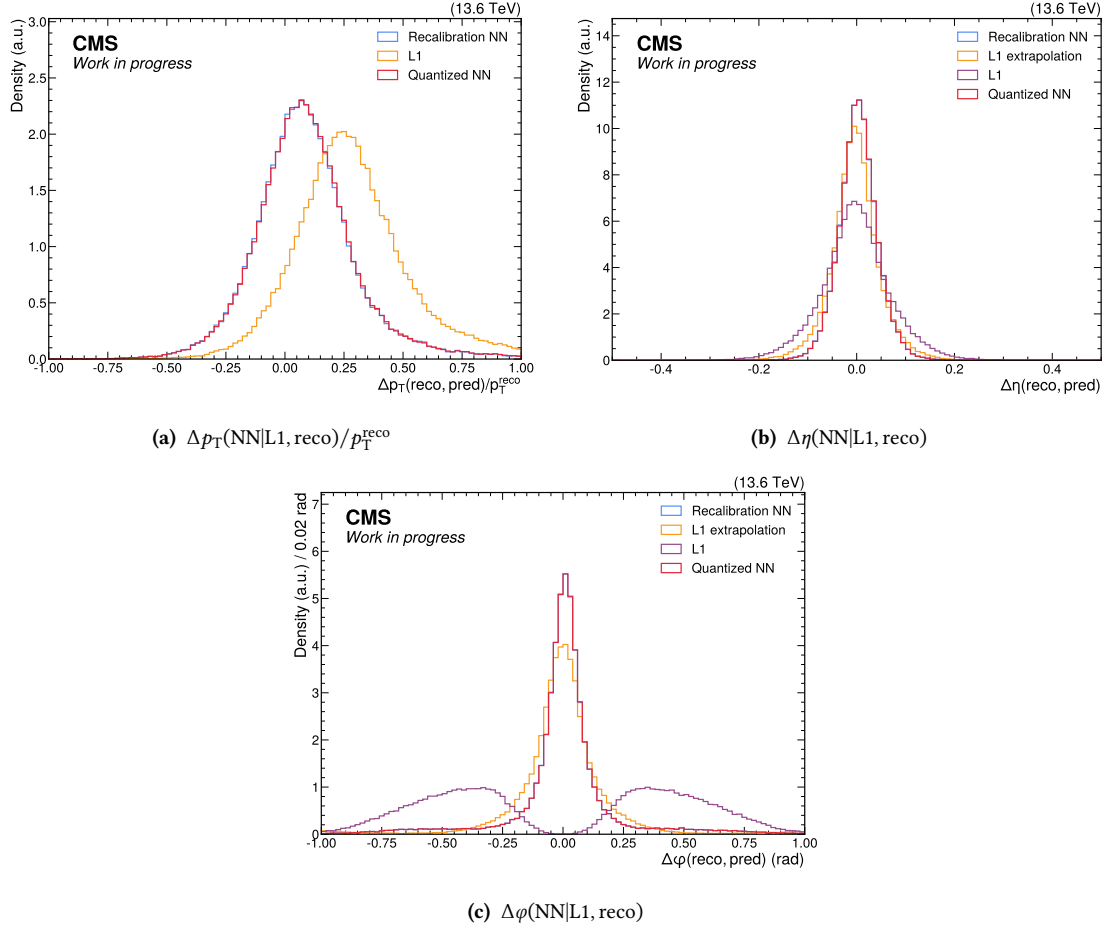


Figure 5.26 • Deep-learning based L1 muon recalibration results, comparing L1 quantities and NN-corrected quantities with the offline reconstruction. The comparison in p_T over the p_T of the offline muon (5.26a), in η_{ext} (5.26b) and ϕ_{ext} (5.26c) are shown. The comparison for the quantized NN using hls4ml is also included.

The resolution peak shows how the recalibration is correcting the central value of the p_T distribution to its nominal value from offline data. Moreover, the NN is correctly learning the propagation at the nominal collision vertex with performance superior to that of the extrapolation done in the μGMT , introduced in Section 4.2.2.

The network parameters are then discretized to a fixed-point format for translation into HLS code for deployment on FPGA. The optimal fixed-point precision is chosen to balance the amount of resources required for hardware implementation and the degradation in resolution compared to the results with floating-point weights: 8 bits are reserved for the integer part, and 16 bits for the fractional one. The results from the quantized network compared to the matched offline quantities are also shown in Fig. 5.26.

5.5. Machine learning inference applications for L1DS

The NN converted to HLS code has been synthesized to evaluate the approximate amount of logic resources required for implementation on hardware. The results, shown in Table 5.7, are acceptable for an implementation reaching timing closure with the VU37P FPGA device, whose available logic resources are listed in Table 5.5. To reduce the usage of DSP components, a single DSP is reused for multiple operations, at the expense of increased inference latency, which amounts to 23 clock cycles at a frequency of 250 MHz. Quantization Aware Training techniques [163, 164] are expected to further reduce the required logic resources with negligible loss in the resolution of the recalibration peaks.

Resource	Utilization	Available	%
LUT	40871	1303680	3.1
FF	16549	2607360	0.6
BRAM	0	4032	0
URAM	0	960	0
DSP	226	9024	2.5

Table 5.7 • Resource usage of deep-learning based L1 muon recalibration application on VU37P FPGA device.

5.5.2 Deep-learning based barrel muon track reconstruction

Muon stub primitives from the barrel muon system undergo pattern recognition and track fitting in the BMTF trigger processors. The track fitting step, using a Kalman filter, is implemented under strict latency and resource constraints. However, within the scouting system, these constraints are lifted, allowing for a custom reconstruction to be performed on the collected muon stubs. To explore potential improvements over the current reconstruction, a deep-learning based approach has been investigated for the track fitting step [165], building upon the existing pattern recognition done in the BMTF algorithms.

The neural network is trained on offline collision data collected during the 2023 data-taking period using random colliding BX triggers. The dataset is complemented by offline collision data collected with muon plus photon triggers to populate the high p_T muon region and ensure uniform network performance across a broader p_T spectrum. The training dataset is prepared by collecting offline muon tracks and propagating them to the second muon station. These propagated tracks are then matched to the tracks produced by the KBMTF algorithm from the muon barrel stubs, following the matching methodology described in Section 5.5.1. From the matched KBMTF tracks, the stubs contributing to the muon candidate are extracted. The network aims to predict the offline muon tracks p_T , η , ϕ at the nominal collision vertex in a regression task, while the charge of the muon candidate is determined by the sign of the predicted p_T . The network can accept up to four stubs, with padding applied when only two or three stubs are associated with the KBMTF track.

The network architecture balances performance and compatibility with hls4ml supported features for potential hardware deployment. It is a “feed-forward” model with three hidden layers, each containing 32 neurons, and uses L2 regularization [161] with $\lambda = 10^{-3}$. Exponential Linear Units (ELU) [166] are chosen for activation functions, and the Mean Absolute Error

(MAE) loss function yields the best performance. Other regularization techniques, such as Batch Normalization [162] and Dropout [167], were tested but showed inferior performance, so they were not pursued further.

The regression performance is evaluated by examining the resolution of each feature, calculated as the difference between the predicted or L1 quantity and the offline reconstructed value. For transverse momentum, the difference is normalized by the offline reconstructed muon p_T . Performance is assessed using a test dataset distinct from the training and validation sets. The charge assignment accuracy of the neural network is comparable to that of the KBMTF, with approximately 97.9% accuracy relative to the offline charge assignment. The resolution of each feature is shown in Fig. 5.27.

The network parameters are discretized into a fixed-point format with 8 bits for the integer part and 16 bits for the fractional part. The optimal fixed-point precision is selected to balance resource usage for hardware implementation and the resolution degradation compared to floating-point weights. The results from the quantized network are also included in the comparison.

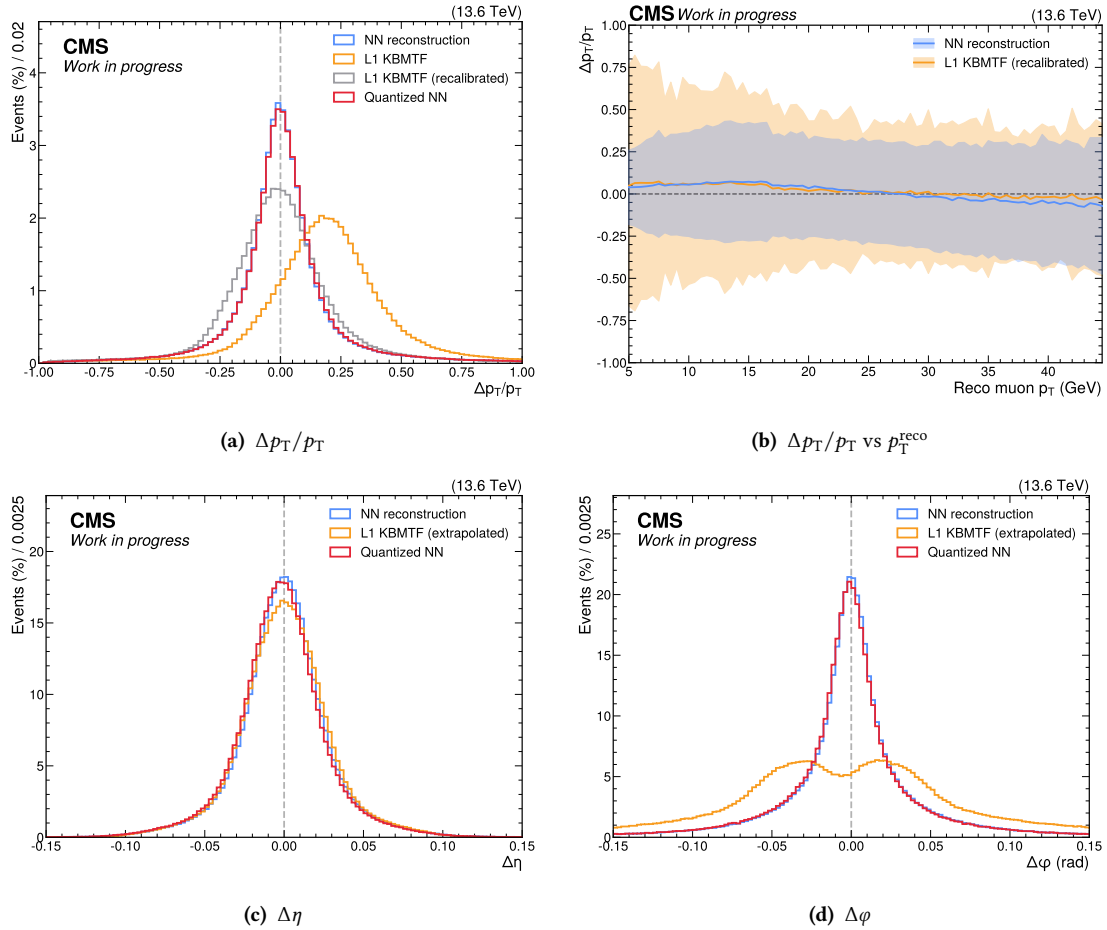


Figure 5.27 • Resolution relative to the offline reconstruction of p_T (5.27a and 5.27b), η (5.27c), and ϕ (5.27d), of the neural network prediction (blue) and the KBMTF assignment (orange) and its recalibrated peak (gray), where the recalibration is performed by dividing the KBMTF assigned p_T by 1.2. The results from the quantized network, with an 8-bit integer part and 16-bit fractional part, are also shown (red).

5.6. Run-3 demonstrator summary and perspectives

The NN converted to HLS code through `hls4ml` has been synthesized to evaluate the approximate amount of logic resources required for hardware implementation. The results, shown in Table 5.8, are acceptable for an implementation reaching the timing closure with the VU37P FPGA device. To minimize the use of DSP components, a single DSP is reused for 16 different multiplicative operations, which increases inference latency to 68 clock cycles at a frequency of 250 MHz.

Resource	Utilization	Available	%
LUT	137383	1303680	10.5
FF	74144	2607360	2.8
BRAM	48	4032	1.2
URAM	0	960	0
DSP	271	9024	3.0

Table 5.8 • Resource usage of deep-learning based barrel muon stubs fit application on VU37P FPGA device.

The performance of this approach generally exceeds that of the standard KBMTF currently in use by the trigger system. This result highlights the feasibility of applying machine learning techniques within the L1DS system, either on the FPGA readout boards in real-time or offline on the collected stubs. Moreover, it demonstrates the potential for improving the physics performance of trigger objects by applying custom reconstruction on lower-level primitives without the constraints of latency and resource limitations. The model described could be extended to predict the transverse momentum without vertex constraint, utilizing simulations of exotic LLP particles decaying to displaced muons, thereby enhancing the L1DS sensitivity to topologies where the current L1 trigger reconstruction and selection are less effective.

5.6 Run-3 demonstrator summary and perspectives

The successful commissioning of the L1DS Run-3 demonstrator marks significant progress towards the development of the Phase-2 system. The characterization and physics studies presented in this chapter have demonstrated that the L1DS is not only a valuable tool for system benchmarking but also holds potential for monitoring purposes and conducting physics searches that are not feasible with standard triggered data or the HLT scouting dataset. Thus, the L1DS experience during Run-3 has the potential to extend beyond its initial demonstration purposes and be used to produce a physics measurement.

Dijet searches, which have been a focus of both the CMS and ATLAS experiments, as well as previous generations of colliders [168], have provided thus far only weak constraints on the coupling of hypothetical exotic bosons with masses ranging from 100 to 300 GeV that decay into a pair of hadrons. The summary of the most recent searches performed by the CMS experiment is illustrated in Fig. 5.28, where the low mass region is primarily explored through HLT scouting analyses.

Simulations have shown that at lower energy regimes, L1 jets reconstructed by the CMS Phase-1 trigger achieve a resolution comparable to that of PF jets and jets with PUPPI pileup subtraction. This comparison is illustrated in Fig. 5.29 for a simulation of leptophobic Z' (with $m_{Z'} = 200$ GeV) decaying to quarks [169]. The L1DS approach could significantly enhance signal efficiency, with the only E_T threshold being the implicit one from the L1 reconstruction.




Figure 5.28 • Limits on the universal coupling g'_q between a leptophobic Z' boson and quarks [169] from various CMS dijet analyses, taken from Ref. [170]. The expected limits are shown in dashed lines, and the corresponding observed limits are shown in solid lines. The hashed areas indicate the excluded regions based on the observed limits. The gray dashed lines represent the g'_q values at fixed ratios $\Gamma_{Z'}/M_{Z'}$.




Figure 5.29 • Simulation of leptophobic Z' boson decay to a pair of jets [169], with $m_{Z'} = 200$ GeV and an $E_T \geq 30$ GeV threshold applied on the jets. The invariant mass peak is compared for generator-level jets, L1 jets, PF jets, and jets with PUPPI pileup subtraction. The resolution for L1 jets is comparable with the HLT online reconstruction and offline jets at low energy regime.

5.6. Run-3 demonstrator summary and perspectives

This measurement poses at the same time considerable challenges that will require extensive investigation. Robust background modeling techniques are essential, especially given that trigger objects might not be immediately calibrated for physics analysis, as evidenced by the structures visible in Fig. 5.21. If a reliable simulation to model QCD background processes is not feasible computation-wise, a data-driven approach must be implemented. An example of this, as used in Ref. [171], leverages the $|\Delta\eta|$ of the dijet system, which is expected to be centered around zero for the signal, in contrast to QCD background processes. The L1DS collected data must also deal with an unprecedented amount of events, presenting a significant computational challenge to address in future work.

6

Study of rare Higgs boson decays with Phase-2 Level-1 trigger Data Scouting

Searching for Higgs boson rare exclusive decays, with branching fractions ranging from 10^{-5} to 10^{-10} , offers a probe to the couplings to first and second generation quarks, as illustrated in Chapter 1. Their study requires high signal selection efficiency, which begins with a high trigger efficiency. As introduced in Chapter 4, triggering on these processes might not be trivial, in particular for final states with hadronic tracks coming from meson decays. When possible to implement, a trigger algorithm that targets one of these specific decays might require a too large amount of logic resources to fit in the trigger menu. The possibility of performing data analysis on all the events, bypassing the L1T accept decision, might overcome these problems and enhance the sensitivity of the search. This approach, known as Level-1 trigger Data Scouting, has been commissioned in the CMS Run-3 L1 trigger with successful results, as described in Chapter 5. Thus, the application in the upgraded Phase-2 L1 trigger might open the possibility to probe physics scenario not accessible before. The CMS upgrade for HL-LHC offers a promising ground for the data scouting approach, given a L1 trigger reconstruction resolution often comparable to the one of the offline reconstruction [115]. This chapter explores the possibility of searching for the rare Higgs decays at the center of this work, using only the information of the Phase-2 L1 trigger reconstruction. The analysis feasibility is evaluated in the optics of deployment inside the Phase-2 L1DS online processing system.

6.1 Study premises and conditions

The projections reported in this chapter are derived using the CMS Phase-2 detector and L1 trigger simulation. The $H \rightarrow QQ$ and $H \rightarrow Q\gamma$ signal samples are generated at LO using PYTHIA for the inclusive H boson production and decay to a pair of mesons or to a meson and a photon. The meson subsequently decays to a pair of charged pions for $Q = \rho$, to a pair of charged kaons for $Q = \phi$, or to a pair of muons for $Q = J/\psi$. For every signal sample, around 25k events are generated at a center-of-mass energy of 14 TeV, with the additional simulation of 200 pileup interactions for a more realistic projection. The expected reducible background yields are estimated from minimum bias collisions simulation at a pileup of 200.

The analysis results are derived for a total integrated luminosity of $\mathcal{L}_{\text{int}} = 400 \text{ fb}^{-1}$, corresponding to one year of data taking at the highest luminosity of $7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. For the reducible background sample, a rate of collision $f_{\text{minbias}} = 31.5 \text{ MHz}$ is considered, corresponding to a maximum number of colliding bunches of 2808 [115]. The integration time is $5.33 \times 10^6 \text{ s}$, corresponding to roughly $t_{\text{DAQ}} = 62$ days of data taking. Thus, the number of estimated background events per year is computed as:

$$n_{\text{bkg}} = \epsilon_{\text{bkg}} \cdot f_{\text{minbias}} \cdot t_{\text{DAQ}} \quad , \quad (6.1)$$

where ϵ_{bkg} is the selection efficiency on the background sample. Conversely, the expected amount of signal events is estimated as:

$$\begin{aligned} n_{\text{sig}}(\text{H} \rightarrow Q\gamma) &= \mathcal{L}_{\text{int}} \cdot \sigma_{\text{H}} \cdot \mathcal{B}(\text{H} \rightarrow Q\gamma) \cdot \mathcal{B}(Q \rightarrow p_1 p_2) \cdot \epsilon_{\text{sig}} \\ n_{\text{sig}}(\text{H} \rightarrow Q_1 Q_2) &= \mathcal{L}_{\text{int}} \cdot \sigma_{\text{H}} \cdot \mathcal{B}(\text{H} \rightarrow Q_1 Q_2) \cdot \mathcal{B}(Q_1 \rightarrow p_1 p_2) \cdot \mathcal{B}(Q_2 \rightarrow p_3 p_4) \cdot \epsilon_{\text{sig}} \end{aligned} \quad , \quad (6.2)$$

where $\sigma_{\text{H}} = 62.5 \text{ pb}$ is the inclusive Higgs production cross section at $\sqrt{s} = 14 \text{ TeV}$ [16], $\mathcal{B}(Q \rightarrow pp)$ is the branching fraction for the meson decay to the final state particle pairs, presented in Table. 1.3, and ϵ_{sig} the signal selection efficiency.

The L1 objects considered for the analysis are the PUPPI candidates, tracker tracks, and standalone e/γ from CL1. Their multiplicities in minimum bias simulated events for different p_{T} thresholds are shown in Fig. 6.1.

The PUPPI candidates have a 3-vector with p_{T} , η and ϕ , with additional information specifying the particle type (charged hadron, neutral hadron, photon, electron, muon) and their charge. Charged particles also contain the track d_z and d_{xy} information, while neutral particles have the output weight from the PUPPI algorithm. The average multiplicity of PUPPI particles in simulated minimum bias events at pileup of 200 is about 31, and increases to about 50 for inclusive $t\bar{t}$ events.

The L1 track candidates come with the same kinematic quantities as for the PUPPI candidates, but they do not provide a discrimination between the different types of particles. Their use is explored as a possible improvement for the meson decay reconstruction, where performance degradation might come from pileup subtraction performed assuming a misassigned primary vertex. The track multiplicity increases with respect to the PUPPI candidates, reaching an average of 156 for simulated minimum bias events at pileup of 200, and over 200 for inclusive $t\bar{t}$ events.

The L1 muons for this study are selected among the PUPPI candidates or identified from the set of tracker tracks in the reconstruction. Future improvements of this work will extract them from dedicated collections, storing all the tracker muons, which have reduced multiplicity per BX.

The L1 photons are available in dedicated collections, produced in the Layer-1 step of the Correlator Trigger. They are provided with their kinematic features and with additional quantities as reconstruction quality and isolation with respect to hadronic activity. Their multiplicity saturates at the maximum of 12 objects almost in every minimum bias simulated event, as shown in Fig. 6.1c.

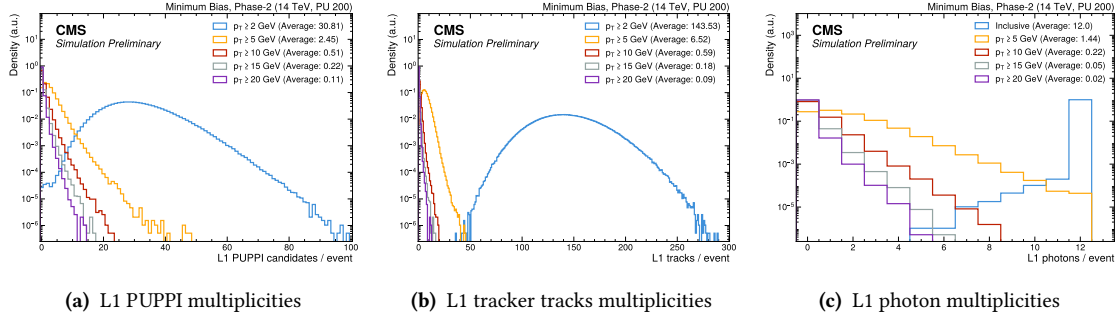


Figure 6.1 • Multiplicity of L1 PUPPI (6.1a), tracker track (6.1b) and photons (6.1c) for different p_T thresholds in minimum bias simulated events.

6.2 Generator level study

A generator-level study is performed first to assess the fraction of events within the L1 scouting acceptance for standard selection criteria on the different objects. Generated photons are required to be within $|\eta| < 2.5$ and to have $p_T > 10$ GeV. Charged particle candidates from the meson decays with $p_T > 2$ GeV are accepted up to $|\eta| < 2.4$. The acceptance for $H \rightarrow Q\gamma$ ($H \rightarrow QQ$) decays is around 74% (73%) for $Q = \phi$, dropping down to around 72% (70%) for $Q = \rho$. For the $H \rightarrow J/\psi\gamma$ ($H \rightarrow \phi J/\psi$) decay it amounts to 68% (67%). The acceptance is calculated considering the LO production of a Higgs boson. Slightly higher numbers could be expected when including higher order production processes.

To perform the matching between the generator level photon from the Higgs boson decay and the L1 photons, the pair with lowest angular distance is selected and the matching is considered successful if $\Delta R < 0.1$. For the charged particles coming from the meson decay, the matching is done between the generator level mesons and the reconstructed candidate mesons from L1 PUPPI charged hadrons or track pairs. All the possible combinations of L1 charged particle pairs are created after preselecting the set of particles with a p_T threshold of 5 GeV, reducing the computational complexity of the matching procedure without significantly affecting the signal efficiency. The pair combinations are further filtered by selecting the pairs with an invariant mass value compatible with the candidate meson mass, namely $m_{\pi\pi} \in [0.4, 1.3]$ GeV for $Q = \rho$, $m_{KK} \in [0.95, 1.25]$ GeV for $Q = \phi$ and $m_{\mu\mu} \in [2.5, 3.5]$ GeV for $Q = J/\psi$.

The matched L1 photon and charged particle pair are further filtered on the basis of the hadronic activity inside a cone $R = 0.25$. A variable quantifying the isolation from hadronic activity is built and it reads:

$$\text{Iso}(\text{obj}) = \sum_{i \mid 0.05 < \Delta R(\text{obj}, i) < 0.25} p_T^i \cdot \frac{1}{p_T^{\text{obj}}} \quad , \quad (6.3)$$

where the particles within $\Delta R < 0.05$ are excluded from the computation to remove possible spurious neutral particles very close to the considered candidate around which the isolation is built. A photon or a meson candidate is considered isolated if the isolation variable in Eq. 6.3 is lower than a threshold of 0.25. The matching efficiencies for $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ are shown in Fig. 6.2 for both the cases of meson candidates reconstructed using L1 PUPPI or tracker tracks.

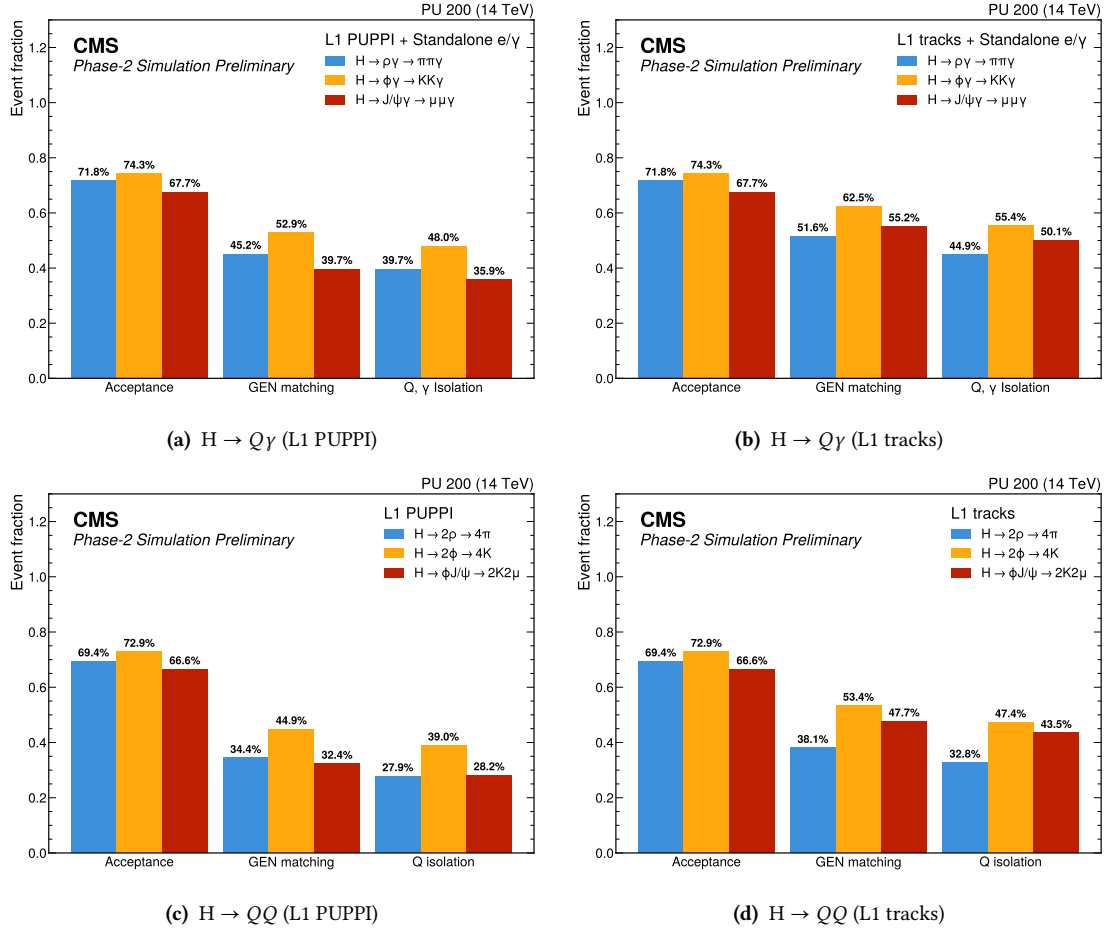


Figure 6.2 • Cut-flow of generator level matching for $H \rightarrow Q\gamma$ (top row) and $H \rightarrow QQ$ (bottom row) decays, where $Q = \rho, \phi, J/\psi$. The results are reported for both L1 PUPPI (left column) and tracker track (right column) candidates used for the meson reconstruction.

Further analysis is needed to assess the impact of photon and meson candidate reconstruction on the H resonance peak resolution. The difference of the candidate Higgs boson invariant mass peak, reconstructed with the L1 information, and the corresponding generator-level mass, considered as reference, is shown in Fig. 6.3a. The differences between the p_T of the L1 photon or L1 reconstructed meson and the generator-level p_T are shown in Figs. 6.3b-6.3c.

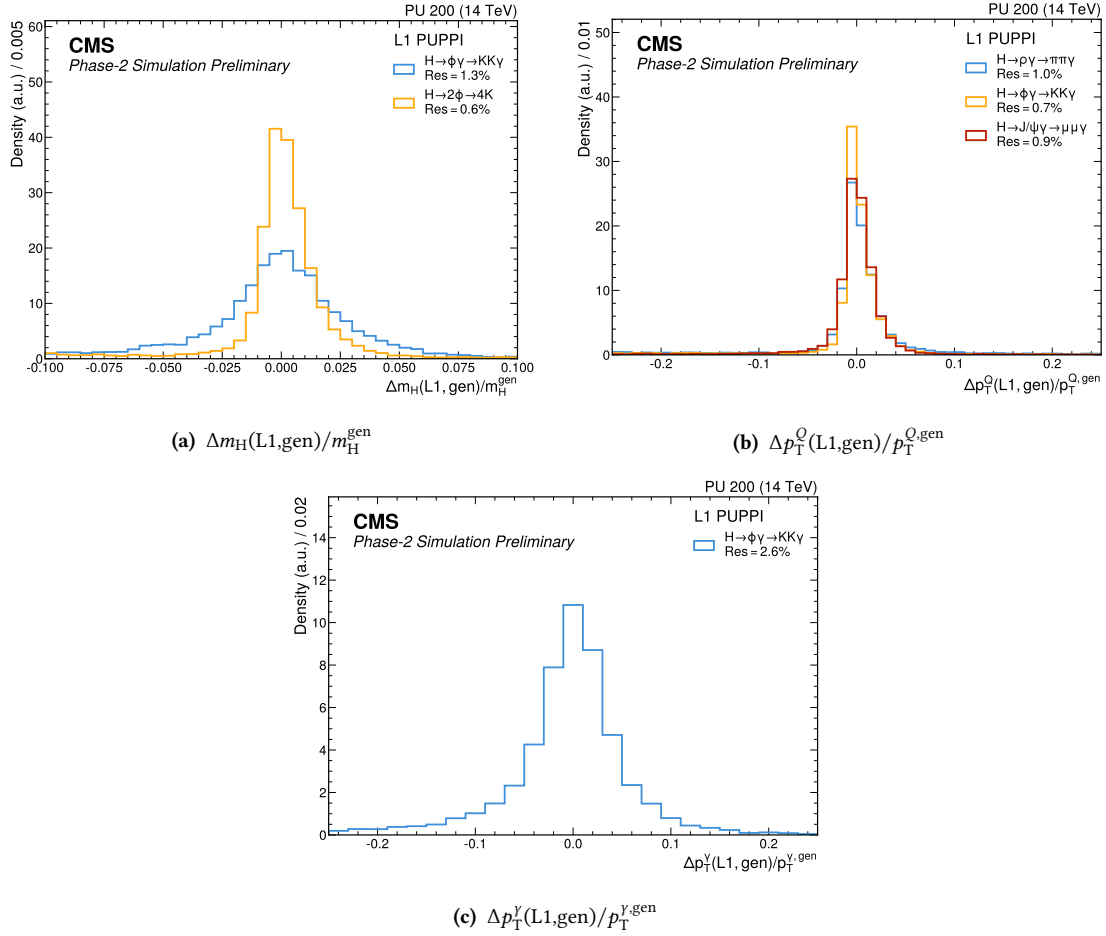


Figure 6.3 • Comparison between L1 reconstruction and generator level quantities for H candidate invariant mass (6.3a), for p_T^Q (6.3b) and for p_T^γ (6.3c). The meson reconstruction is performed with L1 PUPPI candidates.

The results show how the H resolution is driven by the L1 photon rather than the meson reconstruction, thus showing where there is room for improvement with a custom reconstruction in the L1DS, starting from the subdetectors TPs and profiting from relaxed constraints on latency and logic resources with respect to the L1T processors. This observation is confirmed by the Higgs boson mass resolution for the $H \rightarrow 2\phi$ channel, considered as example, where the absence of the photon and the excellent resolution in the reconstruction of the ϕ meson produce a narrower invariant mass peak with respect to the case $H \rightarrow \phi\gamma$. In general, a resolution around the 1% level is observed for the $H \rightarrow QQ$ decay channels and it degrades to around 2% for the $H \rightarrow Q\gamma$

channels, with the resolution quantified using:

$$R = \frac{\sigma_{\text{gaus}}}{\mu_{\text{gaus}}} . \quad (6.4)$$

In Eq. 6.4, σ_{gaus} and μ_{gaus} are the standard deviation and mean, respectively, of the gaussian core of the resonance peak.

6.3 L1 scouting analysis

The event reconstruction with the L1 objects introduced above is performed without the latency and computational resources constraints that need to be taken into account in a trigger algorithm. The study is done in view of deployment inside the online event processing in the Phase-2 L1DS system. A selected set of event records can be saved for offline analysis, where more complex selections and analysis techniques can be applied to further enhance the signal-to-background ratio while keeping a high signal efficiency.

6.3.1 Event selections

The event reconstruction is applied to the signal and background samples in a first preselection filtering step, which keeps enough events to perform the event modelling in Section 6.3.2:

- The L1 PUPPI and tracker track candidates are filtered by keeping only the objects with $p_T > 5$ GeV, reducing the object multiplicity in each event without significant signal efficiency loss.
- For the $H \rightarrow Q\gamma$ decay, the photon objects are preselected by requiring a threshold $p_T > 10$ GeV.
- The meson candidates are reconstructed by building all the pair combinations from the filtered L1 PUPPI and tracker track collections.
- Only the pairs with opposite sign components, $p_T > 10$ GeV and invariant mass compatible with the meson mass are kept. For the mass requirement, the same mass ranges employed for the generator level study are used.
- For the $H \rightarrow Q\gamma$ search, the best meson candidate is selected as the pair with closest components in angular distance ΔR among the set of filtered pairs. For the $H \rightarrow QQ$ search, the same criterion is applied to identify the first meson $Q_1 = \rho, \phi, \phi$, while the second meson $Q_2 = \rho, \phi, J/\psi$ is chosen as the second closest pair in ΔR .
- The H candidate is reconstructed from the highest p_T photon plus meson candidate system for $H \rightarrow Q\gamma$ searches, and from the two mesons system for $H \rightarrow QQ$ searches, and its invariant mass is restricted in the 100-150 GeV range.

The signal and background efficiencies after event preselection are reported in Table 6.1. Notably, the background fraction is reduced of a factor 10^4 when using L1 PUPPI candidates and 10^3 when using tracker tracks. At this step, a significantly higher signal efficiency is observed while using tracker tracks.

Process	Efficiency (L1 PUPPI)		Efficiency (L1 tracks)	
	Signal	Background	Signal	Background
$H \rightarrow \rho\gamma$	4.3×10^{-1}	3.9×10^{-3}	5.1×10^{-1}	7.7×10^{-3}
$H \rightarrow \phi\gamma$	5.1×10^{-1}	2.7×10^{-3}	6.1×10^{-1}	5.3×10^{-3}
$H \rightarrow J/\psi\gamma$	4.0×10^{-1}	5.9×10^{-4}	5.9×10^{-1}	3.5×10^{-3}
$H \rightarrow 2\rho$	3.3×10^{-1}	2.2×10^{-3}	4.6×10^{-1}	7.0×10^{-3}
$H \rightarrow 2\phi$	4.4×10^{-1}	1.0×10^{-3}	5.7×10^{-1}	3.3×10^{-3}
$H \rightarrow \phi J/\psi$	3.2×10^{-1}	2.9×10^{-4}	5.2×10^{-1}	3.4×10^{-3}

Table 6.1 • Signal and background preselection efficiency for $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ searches, where $Q = \rho \rightarrow \pi\pi, \phi \rightarrow KK, J/\psi \rightarrow \mu\mu$, using L1 PUPPI and tracker tracks.

The p_T distribution of the final state candidate kaon and muon tracks, of the photon and of the reconstructed meson after the preselection are shown in Fig. 6.4 for the channel $H \rightarrow Q\gamma$, as their shape is similar for the case of $H \rightarrow QQ$. It is possible to observe how the signal-to-background ratio of the distributions with area normalized to the unit is larger in the high- p_T region, in particular around a value of 35-40 GeV for the photon and candidate meson p_T . A lower p_T threshold of 10-15 GeV is observed for the candidate kaons with respect to pions and muons, for which the signal enriched region starts around 20 GeV. This difference is due to the intrinsic kinematic characteristics of the decay: the K mass is close to half the mass of the ϕ meson. Conversely, products of the ρ and J/ψ mesons have mass significantly lower than half of m_Q .

The results of the meson reconstruction applied in the preselection are shown in Fig. 6.5, where the invariant mass of the best pair of OS candidates is shown for the cases $\rho \rightarrow \pi\pi$, $\phi \rightarrow KK$ and $J/\psi \rightarrow \mu\mu$, taken from the $H \rightarrow Q\gamma$ searches. The distributions highlight the resolution of the resonance peak and they compare how this feature has high discrimination power between signal and background. The resolution of the candidate ϕ meson is significantly higher ($\approx 1 - 2\%$) than the one for the J/ψ ($\approx 3\%$) and the ρ ($\approx 12\%$), for which the natural width of around 145 MeV is dominant. Thus, a better signal-to-background efficiency ratio is expected for the searches involving ϕ reconstruction, translating in better constraints on the upper limit results.

Finer selection criteria are implemented or optimized to define a set of signal region selections for the sensitivity study. The definition is done on the basis of the distribution of the variables with high discrimination power, as discussed above. In particular, the p_T of the selected candidate mesons and photon are requested to be higher than 30 GeV. Additionally, a significant optimization of the signal-to-background ratio comes from the isolation variable for both the candidate meson and the photon. Their distributions after the preselection are shown in Fig. 6.6 for the channel $H \rightarrow \phi\gamma$, where it is possible to observe a large fraction of signal events where the meson and the photon are isolated from hadronic activity.

Given the statistically limited population of the background sample, the application of both meson and photon candidate isolation selections might result in all the events of the sample being filtered out. When present, this limitation is solved by applying only the first isolation requirement and estimating the final efficiency after the second isolation selection as:

$$\epsilon_{\text{bkg}} = \frac{n_{(Q),(\gamma)}}{n_{\text{tot}}} = \frac{n_{(Q),(\cancel{\gamma})} \cdot n_{(\cancel{Q}),(\gamma)}}{n_{(\cancel{Q}),(\cancel{\gamma})}} \cdot \frac{1}{n_{\text{tot}}}, \quad (6.5)$$

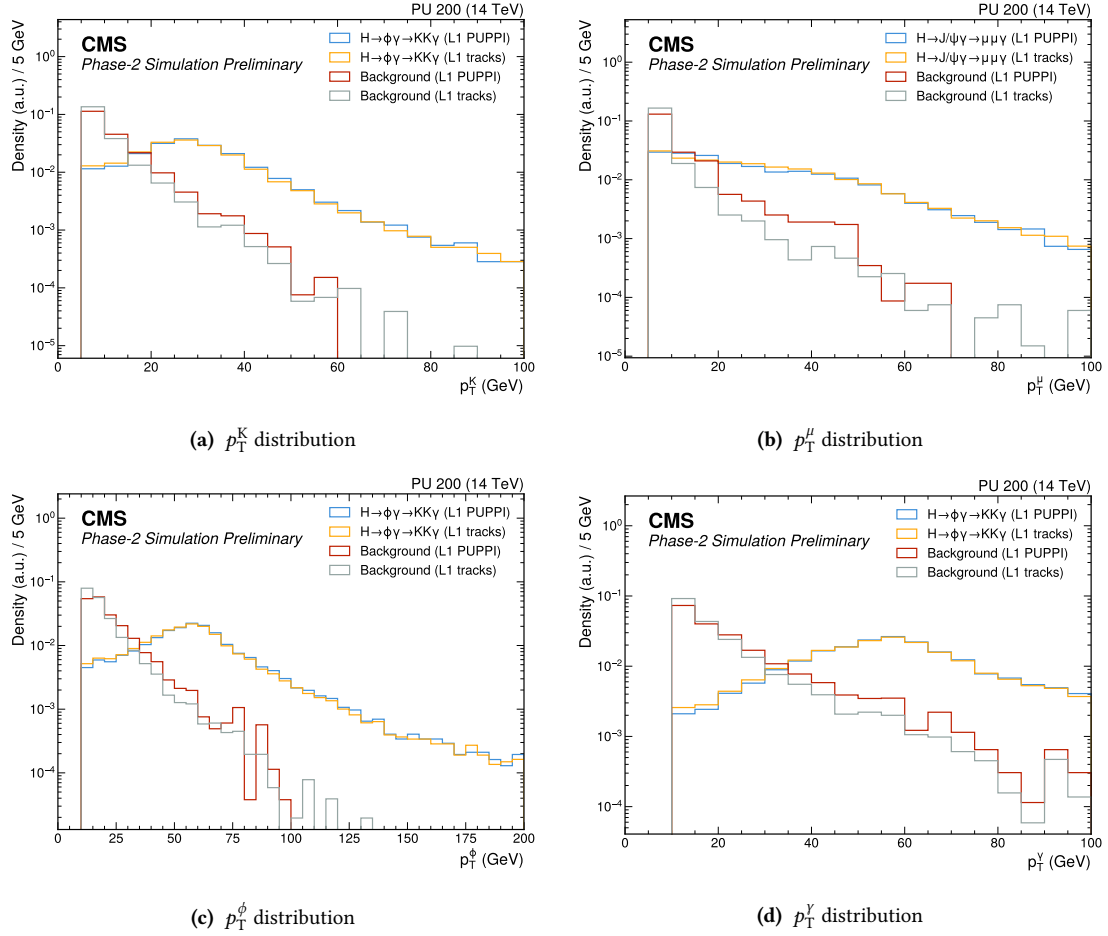


Figure 6.4 • p_T distributions of candidate kaons (6.4a), muons (6.4b), ϕ meson (6.4c) and photon (6.4d). The distributions are shown for both signal and background samples for $H \rightarrow Q\gamma$ channel selections and normalized to the unit for comparison. The results are shown for the cases of L1 PUPPI and tracker tracks used for the reconstruction.

6.3. L1 scouting analysis

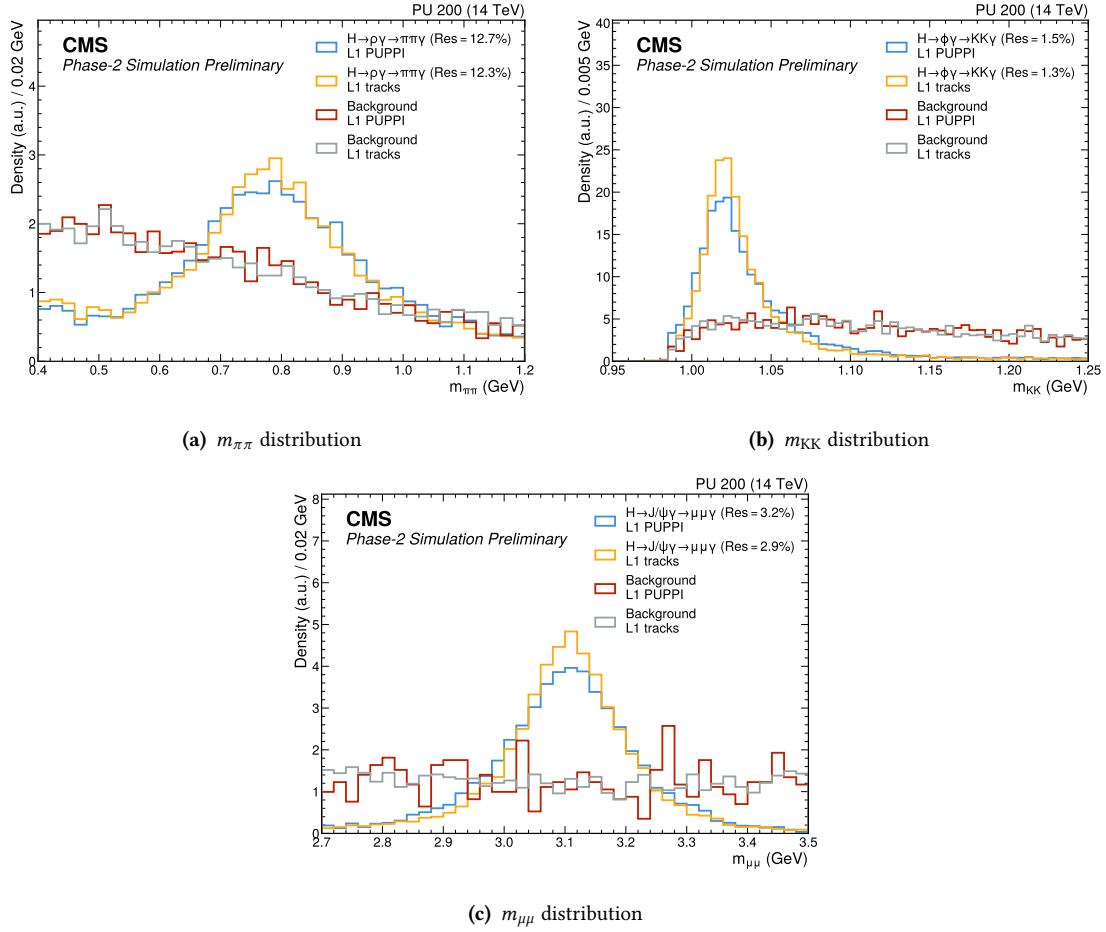


Figure 6.5 • Distribution of candidate ρ (6.5a), ϕ (6.5b) and J/ψ (6.5c) invariant mass, for both signal and background samples in the $H \rightarrow Q\gamma$ channel. The distributions are normalized to the unit for comparison. The reconstruction has been performed using both L1 PUPPI and tracker track candidates.

where the (Q) and (γ) represents the application of the Q and γ isolation requirements, respectively, and (\bar{Q}) and $(\bar{\gamma})$ their negation [172]. The same technique is applied in the case of $H \rightarrow QQ$ decays, where in this case (γ) is replaced by the isolation of the second reconstructed meson candidate. Eq. 6.5 conservatively assumes that the relative efficiency of the second isolation requirement is equal to the first one. Thus, the absolute background efficiency after the application of the first isolation selection is reweighted with this criterion, estimating the final background efficiency of the signal region selections in a more robust and reliable way. The selection efficiencies at each selection step and their final value for the $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ analyses are shown in Tables. 6.2-6.3.

Further selection criteria keeping a high signal efficiency could be applied to enhance the sensitivity. However, the reference background sample is not sufficiently populated to allow further optimization. This limitation affects in particular the $H \rightarrow QQ$ channel where the two m_Q window cuts filter out almost the whole dataset. Thus, this study can be dramatically improved with a larger background simulated sample.

Channel	Selection	Efficiencies (L1 PUPPI)		Efficiencies (L1 tracks)	
		Signal	Background	Signal	Background
$H \rightarrow \rho\gamma$	$p_T^{\text{tracks}} > 10$	8.8×10^{-1}	1.0×10^{-1}	7.5×10^{-1}	1.3×10^{-1}
	$p_T^\gamma > 30$	7.2×10^{-1}	3.4×10^{-3}	6.2×10^{-1}	2.7×10^{-3}
	OS tracks	4.3×10^{-1}	1.5×10^{-3}	5.6×10^{-1}	2.1×10^{-3}
	m_ρ window	3.3×10^{-1}	8.3×10^{-4}	4.1×10^{-1}	1.0×10^{-3}
	$p_T^\rho > 30$ GeV	3.1×10^{-1}	6.5×10^{-4}	3.9×10^{-1}	7.9×10^{-4}
	m_H range	2.5×10^{-1}	6.0×10^{-5}	2.9×10^{-1}	7.6×10^{-5}
	Iso(ρ)	2.4×10^{-1}	1.4×10^{-5}	2.8×10^{-1}	2.4×10^{-5}
	Iso(γ)	2.3×10^{-1}	5.7×10^{-6}	2.8×10^{-1}	8.8×10^{-6}
$H \rightarrow \phi\gamma$	$p_T^{\text{tracks}} > 10$ GeV	8.9×10^{-1}	1.0×10^{-1}	8.4×10^{-1}	1.3×10^{-1}
	$p_T^\gamma > 30$ GeV	7.2×10^{-1}	3.4×10^{-3}	6.9×10^{-1}	2.7×10^{-3}
	OS tracks	5.4×10^{-1}	1.5×10^{-3}	6.6×10^{-1}	2.1×10^{-3}
	m_ϕ window	4.6×10^{-1}	3.9×10^{-4}	5.6×10^{-1}	5.9×10^{-4}
	$p_T^\phi > 30$ GeV	4.4×10^{-1}	2.5×10^{-4}	5.4×10^{-1}	3.7×10^{-4}
	m_H range	3.9×10^{-1}	2.4×10^{-5}	4.6×10^{-1}	3.9×10^{-5}
	Iso(ϕ)	3.7×10^{-1}	9.3×10^{-6}	4.5×10^{-1}	7.2×10^{-6}
	Iso(γ)	3.7×10^{-1}	3.6×10^{-6}	4.4×10^{-1}	3.1×10^{-6}
$H \rightarrow J/\psi\gamma$	$p_T^{\text{tracks}} > 10$ GeV	8.9×10^{-1}	1.0×10^{-1}	7.7×10^{-1}	1.3×10^{-1}
	$p_T^\gamma > 30$ GeV	7.2×10^{-1}	3.4×10^{-3}	6.3×10^{-1}	2.7×10^{-3}
	OS tracks	4.1×10^{-1}	1.5×10^{-3}	5.7×10^{-1}	2.1×10^{-3}
	$m_{J/\psi}$ window	2.6×10^{-1}	6.1×10^{-5}	3.9×10^{-1}	9.3×10^{-5}
	$p_T^{J/\psi} > 30$ GeV	2.5×10^{-1}	4.8×10^{-5}	3.8×10^{-1}	7.3×10^{-5}
	m_H range	2.4×10^{-1}	4.1×10^{-6}	3.6×10^{-1}	5.7×10^{-6}
	Iso(J/ψ)	2.2×10^{-1}	3.1×10^{-6}	3.5×10^{-1}	$< 10^{-7}$
	Iso(γ)	2.2×10^{-1}	1.6×10^{-6}	3.4×10^{-1}	$< 10^{-7}$

Table 6.2 • Signal and background selection efficiencies for $H \rightarrow Q\gamma$ channel, where $Q = \rho \rightarrow \pi\pi$, $\phi \rightarrow KK$, $J/\psi \rightarrow \mu\mu$, using L1 PUPPI and tracker tracks. The listed selections are sequential, from top to bottom.

Channel	Selection	Efficiencies (L1 PUPPI)		Efficiencies (L1 tracks)	
		Signal	Background	Signal	Background
$H \rightarrow 2\rho$	$p_T^{\text{tracks}} > 10 \text{ GeV}$	5.7×10^{-1}	1.7×10^{-2}	5.4×10^{-1}	1.4×10^{-2}
	OS tracks	5.2×10^{-1}	8.9×10^{-3}	5.4×10^{-1}	1.3×10^{-2}
	$Q_{1,2} \text{ mass}$	4.6×10^{-1}	4.8×10^{-3}	5.0×10^{-1}	5.9×10^{-3}
	$p_T^{Q_{1,2}} > 30 \text{ GeV}$	4.5×10^{-1}	3.0×10^{-3}	4.8×10^{-1}	3.4×10^{-3}
	$m_H \text{ range}$	1.5×10^{-1}	8.0×10^{-5}	2.0×10^{-1}	1.2×10^{-4}
	Iso(Q_1)	1.6×10^{-1}	5.8×10^{-5}	2.1×10^{-1}	1.1×10^{-4}
	Iso(Q_2)	1.4×10^{-1}	1.8×10^{-5}	1.9×10^{-1}	3.4×10^{-5}
$H \rightarrow 2\phi$	$p_T^{\text{tracks}} > 10 \text{ GeV}$	6.9×10^{-1}	1.7×10^{-2}	6.9×10^{-1}	1.4×10^{-2}
	OS tracks	6.6×10^{-1}	8.9×10^{-3}	6.9×10^{-1}	1.3×10^{-2}
	$Q_{1,2} \text{ mass}$	6.4×10^{-1}	3.5×10^{-3}	6.8×10^{-1}	4.3×10^{-3}
	$p_T^{Q_{1,2}} > 30 \text{ GeV}$	6.2×10^{-1}	1.9×10^{-3}	6.6×10^{-1}	2.3×10^{-3}
	$m_H \text{ range}$	3.5×10^{-1}	3.9×10^{-5}	4.4×10^{-1}	5.9×10^{-5}
	Iso(Q_1)	3.5×10^{-1}	2.3×10^{-5}	4.4×10^{-1}	4.2×10^{-5}
	Iso(Q_2)	3.3×10^{-1}	7.2×10^{-6}	4.2×10^{-1}	1.1×10^{-5}
$H \rightarrow \phi J/\psi$	$p_T^{\text{tracks}} > 10 \text{ GeV}$	5.8×10^{-1}	1.7×10^{-2}	6.1×10^{-1}	1.4×10^{-2}
	OS tracks	5.5×10^{-1}	8.9×10^{-3}	6.1×10^{-1}	1.3×10^{-2}
	$Q_1 \text{ mass}$	4.5×10^{-1}	3.5×10^{-3}	5.4×10^{-1}	4.3×10^{-3}
	$Q_2 \text{ mass}$	2.2×10^{-1}	9.5×10^{-5}	3.7×10^{-1}	2.1×10^{-4}
	$p_T^{Q_1} > 30 \text{ GeV}$	2.1×10^{-1}	7.4×10^{-5}	3.5×10^{-1}	1.4×10^{-4}
	$p_T^{Q_2} > 30 \text{ GeV}$	2.1×10^{-1}	5.8×10^{-5}	3.4×10^{-1}	1.1×10^{-4}
	$m_H \text{ range}$	2.0×10^{-1}	5.2×10^{-6}	3.2×10^{-1}	1.4×10^{-5}
	Iso(Q_1)	1.9×10^{-1}	9.3×10^{-6}	3.2×10^{-1}	1.8×10^{-5}
	Iso(Q_2)	1.8×10^{-1}	5.2×10^{-7}	3.1×10^{-1}	2.6×10^{-6}

Table 6.3 • Signal and background selection efficiencies for $H \rightarrow QQ$ channel, where $Q = \rho \rightarrow \pi\pi, \phi \rightarrow KK, J/\psi \rightarrow \mu\mu$, using L1 PUPPI and tracker tracks. The listed selections are sequential, from top to bottom.

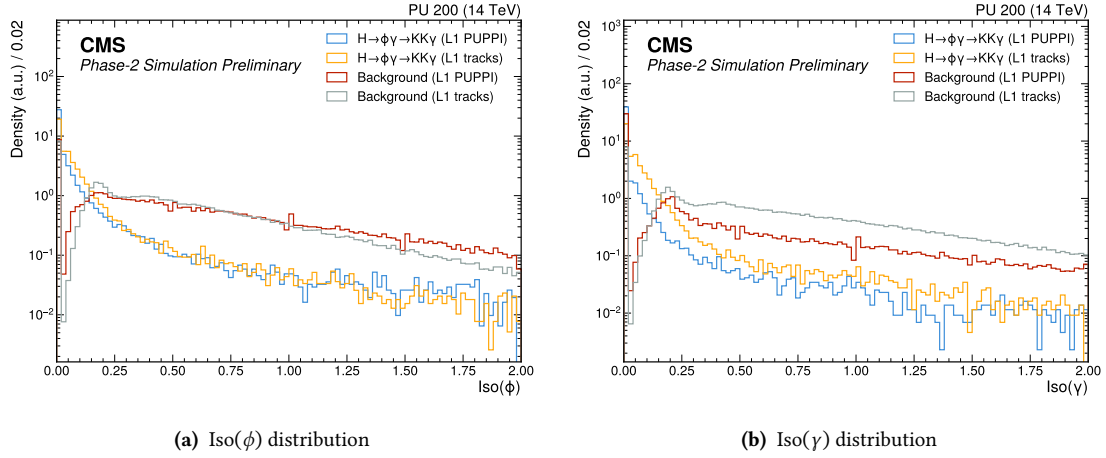


Figure 6.6 • Relative isolation distribution of the reconstructed ϕ (6.6a) and of the highest energy L1 photon (6.6b) in $H \rightarrow \phi\gamma$ search. The distributions are normalized to the unit for comparison. The meson reconstruction and isolation calculation for both the meson and photon candidates have been performed using both L1 PUPPI and tracker track candidates.

6.3.2 Signal and background modelling

The signal contribution is modeled using the resonant DSCB function, presented in Eq. 3.5. The shape of the distribution is obtained from the fit in the range $100 < m_H < 150$ GeV on the signal dataset after the preselection. The signal model normalization is fixed to the final signal efficiency value.

The background contribution is modeled from the minimum bias simulated sample after the preselection, allowing for a significant reduction of statistical fluctuations. A too small amount of events would be available after the signal region selections to allow any reliable modelling of the final state particles invariant mass distribution. This approach conservatively assumes that the signal region selections do not sculpt the m_H distribution in the considered range and that the shape is left unchanged. The chosen functional form of the background distribution is a 2-parameter power-law, as in Eq. 3.8. The power-law function is fitted to the m_H distribution in the considered range and normalized to the expected number of background events after the signal selections from Eq. 6.1. The results of the modelling are shown in Fig. 6.7 just for the case of L1 PUPPI candidates used for the meson reconstruction.

6.3.3 Projected constraints on branching fraction

Expected limits on $\mathcal{B}(H \rightarrow Q\gamma)$ and $\mathcal{B}(H \rightarrow QQ)$ are calculated with the same assumptions as those described in Section 6.1. The limits are extracted with Combine using the signal and background modelling previously described and using the asymptotic approximation in the CL_s method [93]. The expected results are shown in Table 6.4. The upper limits are around 1 order of magnitude less stringent than the ones obtained during Run-2 for the $H \rightarrow Q\gamma$ searches, with better results obtained in general from using L1 tracker tracks for the meson reconstruction. For the case of $H \rightarrow QQ$ searches, no public results are available to date. It should be noted that these are the results from an analysis done uniquely at L1DS level and with a statistically

6.3. L1 scouting analysis

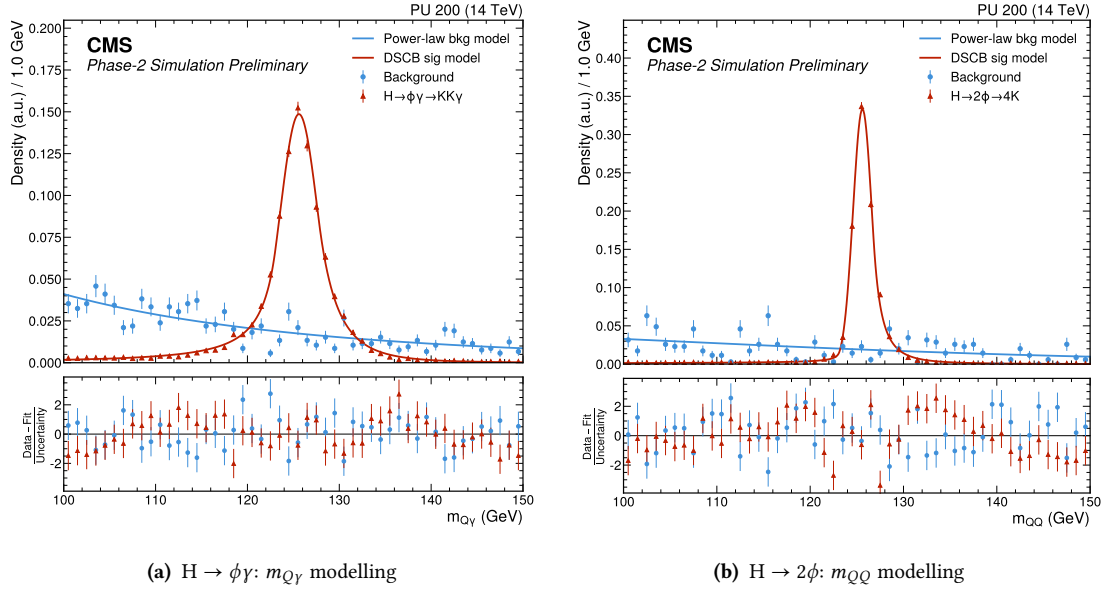


Figure 6.7 • Modelling of $m_{\phi\gamma}$ (6.7a) and $m_{\phi\phi}$ (6.7b) invariant mass distributions after preselection, with background (blue) and signal (red) contributions. The distributions are normalized to the unit and the meson candidate has been reconstructed using L1 PUPPI candidates.

limited background simulated sample, which does not allow for the application of tighter selection criteria, affecting in particular the $H \rightarrow QQ$ searches. The analysis relies on out-of-the-box L1 objects reconstruction performance without further processing or optimization. Moreover, additional handles, such as MVA discriminators, could be used online or offline to further reduce the background contributions and improve the expected limits.

Upper limits on \mathcal{B} at 95% CL			
Process	L1 PUPPI	L1 tracker tracks	Public results
$H \rightarrow \rho\gamma$	1.2×10^{-3}	7.5×10^{-4}	3.74×10^{-4} [31]
$H \rightarrow \phi\gamma$	1.3×10^{-3}	9.8×10^{-4}	2.97×10^{-4} [31]
$H \rightarrow J/\psi\gamma$	6.0×10^{-3}	4.5×10^{-3}	2.0×10^{-4} [32]
$H \rightarrow 2\rho$	9.1×10^{-3}	9.8×10^{-3}	—
$H \rightarrow 2\phi$	5.1×10^{-3}	6.3×10^{-3}	—
$H \rightarrow \phi J/\psi$	1.9×10^{-2}	2.7×10^{-2}	—

Table 6.4 • Upper limit projections on branching fraction for $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ searches, where $Q = \rho \rightarrow \pi\pi, \phi \rightarrow KK, J/\psi \rightarrow \mu\mu$, using L1 PUPPI and tracker tracks. The integrated luminosity of $\mathcal{L}_{\text{int}} = 400 \text{ fb}^{-1}$, corresponding to one year of data taking at the highest luminosity of $7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, is considered for the sensitivity study.

6.4 Online processing timing benchmark

The analyses described in the previous discussion have been optimized for deployment in the online processing system described in Section 5.3 and adapted for the Phase-2 version of the scouting system. After unpacking the needed L1 objects from raw data fragments created by a demonstrator of the Phase-2 L1 trigger readout, the data is given in input to a set of analyzer modules, each implementing an analysis with the same selections described in Section 6.3.1.

The raw data given in input to the online processing infrastructure mimics minimum bias collisions for a realistic benchmark of the time needed to unpack the data, perform the analyses, and save the needed information to retrieve the best triplet $Q\gamma$ or quadruplet QQ and the data from the BXs satisfying the analyses selections. The time needed for each step and analysis is reported in Fig. 6.8 for a single-threaded execution on an AMD EPYC 9654 96-Core CPU.




Figure 6.8 • Timing benchmark per BX of rare Higgs boson decay analyses in Phase-2 L1DS online processing infrastructure, estimated on minimum bias collisions for a single threaded execution on an AMD EPYC 9654 96-Core CPU.

Given an estimate of approximately $\mathcal{O}(1) \mu\text{s}$ to perform the discussed analyses for a single BX, around $\mathcal{O}(30)$ parallel processes could sustain the online processing in realistic HL-LHC collisions. This requirement is well within the specifics of the Phase-2 online processing system, motivated also by a large margin of optimization of the online processing. For instance, the computation of the photon isolation and the reconstruction of a specific set of meson candidates can be performed once, before executing the analyses workflow.

Overall, the timing benchmark demonstrates the feasibility of the Phase-2 L1DS approach, moving closer to the ambitious goal of real-time processing of all collision events. Rare Higgs boson decay analyses can be added to the list of searches that can be performed with the standard L1 trigger menu to extend the physics coverage of the CMS experiment. The similar decays $W, Z \rightarrow Q\gamma$, which are further limited by the L1 threshold on the photon E_T , and $W, Z \rightarrow QQ$ [6] could be searched with a similar analysis strategy as the one presented in this study. The range of analyses is expected to grow, further enriching the Phase-2 L1DS scouting program and demonstrating the full potential of this approach.

Conclusions and outlook

Following the discovery of the Higgs boson by CMS and ATLAS experiments, a series of precision measurements have enhanced our understanding of its properties, particularly its couplings to gauge bosons and third-generation fermions. With the conclusion of LHC Run-2 and the ongoing Run-3 campaign, the focus has shifted towards measuring the couplings to lighter fermions, such as the charm and the strange quarks. Investigating these couplings through rare Higgs boson decays to a meson and a photon, or to a pair of mesons, offers a promising probe due to reduced SM backgrounds, especially compared to the direct searches for more frequent $H \rightarrow q\bar{q}$ processes.

The SM predicts a branching fraction in the range of 10^{-5} to 10^{-10} for decays $H \rightarrow Q\gamma$ or $H \rightarrow QQ$, where Q represents mesons like the ρ , ϕ , $\psi(\text{nS})$ or $Y(\text{nS})$. Thus, searching for these rare processes requires large amounts of data and dedicated trigger strategies, which are particularly challenging to define in cases where mesons decay into pairs of charged hadrons. This thesis presents the search for rare Higgs boson decays to a $\psi(\text{nS})$ meson and a photon, which probes the coupling to the charm quark. The search is facilitated by the subsequent decay of the $\psi(\text{nS})$ to a pair of muons, making it easier to trigger with respect to the cases involving the ρ and ϕ mesons. Additionally, the similar Z boson decays have been explored, as they provide a benchmark for validating theoretical predictions and are likely to be observed first due to a larger expected cross section. The dataset used, comprising 123 fb^{-1} of proton proton collision data at $\sqrt{s} = 13 \text{ TeV}$, was collected by the CMS experiment during Run-2 using single muon plus photon triggers.

No significant excess over the SM expectation was observed in this search and upper limits on the branching fractions of these processes have been set. The previous results from CMS Collaboration have been significantly improved, thanks to a new event selection strategy that categorizes the data based on Higgs boson production modes and decay angular variables, coupled with a more advanced statistical treatment. The limits have been interpreted into constraints on ratio of the coupling modifiers κ_c/κ_γ , set to the interval $(-157, +199)$ at 95% CL. The dataset collected during Run-3 by CMS experiment is expected to further tighten these constraints, as demonstrated by the analysis projection with approximately 300 fb^{-1} of additional proton-proton collision data at 13.6 TeV, expected to be delivered at the end of the campaign.

This analysis, along with other constraints set by the CMS and ATLAS Collaborations, underscores that the current searches for $Z, H \rightarrow Q\gamma$ are statistically limited. Similar limitations apply to decays to a pair of mesons, such as $Z, H \rightarrow 2\rho, 2\phi, \phi J/\psi, 2J/\psi, 2Y$, where the first three decays have never been studied before. Implementing efficient selection algorithms at the first level of the CMS trigger chain is essential for improving, and in some cases enabling, the sensitivity to these processes. Identifying mesons from their hadronic decays and setting a low enough energy threshold for the photon are challenging tasks in L1 trigger. The CMS Phase-2 upgrade for HL-LHC will be essential for advancing the study of these rare decays, introducing the processing of tracks from the tracking system by the L1T system. However, the hardware resources required for implementing meson identification are not negligible. To avoid a loss of signal events due to the trigger selection, the analysis could be performed on all collisions events using solely the L1T

Conclusions and outlook

reconstructed objects. For decays where no muons or photons are available to trigger on, such as $H \rightarrow 2\rho, 2\phi$, this approach may be the only viable solution, enabled by the upgraded L1T objects resolution often comparable to the offline reconstruction level.

In this context, this thesis introduces a novel data taking approach at the first level of the CMS trigger chain, known as Level-1 trigger Data Scouting. The Run-3 demonstrator collects the physics objects reconstructed by the current L1T of Run-3 at the LHC bunch crossing rate, performs real-time analysis on the collected data, and stores a compact event record for subsequent offline analysis. The collected data include high-level objects such as L1 muons and jets, as well as lower-level primitives like the muon segments used in the muon barrel reconstruction algorithm. The commissioning of the L1DS demonstrator system, along with the rationale behind its implementation, have been discussed in detail. The design of the two main components of the L1DS demonstrator has been elaborated upon: an L1T processors readout and concentrator component, based on hardware devices like FPGAs, and an online processing infrastructure, which is software-based and performs the event aggregation and online analysis.

The data collected by the demonstrator during 2024 were validated by searching for well-known SM processes, such as $J/\psi \rightarrow \mu\mu$ and $Z \rightarrow \mu\mu, ee$. Machine learning applications, including a custom reconstruction of barrel muons from the low-level segments without latency and resource constraints, were developed targeting a direct deployment on the scouting readout boards or their use in offline analyses. This effort aimed to enhance the physics reach of the demonstrator and show the potential of the L1DS system as testing ground for ML-based techniques to deploy in the L1T chain. Ultimately, the system has progressed beyond its original demonstration purpose, aiming to perform physics measurements otherwise unattainable with data from the trigger chain or from data scouting at the HLT level. Specifically, dijet searches in the low invariant mass range between 100 and 400 GeV could significantly benefit from the L1DS approach and are currently being explored.

A final projection study was conducted in this thesis to assess the feasibility and sensitivity of the analysis of rare Higgs boson decays to mesons with the Phase-2 L1DS system. Simulations of the CMS Phase-2 detector and L1T were used to study how signals and reducible backgrounds would behave when the events are selected only using the L1T data, under the energy and pileup conditions foreseen for HL-LHC. The L1T objects that will be available in the baseline L1DS system were used for the study. The analyses were repeated including L1 tracker tracks, targeting a potential deployment in an extension of the L1DS baseline system. Thanks to dedicated readout boards with state-of-the-art FPGA devices, it would be possible to aggregate a larger amount of information reconstructed by the L1T and apply more sophisticated processing directly on FPGA, reducing the workload of the subsequent online processing step. The timing benchmarks of the analysis projections demonstrate that real-time analysis in the baseline Phase-2 L1DS system is technically feasible, leaving a large margin for additional exclusive analyses. Although the sensitivity study was limited by the reduced population of the background simulation dataset, it highlights that the analysis feasibility is well within the reach of the L1DS system and represents a synergistic complement to the CMS physics program with standard triggered data.

This thesis, along with three intense years of effort behind it, culminates here, demonstrating the effectiveness of this novel data scouting approach and providing a new tool for searching extremely elusive signatures, otherwise unreachable with the standard trigger selection of the CMS experiment.

List of Figures

1.1	Particle content of the standard model of elementary particles	4
1.2	Main production modes of Higgs boson at the LHC	9
1.3	Higgs boson production cross section as a function of \sqrt{s} or of m_H	9
1.4	Four-lepton invariant mass in the $H \rightarrow ZZ^* \rightarrow 4\ell$ search and combination of results	10
1.5	Higgs boson coupling strength modifiers measurement	11
1.6	Leading-Order Feynman diagrams of Higgs rare decays to $Q\gamma$	12
1.7	Leading-Order Feynman diagrams of Higgs rare decays to a pair of mesons	14
2.1	Schematic view of CERN facilities and LHC accelerator ring	16
2.2	LHC filling scheme overview	17
2.3	Integrated luminosity by CMS experiment during LHC Run-2 and Run-3	18
2.4	Pileup distribution observed by the CMS experiment for each year of data taking	19
2.5	CMS detector overall 3D view	20
2.6	CMS detector octant slice	20
2.7	Silicon tracker of CMS detector	22
2.8	3D models of Electromagnetic and Hadronic Calorimeters of the CMS detector .	23
2.9	Quadrant of the muon system of the CMS detector	23
2.10	Example of application of Particle Flow algorithm for event reconstruction . . .	25
3.1	Distributions $m_{\mu\mu\gamma}$, $m_{\mu\mu}$, $p_T^{\mu_1}$, $p_T^{\mu_2}$, p_T^γ and \mathcal{I}^{μ_1} after preselection	36
3.2	Distribution of $m_{\mu\mu\gamma}$ after the application of signal region selections	37
3.3	Topology of $Z, H \rightarrow \psi(nS)\gamma \rightarrow \mu\mu\gamma$ decay and main angular variables	38
3.4	Distributions of the angular variables $\cos \theta^*$, $\cos \theta_1$, φ_1 after preselection and their correlation with $m_{\mu\mu\gamma}$, $m_{\mu\mu}$	39
3.5	Distributions of $LD_{\text{bkg vs } Z/H \text{ sig}}$ after preselection	40
3.6	Distribution of $m_{\mu\mu\gamma}$ in the control region	41
3.7	Summary of event categorization and number of observed events per category .	42
3.8	Examples of signal and resonant background MC samples $m_{\mu\mu\gamma}$ distribution fit results	43
3.9	Background-only fits for $H, Z \rightarrow \psi(nS)\gamma$ searches (part 1)	46
3.10	Background-only fits for $H, Z \rightarrow \psi(nS)\gamma$ searches (part 2)	47
3.11	Results of the background fit in the Control Region	48
3.12	Pulls distribution of bias studies with discrete profiling method	50
3.13	Efficiency of the $\mu + \gamma$ trigger as a function of p_T and η of the photon and muon legs	52
3.14	Observed and expected exclusion limits on the branching fraction of the $(H, Z) \rightarrow \psi(nS)\gamma$ decays	57
3.15	Likelihood scan over κ_c/κ_γ and over κ_c with $\kappa_\gamma = 1$	58
3.16	Projection to $123+300 \text{ fb}^{-1}$ of likelihood scan over κ_c/κ_γ and over κ_c with $\kappa_\gamma = 1$	59

List of Figures

4.1	Illustration of common BSM signatures in exotic physics searches	62
4.2	Data acquisition strategy with scouting and parking data streams	63
4.3	Level-1 trigger Data Scouting data flow with respect to the standard DAQ and trigger chain	64
4.4	Diagram of the CMS Phase-1 Level-1 trigger system	65
4.5	Scheme of the time-multiplexed architecture of the CMS Phase-1 calorimeter trigger	66
4.6	CMS Phase-1 L1 jets, e/γ , τ_h and E_T^{miss} trigger efficiencies	67
4.7	CMS Phase-1 L1 trigger muon reconstruction efficiencies	69
4.8	Architecture of the CMS Phase-2 Level-1 trigger	71
4.9	Example application of FastHisto algorithm on $t\bar{t}$ simulated events at 14 TeV and 200 pileup	72
4.10	Overview of the PF+PUPPI algorithm	74
4.11	CMS Phase-2 L1 jets, energy sums and τ_h trigger efficiencies	75
4.12	Architecture of the CMS Phase-2 Level-1 trigger Data Scouting system	76
4.13	Baseline architecture of the CMS Phase-2 Level-1 trigger Data Scouting system	78
5.1	Architecture of the CMS Level-1 trigger Data Scouting demonstrator for LHC Run-3	82
5.2	Example of pattern matching scheme in η in TwinMux muon stubs	85
5.3	Scouting VCU128 readout board, mezzanine and crate housing the VCU128 boards	87
5.4	Schematics of FPGA pre-processing pipeline	88
5.5	Schematics of FPGA pre-processing pipeline	89
5.6	Schematics of output stage of FPGA firmware design	90
5.7	Schematic and floorplan of main scouting firmware designs	92
5.8	Scouting Configuration Endpoint (SCONE) software for boards control and monitoring	93
5.9	Grafana monitoring dashboard for Run-3 Level-1 trigger Data Scouting demonstrator	94
5.10	Scouting data acquisition software (SCDAQ)	95
5.11	SCDAQ input and output throughput during normal LHC operations	95
5.12	Scouting Orbit Collection data format	97
5.13	Scouting orbit aggregation and processing schematic	98
5.14	Overview of merging of the chunks produced by the processing units and copy to Tier-0	98
5.15	Online processing selectors rate and throughput	100
5.16	Event display of “ μ -tagged jet” signature with L1 objects	101
5.17	L1 muons, BMTF stubs and jets bunch crossing occupancy in LHC orbit	103
5.18	L1 object multiplicities per colliding bunch crossing	104
5.19	p_T of L1 muons and E_T distribution of L1 e/γ , τ_h and energy sums	105
5.20	E_T distribution of trigger towers versus bunch crossing and comparison with pileup profile	106
5.21	Jet E_T spectrum distribution for data and QCD monte carlo with comparison	107
5.22	2D sector-wheel and station-wheel occupancies for BMTF stubs	108
5.23	Comparison between BMTF and software emulation from stubs collected with the L1DS	108
5.24	Low $m_{\mu\mu}$ invariant mass distribution for dimuon final state	109

5.25	Invariant mass distributions of dilepton events collected by the L1DS	110
5.26	Deep-learning based L1 muons recalibration results	112
5.27	p_T , η and φ resolution of neural network barrel muon stub reconstruction . . .	114
5.28	Summary of dijet searches up to July 2024	116
5.29	Simulation of leptophobic Z' boson decay to a pair of jets	116
6.1	Multiplicity of L1 PUPPI, tracker track and photon candidates in simulated minimum bias events	121
6.2	Cut-flow of generator level matching for $H \rightarrow Q\gamma$ and $H \rightarrow QQ$ decays	122
6.3	m_H , p_T^Q and p_T^γ resolution with respect to generator-level quantities	123
6.4	p_T distributions of reconstructed candidates for $H \rightarrow Q\gamma$ channels	126
6.5	Distributions of candidate mesons invariant mass for the $H \rightarrow Q\gamma$ channels . .	127
6.6	Reconstructed ϕ meson and photon relative isolation distributions	130
6.7	$m_{Q\gamma}$ and m_{QQ} invariant mass distributions modelling after preselection	131
6.8	Timing benchmark of Higgs rare decays analyses on Phase-2 L1DS online processing system	132

Bibliography

- [1] CMS Collaboration, *Search for rare decays of the Z and Higgs bosons to a J/ψ or $\psi(2S)$ meson and a photon in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Submitted to Phys. Lett. B (2024), arXiv:2411.15000. (Cited on pages 3 and 13.)
- [2] 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. (Cited on pages 3 and 15.)
- [3] CMS Collaboration, *Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV*, JHEP **2013**(6), 81 (2013), doi:10.1007/JHEP06(2013)081, arXiv:1303.4571. (Cited on pages 3 and 15.)
- [4] 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. (Cited on pages 3 and 15.)
- [5] CMS Collaboration, *Evidence for Higgs boson decay to a pair of muons*, JHEP **2021**(1) (2021), doi:10.1007/jhep01(2021)148, arXiv:2009.04363. (Cited on page 3.)
- [6] D. d’Enterria and D. V. Le, *Rare and exclusive few-body decays of the Higgs, Z, W bosons, and the top quark*, J. Phys. G (2024), doi:10.1088/1361-6471/ad3c59, arXiv:2312.11211. (Cited on pages 3, 13, 18, 61, and 132.)
- [7] F. Mandl and G. Shaw, *Quantum Field Theory*, A Wiley-Interscience publication. Wiley, ISBN 9780471496830 (2010). (Cited on pages 3 and 5.)
- [8] S. L. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22**, 579 (1961), doi:10.1016/0029-5582(61)90469-2. (Cited on page 3.)
- [9] A. Salam and J. C. Ward, *Electromagnetic and weak interactions*, Phys. Lett. **13**, 168 (1964), doi:10.1016/0031-9163(64)90711-5. (Cited on page 3.)
- [10] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19**, 1264 (1967), doi:10.1103/PhysRevLett.19.1264. (Cited on pages 3 and 5.)
- [11] C. Burgard, *Standard model of physics*, <https://texample.net/tikz/examples/model-physics/>. (Cited on page 4.)
- [12] CMS and ATLAS Collaborations, *Combination of measurements of the top quark mass from data collected by the ATLAS and CMS experiments at $\sqrt{s} = 7$ and 8 TeV*, Phys. Rev. Lett. **132**(26) (2024), doi:10.1103/physrevlett.132.261902, arXiv:2402.08713. (Cited on page 4.)
- [13] Particle Data Group, *Review of Particle Physics*, Phys. Rev. D **110**, 030001 (2024), doi:10.1103/PhysRevD.110.030001. (Cited on pages 5, 8, 9, and 12.)
- [14] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13**, 321 (1964), doi:10.1103/PhysRevLett.13.321. (Cited on page 5.)
- [15] W. N. Cottingham and D. A. Greenwood, *An Introduction to the Standard Model of Particle Physics*, Cambridge University Press, 2 edn., doi:10.1017/CBO9780511791406 (2007). (Cited on page 5.)

Bibliography

- [16] LHC Higgs Cross Section Working Group, *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, CERN Yellow Reports: Monographs, doi:10.23731/CYRM-2017-002 (2018), arXiv:1610.07922. (Cited on pages 7, 9, 11, 17, 31, and 120.)
- [17] CMS Collaboration, *Measurement of the Higgs boson mass and width using the four-lepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Submitted to Phys. Rev. D (2024), arXiv:2409.13663. (Cited on page 10.)
- [18] ATLAS Collaboration, *Evidence of off-shell Higgs boson production from ZZ leptonic decay channels and constraints on its total width with the ATLAS detector*, Phys. Lett. B **846**, 138223 (2023), doi:10.1016/j.physletb.2023.138223, arXiv:2304.01532. (Cited on page 10.)
- [19] G. Ortona, *The Higgs boson couplings: past, present, and future. The relationships between Higgs boson and other known particles as measured by current and future experiments*, Frontiers in Physics **11** (2023), doi:10.3389/fphy.2023.1230737. (Cited on page 10.)
- [20] LHC Higgs Cross Section Working Group, *LHC HXSWG interim recommendations to explore the coupling structure of a Higgs-like particle* (2012), arXiv:1209.0040. (Cited on page 11.)
- [21] CMS Collaboration, *A portrait of the Higgs boson by the CMS experiment ten years after the discovery*, Nature **607**(7917), 60 (2022), doi:10.1038/s41586-022-04892-x, arXiv:2207.00043. (Cited on page 11.)
- [22] ATLAS Collaboration, *A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery*, Nature **607**(7917), 52 (2022), doi:10.1038/s41586-022-04893-w, arXiv:2207.00092. (Cited on page 11.)
- [23] CMS Collaboration, *Search for Higgs boson decay to a charm quark-antiquark pair in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. Lett. **131**, 061801 (2023), doi:10.1103/PhysRevLett.131.061801. (Cited on pages 11 and 29.)
- [24] CMS Collaboration, *Search for Higgs boson and observation of Z boson through their decay into a charm quark-antiquark pair in boosted topologies in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. Lett. **131**, 041801 (2023), doi:10.1103/PhysRevLett.131.041801, arXiv:2211.14181. (Cited on pages 11 and 29.)
- [25] CMS Collaboration, *Observation of $WW\gamma$ production and search for $H\gamma$ production in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. Lett. **132**(12), 121901 (2024), doi:10.1103/PhysRevLett.132.121901, arXiv:2310.05164. (Cited on page 11.)
- [26] G. Bodwin, F. Petriello, S. Stoynev and M. Velasco, *Higgs boson decays to quarkonia and the $H\bar{c}c$ coupling*, Phys. Rev. D **88**, 053003 (2013), doi:10.1103/PhysRevD.88.053003, arXiv:1306.5770. (Cited on pages 12 and 13.)
- [27] G. T. Bodwin, H. S. Chung, J.-H. Ee and J. Lee, *Z-boson decays to a vector quarkonium plus a photon*, Phys. Rev. D **97** (2018), doi:10.1103/PhysRevD.97.016009, arXiv:1709.09320. (Cited on pages 13, 29, and 30.)
- [28] G. T. Bodwin, H. S. Chung, J.-H. Ee and J. Lee, *Addendum: New approach to the resummation of logarithms in Higgs-boson decays to a vector quarkonium plus a photon [Phys. Rev. D 95, 054018 (2017)]*, Phys. Rev. D **96** (2017), doi:10.1103/PhysRevD.96.116014, arXiv:1710.09872. (Cited on page 13.)
- [29] N. Brambilla, H. S. Chung, W. K. Lai, V. Shtabovenko and A. Vairo, *Order v^4 corrections to Higgs boson decay into $J/\psi + \gamma$* , Phys. Rev. D **100** (2019), doi:10.1103/PhysRevD.100.054038, arXiv:1907.06473. (Cited on pages 13 and 57.)

- [30] M. Koenig and M. Neubert, *Exclusive radiative Higgs decays as probes of light-quark Yukawa couplings*, JHEP **2015**(8) (2015), doi:10.1007/jhep08(2015)012, arXiv:1505.03870. (Cited on pages 13 and 30.)
- [31] CMS Collaboration, *Search for the Higgs boson decays to a ρ^0 , ϕ , or K^{*0} meson and a photon in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Submitted to Phys. Lett. B (2024), arXiv:2410.18289. (Cited on pages 13 and 131.)
- [32] ATLAS Collaboration, *Searches for exclusive Higgs and Z boson decays into a vector quarkonium state and a photon using 139 fb^{-1} of ATLAS $\sqrt{s} = 13$ TeV proton-proton collision data*, Eur. Phys. J. C **83**(9) (2023), doi:10.1140/epjc/s10052-023-11869-1, arXiv:2208.03122. (Cited on pages 13, 29, 30, and 131.)
- [33] V. Kartvelishvili, A. V. Luchinsky and A. A. Novoselov, *Double vector quarkonia production in exclusive Higgs boson decays*, Phys. Rev. D **79**(11) (2009), doi:10.1103/physrevd.79.114015, arXiv:0810.0953. (Cited on page 13.)
- [34] R. N. Faustov, A. P. Martynenko and F. A. Martynenko, *Relativistic corrections to paired production of charmonium and bottomonium in decays of the Higgs boson*, Phys. Rev. D **107**(5) (2023), doi:10.1103/physrevd.107.056002, arXiv:2209.12321. (Cited on page 13.)
- [35] D.-N. Gao and X. Gong, *Higgs boson decays into a pair of heavy vector quarkonia*, Phys. Lett. B **832**, 137243 (2022), doi:10.1016/j.physletb.2022.137243, arXiv:2203.00514. (Cited on page 13.)
- [36] R. N. Faustov, A. P. Martynenko and F. A. Martynenko, *Pair quarkonium production in Higgs boson decay*, Phys. Part. Nucl. Lett. **20**(3), 368 (2023), doi:10.1134/S1547477123030287. (Cited on page 13.)
- [37] I. N. Belov, A. V. Berezhnoy, E. A. Leshchenko and A. K. Likhoded, *QCD one-loop correction to Higgs boson decay into quarkonium pairs*, Phys. Rev. D **108**(3) (2023), doi:10.1103/physrevd.108.036013, arXiv:2304.11620. (Cited on page 13.)
- [38] CMS Collaboration, *Search for Higgs boson decays into Z and J/ψ and for Higgs and Z boson decays into J/ψ or Y pairs in pp collisions at $\sqrt{s} = 13$ TeV*, Phys. Lett. B **842**, 137534 (2023), doi:10.1016/j.physletb.2022.137534, arXiv:2206.03525. (Cited on pages 14 and 61.)
- [39] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole and P. Proudlock, *LHC Design Report*, CERN Yellow Reports: Monographs, doi:10.5170/CERN-2004-003-V-1 (2004). (Cited on pages 15 and 16.)
- [40] I. Zurbano Fernandez *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN Yellow Reports: Monographs. CERN, Geneva, doi:10.23731/CYRM-2020-0010 (2020). (Cited on pages 15 and 16.)
- [41] E. Lopienska, *The CERN accelerator complex, layout in 2022* (2022), <https://cds.cern.ch/record/2800984>. (Cited on page 16.)
- [42] CMS collaboration, *Physics Results: Public CMS Luminosity Information*, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults>. (Cited on pages 18 and 19.)
- [43] T. Sakuma and T. McCauley, *Detector and Event Visualization with SketchUp at the CMS Experiment*, J. Phys.: Conference Series **513**(2), 022032 (2014), doi:10.1088/1742-6596/513/2/022032. (Cited on pages 20, 22, and 23.)
- [44] S. R. Davis, *Interactive Slice of the CMS detector* (2016), <https://cds.cern.ch/record/2205172>. (Cited on page 20.)
- [45] CMS Tracker Collaboration, *The CMS Phase-1 Pixel Detector Upgrade*, JINST **16**, P02027 (2021), doi:10.1088/1748-0221/16/02/P02027, arXiv:2012.14304. (Cited on page 21.)

Bibliography

- [46] CMS Tracker Collaboration, *Stand-alone cosmic muon reconstruction before installation of the CMS silicon strip tracker*, JINST **4**(05), P05004 (2009), doi:10.1088/1748-0221/4/05/p05004, arXiv:0902.1860. (Cited on page 22.)
- [47] CMS Collaboration, *CMS Technical Design Report for the Muon Endcap GEM Upgrade*, Tech. rep., <https://cds.cern.ch/record/2021453> (2015). (Cited on page 24.)
- [48] CMS Collaboration, *Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JINST **15**, P10017 (2020), doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165. (Cited on pages 24, 66, 67, and 105.)
- [49] CMS Collaboration, *CMS The TriDAS Project: Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger. CMS trigger and data-acquisition project*, Technical design report. CERN, <http://cds.cern.ch/record/578006> (2002). (Cited on page 24.)
- [50] CMS Collaboration, *The CMS trigger system*, JINST **12**, P01020 (2017), doi:10.1088/1748-0221/12/01/P01020, arXiv:1609.02366. (Cited on pages 24 and 64.)
- [51] CMS Collaboration, *Performance of the cms high-level trigger during lhc run 2*, JINST **19**(11), P11021 (2024), doi:10.1088/1748-0221/19/11/p11021, arXiv:2410.17038. (Cited on page 24.)
- [52] CMS Collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, JINST **12**, P10003 (2017), doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965. (Cited on pages 25 and 33.)
- [53] CMS collaboration, *Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus, and MET*, CMS-PAS-PFT-09-001 (2009), <https://cds.cern.ch/record/1194487>. (Cited on page 25.)
- [54] CMS collaboration, *Commissioning of the Particle-flow Event Reconstruction with the first LHC collisions recorded in the CMS detector*, CMS-PAS-PFT-10-001 (2010), <http://cds.cern.ch/record/1247373>. (Cited on page 25.)
- [55] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, Nucl. Instrum. Meth. A **262**(2), 444 (1987), doi:10.1016/0168-9002(87)90887-4. (Cited on pages 25 and 68.)
- [56] W. Adam, B. Mangano, T. Speer and T. Todorov, *Track Reconstruction in the CMS tracker*, CMS-NOTE-2006-041 (2006), <https://cds.cern.ch/record/934067>. (Cited on page 25.)
- [57] CMS Collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, JINST **9**, P10009 (2014), doi:10.1088/1748-0221/9/10/P10009, arXiv:1405.6569. (Cited on page 26.)
- [58] CMS collaboration, *Track impact parameter resolution for the full pseudo rapidity coverage in the 2017 dataset with the CMS Phase-1 Pixel detector*, CMS-DP-2020-049 (2020), <https://cds.cern.ch/record/2743740>. (Cited on page 26.)
- [59] CMS Collaboration, *Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV*, JINST **13**, P06015 (2018), doi:10.1088/1748-0221/13/06/P06015, arXiv:1804.04528. (Cited on pages 26 and 33.)
- [60] W. Adam, R. Frühwirth, A. Strandlie and T. Todorov, *Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at the LHC*, J. Phys. G **31**(9), N9 (2005), doi:10.1088/0954-3899/31/9/n01, arXiv:physics/0306087. (Cited on page 27.)
- [61] CMS Collaboration, *Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC*, JINST **16**, P05014 (2021), doi:10.1088/1748-0221/16/05/P05014, arXiv:2012.06888. (Cited on pages 27 and 35.)

- [62] CMS Collaboration, *ECAL 2016 refined calibration and Run-2 summary plots*, CMS-DP-2020-021 (2020), <https://cds.cern.ch/record/2717925>. (Cited on page 27.)
- [63] CMS Collaboration, *Performance of photon reconstruction and identification with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV*, JINST **10**, P08010 (2015), doi:10.1088/1748-0221/10/08/P08010, arXiv:1502.02702. (Cited on pages 27 and 53.)
- [64] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP **04**, 063 (2008), doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189. (Cited on pages 27 and 73.)
- [65] M. Cacciari, G. P. Salam and G. Soyez, *FastJet user manual*, Eur. Phys. J. C **72**, 1896 (2012), doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097. (Cited on page 27.)
- [66] CMS Collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, JINST **12**, P02014 (2017), doi:10.1088/1748-0221/12/02/P02014, arXiv:1607.03663. (Cited on page 27.)
- [67] CMS Collaboration, *Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV*, JINST **13**, P10005 (2018), doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816. (Cited on page 27.)
- [68] CMS Collaboration, *Identification of hadronic tau lepton decays using a deep neural network*, JINST **17**, P07023 (2022), doi:10.1088/1748-0221/17/07/P07023, arXiv:2201.08458. (Cited on page 27.)
- [69] ATLAS collaboration, *Searches for exclusive Higgs and Z boson decays into $J/\psi\gamma$, $\psi(2S)\gamma$, and $Y(nS)\gamma$ at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Lett. B **786**, 134 (2018), doi:10.1016/j.physletb.2018.09.024, arXiv:1807.00802. (Cited on pages 29 and 30.)
- [70] CMS collaboration, *Search for rare decays of Z and Higgs bosons to J/ψ and a photon in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79**(2) (2019), doi:10.1140/epjc/s10052-019-6562-5, arXiv:1810.10056. (Cited on pages 29, 30, 56, and 58.)
- [71] 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 **2014**(7) (2014), doi:10.1007/jhep07(2014)079. (Cited on pages 30 and 31.)
- [72] P. Faccioli and C. Lourenço, *Particle polarization in high energy physics: an introduction and case studies on vector particle production at the LHC*, Lecture Notes in Physics. Springer, doi:10.1007/978-3-031-08876-6 (2022). (Cited on page 30.)
- [73] Y. Li and F. Petriello, *Combining QCD and electroweak corrections to dilepton production in FEWZ*, Phys. Rev. D **86**, 094034 (2012), doi:10.1103/PhysRevD.86.094034, arXiv:1208.5967. (Cited on page 31.)
- [74] CMS Collaboration, *Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements*, Eur. Phys. J. C **80**, 4 (2020), doi:10.1140/epjc/s10052-019-7499-4, arXiv:1903.12179. (Cited on page 31.)
- [75] NNPDF Collaboration, *Parton distributions with QED corrections*, Nucl. Phys. B **877**, 290 (2013), doi:10.1016/j.nuclphysb.2013.10.010, arXiv:1308.0598. (Cited on page 31.)
- [76] NNPDF Collaboration, *Unbiased global determination of parton distributions and their uncertainties at NNLO and at LO*, Nucl. Phys. B **855**, 153 (2012), doi:10.1016/j.nuclphysb.2011.09.024, arXiv:1107.2652. (Cited on page 31.)
- [77] S. Agostinelli *et al.*, *GEANT4—a simulation toolkit*, Nucl. Instrum. Meth. A **506**, 250 (2003), doi:10.1016/S0168-9002(03)01368-8. (Cited on page 31.)

Bibliography

- [78] J. D’Hondt, P. Vanlaer, R. Frühwirth and W. Waltenberger, *Sensitivity of Robust Vertex Fitting Algorithms*, CMS-NOTE-2004-002 (2004), <https://cds.cern.ch/record/787499>. (Cited on page 32.)
- [79] T. Speer, K. Prokofiev, R. Frühwirth, W. Waltenberger and P. Vanlaer, *Vertex Fitting in the CMS Tracker*, CMS-NOTE-2006-032 (2006), <https://cds.cern.ch/record/927395>. (Cited on page 32.)
- [80] CMS Collaboration, *Rochester corrections for muon momentum scale and resolution*, <https://twiki.cern.ch/twiki/bin/viewauth/CMS/RochcorMuon>. (Cited on page 33.)
- [81] J. Therhaag, *TMVA Toolkit for multivariate data analysis in ROOT*, PoS **ICHEP2010**, 510 (2010), doi:10.22323/1.120.0510. (Cited on page 35.)
- [82] CMS Collaboration, *Pileup Removal Algorithms*, CMS-PAS-JME-14-001 (2014), <http://cds.cern.ch/record/1751454>. (Cited on page 35.)
- [83] CMS Collaboration, *Jet algorithms performance in 13 TeV data*, CMS-PAS-JME-16-003 (2017), <http://cds.cern.ch/record/2256875>. (Cited on pages 35 and 53.)
- [84] CMS Collaboration, *Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV*, JINST **13**(05), P05011 (2018), doi:10.1088/1748-0221/13/05/p05011, arXiv:1712.07158. (Cited on page 35.)
- [85] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban and D. Whiteson, *Jet flavor classification in high-energy physics with deep neural networks*, Phys. Rev. D **94**(11) (2016), doi:10.1103/physrevd.94.112002. (Cited on page 35.)
- [86] CMS Collaboration, *Identification of b quark jets at the CMS Experiment in the LHC Run 2*, CMS-PAS-BTV-15-001 (2016), <http://cds.cern.ch/record/2138504>. (Cited on pages 35 and 54.)
- [87] CMS Collaboration, *Methods to apply b-tagging efficiency scale factors*, <https://twiki.cern.ch/twiki/bin/view/CMS/BTagSFFMethods>. (Cited on pages 35 and 54.)
- [88] CMS Collaboration, *Performance of the DeepJet b tagging algorithm using 41.9 fb⁻¹ of data from proton-proton collisions at 13 TeV with Phase 1 CMS detector*, CMS-DP-2018-058 (2018), <http://cds.cern.ch/record/2646773>. (Cited on page 35.)
- [89] W. Verkerke and D. P. Kirkby, *The RooFit toolkit for data modeling*, eConf **C0303241**, MOLT007 (2003), arXiv:physics/0306116. (Cited on page 42.)
- [90] M. J. Oreglia, *A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$* , SLAC Report SLAC-R-236, <http://www.slac.stanford.edu/cgi-wrap/getdoc/slac-r-236.pdf> (1980). (Cited on page 42.)
- [91] P. Dauncey, M. Kenzie, N. Wardle and G. Davies, *Handling uncertainties in background shapes: the discrete profiling method*, JINST **10**(04), P04015 (2015), doi:10.1088/1748-0221/10/04/p04015, arXiv:1408.6865. (Cited on page 49.)
- [92] CMS Collaboration, *A measurement of the Higgs boson mass in the diphoton decay channel*, Phys. Lett. B **805**, 135425 (2020), doi:10.1016/j.physletb.2020.135425, arXiv:2002.06398. (Cited on page 49.)
- [93] CMS Collaboration, *The CMS statistical analysis and combination tool: COMBINE*, Comp. Softw. Big Sci. **8**(1) (2024), doi:10.1007/s41781-024-00121-4, arXiv:2404.06614. (Cited on pages 49 and 130.)
- [94] J. Bendavid, *Implication of bias threshold in pull definition*, <https://indico.cern.ch/event/255493/contributions/1584346/attachments/447909/621085/bias-Jun26.pdf>. (Cited on page 50.)
- [95] CMS Collaboration, *Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS*, Eur. Phys. J. C **81**, 800 (2021), doi:10.1140/epjc/s10052-021-09538-2, arXiv:2104.01927. (Cited on page 50.)

- [96] CMS Collaboration, *CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13\text{TeV}$* , CMS-PAS-LUM-17-004 (2018), <http://cds.cern.ch/record/2621960>. (Cited on page 50.)
- [97] CMS Collaboration, *CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13\text{TeV}$* , CMS-PAS-LUM-18-002 (2019), <http://cds.cern.ch/record/2676164>. (Cited on page 50.)
- [98] CMS Collaboration, *Scale factors for muon objects for Run-2 UL campaign*, <https://gitlab.cern.ch/cms-muonPOG/muonefficiencies>. (Cited on page 51.)
- [99] CMS Collaboration, *Scale factors for electron and photon objects for Run-2 UL campaign*, <https://twiki.cern.ch/twiki/bin/view/CMS/EgammaUL2016To2018>. (Cited on page 51.)
- [100] CMS collaboration, *Utilities for Accessing Pileup Information for Data*, <https://twiki.cern.ch/twiki/bin/viewauth/CMS/PileupJSONFileforData>. (Cited on page 51.)
- [101] CMS Collaboration, *Common tools for analyses of Higgs boson decay in the diphoton final state*, URL http://cms.cern.ch/iCMS/jsp/openfile.jsp?tp=draft&files=AN2019_149_v5.pdf, AN-2019/149. (Cited on page 53.)
- [102] CMS and ATLAS Collaborations, LHC Higgs Combination Group, *Procedure for the LHC Higgs boson search combination in Summer 2011*, CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11 (2011), <https://cds.cern.ch/record/1379837>. (Cited on page 56.)
- [103] A. L. Read, *Presentation of search results: The CL_s technique*, J. Phys. G **28**, 2693 (2002), doi:10.1088/0954-3899/28/10/313. (Cited on page 56.)
- [104] T. Junk, *Confidence level computation for combining searches with small statistics*, Nucl. Instrum. Meth. A **434**, 435 (1999), doi:10.1016/S0168-9002(99)00498-2, arXiv:hep-ex/9902006. (Cited on page 56.)
- [105] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C **71**, 1554 (2011), doi:10.1140/epjc/s10052-011-1554-0, [Erratum: doi:10.1140/epjc/s10052-013-2501-z], arXiv:1007.1727. (Cited on page 56.)
- [106] ATLAS Collaboration, *Searches for exclusive Higgs and Z boson decays into a vector quarkonium state and a photon using 139fb^{-1} of ATLAS $\sqrt{s} = 13\text{TeV}$ proton–proton collision data*, Eur. Phys. J. C **83**, 781 (2023), doi:10.1140/epjc/s10052-023-11869-1, arXiv:2208.03122. (Cited on pages 56 and 57.)
- [107] LHC Higgs Cross Section Working Group, *Handbook of LHC Higgs cross sections: 3. Higgs properties*, CERN-2013-004 (2013), doi:10.5170/CERN-2013-004, arXiv:1307.1347. (Cited on page 56.)
- [108] CMS Collaboration, *Measurements of Higgs boson production cross sections and couplings in the diphoton decay channel at $\sqrt{s} = 13\text{TeV}$* , JHEP **07**, 027 (2021), doi:10.1007/JHEP07(2021)027, arXiv:2103.06956. (Cited on page 57.)
- [109] G. Perez, Y. Soreq, E. Stamou and K. Tobioka, *Constraining the charm Yukawa and Higgs-quark coupling universality*, Phys. Rev. D **92**, 033016 (2015), doi:10.1103/PhysRevD.92.033016, arXiv:1503.00290. (Cited on page 57.)
- [110] CMS Collaboration, *Displaced BMTF Efficiency Using 2023 Data*, CMS-DP-2023-056 (2023), <https://cds.cern.ch/record/2868797>. (Cited on page 62.)
- [111] CMS Collaboration, *Level-1 Trigger Algorithm for Long-lived Particle Jets in Run-3*, CMS-DP-2024-058 (2024), <https://cds.cern.ch/record/2904694>. (Cited on pages 62, 66, and 67.)
- [112] CMS Collaboration, *Enriching the physics program of the CMS experiment via data scouting and data parking*, Submitted to Phys. Rep. (2024), arXiv:2403.16134. (Cited on pages 63 and 76.)

Bibliography

- [113] CMS Collaboration, *Observation of the rare decay of the η meson to four muons*, Phys. Rev. Lett. **131**(9) (2023), doi:10.1103/physrevlett.131.091903, arXiv:2305.04904. (Cited on page 63.)
- [114] CMS Collaboration, *Development of the CMS detector for the CERN LHC Run 3*, JINST **19**(05) (2024), doi:10.1088/1748-0221/19/05/p05064, arXiv:2309.05466. (Cited on pages 64, 69, and 82.)
- [115] CMS Collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CMS-TDR-021 (2020), <https://cds.cern.ch/record/2714892>. (Cited on pages 64, 70, 71, 72, 74, 75, 76, 119, and 120.)
- [116] G. Hall, D. Newbol, M. Pesaresi and A. Rose, *A time-multiplexed track-trigger architecture for CMS*, JINST **9**(10), C10034 (2014), doi:10.1088/1748-0221/9/10/C10034. (Cited on page 65.)
- [117] R. Frazier, S. Fayer, G. Hall, C. Hunt, G. Iles, D. Newbold and A. Rose, *A demonstration of a Time Multiplexed Trigger for the CMS experiment*, JINST **7**(01), C01060 (2012), doi:10.1088/1748-0221/7/01/C01060. (Cited on page 65.)
- [118] CMS Collaboration, *Performances of L1 Jets and MET Trigger in early Run-3*, CMS-DP-2023-054 (2023), <https://cds.cern.ch/record/2868796>. (Cited on pages 66 and 67.)
- [119] J. Motta, *Development of machine learning based τ trigger algorithms and search for Higgs boson pair production in the $b\bar{b}\tau\tau$ decay channel with the CMS detector at the LHC*, <https://cds.cern.ch/record/2881939> (2023). (Cited on page 67.)
- [120] CMS Collaboration, *Performance of Level-1 trigger e/γ and tau in Run-3*, CMS-DP-2023-008 (2023), <https://cds.cern.ch/record/2853676>. (Cited on page 67.)
- [121] D. Rabady, *Upgrade of the Global Muon Trigger for the Compact Muon Solenoid experiment at CERN*, <https://cds.cern.ch/record/2315346> (2018). (Cited on page 68.)
- [122] S. Mallios, *The CMS barrel track finder and upgrades for HL-LHC*, <https://oldsite.physics.uoi.gr/sites/default/files/phds/Mallios-PHD-uo-i-2019.pdf> (2019). (Cited on pages 68 and 85.)
- [123] M. Bachtis *et al.*, *Upgrade of the CMS Barrel Muon Track Finder for HL-LHC featuring a Kalman Filter algorithm and an ATCA Host Processor with Ultrascale+ FPGAs*, PoS **TWEPP2018**, 139 (2019), doi:10.22323/1.343.0139. (Cited on page 68.)
- [124] W. Zabolotny and A. Byszuk, *Algorithm and implementation of muon trigger and data transmission system for barrel-endcap overlap region of the CMS detector*, JINST **11**(03), C03004 (2016), doi:10.1088/1748-0221/11/03/C03004. (Cited on page 69.)
- [125] D. Acosta, A. Brinkerhoff, E. Busch, A. Carnes, I. Furic, S. Gleyzer, K. Kotov, J. F. Low, A. Madorsky, J. Rorie, B. Scurlock, W. Shi *et al.*, *Boosted Decision Trees in the Level-1 Muon Endcap Trigger at CMS*, J. Phys.: Conference Series **1085**(4), 042042 (2018), doi:10.1088/1742-6596/1085/4/042042. (Cited on page 69.)
- [126] CMS Collaboration, *Performances of Muons, Jets and MET Level 1 trigger algorithms in Run 3*, CMS-DP-2023-007 (2023), <https://cds.cern.ch/record/2853675>. (Cited on page 69.)
- [127] A. Bocci, P. Demin, R. Ranieri and S. De Visscher, *Tagging b jets with electrons and muons at CMS*, CMS-NOTE-2006-043 (2006), <https://cds.cern.ch/record/951390>. (Cited on page 70.)
- [128] E. Govorkova *et al.*, *Autoencoders on field-programmable gate arrays for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider*, Nat. Mach. Intell. **4**(2), 154 (2022), doi:10.1038/s42256-022-00441-3, arXiv:2108.03986. (Cited on page 70.)
- [129] CMS Collaboration, *Anomaly Detection in the CMS Global Trigger Test Crate for Run 3*, CMS-DP-2023-079 (2023), <https://cds.cern.ch/record/2876546>. (Cited on page 70.)
- [130] CMS Collaboration, *Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm*, CMS-DP-2023-086 (2023), <https://cds.cern.ch/record/2879816>. (Cited on page 70.)

- [131] CMS Collaboration, *2024 Data Collected with AXOL1TL Anomaly Detection at the CMS Level-1 Trigger*, CMS-DP-2024-059 (2024), <https://cds.cern.ch/record/2904695>. (Cited on page 70.)
- [132] A. Zabi, *Design of the CMS calorimeter trigger upgrade from Phase I to Phase II of the LHC*, J. Phys.: Conference Series **1162**(1), 012040 (2019), doi:10.1088/1742-6596/1162/1/012040. (Cited on page 70.)
- [133] CMS Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, CMS-TDR-014 (2017), <https://cds.cern.ch/record/2272264>. (Cited on page 71.)
- [134] C. E. Brown, *Fast Machine Learning in the CMS Level-1 Trigger for the High-Luminosity LHC*, <https://cds.cern.ch/record/2875830> (2023). (Cited on page 72.)
- [135] G. Petrucciani, *Particle Flow reconstruction in the CMS Level-1 trigger for the HL-LHC*, EPJ Web Conf. **214**, 01019 (2019), doi:10.1051/epjconf/201921401019. (Cited on page 73.)
- [136] C. Herwig, *Particle flow reconstruction for the CMS Phase-II Level-1 Trigger*, JINST **18**(01), C01037 (2023), doi:10.1088/1748-0221/18/01/C01037. (Cited on page 73.)
- [137] S. Summers, I. Bestintzanos and G. Petrucciani, *Reconstructing jets in the Phase-2 upgrade of the CMS Level-1 Trigger with a seeded cone algorithm* (2023), arXiv:2310.08062. (Cited on page 73.)
- [138] G. Bortolato, C. Deldicque, D. Gigi, B. Huber, E. Leutgeb, A. Lobanov, D. S. Rabady, A. Racz and H. Sakulin, *Architecture and Prototype of the CMS Global Level-1 Trigger for Phase-2*, JINST **18**(01), C01034 (2023), doi:10.1088/1748-0221/18/01/C01034. (Cited on page 74.)
- [139] D. S. Rabady *et al.*, *A 40 MHz Level-1 trigger scouting system for the CMS Phase-2 upgrade*, doi:10.1016/j.nima.2022.167805 (2023). (Cited on pages 75 and 76.)
- [140] T. O. James *et al.*, *The Level 1 Scouting system of the CMS experiment*, CMS-CR-2023-024 (2023), <https://cds.cern.ch/record/2852916>. (Cited on page 75.)
- [141] R. Ardino *et al.*, *CMS Level-1 trigger Data Scouting firmware prototyping for LHC Run-3 and CMS Phase-2*, JINST **19**(05), C05027 (2024), doi:10.1088/1748-0221/19/05/C05027. (Cited on page 75.)
- [142] P. Ilten, Y. Soreq, J. Thaler, M. Williams and W. Xue, *Proposed Inclusive Dark Photon Search at LHCb*, Phys. Rev. Lett. **116**(25) (2016), doi:10.1103/physrevlett.116.251803. (Cited on page 76.)
- [143] R. T. D’Agnolo and M. Low, *Disorder and mimesis at hadron colliders*, JHEP **2019**(163) (2019), doi:10.1007/jhep08(2019)163. (Cited on page 76.)
- [144] S. Knapen, S. P. Griso, M. Papucci and D. J. Robinson, *Triggering soft bombs at the LHC*, JHEP **2017**(076) (2017), doi:10.1007/jhep08(2017)076. (Cited on page 77.)
- [145] LHCb Collaboration, *Search for the decays $B_s^0 \rightarrow \tau^+ \tau^-$ and $B^0 \rightarrow \tau^+ \tau^-$* , Phys. Rev. Lett. **118**, 251802 (2017), doi:10.1103/PhysRevLett.118.251802, arXiv:1703.02508. (Cited on page 77.)
- [146] X.-Y. Pham, *Lepton flavor changing in neutrinoless τ decays*, Eur. Phys. J. C **8**(3), 513 (1999), doi:10.1007/s100529901088, arXiv:hep-ph/9810484. (Cited on page 77.)
- [147] CMS Collaboration, *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*, CMS-TDR-022 (2021), <https://cds.cern.ch/record/2759072>. (Cited on page 77.)
- [148] Samtec, *Micro Flyover On-Board Optical Engine, FireFlyTM*, <https://www.samtec.com/optics/optical-cable/mid-board/firefly>. (Cited on page 77.)
- [149] P. Zejdl *et al.*, *10 Gbps TCP/IP streams from the FPGA for the CMS DAQ eventbuilder network*, JINST **8**(12), C12039 (2013), doi:10.1088/1748-0221/8/12/C12039. (Cited on pages 77 and 86.)
- [150] P. Zejdl *et al.*, *10 Gbps TCP/IP streams from the FPGA for High Energy Physics*, J. Phys.: Conference Series **513**(1), 012042 (2014), doi:10.1088/1742-6596/513/1/012042. (Cited on pages 77 and 86.)

Bibliography

- [151] AMD Xilinx, Inc, *Virtex UltraScale+ HBM VCU128 FPGA Evaluation Kit*, <https://www.xilinx.com/products/boards-and-kits/vcu128.html>. (Cited on page 86.)
- [152] HiTech Global, *6-Port QSFP28 (6x100G) / QSFP+ (6x40G or 6x56G) FMC+ Module (Vita57.4)*, <https://www.hitechglobal.com/FMCModules/x6QSFP28.htm>. (Cited on page 86.)
- [153] One Stop Systems, *4UV Expansion System*, <https://onestopsystems.com/collections/systems/products/4uv-expansion-system>. (Cited on page 86.)
- [154] CMS Collaboration, *40 MHz Scouting with Deep Learning in CMS*, CMS-DP-2022-066 (2022), <https://cds.cern.ch/record/2843741>. (Cited on page 89.)
- [155] FastML Team, *fastmachinelearning/hls4ml*, doi:10.5281/zenodo.1201549, <https://github.com/fastmachinelearning/hls4ml> (2023). (Cited on pages 89 and 111.)
- [156] J. Duarte *et al.*, *Fast inference of deep neural networks in FPGAs for particle physics*, JINST **13**(07), P07027 (2018), doi:10.1088/1748-0221/13/07/P07027, arXiv:1804.06913. (Cited on pages 89 and 111.)
- [157] D. Valsecchi, *First evidence of VBS in semileptonic decays with $WVjj \rightarrow lvqqjj$ final state and optimization of the CMS electromagnetic calorimeter for Run III*, <https://cds.cern.ch/record/2801848> (2022). (Cited on page 102.)
- [158] CMS Collaboration, *Analysis of muons and calorimeter objects collected by the Level-1 Trigger Data Scouting demonstrator during LHC Run 3*, CMS-DP-2024-056 (2024), <https://cds.cern.ch/record/2904692>. (Cited on page 107.)
- [159] CMS Collaboration, *Measurement of the inclusive cross section of Z boson production in pp collisions at $\sqrt{s} = 13.6$ TeV*, CMS-PAS-SMP-22-017 (2023), <http://cds.cern.ch/record/2868001>. (Cited on page 109.)
- [160] S. Summers *et al.*, *Fast inference of Boosted Decision Trees in FPGAs for particle physics*, JINST **15**(05), P05026 (2020), doi:10.1088/1748-0221/15/05/P05026, arXiv:2002.02534. (Cited on page 111.)
- [161] L. Wan, M. Zeiler, S. Zhang, Y. Le Cun and R. Fergus, *Regularization of Neural Networks using DropConnect*, PMLR **28**(3), 1058 (2013), <http://proceedings.mlr.press/v28/wan13.html>. (Cited on pages 111 and 113.)
- [162] S. Ioffe and C. Szegedy, *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift* (2015), arXiv:1502.03167. (Cited on pages 111 and 114.)
- [163] C. N. Coelho, A. Kuusela, S. Li, H. Zhuang, T. Aarrestad, V. Loncar, J. Ngadiuba, M. Pierini, A. A. Pol and S. Summers, *Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors. Ultra Low-latency, Low-area Inference Accelerators using Heterogeneous Deep Quantization with QKeras and hls4ml*, Nat. Mach. Intell. **3**, 675 (2021), doi:10.1038/s42256-021-00356-5, arXiv:2006.10159. (Cited on page 113.)
- [164] E. Wang, J. J. Davis, D. Moro, P. Zielinski, J. J. Lim, C. Coelho, S. Chatterjee, P. Y. K. Cheung and G. A. Constantinides, *Enabling Binary Neural Network Training on the Edge* (2023), arXiv:2102.04270. (Cited on page 113.)
- [165] N. Lai, R. Ardino and T. James, *Barrel muon track reconstruction with deep learning for Level-1 trigger data scouting in the CMS experiment*, CERN-STUDENTS-NOTE-2023-245 (2023), <https://cds.cern.ch/record/2883887>. (Cited on page 113.)
- [166] D.-A. Clevert, T. Unterthiner and S. Hochreiter, *Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)* (2016), arXiv:1511.07289. (Cited on page 113.)

- [167] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever and R. R. Salakhutdinov, *Improving neural networks by preventing co-adaptation of feature detectors* (2012), arXiv:1207.0580. (Cited on page 114.)
- [168] R. M. Harris and K. Kousouris, *Searches for Dijet Resonances at Hadron Colliders*, Int. J. Mod. Phys. A **26**(30n31), 5005 (2011), doi:10.1142/s0217751x11054905, arXiv:1110.5302. (Cited on page 115.)
- [169] CMS collaboration, *Search for dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter and other models*, Phys. Lett. B **769**, 520 (2017), doi:10.1016/j.physletb.2017.02.012, arXiv:1611.03568. (Cited on page 116.)
- [170] CMS collaboration, *CMS Exotica Summary plots for 13 TeV data*, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEX013TeV>. (Cited on page 116.)
- [171] CMS Collaboration, *Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JHEP **2020**(5) (2020), doi:10.1007/jhep05(2020)033, arXiv:1911.03947. (Cited on page 117.)
- [172] P. Cappelli, *On the measurement of $W \rightarrow 3\pi$ with the Phase 2 L1 scouting system at CMS*, doi:20.500.12608/59324 (2022). (Cited on page 127.)

Acronyms

AK4	Anti- k_T algorithm with $R = 0.4$ clustering parameter	35, 66, 73
AOD	Analysis Object Data	83, 97
API	Application Programming Interface	91
ATCA	Advanced Telecommunications Computing Architecture	70
ATLAS	A Toroidal LHC Apparatus	1, 3, 10, 11, 15, 29, 56, 57, 115, 133
bbH	Higgs production mode in association with a bottom quark pair	8, 9, 31, 54
BC	(Phase-2) Barrel Calorimeter	70
BCT	(Phase-2) Barrel Calorimeter Trigger	70
BDT	Boosted Decision Tree	35, 69, 73, 74
BMTF	Barrel Muon Track Finder	68, 69, 71, 82–84, 86, 88–92, 99, 100, 102–104, 106–108, 110, 111, 113, 136
BSM	Beyond the Standard Model	1, 3, 14, 62, 136
BX	Bunch Crossing	16, 74, 79, 87, 88, 94, 96, 97, 99–106, 111, 113, 120, 132
CERN	European Organisation for Nuclear Research	15, 16
CHS	Charged Hadron Subtraction	35
CL	Confidence Level	10, 11, 13, 14, 56–59, 133
CL1	Correlator Layer-1	72–74, 120
CL2	Correlator Layer-2	73, 74
CL_s	Statistical method for setting upper limits	56, 130
CMS	Compact Muon Solenoid	1–3, 10, 11, 15, 17–24, 26, 29–31, 33, 49, 56, 58, 61–71, 75–79, 81, 83, 86, 91, 95–99, 105, 106, 115, 116, 119, 132–134
CMS-DAS	CMS Data Aggregation System	83
CMSSW	CMS offline software	83, 96, 97, 99, 107
COTS	Commercial-Off-The-Shelf	77
CPU	Central Processing Unit	79, 94, 97, 132
CR	Control Region	37, 40–43, 45, 48, 55–57

Acronyms

CRC	Circular Redundancy Check (Level-1 trigger link protocol)	87
CSC	Cathode Strip Chambers	23, 24, 65, 68, 69, 71
CSP	CMS Standard Protocol	70
CSV	Combined Secondary Vertex	35
CT	Calorimeter Trigger	65, 69–73, 110
CTF	Combinatorial Track Finder	25
CTP7	Calorimeter Trigger Processor (Virtex) 7	65, 66
DAQ	Data Acquisition	63, 77, 79, 86, 89–91, 96, 97
DeMux	Calorimeter Trigger Demultiplexer	81, 82, 85, 90–92, 102
DM	Dark Matter	62, 76
DMA	Direct Memory Access	91
DQM	Data Quality Monitoring	79
DSBU	Data Scouting Buffer Unit	82, 83, 93, 94, 96–98
DSCB	Double Sided Crystal Ball	42, 43, 53, 110, 111, 130
DSP	Digital Signal Processor	91, 113, 115
DSPU	Data Scouting Processing Unit	82, 83, 97, 98
DT	Drift Tubes	23, 65, 68, 69, 71, 84
DY	Drell-Yan	109
EB	Electromagnetic Calorimeter Barrel	21
ECAL	Electromagnetic Calorimeter	21, 23, 25–27, 34, 41, 53, 54, 65, 66, 70, 72, 73, 84, 110
EE	Electromagnetic Calorimeter Endcap	21
ELU	Exponential Linear Unit	113
EMTF	Endcap Muon Track Finder	68, 69, 71, 102, 103
EoLS	End-of-Lumisecion (JSON file)	96
EoR	End-of-Run (JSON file)	96
EW	ElectroWeak	8
FIFO	First-In First-Out	87, 88
FM	Function Manager	83, 91, 93, 96
FPGA	Field Programmable Gate Array	1, 24, 64, 65, 68–71, 77, 79, 82, 83, 86, 89–93, 107, 111–113, 115, 134
FSM	Finite State Machine	88, 90
FSR	Final State Radiation	30, 40, 41
GbE	Gigabit Ethernet	77, 83, 86, 90, 91, 93
GCT	Global Calorimeter Trigger	70–73
GEM	Gas Electron Multiplier	24, 69, 71
ggF	Gluon-gluon Fusion Higgs production mode	7, 9, 31, 40, 43, 46

μGMT	(micro) Global Muon Trigger	68–72, 81, 82, 84, 90–92, 102, 107, 110–112
GPU	Graphics Processing Unit	79
GSF	Gaussian Sum Filter	27
μGT	(micro) Global Trigger	65, 66, 68–71, 74–76, 82, 83, 88, 90–92
GTT	Global Track Trigger	70–74
HB	HCAL Barrel detector	22, 102
HBM	High-Bandwidth Memory	77, 83, 86–91, 93
HCAL	Hadronic Calorimeter	21–23, 25, 26, 65, 66, 72, 73
HE	HCAL Endcap detector	22, 102
HF	HCAL Hadron Forward detector	22, 40, 46, 54, 70, 73, 102
HGCAL	High-Granularity Calorimeter	70, 72, 73
HL-LHC	High-Luminosity Large Hadron Collider	1, 2, 15, 17–19, 29, 61, 62, 70, 119, 132–134
HLS	High-Level Synthesis	89, 112, 113, 115
HLT	High-Level Trigger	24, 62, 63, 65, 70, 81, 97, 115, 116, 134
HO	HCAL Outer barrel detector	22
HP	High Purity	39, 40, 43, 46, 47
HPC	High Performance Computing	78
HSCP	Heavy Stable Charged Particle	2, 62, 74, 77
HT	Hadron Transverse Energy	68
I/O	Input/Output	76, 77
ID	Identification Criterium	32, 33, 35
IP	(Semiconductor) Intellectual Property (core)	86
iRPC	Improved Resistive Plate Chambers	71
ISR	Initial State Radiation	30
IU	Ingestion Units	77, 78
JER	Jet Energy Resolution	27, 53, 54
JES	Jet Energy Scale	27, 53, 54
JSON	Java Script Object Notation	93
KBMTF	Kalman Muon Barrel Track Finder	81, 101, 107, 108, 113–115
KF	Kalman Filter	25, 68, 71
L1	Level-1	1, 2, 24, 54, 62–72, 75–78, 81, 83, 87–89, 96, 97, 99, 101–109, 111–116, 119–134, 136

Acronyms

L1DS	Level-1 trigger Data Scouting	2, 62–64, 75–77, 81, 82, 90, 91, 94, 96, 97, 100–102, 105–108, 110, 111, 115–117, 119, 123, 124, 130, 132, 134
L1T	Level-1 trigger	1, 2, 24, 61–64, 66–68, 70, 75, 76, 79, 119, 123, 133, 134
LCDA	Light-Cone Distribution Amplitude	13
LD	Likelihood Angular Discriminator	38–40
LDME	Long-Distance Matrix Elements	13
LEP	Large Electron–Positron Collider	15
LHC	Large Hadron Collider	1–3, 7, 9–11, 15–18, 24, 29, 64, 71, 74, 75, 77, 79, 81, 83, 87, 94–96, 99, 102, 103, 133, 134
LID	Link Identification word (Level-1 trigger link protocol)	87
LINAC4	Linear accelerator 4	15
LLP	Long-Lived Particles	62, 66, 69, 74, 115
LO	Leading-Order	7, 30, 119, 121
LP	Low Purity	39, 40, 46, 47
LS	Lumi Section	96, 97
LUT	Look-Up-Table	68
MAE	Mean Absolute Error	114
MC	Monte Carlo	30, 31, 43, 44, 49, 54–56, 68, 106, 109, 111
MET	Missing Transverse Energy	21, 67, 68, 101
ML	Maximum Likelihood	42, 44, 74, 82, 134
MP7	Master Processor (Virtex) 7	65, 66, 68, 69, 81
MTF7	Modular Track Finder (Virtex) 7	69
MVA	Multivariate Classifier	34, 35, 51, 131
N³LO	Next-to-Next-to-Next-Leading-Order	8, 31
NLL	Negative Log-Likelihood	49
NLO	Next-to-Leading-Order	8, 31
NN	Neural Network	35, 74, 112, 113, 115
NNLO	Next-to-Next-Leading-Order	8, 31
NRQCD	Non-Relativistic Quantum Chromodynamics	13
OMS	CMS Online Monitoring System	105, 106
OMTF	Overlap Muon Track Finder	68, 69, 71, 102, 103, 106
OS	Opposite Sign pair	109, 110, 125
OT	Outer Tracker	71, 72
PCB	Printed Circuit Board	86

PCIe	Peripheral Component Interconnect Express	86–88, 90, 91
PDF	Parton Distribution Function	8, 31, 45, 53
PF	Particle Flow	25–27, 33, 34, 61, 70–74, 116
POG	Physics Object Group	49, 51, 54
PS	Proton Synchrotron	15
PSB	Proton Synchrotron Booster	15
PSV	Pixel Seed Veto	35
PU	Pileup	18, 78, 79
PUPPI	Pileup per Particle Identification	73, 74, 116, 120–131
PV	Primary Vertex	25, 32, 73
QCD	Quantum Chromodynamics	5, 8, 12, 13, 30, 31, 39, 44, 45, 49, 55, 66, 67, 106, 107, 117
QSFP	Quad Small Form-factor Pluggable	77, 86
RAM	Random Access Memory	94, 96–98
REST	Representational State Transfer	91
RPC	Resistive Plate Chambers	23, 24, 65, 68, 69, 71
RSS	Residual Sum of Squares	45
SCDAQ	Scouting Data Acquisition software	94–96
SCET	Soft-Collinear Effective Theory	13
SCONE	Scouting Configuration Endpoint software	91, 93, 94
sDS	Scouting Decision System	76
SF	Scale Factor	49, 51, 52, 54, 109, 111
sGS	Scouting Global System	76
SL	Super Layer	84, 85
sLS	Scouting Local System	76
SM	Standard Model of particle physics	1–8, 11, 12, 14, 15, 21, 29–31, 45–47, 55–58, 61, 63, 77, 81, 101, 108, 109, 133, 134
SPS	Super Proton Synchrotron	15
SR	Signal Region	37, 40, 42, 43, 45, 55
SR1	Signal Region 1 ($Z, H \rightarrow J/\psi \gamma$ searches)	37, 43
SR2	Signal Region 2 ($Z, H \rightarrow \psi(2S)\gamma$ searches)	37, 39, 40
SS	Same Sign pair	109
SSB	Spontaneous Symmetry Breaking	5–7
sTS	Scouting Track System	76
SU	Storage Units	78, 79, 83
SUEP	Soft Unclustered Energy Patterns	77
SV	Secondary Vertex	25
TBB	Intel Threading Building Blocks	94, 95

Acronyms

μTCA	Micro Telecommunications Computing Architecture	64
TCDS	Timing and Control Distribution System	75
TCP/IP	Transmission Control Protocol/Internet Protocol	77, 86, 88, 90, 91, 93, 94
TEC	Tracker Endcap	21
TF	Track Finder	83
TIB	Tracker Inner Barrel	21
TID	Tracker Inner Disk	21
TMT	Time Multiplexed Trigger	65
TMVA	Toolkit for Multivariate Data Analysis	35
TOB	Tracker Outer Barrel	21
TP	Trigger Primitive	64, 65, 70, 71, 75, 102, 123
TT	Trigger Tower	65–68, 70, 72, 105, 106
ttH	Higgs production mode in association with a top quark pair	8, 9, 31, 54
VBF	Vector Boson Fusion	8, 9, 31, 35, 40, 46, 54, 74
VH	Higgs-strahlung production mode	8, 34
WH	Higgs production mode in association with a W boson	8, 9, 31, 34, 40
ZH	Higgs production mode in association with a Z boson	8, 9, 31, 40