

Contents

1	Introduction to data scouting and data parking	2
1.1	Report structure	2
1.2	Physics motivations	2
1.3	Challenges and solutions	3
1.4	Origin and evolution of the data-scouting model	6
1.5	Evolution of the data-parking program	8
2	The CMS detector, trigger, and event reconstruction	10
2.1	The CMS detector	10
2.2	Trigger and data acquisition	10
2.3	Online and offline event reconstruction	11
3	Data scouting in Run 1 and Run 2	19
3.1	General strategy of data scouting	19
3.2	Trigger definitions and event content	19
3.3	Jets	22
3.4	Muons	26
3.5	Physics results	31
4	Data scouting in Run 3	45
4.1	The new Run 3 scouting strategy at the HLT	45
4.2	Event size and content	48
4.3	Jets	49
4.4	Muons	53
4.5	Electrons and photons	56
5	Data parking in Run 1 and Run 2	59
5.1	Data parking in Run 1	59
5.2	Data parking for B physics in Run 2	62
5.3	Physics results with the B parking data set	73
6	Data parking in Run 3	78
6.1	Data parking for B physics in Run 3	80
6.2	Alternative data parking strategies in Run 3	88
7	Summary	93
	Glossary	98
8	The CMS Collaboration	113

1 Introduction to data scouting and data parking

The Compact Muon Solenoid (CMS) experiment at CERN’s Large Hadron Collider (LHC) has achieved remarkable success in its mission to probe the fundamental structure of the universe.

Results from CMS and other experiments have considerably constrained the available parameter space for physics beyond the standard model (BSM), excluding the possibility of new states with masses up to several TeV predicted by a wide range of models of new physics. It has also scrutinized the realm of strong and weak interactions with great precision, including the discovery of the Higgs boson (H) and the measurements of its couplings [1–4]. As we delve deeper into the extensive data set afforded by the LHC, the absence of clear signals for new BSM physics prompts us to explore further avenues of investigation.

In this report, we describe the data-scouting and data-parking techniques, which involve the nonstandard use of the trigger, data acquisition (DAQ), and offline computing and software environments of CMS. Data scouting and data parking can overcome the limits of the conventional data processing strategies employed within CMS, by leveraging the capability and flexibility of the DAQ and offline computing systems. These techniques also exploit the advanced capabilities of the smart algorithms embedded within the level-1 (L1) trigger firmware and the sophisticated software-based event reconstruction algorithms used by the high-level trigger (HLT). Data scouting and data parking were introduced during the early running period of proton-proton (pp) collisions at the LHC, have been employed ever since, and equip CMS with the ability to substantially extend its sensitivity to low-mass and rare phenomena.

1.1 Report structure

This report is organized as follows. Section 1 introduces the physics motivations, the common data processing challenges and the solutions adopted to mitigate them, and the evolution of data scouting and data parking since they were initially introduced in CMS. Section 2 describes the CMS detector and trigger system, and details the typical event reconstruction workflow used in the experiment. In Section 3, the scouting strategy adopted in 2010–2012 (Run 1 period) and 2015–2018 (Run 2 period) is discussed, along with the main physics results obtained with this technique. Section 4 describes new scouting developments for the ongoing Run 3 (started in 2022, and planned to continue through 2025). Section 5 introduces the original data-parking implementation in 2012 and then focuses on the B parking strategy developed in Run 2 to increase the CMS sensitivity to flavor physics processes. In Section 6, new parking improvements designed for Run 3 are discussed, which are meant to complement the existing standard triggers with a large variety of physics goals in mind. Finally, Section 7 summarizes the main features and achievements of the data-scouting and data-parking strategies in CMS.

1.2 Physics motivations

The search for new physics often leads to scenarios where hypothetical particles have low masses and feeble couplings. Processes involving such particles are difficult to detect, given the large rate of standard model (SM) backgrounds at the LHC [5, 6]. In order to maintain a manageable overall trigger rate, traditional data acquisition protocols frequently necessitate relatively high thresholds to mitigate SM backgrounds. Consequently, intriguing signal events characterized by lower energy and momenta may inadvertently be discarded. Low-mass BSM particles that decay into final states involving low-energy jets or lepton pairs therefore present considerable challenges at the LHC. These challenges stem from the huge cross sections associated with jet production. Analogously, the large quantum chromodynamics (QCD) cross section and the subsequent (semi)leptonic decays of hadrons pose similar problems to searches

for new physics in light dilepton final states.

In addition to direct searches for new physics, we pursue indirect strategies where new physics may manifest as significant deviations between precise SM predictions and experimental measurements. One example is the study of rare B meson decays, involving particles with momenta in the few GeV range. However, online selection of these events presents formidable challenges, which are compounded by the need to collect a substantial amount of data to achieve sufficient statistical precision.

1.3 Challenges and solutions

The LHC facility features two adjacent parallel circular beamlines, each containing a bunched beam of protons traveling in opposite directions around the 27 km ring [7]. Each proton bunch orbits the ring at close to the speed of light 11,245 times per second. The proton beams are directed by superconducting magnets, and made to intersect at various points around the ring, where the pp collisions take place. In 2011 and 2012, the protons were accelerated to energies of 3.5 and 4 TeV, respectively. Starting from 2015, the energy was raised to 6.5 TeV and then, from 2022, further increased to 6.8 TeV.

The LHC orbit is divided into a total of 3564 time windows, each 25 ns in duration (bunch crossing slots) and potentially containing a colliding proton bunch. The actual collision rate depends on the number of colliding bunches and the structure of the filling scheme, which varies with time. Bunches are grouped into “trains” with 25 ns spacing (50 ns before 2015), and larger gaps between trains. The largest number of colliding bunches, 2544, was reached in 2017 and 2018, corresponding to an average collision rate of almost 30 MHz. Moreover, the number of multiple pp interactions within the same or adjacent bunch crossings, termed “pileup”, has also varied in time, ranging between ≈ 20 on average in Run 1, ≈ 40 in Run 2, and finally ≈ 50 in Run 3.

Protons are delivered in “fills” and, in 2023, the peak instantaneous luminosity ($\mathcal{L}_{\text{inst}}$) at the beginning of each fill was above $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This level was typically maintained (“luminosity leveling”) for six hours, after which the $\mathcal{L}_{\text{inst}}$ slowly decayed to lower values for the remainder of the fill (usually lasting several hours). The process of luminosity leveling entails deliberately diminishing the instantaneous luminosity from its maximal potential by slightly defocusing and/or separating the beams. This adjustment is crucial to prevent excessive pileup in experiments. Starting in 2024, the LHC aims to further increase the integrated luminosity (\mathcal{L}_{int}) delivered by extending the duration of the luminosity-leveling period. This continuous push for improvements in the performance of the LHC operations requires the experiments to develop innovative trigger and DAQ strategies in order to continue recording data sets rich with physics potential.

The traditional paradigm for data analysis at the LHC is that pp collision events are selected online by a trigger system, stored to disk in raw data format, and finally reconstructed and analyzed. The offline reconstruction aims to provide the best physics objects for analysis and, since it is not bound to be executed at the same pace of data acquisition and with the same low latency, as opposed to the trigger-level reconstruction, it achieves this goal at the cost of being computationally expensive.

The CMS experiment uses a two-tiered trigger system to filter the interesting collision events. The first level, L1, composed of custom hardware processors, relies on information from the calorimeters and muon detectors to select events up to a rate of around 100 kHz within a fixed latency of about $4 \mu\text{s}$ [8]. The second level, HLT, consists of a farm of processors running a

version of the event reconstruction software optimized for fast processing [9]. The HLT reduces the event rate to several kHz before data storage.

There are various constraints imposed on the trigger system and on the data processing framework that limit the number of events that can be selected, recorded and analyzed in this way:

- **L1 acquisition rate.** The rate of events that are accepted by the L1 system is limited to ≈ 100 kHz, determined from the finite bandwidth of the detector readout systems and the amount of raw detector information transmitted per event [10–12]. This is a hard constraint dictated by the detector design. Operating the system at rates beyond this threshold would result in dead time (the recording time lost because the readout system is not ready to transmit data for a new event) [8] and, effectively, no additional DAQ capability.
- **Event-processing time at the HLT.** The processing capacity of the HLT farm is proportional to the number of computing cores available and to the speed of such cores. The maximum processing time per event is therefore determined by the rate of L1-accepted events passed on to the HLT and by the total capacity of the farm. In 2018, it corresponded to a limit of about 600 ms per event assuming 100 kHz of L1 throughput. This rate is somewhat less of a hard constraint, as the HLT computing farm can be and is continuously being expanded via new acquisitions or via the replacement of older machines.
- **DAQ output bandwidth.** The DAQ throughput, increased from a few GB/s in Run 2 to about 20 GB/s in Run 3, is not considered to be a limiting factor. More relevant are the restrictions on the output bandwidth from the DAQ system, imposed by the size of the temporary raw data storage buffer at the site hosting the CMS experiment and by the bandwidth of the link transferring the raw data from the temporary to the permanent storage at the main CERN site. These limit the product of the HLT output rate and the event size, which in turn opens the possibility of collecting data at higher rates in exchange for reduced event sizes.
- **Prompt reconstruction of recorded data.** Normally, the full offline reconstruction of freshly recorded data, called “prompt reconstruction”, starts with only a short delay of about 48 hours once various detector calibration and alignment data are available [13]. Routine performance measurements of high-level physics objects and simulation-to-data corrections are often essential requirements for analyses and therefore time critical. The available computing resources allow for the prompt reconstruction of data with an approximately constant 48 hours turnaround time for HLT rates up to a few kHz on average.
- **Finite permanent data storage.** Ultimately, data storage is the remaining potential bottleneck to consider in the DAQ chain. Data can be stored on disks as well as on tapes. Disks are faster to access, but offer reduced storage relative to tapes. Very large data sets may stay on disk only for short periods of time until they are processed and stored in higher-level, smaller-sized data formats. After that, they must be moved to tape, where their retrieval is not immediate. However, this is also a soft constraint, as the purchase of additional disk storage is less costly compared to the purchase of computing cores.

The trigger system selects interesting events for physics analysis at a rate that is four orders of magnitude smaller than the bunch crossing rate. As the LHC performance improves over time, the higher $\mathcal{L}_{\text{inst}}$ values delivered impact the operations of the trigger, DAQ, and computing systems. Higher $\mathcal{L}_{\text{inst}}$ values imply higher pileup, which can degrade the performance of the

trigger algorithms and increase both the event size and the computational load from the event reconstruction.

The data-scouting and data-parking strategies [14] overcome two of the main limitations in the CMS data acquisition chain, namely the finite bandwidth available to write data to permanent storage and the finite ability to promptly process (i.e., reconstruct) the data as they are recorded. These two techniques are illustrated schematically in Fig. 1 for a representative year of data taking. Data scouting is a novel concept that CMS first prototyped in 2011 at the end of Run 1, used throughout Run 2 [15], and developed substantially for Run 3. Data parking is novel at the LHC, borrowing from a frequently used strategy by fixed-target experiments in which raw data are recorded and subsequently processed for analysis much later in time.




Figure 1: 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 $\mathcal{L}_{\text{inst}} = 1.2 \times 10^{34} \text{ cm}^{-2} \text{s}^{-1}$ over a typical 2018 fill, corresponding to an average pileup of 38, is considered.

The data-scouting strategy enhances sensitivity to low-energy physics processes by significantly lowering the HLT thresholds and storing a reduced event content on disk. Events reconstructed at the HLT are selected based on the kinematic quantities of their reconstructed objects using looser trigger thresholds than those applied in the standard trigger paths. For each event passing these looser selections, only high-level physics objects (such as jets or leptons) reconstructed at the HLT are stored on disk. No raw data from detector channels are stored for later offline analysis. These dedicated data samples are then used offline to perform physics analysis. The excellent performance of the HLT online reconstruction, which closely approximates the performance of the standard offline reconstruction, is the basis of the success of this strategy.

The data-parking strategy also lowers the thresholds used by the trigger algorithms, thereby increasing the experimental acceptance to low-mass physics processes. The event collection rate is thus substantially increased, potentially beyond the capacity of the computational resources available to promptly reconstruct the events as they are acquired. In this case, the data parking stream is transferred, unprocessed, to tape storage and is kept in a raw format until sufficient computational resources are available for the events to be reconstructed, such as between data-taking periods.

Figure 2 shows the time evolution of the output HLT rates for the standard, data-scouting, and

data-parking streams, averaged over one typical fill of a given data-taking year.

The CMS efforts have helped set a now-established trend in our field. Similar to the CMS data scouting, the LHCb and ATLAS Collaborations have “turbo” [16] and “trigger-level analysis” [17] streams, respectively, which were implemented during Run 2. Concurrently with the CMS data-parking developments in 2012, the ATLAS Collaboration developed a comparable “delayed stream” [18] approach. Finally, one additional DAQ technique at the LHC that circumvents limitations in the “standard” infrastructures is the ALICE “triggerless readout system” [19, 20].




Figure 2: Comparison of the typical HLT rates of the standard, parking, and scouting data streams from Run 1 to Run 3. The $\mathcal{L}_{\text{inst}}$ averaged over one typical fill of a given data-taking year is shown in pink.

1.4 Origin and evolution of the data-scouting model

The history of data scouting starts in the last weeks of the 2011 data-taking period. The CMS Collaboration aimed to preserve the physics sensitivity of searches for resonances with sub-TeV masses decaying to a pair of jets (dijet). The low-mass regime had become inaccessible because of the more stringent trigger selections applied to multijet event topologies in response to increases in the peak $\mathcal{L}_{\text{inst}}$ values delivered by the LHC. These additional selections were required because the cross section for the production of two jets mediated by the strong interaction grows substantially (roughly according to a power law) as the dijet mass decreases. During Run 1, the particle-flow (PF) algorithm [21] was introduced in the HLT online reconstruction. The PF algorithm, described in more detail in Section 2.3.4, aims to reconstruct and identify each individual particle (called a PF candidate) in an event with an optimized combination of information from the various CMS subdetectors. In 2011, the HLT jet algorithms

adopted the PF reconstruction providing a better jet momentum and spatial resolutions. Lowering the thresholds of the jet-based HLT triggers to select interesting low-mass events led to considerably higher trigger rates and to a higher data volume. The standard trigger and DAQ pipelines were not designed to handle such large amounts of data, which would have exceeded the resources allocated for the entire physics program of the experiment.

The proposed solution was to maintain low thresholds in the HLT PF jet algorithms and mitigate the high trigger bandwidth issue by permanently recording only a reduced data format (about a hundred times smaller than the standard one) in order to satisfy the design constraints of the DAQ system at the time. The data consisted essentially of the four-momenta of the jets reconstructed at the HLT and little additional information. This new data-scouting approach, which aimed to explore previously inaccessible regions of the mass-coupling model parameter space, was successfully tested in the last days of the data-taking period in 2011 and employed to produce a first preliminary result in a search for dijet resonances. It was the first attempt of its kind at the LHC.

Since its inception, data scouting has evolved, becoming a well-established approach in CMS, as described in Section 3. In 2012, the final year of Run 1, the data-scouting stream was used to search for dijet resonances with jets reconstructed from the calorimeter energy deposits alone [15]. In 2015, after the first long shutdown (LS1) of the LHC in 2013–2014, the data-scouting approach was consolidated by introducing a comprehensive event record based on the PF algorithm. The jet, muon, and electron candidates provided by the PF algorithm were all added to the scouting event record, which in turn allowed complex analyses to be performed, similar to what is possible with standard CMS data. An example is the study of jet substructure originating from the hadronic decay of a Lorentz-boosted resonance, as described in Section 3.3. New muon-based trigger algorithms were introduced to select events containing a pair of muons (dimuon) with transverse momenta of only a few GeV. These algorithms allowed extended searches for new dimuon resonances below 40 GeV, almost down to the kinematic threshold of twice the muon mass. The excellent performance of scouting muons in Run 2 is presented in Section 3.4. The scouting strategy in Run 2 enabled CMS to embark on pioneering searches for low-mass resonances, including pairs of jets or muons, promptly produced or displaced with respect to the primary pp interaction vertex, and complex decay chains involving multiple jets in the final state. An overview of these results is presented in Section 3.5.

The primary constraint in implementing the scouting strategy was found to be the HLT event processing time for the PF reconstruction algorithm. By the end of the second long shutdown of the LHC (LS2), in 2019–2021, the computing capabilities of the HLT system were greatly improved, thanks to the new computing farm equipped with graphics processing units (GPUs). Within this new GPU-based model, events are reconstructed at the HLT with a novel scouting PF algorithm that exploits charged-particle tracks built solely from information provided by the innermost silicon-pixel tracker. The substantial reduction in the average HLT event processing time (by over 40%) contributes to the increase in the maximum event rate that can be processed by the scouting stream.

At the beginning of Run 3, a single, unifying data-scouting stream has been available, comprising a complete PF-based event record for all events that satisfied the requirements imposed by at least one of several L1 algorithms based on jets, muons, and electrons or photons. The development of a suitably compact, yet complete, event record for scouting relied on the substantial experience developed within CMS [22, 23]. The total rate of events accepted by the combination of L1 algorithms fed to the scouting stream has increased to ≈ 30 kHz in 2022. A minimal event selection is then applied based on the PF candidates reconstructed at the HLT before scout-

ing events are recorded permanently for analysis. The event output rate of the data-scouting stream reached a maximum value of ≈ 30 kHz in 2022 and ≈ 26 kHz in 2023, roughly a factor of 10 higher than the standard data stream. In addition to jets, muons, and electrons, photons and individual PF candidates (such as hadron candidates) are now stored in the reduced scouting data format. More complex objects, such as hadronically decaying tau leptons or jets coming from the decay of heavy quarks, can in principle be reconstructed later from the constituent PF candidates. More information on the data-scouting strategy and event content in Run 3 is provided in Sections 4.1 and 4.2, respectively.

Data scouting made it possible to significantly reduce trigger thresholds with respect to the standard data stream, thereby increasing the sensitivity to previously unexplored new physics domains. As an example, scouting analyses with jet-based final states can select events with H_T greater than about 300 GeV, where H_T is defined as the scalar sum of the jet p_T momenta reconstructed in the event. This can be compared to analyses using standard triggers, which require $H_T \gtrsim 1000$ GeV. Similarly, the scouting dimuon triggers require two muons with $p_T > 3$ GeV, compared to $p_T > 17$ and 8 GeV for the leading and subleading p_T muons, respectively, for the standard inclusive triggers. One of the key aspects of the success of data scouting is the outstanding quality of the HLT online reconstruction. The muon reconstruction at the HLT is very similar to the offline one, which guarantees excellent performance in terms of identification efficiency and momentum resolution. Jets reconstructed from PF candidates and electron and photon objects also show comparable online and offline performance, in terms of energy scale and resolution, and particle tagging capabilities. The Run 3 performance of jets and muons, and initial studies on electron and photon objects, are presented in Sections 4.3, 4.4, and 4.5, respectively.

1.5 Evolution of the data-parking program

The LHC delivered pp collisions at center-of-mass energies of 7 and 8 TeV during Run 1. Towards the end of that data-taking period, CMS was accumulating data for the core pp physics program with an HLT rate of 300–350 Hz. In 2012, the data-parking technique was first deployed in CMS. New trigger algorithms, or existing ones with relaxed kinematical thresholds, were introduced to accumulate an additional 350 Hz of data, which were subsequently parked in raw format and later reconstructed during 2013. The triggers targeted a range of SM and BSM physics scenarios, including vector boson fusion (VBF) topologies and Higgs boson measurements, B physics measurements, and searches for models of compressed supersymmetry (SUSY) and dark matter (DM). Section 5.1 summarizes the Run 1 activities and the physics analyses served by these data-parking streams.

Early in Run 2, a data parking stream was enabled as a monitoring tool for the PF-based data scouting stream. As the LHC approached the end of Run 2, CMS initiated a powerful data-parking program to enable measurements of observables connected to the “flavor anomalies”. This collective term refers to several measurements of rare b hadron decays that exhibit some level of discrepancy with respect to the SM predictions [24]. These measurements have been the subject of substantial interest in the field since 2015.

In early 2018, a new trigger strategy was designed and implemented to identify muons originating from b hadron decays and thus accumulate a high-purity sample of b quark-antiquark ($b\bar{b}$) pairs. Kinematical requirements on the muon were progressively relaxed in the L1 and HLT algorithms during an LHC fill: the L1 algorithms were adjusted such that the system operated at or near its design limit throughout the entire LHC fill; the higher trigger rates from relaxing thresholds in the HLT algorithms were mitigated by the parking strategy. This ap-

proach minimized the impact on the core physics program while maximizing the sample size of $b\bar{b}$ events: Around 10^{10} $b\bar{b}$ events were recorded in 2018, which enabled a new program of measurements involving b hadron decays. The “tag-side” b hadron decays to a muon (responsible for the positive trigger decision) and other particles allowed for precision measurements of rare and low-mass signatures. Furthermore, the sample also crucially provided an unprecedented sample of 10^{10} unbiased decays from the other b hadron in each event, which allowed the studies of final states involving low- p_T leptons and hadrons that previously could not be probed with existing triggers. In comparison, other data sets highly enriched in b hadrons collected during the same period comprise at most 5×10^8 unbiased b decays. The data sample also has rich potential for BSM searches involving, e.g., low-mass states and very rare decays, which is complementary to the data sets that serve the high- p_T searches typical at the LHC, and thus substantially extends the reach of the CMS physics program. This data sample was parked and subsequently reconstructed in 2019 during LS2. Section 5.2 motivates and describes the strategy in detail, and summarizes some key physics results based on the analysis of the B parking data set.

The evolution of the B parking physics program during the LHC Run 3 was facilitated by the L1 trigger and DAQ and HLT systems operating at capacities beyond their original design specifications. For instance, the L1 system routinely operated at ≈ 110 kHz in 2023, which is exploited by the data-parking programs. Perhaps most crucially, additional computing resources are opportunistically available that allow CMS to reconstruct more data, by accommodating higher trigger rates from the HLT system. These operational developments directly and significantly enhance the scope of the CMS physics program. The improvements in LHC performance in Run 3 provide exciting new opportunities as well as challenges for data-parking strategies. The luminosity leveling periods impose constraints on available resources while enabling improved sensitivity to a wide variety of new physics searches and precision SM measurements. Section 6.1 describes the B parking strategy for Run 3.

The aims of the B parking strategy in Run 3 are twofold: collecting events with dimuon final states inclusively, and collecting events with dielectron final states inclusively. The dimuon approach simplifies the array of exclusive dimuon-based triggers that served much of the B physics program in Run 2. The new dimuon trigger consolidates the B physics program in many ways: a more efficient use of allocated trigger, DAQ, and computing resources; a common trigger strategy for the B physics group as a whole; and substantial gains in yields for b hadron decay modes (e.g., by more than a factor of 10 for $B^0 \rightarrow J/\psi K_S^0$) that were poorly served during Run 2. The dimuon trigger logic and physics performance are described in Section 6.1.3. The dielectron trigger primarily targets a measurement of the R_K observable [25–28] with a precision that is substantially improved with respect to that achieved using the single-muon trigger strategy of Run 2. However, the dielectron trigger logic is sufficiently inclusive to provide a data set that is also rich in possibility with regards to low-mass BSM searches. The dielectron trigger adopts the same approach as the single-muon trigger algorithms in 2018, by progressively lowering kinematical thresholds at L1 and HLT during the LHC fill; transverse energy (E_T) thresholds for each electron candidate as low as 5 GeV are deployed in the L1 system towards the end of an LHC fill. Section 6.1.4 describes the trigger logic and characterizes the data set recorded in 2022.

Improvements to the dielectron trigger strategy in 2023 opened up possibilities to further diversify the data-parking program to cover a wider range of physics topics, as it was originally conceived in 2012. Several triggers were added to the data-parking streams, with changes to their algorithms in both the L1 and HLT systems, providing improved sensitivity to a range of interesting physics processes beyond the scope of the B physics program. The VBF produc-

tion mode for the Higgs boson is covered by a suite of triggers that identify pairs of jets in the forward regions of the CMS detector. The Higgs boson self-coupling is a key parameter of the Higgs potential that remains unmeasured. Thus, optimized triggers that provide sensitivity to the pair production of Higgs bosons via final states containing pairs of τ leptons and jets from b quark decays were developed. Finally, new triggers were added to provide sensitivity to the distinctive experimental signatures of long-lived particles (LLPs), predicted by many BSM models. Examples include triggers that identify displaced dijets, or make use of timing information from the electromagnetic calorimeter (ECAL) subdetector. These additional data-parking strategies are discussed in Section 6.2.

2 The CMS detector, trigger, and event reconstruction

This section introduces the CMS detector and describes in more detail the standard CMS trigger and event reconstruction workflow. These topics are relevant for the subsequent discussion of the data-scouting and data-parking techniques developed in the remainder of the Report.

2.1 The CMS detector

The central feature of the CMS apparatus is a superconducting solenoid of 6 m internal diameter, providing a magnetic field of 3.8 T. Within the solenoid volume are a silicon pixel and strip tracker, a lead tungstate crystal ECAL, and a brass and scintillator hadron calorimeter (HCAL), each composed of a barrel and two endcap sections. Forward calorimeters extend the pseudorapidity coverage provided by the barrel and endcap detectors. Muons are measured in gas-ionization detectors embedded in the steel flux-return yoke outside the solenoid. At the start of 2017, a new pixel detector was installed [29] to provide four-hit pixel coverage in the pseudorapidity range $|\eta| < 2.5$. A more detailed description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in Refs. [30, 31].

2.2 Trigger and data acquisition

All CMS analyses rely heavily on an efficient trigger system that is able to separate the interesting processes from the huge number of background events produced in pp collisions at the LHC. The trigger system in CMS is split into two levels, the first one relying on custom-design hardware boards that use a minimal amount of information from the subdetectors with the fastest response, and the second one exploiting the complete event information for trigger decisions.

The L1 trigger [8] utilizes high-bandwidth optical links and large field programmable gate arrays (FPGAs) to process the information from the calorimeters and the muon trigger system to build trigger primitives (TPs). These trigger objects along with their kinematic features are used in a large set of algorithms for the final decision of the global trigger (GT). The products of these algorithms are referred to as L1 seeds. The L1 calorimeter trigger operates in two stages (layers). The calorimeter TPs (local energy deposits) from the electromagnetic, hadronic, and hadronic forward calorimeters are received by the first layer and calibrated. Then the ECAL and HCAL TPs are combined into single trigger towers and transmitted to the second layer for further processing. Jet, electron, photon, and tau lepton candidates are reconstructed and calibrated by the second layer and then fed to the GT together with the computed energy sums. The L1 muon trigger receives TPs from the overlapping muon subdetectors and feeds the reconstructed muon tracks into the GT. Finally, the inputs received by the GT are evaluated with a suite of algorithms and selection criteria, collectively called the trigger menu. The flexibility

of the GT hardware allows regular updates of the L1 menu in response to physics program choices and to changes in the LHC beam conditions.

The HLT [9] operates on fully assembled events that contain the entire event information, for example reconstructing tracks of particle trajectories and providing precise energy measurements from various subdetectors, making use of the full detector granularity and resolution. The HLT uses software algorithms running asynchronously on commercial computing hardware. Compared to the L1, the HLT menu defines a set of more complex algorithms and selection criteria, reconstructing online physics objects and filtering events. The HLT enables a more refined event selection than the L1 for storage and posterior analysis. It is also more computationally intensive, requiring longer processing times compared to the L1 trigger.

The data processing in the HLT is based on the concept of a trigger “path”, which is a set of algorithmic processing steps run in a predefined order that both reconstruct physics objects and apply selections on these objects based on physics requirements. The HLT paths targeting similar physics processes are grouped into common primary data sets (PDs). The PDs are defined to keep the total event rate balanced and within the limits imposed by the available data offline processing resources. While events can end up in more than one PD because of different trigger selections, significant effort is made to keep the event overlap to a minimum. Collections of PDs are organized into data “streams”, for efficient data handling. A data stream consists of a set of HLT paths and a well-defined event content.

In addition to the algorithms used to record events for physics analyses, the HLT also contains specific paths and data streams to gather information for detector calibrations and to conduct online data quality monitoring during data taking. Over the years, the HLT processing requirements increased notably in response to the evolving LHC and detector conditions, and the HLT computing capacity has been gradually scaled up to reflect the needs of the experiment. The Run 3 system was largely renewed by including general-purpose GPUs to provide cost-effective computing acceleration. Further details about the main Run 3 changes relevant to this Report are described in Section 4.

The DAQ system provides the data pathway and time decoupling between the synchronous detector readout and data reduction, the asynchronous selection of interesting events in the HLT, their intermediate or temporary local storage at the experiment site, and the transfer to Tier-0 for offline permanent storage and analysis. The DAQ system includes software to perform data handling, a hierarchical system to control the electronics components, monitoring systems to collect relevant metrics, and several monitoring clients to interpret those metrics. More details on the CMS DAQ and offline computing systems can be found in Refs. [30, 32, 33].

2.3 Online and offline event reconstruction

This section describes the physics event reconstruction workflow of CMS, both online (at the HLT) and offline. We focus on the physics objects that are also used in scouting analyses. A more complete description of the event reconstruction in CMS can be found in the references provided in the next sections.

2.3.1 Tracks and primary vertices

The tracking and vertex reconstruction algorithms employed by CMS [34] aim to precisely reconstruct the trajectories of charged particles and pinpoint the locations of pp interaction vertices within collision events.

In the initial stages of offline tracking, raw detector signals are converted into “hits” represent-

ing particle interactions with the various layers of the CMS detector, including the silicon pixel tracker and the silicon strip tracker. These hits are then utilized in a multi-step track reconstruction process: track seed generation, track finding, and track fitting. During track seed generation, potential track candidates are identified using subsets of hits, and various algorithms evaluate their compatibility with charged-particle trajectories. The subsequent track-finding stage propagates these candidates through the detector layers iteratively, refining their parameters to best fit the observed hit positions. Finally, track fitting optimizes the track parameters, such as position, direction, and momentum, by minimizing the discrepancies between the predicted and measured hit positions. The HLT uses track reconstruction software that is identical to that used for the offline reconstruction summarized above but configured to meet the constraints of the available central processing unit (CPU) resources at the HLT. This primarily involves reducing the number of iterations in the iterative tracking and/or only running iterations around objects of interest. Additionally, for some purposes at the HLT, the performance of pixel-only tracks is sufficiently robust to omit the time-consuming pixel-plus-strip tracking step, thus increasing the tracking rate for the same CPU budget. In general, the simplified HLT tracking reduces the efficiency and resolution of track parameters in some regions of phase space compared to the standard offline reconstruction.

The vertex reconstruction aims to measure the location and associated uncertainty of all pp interaction vertices in each event, including the vertex from the hard parton scattering and any additional pileup vertices, using the available reconstructed tracks. It consists of three steps: (i) selection of the tracks, (ii) clustering of the tracks that appear to originate from the same interaction vertex, and (iii) fitting for the position of each vertex using its associated tracks.

The first stage is to select high-quality tracks that are likely to be associated with the primary interaction. This involves applying track quality criteria to filter out noise and low-quality tracks. Track clustering is then performed using a “deterministic annealing” algorithm [34] converging towards a set of vertex candidates. Once the initial vertex seeds are found, a vertex fitting algorithm is employed. This algorithm iteratively refines the vertex positions and uncertainties by considering the selected tracks associated with each vertex candidate. The primary vertex (PV) is taken to be the vertex corresponding to the hardest scattering in the event, evaluated using tracking information alone, as described in Section 9.4.1 of Ref. [35].

At the HLT, pixel tracks can be used in the reconstruction of the vertex position. A “gap” algorithm [34] and a “density-based” algorithm [36] are used in Run 2 and Run 3, respectively. The pixel vertex reconstruction improves the overall speed while sacrificing some of the efficiency and resolution.

2.3.2 Calorimeters

Calorimeters are used to measure the energies of the various particles produced in each collision. The ECAL measures the energy of electrons and photons by absorbing them completely. Hadrons typically pass through the ECAL and are absorbed and measured by the HCAL. The local reconstruction of energy deposits in ECAL and HCAL is described in the following paragraphs.

2.3.2.1 The ECAL The ECAL consists of lead tungstate (PbWO_4) crystals emitting scintillation light when particles interact within their volumes. The 75,848 crystals are arranged in a central, cylindrical barrel section (EB), with pseudorapidity coverage up to $|\eta| = 1.48$, closed by two flat endcap sections (EE), extending the coverage to $|\eta| = 3.0$. The scintillation light produced inside the crystals is collected by photodetectors, creating an electrical signal amplified and shaped using a multigain preamplifier, which provides analog outputs that are converted

into digital signals by analog-to-digital converters. Because of the increased $\mathcal{L}_{\text{inst}}$ provided by the LHC and thus the higher number of overlapping signals from neighboring bunch crossings in Run 2 compared with Run 1 (resulting from the LHC bunch spacing changing from 50 to 25 ns), a novel ECAL amplitude reconstruction algorithm was developed in Run 2. The algorithm is based on a template fit called “multifit”, introduced since 2017, which attempts at resolving the many overlapping signals coming from pulses emitted in different bunch crossings, and has replaced the Run 1 method that was based on a digital filtering technique [37]. This “multifit” algorithm is robust and fast enough to be used both in the offline CMS reconstruction and at the HLT.

The energy response of the ECAL changes with time due to ageing of the crystals and of the photodetectors, caused by the high radiation levels at the LHC [38]. A dedicated monitoring system, using lasers that inject light during the LHC orbit gap, which is without proton collisions, is used to measure the transparency of each crystal and the photodetector response. For energy measurements at the HLT level, correction factors for the change in transparency are derived using measurements from the light monitoring system recorded in the preceding hours or days. In Run 3, these corrections are updated once per LHC fill, which is deemed sufficient given the existing running conditions. The finer time granularity of these offline corrections, which enables an accurate monitoring of the evolution of the detector response during an LHC fill, introduces some differences between online and offline reconstructed ECAL energy deposits.

A clustering algorithm is required to sum the energy deposits of adjacent channels that are associated with a single electromagnetic shower [39]. Corrections are applied to rectify the cluster partial containment effects. The ECAL clusters are dynamically combined into larger clusters to capture the full energy deposit from an electron or photon that might have undergone bremsstrahlung emission or conversions in the inner pixel tracker.

2.3.2.2 The HCAL The HCAL system includes several sections: the barrel (HB), endcaps (HE), outer (HO), and forward (HF). The HB and HE are sampling calorimeters made of interleaved brass and scintillating material, stationed outside the ECAL and inside the solenoid magnet. The HO is a plastic scintillator placed outside the solenoid and designed to catch highly-energetic hadrons. Finally, the HF is a quartz fiber Cherenkov calorimeter with steel absorbers also located outside the solenoid. Scintillation light produced inside the HB and HE are collected with wavelength-shifting fibers, optically summed, and sent to photodetectors to form analog electric signals. These signals are digitized by a charge integrator over a 25 ns interval, the latter known as a time sample (TS). Each recorded pulse shape consists of 10 TSs (8 since 2018 to reduce the data volume).

In the HB and HE, approximately 85–90% of the integrated energy occurs in a 50 ns window (2 TSs), while the LHC has delivered proton bunches every 25 ns since Run 2. The overlapping signals from nearby bunch crossings required the development of dedicated algorithms to estimate energy deposition in the HCAL. Used prior to 2015, a method based on the simple corrected sum of charges deposited in 2 TSs became unsuitable with the 25 ns bunch spacing. Consequently, several algorithms [40] were developed, based on fitting pulse-shape templates similar to the ECAL local reconstruction. Since 2018, the “minimization at HCAL, iteratively” (MAHI) algorithm, based on a fast chi-square minimization, has been used. This algorithm, deployed both offline and online, leads to a smaller difference between the offline and online reconstruction performance compared to the previous methods developed in 2016–2017.

2.3.3 Muon detectors

The CMS detector was designed with subdetectors dedicated to muon identification and to muon triggering, as well as to the measurement of the muon momentum and charge over a broad range of kinematic parameters [41]. The drift tubes and cathode strip chambers are located in the regions $|\eta| < 1.2$ and $0.9 < |\eta| < 2.4$, respectively, and are complemented by resistive plate chambers in the range $|\eta| < 1.9$. Three regions are distinguished, naturally defined by the cylindrical geometry of CMS, referred to as the barrel ($|\eta| < 0.9$), overlap ($0.9 < |\eta| < 1.2$), and endcap ($1.2 < |\eta| < 2.4$) regions. The chambers are arranged to maximize the coverage and to provide some overlap where possible.

Muons and other charged particles that traverse a muon subdetector ionize the gas in the chambers, which eventually causes electric signals to be produced on the wires and strips. These signals are read out by electronics and are associated with well-defined locations, generically called “hits”, in the detector. The precise location of each hit is reconstructed from the electronic signals using different algorithms depending on the detector technology.

2.3.4 Particle flow

The PF algorithm [21] aims to reconstruct and identify each individual particle (called a PF candidate) in an event, with an optimized combination of information from the various elements of the CMS detector. The energy of photons is obtained from the ECAL measurement. The energy of electrons is determined from a combination of the electron momentum at the primary interaction vertex as measured by the tracker, the energy of the corresponding ECAL cluster, and the energy sum of all bremsstrahlung photons spatially compatible with originating from the electron track. The energy of muons is obtained from the curvature of the corresponding track. The energy of charged hadrons is determined from a combination of their momentum measured in the tracker and the matching ECAL and HCAL energy deposits, corrected for the response function of the calorimeters to hadronic showers. Finally, the energy of neutral hadrons is obtained from the corresponding corrected ECAL and HCAL energies.

Offline PF reconstruction is used in the vast majority of physics analyses in CMS, and has also been deployed at the HLT. To cope with the stringent timing constraints, the HLT relies on a simplified PF algorithm. Offline, most of the processing time is spent reconstructing the inner tracks for the PF algorithm. The online version of the PF algorithm runs with two minor differences compared to its offline counterpart: the electron and isolated photon identification and reconstruction tasks are not included, and the reconstruction of tracks arising from nuclear interactions in the tracker material is not performed.

2.3.5 Jets

One important aspect of event reconstruction in CMS is the identification and reconstruction of jets. Jets are collimated streams of particles that arise from the fragmentation and hadronization processes of quarks and gluons produced in high-energy collisions. Reconstructing jets is crucial for understanding the properties of the particles involved in the collision and for identifying potential new physics phenomena.

The offline jets considered in this Report are reconstructed using the infrared- and collinear-safe anti- k_T (AK) algorithm [42, 43]. The default distance parameters used by the algorithm are 0.4 or 0.8, to reconstruct AK4 jets from single quarks/gluons or AK8 jets from the decay of Lorentz-boosted hadronic resonances, respectively. The inputs to the clustering algorithm are the four-momentum vectors of calorimeter energy deposits or PF reconstructed particles, which result in a calorimeter (Calo) jet or a PF jet, respectively.

2.3.5.1 Calo and PF jets

Calo jets are reconstructed from energy deposits in the calorimeter towers. A calorimeter tower consists of one or more HCAL cells and the geometrically corresponding ECAL crystals. In this process, the contribution from each calorimeter tower is assigned a momentum, the absolute value and direction of which are given by the energy measured in the tower and by the coordinates of the tower. The jet energy is obtained from the sum of the tower energies, and the jet momentum by the vector sum of the tower momenta. The jet energies are then corrected to establish a relative uniform response of the calorimeter in η and a calibrated absolute response in p_T .

In contrast, PF jets are reconstructed by clustering the four-momentum vectors of PF candidates. The jet momentum is determined as the vector sum of all the particle momenta in the jet. Pileup interactions can contribute extra tracks and calorimetric energy depositions to the jet momentum. To mitigate this effect in offline analysis, the jets are subject to the charged-hadron subtraction (CHS) or the pileup-per-particle identification (PUPPI) [44, 45] algorithms. In CHS, charged particles identified as originating from pileup vertices are discarded and an offset is applied to correct for remaining contributions. In PUPPI, the effect of pileup is mitigated at the reconstructed particle level, making use of local shape information, event pileup properties, and tracking information. While pileup charged particles are discarded, the momenta of neutral particles are rescaled according to their probability to originate from the PV, which is deduced from a local shape variable.

Calo jets result from a relatively simple yet robust approach and were widely used in early CMS publications. However, as the performance of the PF reconstruction has proven reliable and more powerful, PF jets have become the norm in CMS analyses. The advantages of using PF jets over Calo jets include a more complete event description as well as improved jet momenta and spatial resolutions, stemming from the combined use in PF of tracking detectors and of the high granularity of the ECAL.

2.3.5.2 Jet calibration

Jet energy corrections are derived from simulated samples to bring the measured response of jets to that of particle-level jets on average. In situ measurements of the momentum balance variable 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 [46]. The PF jet energy resolution (JER) typically amounts to 15–20% at 30 GeV, 10% at 100 GeV, and 5% at 1 TeV [46]. Additional selection criteria are applied to each jet to remove jets potentially dominated by anomalous contributions from various subdetector components or reconstruction failures [47].

The HLT reconstruction of jets uses the same clustering algorithm as its offline counterpart but differs in the calorimeter energy deposits or the PF candidates provided as input, as discussed in previous sections. Similarly to offline jets, jet energy corrections are derived from simulation to correct the response of HLT reconstructed jets. Dedicated studies to quantify and account for residual differences in jet energy scale and resolution between online (with scouting) and offline reconstructed jets are presented in Sections 3.3.2 and 4.3.

2.3.5.3 Jet substructure

The collisions at the LHC can produce heavy particles with large transverse momenta. In events that contain W and Z gauge bosons, Higgs bosons, top quarks, or even new resonances predicted in new physics scenarios, it is possible to achieve a high selection efficiency through the use of their hadronic decay channels. At sufficiently large Lorentz boosts (typically with p_T

of a few hundreds of GeV), the final-state hadrons from decays of such resonances merge into a single large-radius jet. In these cases, the analysis of jet substructure can be used to distinguish between those jets arising from a resonance decay and those arising from the numerous SM events composed uniquely of jets produced through the strong interaction, referred to as QCD multijet events [48].

The jet mass is one of the most powerful observables to discriminate resonance jets from background jets (i.e., jets stemming from the hadronization of light-flavor quarks or gluons). Contributions from initial-state radiation, the underlying event, and pileup can strongly impact the jet mass. Jet “grooming” techniques (such as the jet trimming [49] employed at the HLT) are applied to remove low-energy or uncorrelated radiation contributions from jets, thus improving the jet mass scale and resolution. Powerful machine learning (ML) techniques based on particle-level information have been recently used in offline analyses to identify and classify hadronic decays of highly Lorentz-boosted resonances [50]. In the analysis of the scouting data described in Sections 3.3.2 and 3.5.1.2, techniques without ML were used to identify the three-prong substructure from boosted top quarks or trijet resonances decays (the N -subjettiness ratio τ_{32} [51]), and the two-prong substructure of boosted W and Z bosons or dijet resonance decays (the N_2^1 variable based on energy correlation functions [52]).

2.3.5.4 Tagging of b jets

Jets from the hadronization and subsequent decay of bottom quarks (or b quarks) are called b jets. The hadronization of a b quark produces a b hadron that traverses the detector before decaying within the tracker volume. This phenomenon results in distinctive attributes within the emerging b jet, exemplified by the presence of a displaced secondary vertex (SV) that exhibits a displacement from the PV exceeding the CMS tracker resolution. The tracks stemming from this secondary vertex have a large impact parameter. Occasionally, the b jet is accompanied by a tertiary vertex (an outcome of the decay of the b hadron into a charm hadron), or by a lepton via the semileptonic decay of the b hadron or the charm hadron from a b cascade decay.

Physics analyses with b jets in the final state rely greatly on the identification, or tagging, of b jets. The b jets can be discriminated from jets produced by the hadronization of light quarks based on characteristic attributes of b hadrons, such as those described above. The CMS experiment employs a variety of b tagging algorithms. During Run 1, the principal tools employed for b jet identification consisted of likelihood-based discriminators [53]. Subsequently, in Run 2 and Run 3, the evolution of b tagging algorithms led to the adoption of multilayer perceptrons [54], deep neural network multiclassifiers [55, 56], and graph convolutional neural networks [57]. Each successive algorithm yielded notable enhancements in the efficiency of b jet identification. Similar algorithms, trained with the online reconstructed objects as input, were employed at the HLT in Run 2 and Run 3 to increase the online selection of events containing b jets. Tagging of b jets has not been employed so far in scouting-based analyses, but information that would allow such an analysis has been stored in the scouting data set since the beginning of Run 3.

2.3.6 Muons

Muons are crucial objects for the physics program of CMS since the original design of the detector. Thanks to their very clean experimental signature as they pass through the detector, muons are excellent probes to study known SM processes and to search for the production of new particles at colliders.

Following the hardware-driven reconstruction steps within the L1 trigger system, the standard

reconstruction of muon objects and their trajectories takes place via a two-step process at the software level. First, muons are reconstructed within the muon system only, which produces level-2 (L2) muons. Then, tracks produced in the pixel tracker are combined with the information from the muon spectrometer to reconstruct the full trajectory of the muon through the detector, which are termed level-3 (L3) muons.

The L2 muon reconstruction can refine the initial estimate of the muon trajectory by applying more accurate algorithms that are not feasible at L1 trigger. Standalone muon tracks are constructed by combining information from all muon subdetectors along a muon trajectory with a Kalman filter technique [58]. This iterative algorithm executes pattern recognition on a detector layer basis while concurrently refining the trajectory parameters.

The L3 muon reconstruction uses all available information about the muon trajectory from both the muon detectors and the tracker. Different L3 algorithms were used over the data-taking years. Generally, global muon tracks are built via an outside-in (OI) matching between a standalone muon track and a tracker track. The information from both tracks is used to perform a combined fit with the Kalman filter. Tracker muon tracks are instead constructed with an inside-out (IO) extrapolation by looking for a loose match between the tracker tracks and at least one muon detector segment. With the installation of a new pixel detector before the beginning of the 2017 data taking, a new iterative algorithm was adopted. It works in three steps. The first two, the OI and the IO steps, are both seeded by L2 muons. The second step considers only muons that were not already reconstructed by the previous step. Then, an additional IO step seeded by L1 muons is performed to recover candidates that could not be matched to an already reconstructed L3 muon. This IO step recovers some of the efficiency loss observed in previous steps, ensuring excellent performance for high- p_T muons and for muons in high pileup scenarios.

At the HLT [59], the procedure to build L2 muons as seeds for the track reconstruction in the inner tracker is identical to the one used for offline standalone muons. However, the HLT computing time constraints preclude conducting the full track reconstruction based on the multi-iteration approach across the complete volume of the inner tracker. The L3 reconstruction algorithms are performed only in smaller regions of the detector based on the presence of L1 or L2 muons. As a result, high reconstruction efficiency is achieved while minimizing computing resources. Differences between online and offline muons are typically small in terms of muon momentum scale and resolution, as described in Section 4.4.

2.3.7 Electrons and photons

Electrons and photons in the CMS detector are reconstructed with high purity and efficiency, and excellent resolution, making them ideal to use both in SM precision measurements and in BSM searches. Electrons and photons deposit almost all of their energy in the ECAL. In addition, electrons produce hits in the tracker layers. As electrons and photons propagate through the material in front of the ECAL, they may interact with the medium, with electrons emitting bremsstrahlung photons and photons converting into electron-positron pairs. Thus, by the time they reach the ECAL, they could consist of a shower of multiple electrons and photons. Their resulting clusters are combined into a single supercluster (SC) object to recover the energy of the primary electron or photon. Additionally, for an electron that loses momentum by emitting bremsstrahlung, the curvature of its trajectory changes in the tracker. A tracking algorithm based on a Gaussian sum filter (GSF) [60] is used to estimate the track parameters of electrons even in the presence of such emissions.

In CMS, there are three main ways to reconstruct an electron: seeded by the ECAL, seeded

by the tracker, and with a special low- p_T electron reconstruction. The ECAL-driven approach starts by combining ECAL clusters into a SC. For each SC found, compatible pixel hits in the inner tracker are sought, and any matches are used to seed the GSF tracking algorithm that builds the electron candidate. The tracker-driven approach takes the standard track collection and looks for one track that is compatible with ECAL energy clusters after applying some preselection. It then uses that track to seed the GSF tracking step. All ECAL clusters compatible with the track are associated with a single SC. Finally, the low- p_T electron reconstruction is a variant of the tracker-driven one and optimized for very low track momenta. Because of the high CPU cost to reconstruct all tracks in the event, only the ECAL-driven algorithm is available at the HLT and thus all electrons at the HLT require at least two hits in the inner tracker.

The differences between the ECAL-driven HLT and offline reconstruction algorithms are minimal and primarily driven by the limited CPU time available at the HLT and by the lack of final calibrations, which are not promptly computed during the data-taking period. The main distinction is in the GSF tracking algorithm, which is applied with fewer iterations compared to the offline reconstruction. Additionally, the formation of SCs is purely calorimeter-based, and not refined with tracking information, which would more accurately account for energy deposits that may be compatible with bremsstrahlung interactions.

More details on electron and photon reconstruction in CMS can be found in Ref. [39]. The dedicated offline reconstruction and identification algorithm optimized for electrons with $p_T < 10$ GeV is described in Section 5.2.8.

2.3.8 Missing transverse momentum

The presence of particles that do not interact with the detector material is indirectly measured by the missing transverse momentum (p_T^{miss}). The measurement captures the momentum carried away by undetected or invisible particles, such as neutrinos or other weakly interacting particles. The \vec{p}_T^{miss} vector is computed as the negative vector sum of the transverse momenta of the input objects in an event. The inputs can be calorimeter towers, PF candidates or jets (in the latter case the symbol H_T^{miss} is used). Similar to jets, the offline and online missing transverse momentum reconstruction algorithms mainly differ in the inputs fed to the algorithm.

More details on the reconstruction and calibration of these objects are provided in Ref. [61].

2.3.9 Tau leptons

The tau lepton (τ), with a mass of about 1.78 GeV, is the only lepton sufficiently massive to decay into hadrons. About one third of the time, tau leptons decay into an electron or a muon, plus two neutrinos. The neutrinos escape undetected, but the electron and muon are reconstructed and identified through the usual techniques available for such leptons, as described in previous sections. Almost all of the remaining decay final states of tau leptons contain hadrons, typically with a combination of charged and neutral mesons, and a tau neutrino.

Hadronic τ lepton decays (τ_h) are reconstructed from jets, using the hadrons-plus-strips (HPS) algorithm [62], which combines one or three tracks with energy deposits in the calorimeters to identify the tau lepton decay modes. Neutral pions are reconstructed from electrons and photons as strips with dynamic size in the η - ϕ plane, where the strip size varies as a function of the p_T of the electron or photon candidate.

To distinguish genuine τ_h decays from jets originating from the hadronization of quarks or gluons, and from electrons and muons, the DEEPTAU algorithm is used [63]. Information from all individual reconstructed particles near the τ_h axis is combined with properties of the τ_h

candidate and of the event.

The HLT system runs a version of the τ_h reconstruction that is slightly different from the one used offline. This is achieved via specialized, fast, and regional versions of the reconstruction algorithms, and via the implementation of a multistep selection logic, designed to reduce the number of events processed by the more complex, and therefore more time-consuming, subsequent steps. Reconstructed tau leptons have not been employed so far in scouting-based analyses, but information that would allow such an analysis is stored in the Run 3 scouting data set.

3 Data scouting in Run 1 and Run 2

This section details the development and application of the scouting technique by the CMS Collaboration during the first two periods of LHC operation. Two scouting data streams were defined, one based on jets and the other on muons. First, we describe in Section 3.1 the general trigger and reconstruction strategy for the scouting streams throughout the Run 1 and Run 2 data-taking periods. In Section 3.2 we focus on the definition of the triggers used to select interesting collision events and describe the corresponding event content of data stored with those triggers. In Sections 3.3 and 3.4, we report efficiency measurements of the scouting triggers and of the reconstruction performance for jet and muon objects, respectively. Finally, Section 3.5 showcases the physics results obtained with scouting-based analyses.

3.1 General strategy of data scouting

The scouting strategy at the HLT was originally designed and tested in 2011 to improve access to the enormous amount of data collected by the CMS detector, totaling over a hundred million individual readout channels. Scouting events are selected with a dedicated set of L1 algorithms at a higher HLT rate with respect to the standard streams to provide additional sensitivity to specific parts of the CMS physics program. These events are then processed in real time by the HLT computer farm and written on disk with reduced content. The majority of scouting events are reconstructed as part of the standard HLT event selection workflow. The CPU count dedicated to scouting thus constituted less than 5% of the total HLT farm resources, which in 2018 featured approximately 30,000 CPU cores. In Run 2, the scouting event rate accepted by the HLT was approximately 5 kHz on average and 6 kHz at the highest value of $\mathcal{L}_{\text{inst}}$, while the total allocated rate for the standard CMS physics program was approximately 1 kHz.

A comparison of the typical rates for each data stream during Run 1 and Run 2 operation is reported in Table 1, ranging from the initial tests performed in 2011 to the final configuration reached in 2018.

3.2 Trigger definitions and event content

The CMS trigger system is a dynamical entity, with operational parameters that are adjusted frequently to adapt to changing data-taking conditions in the short term, and less frequently to adjust to different physics goals in the long term. This section describes the specific event content and the algorithms designed for each scouting stream, as well as the dedicated rate budget available for data scouting. Most of the information is reported for the 2018 data-taking scenario, because it represents the final configuration achieved after various developments in Run 1 and Run 2, thus serving as a useful benchmark reference.

The initial scouting development in Run 1 focused on dijet triggers to search for low-mass hadronic resonances. Dedicated trigger paths based on calorimeter jets and on PF jets were

Table 1: Comparison of the typical HLT trigger rates of the standard, parking, and scouting data streams during Run 1 and Run 2. The average $\mathcal{L}_{\text{inst}}$ over one typical fill of a given data-taking year and the average pileup (PU) are also reported, consistent with the scenarios reported in Fig. 2.

Year	$\mathcal{L}_{\text{inst}} [\text{cm}^{-2}\text{s}^{-1}]$	PU	Standard rate [Hz]	Parking rate [Hz]	Scouting rate [Hz]
2012	0.5×10^{34}	28	420	400	1000
2016	0.9×10^{34}	35	1000	500	4500
2017	1.0×10^{34}	43	1000	400	4500
2018	1.2×10^{34}	38	1000	3000	5000

successfully commissioned in the final months of 2011, leading to the first preliminary results from dijet resonance searches. In Run 2, a new set of dimuon algorithms was employed to feed the scouting reconstruction in addition to the existing hadronic algorithms. Two versions of the hadronic trigger path were still in place: one using the calorimeter information and the other an optimized version of the PF reconstruction, which relied on additional tracking algorithms needed to improve the momentum resolution. As a result, two scouting data sets were produced and stored on disk: one from the “Calo” scouting stream, including both the muon and the hadronic triggers, and one from the “PF” scouting stream. This notation will be used in the following sections to identify the various groups of triggers. The event content of the PF scouting stream includes all PF candidates, resulting in a significant event size increase relative to the Calo scouting stream. Finally, a complementary data set that includes both the scouting event content and the complete CMS raw detector output was also defined, and used to collect events at a much lower rate. This data set is used to fully reconstruct a subset of scouting events offline, providing a useful way to validate the scouting reconstruction performance.

Table 2 lists the most important L1 and HLT triggers deployed in 2018 to collect scouting events. The dimuon scouting triggers were fully commissioned during 2017 with the aim of substantially lowering the muon p_{T} thresholds compared to the standard triggers. The L1 requirements on the dimuon invariant mass $m_{\mu\mu}$ and angular separation $\Delta R_{\mu\mu}$ help reduce the trigger rates. Lower-mass resonances are typically produced with considerable Lorentz boosts at the LHC, leading to final-state muons with significant momentum vector collimation (or low values of $\Delta R_{\mu\mu}$). The hadronic triggers are based on the H_{T} content of the event. In the Calo and PF scouting streams, only jets with $p_{\text{T}} > 40 \text{ GeV}$ are considered in the H_{T} sum. In the hadronic PF trigger, the L1 threshold was below 300 GeV in 2016 but subsequently raised to 360 GeV because of the increased pileup in 2017 and 2018. In parallel, new single-jet and double-jet L1 algorithms were added to better serve low-mass dijet analyses.

To maintain the event rate, data set size, and processing time within the allocated resources, minimal additional selection criteria are implemented in the scouting paths at the HLT. The dimuon and triple-muon L1 algorithms require each muon to have $p_{\text{T}} > 3 \text{ GeV}$, without imposing the need for muon tracks to point back to the nominal interaction point. The Calo scouting stream affords an H_{T} threshold at the HLT as low as 250 GeV , while maintaining a reasonable rate and good energy scale and resolution. The PF scouting stream, in contrast, requires a higher threshold of $H_{\text{T}} > 410 \text{ GeV}$ because of the larger event content compared to the Calo stream. A summary of the typical trigger rates achieved for each stream is reported in Table 3, for a scenario corresponding to the end of the 2018 data taking.

Table 4 summarizes the event content of the Run 2 scouting streams. Since there is no offline reconstruction in the scouting streams, the scouting event content comprises physics objects

Table 2: List of L1 and HLT thresholds for the most relevant scouting triggers in Run 2. The list corresponds to the 2018 thresholds that were valid for the overall Run 2 data-taking period. Differences with respect to the 2016 or 2017 scenario are reported in parentheses. Muons and photons are annotated as μ and γ , respectively, while OS stands for opposite-sign muon pairs. In cases where the same threshold is applied to all selected objects in an event, a single number is shown, while if different thresholds are applied to the objects, they are shown separated by slashes from the highest to the lowest.

Stream	L1 thresholds	HLT thresholds
Calo	$1\mu, p_T > 22 \text{ GeV}$ (not in 2017)	$p_T > 3 \text{ GeV}$
	$2\mu, p_T > 15/7 \text{ GeV}$	
	$2\mu, p_T > 4.5 \text{ GeV}, \eta < 2.0, \text{OS}, 7 < m_{\mu\mu} < 18 \text{ GeV}$	
	$2\mu, p_T > 4.5 \text{ GeV}, \eta < 2.0, \text{OS}, m_{\mu\mu} > 7 \text{ GeV}$ (not in 2017)	
	$2\mu, p_T > 0 \text{ GeV}, \eta < 1.5, \text{OS}, \Delta R < 1.4$	
	$2\mu, p_T > 4 \text{ GeV}, \eta < 2.5, \text{OS}, \Delta R < 1.2$	
PF	$3\mu, p_T > 5/3/3 \text{ GeV}$	$H_T > 250 \text{ GeV}$
	$H_T > 360 \text{ GeV}$ (200 GeV in 2016)	
PF	$H_T > 360 \text{ GeV}$	$H_T > 410 \text{ GeV}$
	$1 \text{ jet}, p_T > 180 \text{ GeV}$	-
	$2 \text{ jets}, p_T > 30 \text{ GeV}, \eta < 2.5, \Delta\eta < 1.5, m_{jj} > 300 \text{ GeV}$	-
	$1\gamma, p_T > 60 \text{ GeV}$	$p_T > 200 \text{ GeV}$

Table 3: Comparisons of the event rate, event size, and total bandwidth between the standard and scouting trigger strategies, for an LHC fill corresponding to data collected in 2018 with $\mathcal{L}_{\text{inst}} \approx 1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ at the start of the fill, one of the highest at the LHC in Run 2, and pileup around 50.

Data stream	Event rate [Hz]	Event size	Total bandwidth [MB/s]
Standard muons	600	0.86 MB	485
Standard jets/ H_T	400	0.87 MB	385
Scouting Calo muons and Calo H_T	5970	8.9 KB	45
Scouting PF jets and PF H_T	1766	14.8 KB	25

reconstructed online by the HLT. In the Calo stream, the jet information includes the kinematic observables of jets reconstructed with the calorimeter, which are stored if they satisfy $p_T > 20 \text{ GeV}$ and $|\eta| < 3$. In addition, the p_T^{miss} and the average energy density per unit area in the event (ρ) [64] are also stored. This stream also includes muon objects in events with at least two reconstructed muons accepted by the muon scouting triggers. Muon information includes kinematic and identification observables, such as the muon track momentum and the number of hits in the tracker and muon detectors, and information about the dimuon vertices such as the three-dimensional (3D) vertex position and corresponding uncertainty. These objects add up to about 10 KB per event, compared to roughly 1 MB in a typical standard event. In the PF scouting stream, the information stored per event consists of all PF candidates with $p_T > 0.6 \text{ GeV}$, as well as PF jets, leptons, and photons as reconstructed at the HLT. In addition, the p_T^{miss} object reconstructed with all PF candidates and the collection of primary vertices

along with ρ are also stored.

Table 4: List of observables saved in the scouting output during Run 2. The upper part of the table lists the observables present in the Calo stream and the lower part lists the contents of the PF stream. The PF candidates are sorted into charged and neutral hadrons, muons, electron and photons, hadronic and electromagnetic deposits in HF.

Observable	Definition
<i>Calo scouting stream</i>	
$(m^{\text{jet}}, p_T^{\text{jet}}, \eta^{\text{jet}}, \phi^{\text{jet}})$	Calo jet four-momentum
A^{jet}	Jet area
$E_{\text{max.}}^{\text{EM}}$	Maximum energy in electromagnetic towers
$E_{\text{max.}}^{\text{had.}}$	Maximum energy in hadronic towers
$E_{\text{HB,HE,HE}}^{\text{EM}}$	Electromagnetic energy in the HB, HE, and HF
$E_{\text{HB,HE,HE}}^{\text{had.}}$	Hadronic energy in the HB, HE, and HF
A^{towers}	Area of the EM and hadronic towers
$p_T^{\text{miss}}, \phi^{\text{miss}}, \rho$	Missing transverse momentum, angle, energy density
$(E^\mu, p_T^\mu, \eta^\mu, \phi^\mu)$	Muon four-momentum
$d_0 \pm \sigma_{d_0}, d_z \pm \sigma_{d_z}$	Muon impact parameters and uncertainties
I_E, I_H, I_T	ECAL, HCAL, and tracker isolation
N_P, N_S, N_M	Number of pixel, strip, and muon detector hits
N_L^S, N_L^T	Number of muon stations and tracker layers with hits
$(p_T^{\text{track}}, \eta^{\text{track}}, \phi^{\text{track}})$	Track three-momentum
χ^2, dof	Track χ^2 and number of degrees of freedom
$(q/p \pm \sigma_{q/p}, \lambda \pm \sigma_\lambda, \phi \pm \sigma_\phi, d_{sz} \pm \sigma_{d_{sz}})$	Fitted track parameters and uncertainties
i_{vertex}	Reference to the corresponding dimuon vertex
$(x \pm \sigma_x, y \pm \sigma_y, z \pm \sigma_z)$	List of 3D positions and uncertainties of dimuon vertices
<i>PF scouting stream</i>	
$(m^{\text{jet}}, p_T^{\text{jet}}, \eta^{\text{jet}}, \phi^{\text{jet}})$	PF jet four-momentum
A^{jet}	Jet area
E_i, N_i	Energy fractions and multiplicity for i^{th} particle type in jet
$p_T^{\text{miss}}, \phi^{\text{miss}}, \rho$	Missing transverse momentum, angle, energy density
$(m, p_T, \eta, \phi), \text{id}, i_{\text{vertex}}$	PF candidate four-momentum, type, vertex index
$(x \pm \sigma_x, y \pm \sigma_y, z \pm \sigma_z)$	List of 3D positions and uncertainties of primary vertices

The next sections demonstrate the feasibility of using scouting jet and muon objects with a reduced event content, and without applying the offline reconstruction algorithms, making scouting a valuable technique for several physics analyses.

3.3 Jets

Jets are the experimental signature of quarks and gluons produced in high-energy collisions such as the pp interactions at the LHC. The understanding of jet properties is a key ingredient

of several physics measurements and searches for BSM physics. Jets have been extensively employed in past CMS searches for new hadronic resonances with the data-scouting technique. This section presents the performance of the scouting jet triggers, showing the large increase in trigger efficiency for low-energy signals compared to the standard data stream. The reconstruction performance of jets in data scouting is also analyzed, demonstrating the feasibility of constructing and applying jet substructure variables with data scouting.

3.3.1 Jet trigger performance

The Calo scouting stream was active during both Run 1 and Run 2, and included jets from energy deposits in the ECAL and HCAL. The main trigger selection requires H_T larger than 250 GeV at the HLT, compared to $H_T > 800\text{--}900$ GeV for the triggers in the standard data stream. Although designed to select generic collision events that include jets in the final state, this trigger was primarily used to perform searches for new resonances decaying to pairs of jets, as described in Section 3.5.1.1. This analysis searches for a resonance peak in the invariant mass distribution of the two leading jets (the dijet mass m_{jj}) and it provides a benchmark for testing the performance of the data-scouting approach. Figure 3 shows the total trigger efficiency as a function of m_{jj} for the scouting (left) and the standard (right) triggers. While the standard trigger becomes fully efficient only for $m_{jj} > 1.25$ TeV, the scouting trigger efficiency reaches the 100% plateau at around 500 GeV, thus significantly extending the sensitivity of searches for low-mass resonances. Given the generic design of the H_T trigger, a similar improvement from scouting compared to the standard triggers is also expected for other new-physics signatures with final states dominated by the presence of high- p_T jets.




Figure 3: The efficiency of the Run 2 Calo scouting (left) and standard (right) jet triggers as a function of the reconstructed mass of the dijet system. Figures taken from Ref. [65].

The PF scouting data stream was introduced in Run 2 and was primarily used to perform searches for new resonances decaying to multijet final states. The data collected in Run 2 and used for physics analysis correspond to $\mathcal{L}_{\text{int}} = 128\text{ fb}^{-1}$. The main trigger selection requires $H_T > 410$ GeV, where here H_T is calculated with jet $p_T > 40$ GeV. Analyses using this trigger typically estimate the trigger efficiency as a function of H_T . The standard jet triggers require a threshold at the HLT of $H_T > 1080$ GeV. Figure 4 (left) shows that the scouting PF H_T trigger is fully efficient at around $H_T > 500$ GeV, offering a significant improvement in signal efficiency for low-energy multijet signals compared to standard triggers. The PF scouting data stream

also contains information about the individual particles as reconstructed by the PF algorithm. Their availability enables the reconstruction of jets with different cluster radii, for example large-radius jets with distance parameter of 0.8, which is useful for identifying resonances with high Lorentz boost that decay to jets. In the case of signals featuring merged decays of individual quarks, the trigger efficiency is measured as a function of the p_T of the leading large-radius jet and the jet mass, the latter being related to the resonance mass. Figure 4 (right) indicates that the PF scouting H_T trigger is fully efficient when $p_T > 300$ GeV, for any trimmed jet-mass (described in Section 2.3.5.3), while the standard triggers are fully efficient for jet momenta that are twice as high. These properties make the trigger suitable for new-resonance searches with a wide range of mass hypotheses.




Figure 4: The efficiency of the Run 2 PF scouting jet triggers as a function of H_T (left) and as a function of the leading large-radius jet p_T and trimmed jet mass (right).

3.3.2 Jet reconstruction performance

Jets in events collected by scouting triggers are formed from input calorimeter energy deposits or from PF candidates reconstructed at the HLT. To meet the stringent HLT time constraints, the online algorithms used to construct these inputs are in general simplified versions of those applied in the standard offline reconstruction. This can cause differences in the JES and JER between the online and offline jet objects. These effects are studied in this section, focusing on the performance of both Calo and PF jet reconstruction in Run 2 data scouting.

The JES of scouting Calo jets that are reconstructed online is calibrated to the one obtained with PF jets reconstructed offline. A monitoring data set has been defined, including both Calo jets at the HLT and the offline reconstructed PF jets, to measure the p_T difference between the two types of jets. A tag-and-probe method [46] is used to obtain these measurements. Figure 5 shows the observed p_T difference between the two collections as a function of jet p_T . The measured points are fitted with a smooth function and the resulting curve is used to calibrate the Calo jets collected by scouting triggers. The JES of Calo jets at the HLT is slightly smaller, by around 4% at low p_T and 1% at high p_T , compared to the PF jets reconstructed offline. With the dijet asymmetry method [46], we estimate that the scouting Calo JER is only about 10% worse compared to offline PF jets. These results confirm the good performance of Calo jet reconstruction in data scouting in the high jet p_T range considered.




Figure 5: The observed percent difference between the p_T of Calo jets at the HLT and the p_T of PF jets reconstructed offline (points), fitted to a smooth function (curve), vs. the Calo jet p_T . Both Calo and PF jets are calibrated with corrections derived from simulation. Figure taken from Ref. [65].

Studies similar to those presented above, which concern the comparison between PF jets reconstructed at HLT and offline, are described in detail in Section 4.3 based on Run 3 data. In this section we focus on studying the calibration level of PF jets at the HLT by reconstructing SM resonances. The HLT PF jet energies are corrected with the standard methods described in Section 2.3.5 using both simulation and data. The corrections are derived from data sets that have undergone full offline reconstruction, rather than from the scouting data set. The top quark, which can be reconstructed as a resonance in three-quark final states, is clearly visible in Fig. 6. The figure shows the invariant mass distribution of three PF jets and the mass distribution of single PF jets selected using jet substructure techniques that indicate they are likely to contain decays of three separate partons. The observed top quark mass peak positions agree with the expected ones from $t\bar{t}$ simulation within less than 2%, while the resolution in data is only 5% worse.

To increase the sensitivity of multijet searches in scouting, new techniques are employed. One innovation is the use of the information in PF candidates within a jet to construct a quark-gluon discriminator (QGD) that enhances signals featuring decays to quarks while suppressing QCD backgrounds consisting largely of gluon jets. The QGD is constructed from observables sensitive to fundamental differences in the fragmentation properties of quarks and gluons, such as the number of constituents and the jet radius. It uses a neural network (NN) architecture based on the DEEPSETS technique [66]. The NN inputs are the normalized four-momenta information along with the particle type of each jet constituent (PF candidates).

The QGD NN selects quarks and rejects gluons with better performance than traditional methods that rely on jet multiplicity and jet mass [47]. Figure 7 (left) shows the QGD score distributions obtained for quarks and gluons, indicating a clear separation. Different working points are considered in analyses, corresponding to quark signal efficiency (gluon background rejection rates) of 98% (31%), 83% (70%), and 61% (87%), respectively, for loose, medium, and tight selections on the QGD score. In Fig. 7 (right), the tight selection on the QGD score is applied to the invariant mass of jet triplets in the search for R -parity violating (RPV) gluinos, described in Section 3.5.1.2. The top quark peak is clearly seen above the QCD background.

The jets in the top quark peak are mostly quarks, while the continuum background contains a large component of jets originating from gluons. The QGD significantly reduces the continuum background, in comparison to the inclusive selection (without QGD), while preserving the top quark signal. This demonstrates the power of the scouting technique in advancing the jet-based physics program of CMS, despite the limited event content stored on disk.




Figure 6: The distribution of m_{3jj} for the resolved three-jet search (left), and average jet mass ($\tilde{m} = (m_1 + m_2)/2$) for the merged three-parton search (right), adapted from Ref. [67]. Both analyses use PF jets. The peak around 170 GeV in both distributions corresponds to the all-hadronic decay of the top quark. The data (points) are compared to the background-only prediction (blue) and the full background fit including simulations of the top quark resonance (red).

3.4 Muons

Muons are often indicators of interesting physics because they can be produced via the electroweak interaction, where Z or W bosons are involved, or via new hypothetical interactions featuring unknown gauge bosons. The CMS detector is particularly well suited to the task of reconstructing and identifying muons, as explained in Section 2. The clean signature of a pair of muons in the final state is exploited in a large number of analyses, from SM precision measurements to searches for new physics up to the TeV scale. Thus, increasing the number of collected dimuon events was a key motivation in the development of the scouting strategy.

3.4.1 Muon trigger performance

Dedicated trigger algorithms targeting dimuon events with significantly lower muon p_T thresholds than those of the standard trigger paths were implemented and fully commissioned in 2017. The trigger definitions are discussed in Section 3.2. The number of selected events at low dimuon masses ($m_{\mu\mu} < 40$ GeV) is substantially increased by reducing the size of the event content at the HLT. The data collected during the last two years of Run 2 (2017 and 2018) correspond to a total \mathcal{L}_{int} value of 101.1 fb^{-1} , 96.6 fb^{-1} of which are used for analysis. The scouting dimuon triggers provided an overall rate of approximately 2 kHz for $\mathcal{L}_{\text{inst}} \approx 1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, about a hundred times higher than the standard dimuon triggers.




Figure 7: Left: output of the QGD for quark (orange) and gluon (blue) jets. The corresponding receiver operating characteristic (ROC) curve is also shown. Right: observation of fully hadronic top quark decays in the invariant mass of three jets with a QCD multijet background, for an inclusive selection (no QGD), and for a selection including the QGD score. Figure adapted from Ref. [67].

Dimuon invariant mass spectra obtained using data collected with the standard and scouting triggers are compared in Fig. 8. The two curves are normalized to the amount of data collected with the scouting triggers in 2017 and 2018. The standard triggers show significant acceptance losses below about 40 GeV because of the higher p_T thresholds on the leading and subleading p_T muons. For the standard trigger strategy, these thresholds are 17 and 8 GeV, respectively. The acceptance is considerably recovered by the scouting triggers thanks to the looser HLT selections. These selections include the reconstruction of at least two muons at the HLT, each with $p_T > 3$ GeV. These requirements are minimal relative to the L1 selections, which are described in Table 2 for 2017 and 2018. The distributions in Fig. 8 are obtained by selecting events with offline muon p_T requirements of 20 and 10 GeV in the standard dimuon triggers, and 4 GeV for both muons in the scouting dimuon trigger. These selections ensure operation on the plateau of the trigger efficiency curves. We note that the comparison between the two data sets has some limitations. The standard thresholds of the dimuon selections in 2016 were lower than the ones adopted in 2017. In addition, the single-muon L1 path was only added to the dimuon scouting stream in 2018. This explains the discrepancy in event yields near the Z boson mass peak and at higher masses, observed when normalizing both curves to the same \mathcal{L}_{int} value.

The high-rate dimuon scouting stream with lower transverse momentum selections at the L1 (as low as the ones at the HLT) enables the exploration of an otherwise inaccessible phase space at low dimuon masses, down to about twice the muon mass at ≈ 210 MeV, which is the dimuon kinematic threshold. Figure 9 shows the dimuon invariant mass distribution obtained with the various L1 algorithms in 2018 and reconstructed at the HLT.

To extend the physics use case of the scouting stream, the scouting triggers utilize HLT reconstruction algorithms that lack any association between muons and the PV. This enables scouting searches for resonances that have nonzero displacement from the PV. The dimuon invariant mass distributions for different values of the dimuon transverse displacement from the interaction point (referred to as l_{xy}) are shown in Fig. 10. The maximum transverse displacement of about 11 cm is determined by the requirement that muon tracks deposit energy in at least




Figure 8: Dimuon invariant mass distribution of events selected with the standard muon triggers (blue, dashed) and scouting muon triggers (pink, solid) in the mass range 11–240 GeV, normalized to $\mathcal{L}_{\text{int}} = 96.6 \text{ fb}^{-1}$, corresponding to the scouting data collected in 2017 and 2018. The selection applied to obtain each distribution is described in Ref. [68].

two layers of the CMS pixel tracker. The definition of the muon scouting triggers in Run 3 have been updated to remove this requirement, providing sensitivity to resonances with even higher transverse displacements, as described in Section 6.1.1.

In Fig. 11, the efficiency of the dimuon scouting triggers, including the HLT selection and the main L1 paths, is presented in two-dimensional (2D) maps as a function of the angular separation $\Delta R_{\mu\mu}$ between the two muons and the dimuon invariant mass $m_{\mu\mu}$ (left plot), and as a function of $\Delta R_{\mu\mu}$ and the subleading muon p_T (right plot). These maps are needed to account for correlations between muons in the L1 algorithms. For very low mass resonances, the muon momenta are typically collinear, with low values of $\Delta R_{\mu\mu}$. Therefore, L1 requirements on $\Delta R_{\mu\mu}$ are applied as the p_T threshold is lowered to focus on the relevant physics expected in this kinematic region, as reported in Table 2. The efficiency measurements are performed with orthogonal data sets that are independent of the presence of muons. All muons with $p_T > 3 \text{ GeV}$ and $|\eta| < 1.9$ are considered. The additional selection $0.45 < m_{\mu\mu} < 0.65 \text{ GeV}$ is applied to the right plot to focus on events likely to contain $\eta \rightarrow \mu^+ \mu^-$ decays. The η meson, with a mass around 0.55 GeV, is one of the lightest resonances decaying to pairs of muons for which CMS has sensitivity, and thus serves as a useful proxy to study the performance of low-mass dimuon reconstruction.

The left plot in Fig. 11 shows that the efficiency to trigger on muon pairs with small angular separation ($\Delta R_{\mu\mu} < 1.0$) ranges from 60 to 80%. The efficiency gradually decreases for higher $\Delta R_{\mu\mu}$ and lower $m_{\mu\mu}$. Near the η meson mass, it ranges from 50 to 90% vs. subleading muon p_T , within the region $\Delta R_{\mu\mu} < 0.06$, which contains most of the $\eta \rightarrow \mu^+ \mu^-$ decays as determined from simulation. For a fixed value of $m_{\mu\mu}$, higher values of $\Delta R_{\mu\mu}$ imply lower muon p_T . This explains the lack of events in the upper left region of the left plot. In the right plot, for fixed muon p_T , higher $\Delta R_{\mu\mu}$ implies higher $m_{\mu\mu}$. Since the selection in this plot includes $0.45 < m_{\mu\mu} < 0.65 \text{ GeV}$, no events are found with $\Delta R_{\mu\mu}$ values above 0.12.




Figure 9: Dimuon invariant mass spectrum and event rate of each L1 seed (legend) obtained with the scouting stream reconstructed at the HLT, using data collected in 2018 corresponding to $\mathcal{L}_{\text{int}} = 60 \text{ fb}^{-1}$. Well-known dimuon resonances from various meson decays or from Z boson decays are indicated above each peak.




Figure 10: Dimuon invariant mass distributions in bins of transverse displacement from the PV (l_{xy}). Figure adapted from Ref. [69].




Figure 11: Efficiency of the dimuon scouting trigger and logical “OR” of all L1 triggers measured with 2017 and 2018 data. The efficiency is shown as a function of the angular separation between the two muons $\Delta R_{\mu\mu}$ and the dimuon mass $m_{\mu\mu}$ (left), and as a function of $\Delta R_{\mu\mu}$ and the subleading muon p_T (right). The selection in the right plot also requires $0.45 < m_{\mu\mu} < 0.65$ GeV to focus on the η meson resonance region. The statistical uncertainty in the measured values is generally less than 3% per bin on the left plot and less than 15% per bin on the right plot.

3.4.2 Muon reconstruction performance

The process of reconstructing muon objects within the scouting stream mirrors that of the standard stream, differing only in the removal of the vertex constraint. However, recorded events contain only a limited amount of information compared to the offline muons. Therefore, dedicated identification (ID) criteria were developed to select muons from the scouting data set.

A customized selection based on standard physics variables (cut-based selection) was initially designed in the context of a search for dimuon prompt production in the mass range 11–45 GeV, as described in Section 3.5.2.1. It relies on some requirements applied to the muon track, such as the number of tracker pixel hits, the total number of tracker layers containing energy deposits, and the quality and relative isolation of the muon track. This selection is not ideal for lower mass resonances, however. The angular separation between the muons is small when $m_{\mu\mu} < 10$ GeV. The muon isolation is less efficient for the boosted system because the isolation cone of the two muons may partially overlap. An optimized selection based on a multivariate analysis (MVA) technique was therefore developed to improve sensitivity to lower mass signals, by increasing the signal muon efficiency and suppressing the rate of background muons, mainly coming from decays in flight of hadrons.

The set of input variables for training the MVA classifier contains a combination of muon and vertex variables, such as the number of pixel hits and tracker layers, the muon track and vertex χ^2 , the track isolation, and the vertex transverse displacement from the interaction point. These parameters are combined into a single discriminator using a boosted decision tree (BDT) [70].

The MVA is optimized separately for the higher mass (4–10 GeV) and lower mass (< 4 GeV) regions. The signal samples used for the MVA training and validation are extracted from events in data containing $Y(1S)$ and J/ψ meson decays, for the high-mass and low-mass regions respectively, while same-charge muon events are used as background samples. The background rejection vs. signal efficiency of all IDs are summarized in Fig. 12, demonstrating the improved

sensitivities of the MVA-based selections compared to the cut-based ones in comparable mass regions. Considering a similar signal efficiency for $\Upsilon(1S)$ and J/ψ signals, the new MVA ID achieves significantly higher background rejection.




Figure 12: Background rejection vs. signal efficiency of the new MVA-based muon ID strategies evaluated on scouting data: $\Upsilon(1S)$ -trained MVA (black line), J/ψ -trained MVA (red line). A comparison with the performance of the previous cut-based selection, which was optimised for signals at the mass range higher than 11 GeV, is also shown for the $\Upsilon(1S)$ (blue triangle) and J/ψ (brown square) signals. Figure adapted from Ref. [71].

The performance of the muon scouting stream was also measured by means of the dimuon mass resolution. Figure 13 shows mass resolutions measured for various SM resonances. The resolution was determined by fitting a signal-plus-background model to the dimuon mass spectrum in 2017 and 2018 data around known resonances, such as the η , ϕ , J/ψ , and Υ , and then extracting the relative width of the peaks. The signal is modeled with the sum of a double-sided Crystal Ball (CB) function [72, 73] and a Gaussian function in a mass window of $\pm 5\%$ around the mean peak value, while the background is described by a third-order Bernstein polynomial. The signal resolution, estimated with the σ_{CB} parameter of the Gaussian core of the CB function, is found to be $\approx 1.3\%$ and roughly independent of year, mass hypothesis, or detector region. The uncertainty in the resolution is evaluated as the variation introduced by alternative signal models, such as a double-Gaussian function, and measured to be 13–18% depending on the resonance. The mass resolution of scouting muons in the low- p_T regime (the region of interest) is within 1% of the one obtained with offline muons, demonstrating the remarkable capabilities of data scouting in boosting the muon-based physics program of CMS.

3.5 Physics results

We now describe some of the CMS physics results obtained with the scouting data sets collected during Run 2. The results are organized in three sections: searches for new physics in hadronic final states, searches for new physics in leptonic final states, and observations of rare SM decays.

3.5.1 Searches in hadronic final states

Searches for direct production of hadronic resonances are particularly important at the LHC, as any hypothetical particle produced via the strong interaction in pp collisions can also decay to quarks and gluons, which hadronize to form jets. The main source of backgrounds consists of




Figure 13: Relative width of dimuon resonances as a function of mass, measured in 2017 and 2018 scouting data. The fits are performed separately for the 2017 and 2018 data sets. The values shown are the average width of each fit, weighted by the \mathcal{L}_{int} value corresponding to the data accumulated in each year. From left to right, the η , ϕ , J/ψ , $\psi(2S)$, $Y(1S)$, $Y(2S)$, and $Y(3S)$ resonances are shown. The inserts display fits of the J/ψ and Y peaks obtained with scouting data (black markers) separately for the signal (green) and background (red) components, and for their sum (blue).

SM multijet QCD backgrounds. The rates for these processes are typically very large compared to those of the potential signals of new physics, increasing substantially for lower resonance masses. Consequently, the event rate and the amount of data that must be recorded in order to carry out physics analyses in this low-energy regime rapidly increases. In this context, the data-scouting approach plays a crucial role in probing the low-mass region in final states with jets.

The CMS hadronic physics program covers a wide range of experimental signatures, such as resonances decaying to a pair of jets (dijet), dijet resonances in association with initial-state radiation, resonances decaying to three jets, and pair-produced resonances resulting in final states with four or more jets. In this section, we describe CMS searches that apply the data scouting technique to extend sensitivity to new physics in the mass region below the TeV scale.

3.5.1.1 Dijet resonances

Proton-proton collisions can produce two or more energetic jets when the constituent partons scatter with large momentum transfer. The invariant mass distribution of the two jets with the largest p_T is predicted to fall steeply and smoothly with increasing mass, based on known multijet processes. Many proposed extensions of the SM predict the existence of new states coupling to quarks and gluons, which would appear as resonances on top of this smooth background in the dijet mass spectrum.

A review of searches for dijet resonances at hadron colliders can be found in Ref. [74]. The first searches for dijet resonances were presented by the UA1 and UA2 experiments after collecting data at the SppS accelerator with $\sqrt{s} = 630$ GeV. These results were later extended to higher resonance masses by the CDF and D0 experiments, using the Fermilab proton-antiproton Tevatron collider, which operated with center-of-mass energies of 1.8 and 1.96 TeV. Finally, the mass reach was increased further by ATLAS and CMS, relying on pp collisions at the LHC with $\sqrt{s} = 7, 8$, and 13 TeV.

Results obtained with different collider energies were compared by translating the upper limits on the quark-quark resonance cross sections into upper bounds on the coupling constant g between the new resonance and a pair of partons as reported in Ref. [75]. That study demonstrated that the existing searches, up until early Run 1, were not sensitive to the presence of low-mass (< 1 TeV) resonances with small couplings to quarks ($g < 1$). In particular, LHC experiments were affected by the aforementioned limitations in the conventional DAQ approaches for triggering, processing and storing data, resulting in decreased sensitivity to lower mass resonances.

To address this issue, in the last days of the 2011 data-taking period, the CMS Collaboration tested the new data scouting approach for the first time [14]. A preliminary search for dijet resonances using this special data set was performed with pp collision data corresponding to 0.13 fb^{-1} at $\sqrt{s} = 7\text{ TeV}$. This search improved the limits on the production cross section of new dijet resonances in the $0.6\text{--}0.9\text{ TeV}$ range, a region otherwise inaccessible with standard triggers. Since 2012, data scouting has become a well-established approach in CMS, leading to the publication of several physics results on searches for dijet resonances, which are summarized in the following paragraphs.

The basic strategy of these searches consists in reconstructing the two jets that correspond to the pairs of quarks or gluons arising from the decay of a new particle. We therefore look for a peak in the invariant mass distribution of the reconstructed dijet system (m_{jj}), with characteristic shape compatible with the one expected from a resonance decay. The main background from QCD multijet production is predicted by fitting the m_{jj} distribution with an empirical functional form that describes well the QCD simulation and the data in absence of a new physics signal. The main trigger for dijet resonance searches requires H_T to exceed a predefined threshold. As discussed in Section 3.3.1, the scouting trigger has a lower threshold than standard triggers and becomes fully efficient for $m_{jj} > 500\text{ GeV}$, compared to $m_{jj} > 1.25\text{ TeV}$ required by the standard triggers. The data scouting approach is thus able to extend the search for resonances down to the sub-TeV mass range.

The analyses combine calorimeter jets originally reconstructed with the standard anti- k_T algorithm at the HLT with distance parameter $R < 0.4$ or 0.5 (AK jets) into “wide jets”, which are then used to measure the mass spectrum and to search for new physics. The partons from the decay of heavy objects can radiate additional partons, which are often produced at large angles with respect to the original parton direction and thus clustered into a separate AK jet. To reduce this effect, the two p_T -leading AK jets are used as seeds and the four-momenta of all other jets, if within $\Delta R < 1.1$ of the seed jet, are added to it to obtain two wide jets, which then form the dijet system. Wide jets collect more of the final-state radiation compared to AK jets and therefore improve the mass resolution of dijet resonances.

Inclusive searches for dijet resonances have been published using data scouting at both $\sqrt{s} = 8\text{ TeV}$ [15] and 13 TeV [65, 76]. The dijet mass spectrum for the most recent analysis at 13 TeV , shown in Fig. 14 (left), is well described by a smooth background parametrization, and no evidence for the production of new particles is observed. The spectrum is only shown up to about 2 TeV , as standard offline reconstructed data is used for higher dijet masses. Upper limits at 95% confidence level (CL) are reported on the production cross section for narrow resonances, with masses between 0.6 and 1.6 TeV . The limits range between about 0.1 and 50 pb depending on the final state considered for the signal model and the resonance mass hypothesis.

To probe resonance masses below 600 GeV , we focus on events where at least one additional high- p_T jet is produced in association with a dijet resonance, resulting in a three-jet final state.




Figure 14: Left: dijet mass spectra (points) compared to a fitted parametrization of the background (solid curve) for the inclusive search performed in Ref. [65]. Right: dijet mass spectrum (points) compared to a fitted parametrization of the background (solid curve) for the three-jet analysis performed in Ref. [77]. The lower panel shows the difference between the data and the fitted parametrization, divided by the statistical uncertainty of the data. Examples of predicted signals from narrow gluon-gluon, quark-gluon, and quark-quark resonances are shown with cross sections equal to the observed upper limits at 95% CL. Figures taken from Refs. [65] (left) and [77] (right).

The requirement of an additional jet (for example coming from initial state radiation) provides enough energy in the event to satisfy the trigger selection. A search for a dijet resonance decaying to a pair of jets with mass between 350 and 700 GeV is performed using events containing at least three jets [77]. The dijet invariant mass spectrum, calculated from the two jets with the largest transverse momenta in the event, is used to search for a resonance. No significant excess over a smoothly falling background is found, as shown in Fig. 14 (right). Limits at 95% CL are set on the production cross section of a narrow resonance in the range between 1 and 20 pb, depending on the resonance mass. The three-jet final state provides sensitivity to even lower resonance masses than in previous searches with the data-scouting technique.

Following the method presented in Ref. [65], the model-independent upper limits on the cross section of dijet resonances are translated into 95% CL upper limits on the coupling g'_q of a hypothetical leptophobic resonance $Z' \rightarrow q\bar{q}$ as a function of its mass. Figure 15 shows the upper limits obtained by various CMS searches for dijet resonances. The aforementioned analyses with the data scouting technique provide the best limits in the mass region from 400 to 1600 GeV.

The mass region below 350 GeV is probed by analyses that rely on standard triggers and study events where the hypothetical resonance is produced with sufficiently high transverse momentum such that its decay products are merged into a single jet, with a two-prong substructure (boosted dijet) [78, 79]. A signal would be identified as a peak over a smoothly falling background in the distribution of the invariant mass of the jet, using jet substructure techniques. These analyses study resonances produced in association with a high- p_T photon or jet and

probe the resonance mass range 10–125 GeV and 50–500 GeV, respectively. For full efficiency with respect to the standard trigger requirements, events are selected by demanding the presence of a photon with $p_T > 200$ GeV in the first case or a jet with $p_T > 500$ GeV in the second. Future developments for these analyses include exploiting the scouting triggers, which would significantly reduce the photon and jet p_T trigger thresholds and hence improve the signal efficiency and sensitivity of these analyses. A similar approach has been investigated for the study of boosted $H \rightarrow b\bar{b}$ decays, as reported in Section 4.3. As demonstrated in Section 3.3.2, the use of jet substructure techniques is now established in data scouting and could be also applied to the case of boosted dijet resonance searches to reconstruct the jet mass and identify the two-prong jet substructure.




Figure 15: Observed limits on the universal coupling g'_q between a leptophobic Z' boson and quarks [76] from various CMS dijet analyses. Regions above the lines are excluded at 95% CL. The grey dashed lines show the g'_q values at fixed values of $\Gamma_{Z'}/m_{Z'}$. Limits from scouting-based analyses are indicated with bold lines.

3.5.1.2 Multijet resonances

Scouting data is ideal for new physics searches in regions of parameter space dominated by large backgrounds, such as resonances decaying to multiple jets. New physics signatures with multijet final states can be produced in several ways, including from the decay of new colored particles. Relevant SUSY models include RPV squarks and gluinos, which can produce paired dijets and paired jet triplets, respectively [80]. For low masses, the partons arising from the decays of these squarks and gluinos can merge into a single jet. To probe such low masses, scouting analyses utilize jet substructure techniques [51, 52]. Here we highlight the results of searches for pairs of two- and three-parton resonances [67], interpreted as RPV squarks and RPV gluinos, respectively. The analysis relies on a data set collected with the PF scouting triggers, which require thresholds as low as $H_T > 410$ GeV at the HLT. The data set also stores relevant jet substructure variables, allowing us to probe low-mass resonances in which the partons merge into single jets, as discussed in Section 2.3.5.

The investigation of the multijet phenomena follows three different paths within the analysis:

pairs of large-radius jets with substructure consistent with three underlying quarks (*merged trijets*), pairs of large-radius jets with two-quark substructure (*merged dijets*), and pairs of well-resolved triplets of jets (*resolved trijets*). Given the characteristic decay of RPV gluinos into three quarks in the final state, high-mass gluino pair production is studied with resolved triplets of jets while gluinos of lower mass are studied with dijet events wherein each jet exhibits substructure indicative of the merging of three partons into a single large jet. The RPV squarks undergo decay into pairs of quarks. These decays manifest as events with two jets, where the substructure of each jet aligns with the fusion of two partons. The merged dijet study centers on squark pair production scenarios characterized by masses below 200 GeV. The resolved trijet analysis exhibits sensitivity to RPV gluinos across the mass spectrum of 200–2000 GeV. In contrast, the merged trijet analysis leverages the τ_{32} jet substructure variable introduced in Section 2.3.5, formulated utilizing the designed decorrelated tagger (DDT) technique [81]. With this analysis, the search sensitivity is extended to resonance masses as low as 70 GeV.

For the resolved trijet analysis, all pairs of jet triplets are analyzed using kinematic variables that differentiate between multijet backgrounds and signal triplet pairs. The jets within the triplets are subjected to QGD methods, as described in Section 3.3.2. Figure 6 shows the invariant three-jet mass for triplets that pass all selection criteria. In that figure, we show the lowest mass range used in the analysis, as well as the jet mass distribution for the low mass gluino search where the three partons merge into a single jet. The top quark mass is clearly discernible in both distributions. The sensitivity to gluinos in searches using scouting data is better than searches performed by the CDF Experiment at the Tevatron [82] and other searches by LHC experiments [83–87], achieving both a lower mass reach and lower cross section limits as shown in Fig. 16.

Searches for final states consisting of pairs of two merged partons were previously performed using standard CMS triggers [88, 89], with limits of roughly 500 GeV set on RPV squarks. Limits on the production cross section of RPV squarks with the scouting data are shown in Fig. 17. Using data scouting and $N_{2,DDT}$ jet substructure techniques, we have extended the sensitivity to RPV squark masses down to 70 GeV.

In summary, with the scouting technique CMS has achieved unprecedented sensitivity to hadronic resonances with low masses. In the case of new particles decaying to three partons, we are sensitive to weak production (Higgsino) cross sections.

3.5.2 Searches in muon final states

Searches for resonant pair production in dilepton final states played a crucial role in the development of the SM, leading to, e.g., the discovery of the charm and bottom quarks via $J/\psi \rightarrow \mu\mu$ and $Y \rightarrow \mu\mu$ decays, respectively, and of the Z boson via $Z \rightarrow ee/\mu\mu$ decays. Today the same approach is used to investigate still unexplored regions of phase space and to look for new particles in dilepton mass spectra. Here we present searches for the production of dimuon resonances in the mass range below 200 GeV. Both scenarios where the resonance decays promptly or displaced from the interaction point are considered.

3.5.2.1 Prompt dimuon resonances

The availability of the scouting data set makes it possible to focus on searches for prompt dimuon resonances in the mass range below 40 GeV. Several sources of cosmological evidence point to the existence of dark matter, and a clear understanding of the origin and nature of dark matter is one of the open problems of particle physics. Dark matter is expected to interact very weakly, if at all, with SM particles. This raises the possibility of a hidden, dark sector of




Figure 16: Comparison of limits from searches for RPV gluinos decaying to three partons. Regions above the lines are excluded at 95% CL. The two CMS analyses that use data scouting are also indicated with bold lines.




Figure 17: Observed (points) and expected (dashes) limits on the product of production cross section, branching fraction, and acceptance for pair-produced merged two-quark resonances. The variations at the one and two standard deviation levels in the expected limits are displayed with shaded bands. A comparison with the theoretical predictions for top squark production (red) is also shown. Figure taken from Ref. [67].

particles. As a prototypical hidden sector, we consider the compelling possibility of a spontaneously broken dark $U(1)_D$ gauge symmetry, mediated by a vector boson called the dark photon, Z_D . The dark photon's only renormalizable interaction with the SM is through kinetic mixing with the hypercharge gauge boson, controlled by the kinetic mixing coefficient ϵ . If there are no other hidden-sector states below the Z_D mass, this mixing causes the dark photon to decay exclusively to SM particles, with a sizable branching fraction to leptons. Given the extraordinary capabilities of CMS to reconstruct and identify muons down to a p_T of just a few GeV, searching for dark photons in the dimuon channel is a clear natural target.

Two searches were performed using the full muon scouting data set collected during Run 2, corresponding to $\mathcal{L}_{\text{int}} = 96.6 \text{ fb}^{-1}$. The mass spectrum up to 200 GeV was scanned to search for a narrow resonance with subsequent prompt decay to a pair of oppositely charged muons. Different strategies were adopted according to the mass window. The regions around the known resonances, namely the J/ψ , $\psi(2S)$, $Y(1S)$, and Z , were excluded because of the difficulty in looking for a new particle in the vicinity of existing resonances with the same final-state signature.

The first search [68] investigated the resonance mass ranges of 45–75 and 110–200 GeV by exploiting conventional trigger paths and event reconstruction techniques. The coverage was extended down to 11.5 GeV by exploiting the scouting triggers. The dimuon mass resolution depends strongly on the pseudorapidity of the muons. The p_T resolution of muons with $p_T < 50 \text{ GeV}$ is around 1% in the barrel region of the detector and 3% in the endcaps. Therefore, events are divided in two categories based on the pseudorapidity of the muons. In the search performed with the scouting triggers, events are required to contain two muons of opposite charge that are consistent with same-vertex production. The muons are required to be well isolated and to pass selection requirements based on the track quality information available in the scouting event content. To suppress sources of background involving muons originating from heavy flavor decays that typically have low p_T , the muons with the largest and second largest p_T are required to have $p_T > m_{\mu\mu}/3$ and $p_T > m_{\mu\mu}/4$, respectively.

The data are found to be consistent with the background prediction. The results of this search are interpreted in the context of the dark photon model introduced earlier. Upper limits are provided at 95% CL on the product of the signal cross section, branching fraction to a pair of muons, and kinematic and geometrical CMS acceptance of a narrow resonance. Moreover, expected and observed upper limits at 90% CL on ϵ^2 as a function of Z_D mass are obtained and compared with the existing results by the LHCb Collaboration [90], as shown in Fig. 18. The search using scouting data sets stringent constraints on dark photon production in the 11–45 GeV mass range.

The second search [71] is an extension of the first, focused on the mass window below the Y resonance peak, in the 1.1–2.6 and 4.2–7.9 GeV mass ranges. The region around the J/ψ peak was excluded. A dedicated MVA muon identification technique trained on control samples in data was used to enhance the sensitivity to this very low mass region. This strategy allowed the optimization of the selection of a promptly produced dimuon resonance while minimizing the rate of muon misidentification. Details are provided in Section 3.4.2. The training of the algorithm was performed on J/ψ and $Y(1S)$ events for the lower and higher mass windows, respectively. The algorithm trained on J/ψ events recovered the selection efficiency in the very low mass region, compared to the one trained on Y events, as demonstrated by the dimuon spectra obtained with the two different selections in Fig. 19.

Event candidates are selected by requiring at least one interaction vertex, as reconstructed by the HLT system, and a pair of oppositely charged muons originating from this vertex. A dedi-




Figure 18: Expected and observed upper limits at 90% CL on the square of the kinetic mixing coefficient (ϵ^2) as a function of dark photon mass. Results obtained with scouting triggers are displayed to the left of the vertical purple line, while those obtained with standard triggers are shown to the right. Limits at 90% CL obtained from the search performed by the LHCb Collaboration [90] are shown in red, and constraints at 95% CL from the measurements of electroweak observables are shown in light blue [91]. Figure taken from Ref. [68].




Figure 19: The $m_{\mu\mu}$ distribution obtained with the scouting data collected during 2017 and 2018 with two sets of selections: the J/ψ -trained (red) and the $Y(1S)$ -trained (blue) MVA-based muon identification algorithms. Figure taken from Ref. [71].

cated mass-dependent vertex displacement criterion is applied to focus on promptly produced dimuon resonances. Two selections are defined: an inclusive one and a high- p_T one, focusing on Drell–Yan and gluon fusion production, respectively. The two signal categories, referred to as inclusive and boosted, are used to obtain an interpretation of the results in the context of two specific models: a minimal dark photon model (as in the previous search), and a model with two Higgs doublets as well as an extra complex scalar singlet (2HDM+S). The inclusive selection requires muon $p_T > 4$ GeV in the pseudorapidity region $|\eta| < 1.9$, while the boosted selection requires muon $p_T > 5$ GeV and a dimuon p_T larger than 35 (20) GeV in the mass region below (above) 4 GeV.

The signal is extracted from maximum likelihood fits to the $m_{\mu\mu}$ distribution in data corresponding to selected events. The fit relies on a signal-plus-background model under the assumption that the natural width of the new resonance is much smaller than the detector dimuon mass resolution. Various empirical functions are investigated and used to model the background shape and to estimate the associated systematic uncertainties. For each mass hypothesis, the fit is performed over a mass window spanning ± 5 or ± 8 times the mass resolution around the hypothesized resonance mass for the inclusive and boosted scenario, respectively. Model-independent limits on the product of the resonance production cross section, branching fraction to muons, and geometrical acceptance are computed. The results are also interpreted as constraints on the parameters of the two models introduced earlier. The model-specific limits rely on the theoretical calculation of cross sections and branching fractions, and on the experimental acceptance derived with simulation. The sensitivity to the kinetic mixing coefficient ϵ is significantly improved or comparable to the one obtained by the LHCb [92] and BaBar [93] Collaborations, as shown in Fig. 20.




Figure 20: Upper limits at 90% CL on the square of the kinetic mixing coefficient (ϵ^2) in the minimal dark photon model obtained as a recast of model-independent limits on the production rates of dimuon resonances for the inclusive category. The CMS limits (pink) are compared with the existing limits at 90% CL provided by LHCb [92] (blue) and BaBar [93] (gray). Figure taken from Ref. [71].

3.5.2.2 Displaced dimuon resonances

The scouting triggers do not require muons to be associated with the reconstructed PV, as described in Section 3.4.1. Therefore, the muon scouting data stream can also be used to search for displaced muon signatures. A search for narrow, long-lived dimuon resonances [69] was performed based on dimuon data collected with the CMS experiment during the LHC Run 2 in 2017 and 2018 using the dimuon scouting trigger stream, with muon $p_T > 3$ GeV and $|\eta| < 2.4$, and probing resonance masses down to $\approx 2m_\mu$. The selected data correspond to $\mathcal{L}_{\text{int}} = 101 \text{ fb}^{-1}$.

The search targets narrow, low-mass, long-lived resonances decaying into a pair of oppositely charged muons, where the lifetime of the long-lived particle is such that the transverse displacement l_{xy} of its decay vertex is within 11 cm of the PV. As mentioned in Section 3.4.1, this constraint arises from the requirement for muon tracks to leave a hit in at least two layers or disks of the pixel tracker. The selected muons are used in pairs to form dimuon vertices, considering all possible pairs. These vertices are hereafter referred to as SVs, and they may or may not be displaced from the PV. The signal is expected to appear as a narrow peak on top of

the dimuon mass continuum, with an intrinsic resonance width smaller than the experimental mass resolution. Such a signal is predicted in BSM frameworks with the Higgs boson decaying into a pair of long-lived dark photons, as shown in Fig. 21 (left), or with a long-lived scalar resonance arising from a decay of a b hadron, as shown in Fig. 21 (right).




Figure 21: Left: diagram illustrating an SM-like Higgs boson (H) decay to four leptons (ℓ) via two intermediate dark photons (Z_D). Right: diagram illustrating the production of a scalar resonance ϕ in a b hadron decay, through mixing with an SM-like Higgs boson. Figure taken from Ref. [69].

Events are required to contain at least one pair of opposite-sign muons associated with a selected SV. The events that contain a single muon pair are then categorized according to the decay length l_{xy} , the p_T of the muon pair ($p_T^{\mu\mu} < 25$ and $p_T^{\mu\mu} > 25$ GeV), and isolation (to distinguish isolated, partially isolated, and nonisolated topologies). The l_{xy} categorization is intended to maximize the search sensitivity to a range of potential BSM signal lifetimes, and is based on the CMS pixel tracker geometry. The categorization in $p_T^{\mu\mu}$, on the other hand, improves the sensitivity to signal models with different production modes and boost distributions.

In each (l_{xy} , $p_T^{\mu\mu}$, isolation) bin, we define mass windows sliding along the dimuon invariant mass spectrum, and perform a search for a resonant peak in each mass window. The size of the sliding windows is set to $\pm 5 \sigma_{\mu\mu}^{\text{mass}}$ around the signal mass hypothesis, where $\sigma_{\mu\mu}^{\text{mass}}$ is determined from simulation and equals about 1% of the mass. Mass regions corresponding to known resonances decaying either to a pair of muons or to a pair of charged hadrons are not considered, i.e., they are “masked” for the purpose of this search. In events with two muon pairs, each associated with an SV, we further require the difference between the two dimuon masses to be within 5% of their mean, and the four-muon mass to be consistent with the mass of the SM Higgs boson ($115 < m_{4\mu} < 135$ GeV). The selected four-muon events are treated as an exclusive independent category, with a single four-muon mass window centered around the known Higgs boson mass (i.e., $115 < m_{4\mu} < 135$ GeV). This additional region is aimed at improving the search sensitivity to models of BSM physics where an SM Higgs boson decays to a pair of Z_D bosons, each decaying to two muons.

In each mass window, the signal is parametrized using the sum of a Gaussian function and a double-sided CB function. The SM background is modeled by means of different functional forms, that include Bernstein polynomials, exponential functions, and combinations of the two. Then, binned maximum likelihood fits to the data are performed simultaneously in all search bins, under either background-only or background-plus-signal hypotheses.

No significant peak-like structures are observed in data. The background-plus-signal fits are used to set upper limits at 95% CL on a wide range of mass and lifetime hypotheses for models

of BSM physics where a Higgs boson decays to a pair of long-lived dark photons, or where a long-lived scalar resonance arises from the decay of a b hadron. These constraints are the most stringent to date for substantial regions of the explored parameter space. For illustrative purposes, the exclusion limits obtained for the $H \rightarrow Z_D Z_D$ model of Ref. [91] and for the $h_b \rightarrow \phi X$ inflaton model of Ref. [94] are shown in Fig. 22. The search is mostly sensitive to signatures with a dimuon resonance produced at nonzero displacement from the PV. At large displacement values, the sensitivity is degraded because the transverse displacement must be within the first three layers of the pixel detector, namely $l_{xy} < 11$ cm.

3.5.3 Observation of the rare $\eta \rightarrow 4\mu$ decay channel

The scouting data set also enables the study of rare light-meson decays to be considered. The lower muon momentum thresholds considerably expand the mass range of particles that can be probed. The power of the data set to measure rare SM decays was demonstrated by the first observation of the four-muon decay of the η meson. The production rate of the η meson falls quickly vs. p_T , so lowering the muon momentum thresholds is essential to enhance the collection of events involving the η meson.

The $\eta \rightarrow \mu^+ \mu^- \mu^+ \mu^-$ decay was observed by the CMS Collaboration using a four-muon selection in the scouting data set. Figure 23 shows the measured distribution of the four-muon invariant mass, after requiring four muons that are compatible with same-vertex production. About 50 signal events ($N_{4\mu}$) are observed on top of a background of roughly 17 events, corresponding to a statistical significance much greater than 5 standard deviations.

The branching fraction of the newly observed decay channel was measured relative to the reference channel $\eta \rightarrow 2\mu$, which is known with a precision of about 14%. The additional ingredients needed for this measurement are the yields of the two-muon reference channel, $N_{2\mu}$, and the products A of the CMS detector geometric acceptance and reconstruction efficiency for both channels, which were determined with simulation. The relative branching fraction is computed via

$$\frac{\mathcal{B}_{4\mu}}{\mathcal{B}_{2\mu}} = \frac{N_{4\mu}}{\sum_{i,j} N_{2\mu}^{i,j} \frac{A_{4\mu}^{i,j}}{A_{2\mu}^{i,j}}},$$

where $N_{4\mu}$ is the total four-muon signal yield, and $A_{4\mu}^{i,j}$, $A_{2\mu}^{i,j}$, and $N_{2\mu}^{i,j}$ are the four-muon A , two-muon A , and two-muon yields in bin i of the candidate η meson p_T and bin j of the η meson rapidity, respectively. We define 32 bins in p_T , in the range 7–70 GeV, and 2 bins in $|y|$.

Figure 24 shows A for both channels as determined by simulation. The efficiency of the two-muon channel is limited by the trigger efficiency, while that of the four-muon channel is constrained by the efficiency to reconstruct all four signal muons. This is more challenging for higher η meson p_T , since the muons are more collimated, leading to overlapping tracks and decreased reconstruction efficiency.

The measured relative branching fraction is

$$\frac{\mathcal{B}_{4\mu}}{\mathcal{B}_{2\mu}} = [0.86 \pm 0.14 \text{ (stat)} \pm 0.12 \text{ (syst)}] \times 10^{-3},$$

where the systematic and statistical uncertainties are roughly balanced. Using the world average value for the reference branching fraction $\mathcal{B}(\eta \rightarrow 2\mu)$, the absolute branching fraction of the four-muon decay is measured as

$$\mathcal{B}(\eta \rightarrow 4\mu) = [5.0 \pm 0.8 \text{ (stat)} \pm 0.7 \text{ (syst)} \pm 0.7 (B_{2\mu})] \times 10^{-9},$$




Figure 22: Observed limits at 95% CL on (upper) the branching fraction $\mathcal{B}(H \rightarrow Z_D Z_D)$ and (lower) the branching fraction product $\mathcal{B}(h_b \rightarrow \phi X) \cdot \mathcal{B}(\phi \rightarrow \mu\mu)$ as contours in the parameter space containing the signal mass (m_{Z_D} or m_ϕ , respectively) and the signal lifetime $c\tau_0$. The vertical gray bands indicate mass ranges containing known SM resonances, which are masked for this search. The limits are obtained using the combination of all dimuon and four-muon event categories. Figures taken from Ref. [69].

where the last term is the uncertainty in the branching fraction of the reference channel. This result is in agreement with theoretical predictions, e.g., $\mathcal{B}(\eta \rightarrow 4\mu) = (3.98 \pm 0.15) \times 10^{-9}$ from Ref. [96], and improves on the precision of previous upper-limit measurements [97] by more than 5 orders of magnitude.




Figure 23: The four-muon invariant mass ($m_{4\mu}$) distribution in the range 0.46–0.90 GeV, obtained with pp collision data collected during 2017–2018. The observed distribution (points) is compared to the background-only prediction (green dashed) and to the full background fit including simulations of the signal (solid blue). The peak observed in the mass window 0.53–0.57 GeV corresponds to the η meson. The pull distribution in the lower panel is shown relative to the background component of the fit model and defined as $(\text{Data} - \text{Fit}) / \text{Uncertainty}$, where the uncertainty is statistical only. Figure taken from Ref. [95].




Figure 24: The product of acceptance (A) and efficiency as a function of the generated η meson p_T for the two-muon (red circles and blue squares) and four-muon (orange up triangles and green down triangles) η meson decays. The product is evaluated using simulated samples. Figure taken from Ref. [95].

4 Data scouting in Run 3

The CMS data scouting technique was explored, developed, and brought to maturity during Run 1 and Run 2, proving to be an innovative trigger strategy and a successful paradigm for data analysis. During this period, the primary constraint in implementing the scouting strategy was found to be the HLT event processing time. In Run 3, the computing capabilities of the HLT were greatly improved, as described in Ref. [31]. The availability of a new GPU-equipped HLT farm, provided as a way to gain expertise for the next phase of operation of the CMS detector during HL-LHC, and the subsequent improvement of the HLT reconstruction (as detailed in the next section) facilitated a significant broadening of the scouting scope. It is interesting to note that comparable performance based on an HLT farm equipped only with CPUs would require a higher overall cost (by approximately 15%) and higher power consumption (by about 30%). This expansion of data scouting is illustrated in Fig. 2, where the average rates allocated to the standard, parking, and scouting streams are reported from Run 1 to Run 3.

This chapter outlines the improvements to the data scouting strategy in Run 3, providing insights into the relevant updates with respect to the Run 2 approach described previously in Chapter 3. First, the trigger rates and event content of the 2022 and 2023 data-taking periods are discussed in Sections 4.1 and 4.2. Then, we investigate the performance of data scouting in the context of jets (Sections 4.3), muons (Section 4.4) and electrons and photons (Section 4.5).

4.1 The new Run 3 scouting strategy at the HLT

Following the hardware upgrade and increased usage of GPUs during Run 3, the HLT algorithms were redesigned to harness the capabilities of parallel architectures. The primary focus was to offload HLT reconstruction steps to GPUs, particularly for complex tasks. As a result, during 2021 and 2022 a new GPU-based approach was developed and fully commissioned for the calorimeter reconstruction, the pixel local reconstruction, and the pixel-based tracking.

Within this new GPU-based paradigm, events are reconstructed with a novel scouting PF algorithm that exploits tracker tracks built solely with pixel hits, using the PATATRACK algorithm [36]. These pixel-only tracks replace the tracks reconstructed with the combined information from the pixel and strip trackers, as done with the standard PF algorithm described earlier in Section 2.3.4. The main advantage of the PATATRACK pixel-only tracking is the possibility of offloading the track reconstruction to GPUs, thereby notably accelerating event processing at the cost of a slightly worse track resolution compared to standard tracks. The degradation is more significant in high- p_T tracks [98]. As low- p_T tracks are most relevant to scouting, this acceleration particularly benefits the scouting strategy. Quality criteria based on the momentum resolution and on the distance in the longitudinal plane from the two leading vertices are applied to the pixel-only tracks before offering them as input to a modified PF algorithm. Vertices are reconstructed using pixel-only tracks and the measured transverse coordinates are computed relative to the online measurement of the interaction point.

The processing time required for the full reconstruction of scouting events and the application of the selection criteria are shown in Table 5 for the scouting paths active during 2023. The difference between scenarios with and without outsourcing of certain steps to GPUs is also presented. The significant speed-up provided by the GPUs is clearly seen.

In the Run 3 scouting strategy, unlike in Run 2, only jets clustered from PF inputs are reconstructed and stored. The reconstruction of PF jets follows the same method employed in Run 2 (see Section 2.3.5), but uses PF candidates reconstructed from pixel-only tracks. All PF candidates are used to calculate the missing transverse momentum. More details on the performance

Table 5: Time required for the object reconstruction and selection criteria in the scouting paths seeded by different L1 algorithms, using only a CPU or accelerating certain steps with a GPU. For the comparison, we pick a representative run recorded in 2023. The time needed to run the full HLT menu reconstruction including non-scouting paths is also shown for reference.

Scouting path	CPU-only [ms]	CPU+GPU [ms]
1 electron/photon	76.0	49.5
≥ 2 electrons/photons	9.3	6.8
≥ 2 muons	69.0	41.6
Jets or MET	83.3	52.1
Full HLT menu	578.4	377.7

of jet objects and algorithms is presented in Section 4.3.

Muon candidates are reconstructed based on information from the silicon tracker and the muon system, as described in Section 3.4. The reconstruction of scouting muons benefited from the integration of ML-based outside-in and inside-out seeding improvements in the standard muon reconstruction at HLT for Run 3 [31]. A relevant difference between the online and offline reconstruction is represented by the removal of the requirement on the minimum number of hits in the pixel layers, which was introduced at the beginning of Run 3. This adjustment enhances the sensitivity to signals with displaced muons in the final state, since, depending on their displacement from the collision region, they may not create many hits in the pixel layers. More details are available in Section 4.4 where we discuss the Run 3 muon performance.

The electron and photon reconstruction is a novelty in Run 3 scouting, made possible by the newly available resources arising from the offloading of pixel-only tracking to GPUs. The reconstruction of electrons and photons themselves is identical to the standard online reconstruction, as discussed in Chapter 2. The scouting path runs reconstruction over the full ECAL volume. Since 2023, an HLT preselection is applied, as shown in the third column of Table 6, in the scouting trigger paths seeded by the L1 e/γ triggers. The preselection reduces the number of events on which the scouting reconstruction is run and frees up resources to enable reconstruction over the full ECAL volume irrespective of the L1 seed. Most other HLT paths reconstruct SCs in a geometrical region matching an L1 ECAL trigger tower. The full ECAL reconstruction results in larger efficiency for low-energy e/γ in the scouting events. The performance of electron and photon reconstruction is further discussed in Section 4.5.

The substantial improvement in reconstruction speed enables an expansion of the maximum L1 event rate that can be processed by the scouting stream and thereby a deeper exploration of physics processes at lower masses and weaker couplings. For events to be reconstructed in the scouting data stream, they must be selected by one of several L1 seeds targeting signatures such as one or two photons or electrons, muons with low p_T , one or two jets, or a moderate amount of H_T . The scouting streams in Run 3 sustained an L1 input rate of approximately 30 kHz for high pileup data-taking scenarios. Figure 25 illustrates the fractional rate relative to the scouting L1 input rate of each L1 category for two reference runs in 2022 and 2023, with average pileup of 60. The list of L1 seeds used in Run 3 is summarized in Table 6. The most notable change with respect to the list of L1 seeds previously reported for Run 2 in Section 3.2 is the inclusion of dedicated algorithms targeting events with one or two electrons or photons. The single e/γ trigger has been included to target single-photon signatures that exploit the notable lowering of the photon p_T threshold in scouting. The list of algorithms for muons and

jets/ H_T remained largely unchanged, except for the temporary removal of the lowest threshold dimuon seed in 2022 (which was added back in 2023) and the lowering of the dijet invariant mass and H_T requirements in 2023. The latter is evident in the increased proportional rate of the hadronic category. The rate variation of the e/ γ categories is due to changes in the online data-taking conditions at L1 and not due to updates of the algorithms.




Figure 25: Relative rate of each L1 algorithm category, shown as the fraction of the total rate based on the 2022 (orange) and 2023 (blue) configurations. The proportional rate of each category with respect to the total is shown on the right. The values are computed using reference runs with average pileup of 60.

As discussed in Section 3, due to the limited computational resources in Run 2, the scouting data was split into two separate streams: PF scouting, which involved the computationally intensive PF reconstruction, and Calo scouting, which used fewer resources. But with the increased computing power available in Run 3, a transition was made to a single scouting data stream where the PF reconstruction is run for all events passing one of the input L1 algorithms. After reconstruction, a minimal event selection is applied before storing events on disk. Therefore, the HLT scouting data stream achieved a data output rate of approximately 26 kHz for high pileup data-taking scenarios during the beginning of Run 3.

In 2023, optimized selections were added to improve the purity of certain final-state topologies. These selections vary based on the L1 algorithm seed used to collect a given event. The selection criteria are listed in Table 6. No further selection was applied to scouting events seeded by the jet and H_T algorithms. In the dimuon path, events were only stored if they contained at least two scouting muons, each with $p_T > 3$ GeV. Events seeded by the single e/ γ path required an e/ γ candidate with $p_T > 30$ GeV, and the L1 double e/ γ path required at least two e/ γ candidates, each with $p_T > 12$ GeV.

A comparison of the average HLT output rates for the standard, parking, and scouting streams during Run 2 and Run 3 is reported in Table 7 and also shown in Fig. 2. The average scouting rate decreased from 2022 to 2023 as a result of the additional object and event selections listed earlier. Table 8 shows the peak rates of the scouting paths, corresponding to two reference runs with an average pileup of 60 recorded during 2022 and 2023. The rates are consistent with the

Table 6: List of L1 and HLT thresholds for the lowest unprescaled scouting triggers during Run 3. The list corresponds to the 2022 and 2023 data-taking periods. Conditions that changed between the two years are annotated in bold face. The four separate HLT paths and corresponding thresholds were only present in 2023. No thresholds were applied at the HLT in 2022. In cases where the same threshold is applied to all selected objects in an event, a single number is shown, while if different thresholds are applied to the objects, they are shown separated by slashes from the highest to the lowest. The notation “OS” stands for opposite-sign muon pairs and “SC” for calorimeter superclusters.

Type	L1 threshold	HLT threshold (2023)
e/γ	$1 e/\gamma, p_T > 30 \text{ GeV}, \eta < 2.1$ $2 e/\gamma, p_T > 18/12 \text{ GeV}, \eta < 1.5$	$1 \text{ SC (loose)}, p_T > 30 \text{ GeV}$ $2 \text{ SC (loose)}, p_T > 12 \text{ GeV}$
μ	$2\mu, p_T > 15/7 \text{ GeV}$ $2\mu, \text{OS}, p_T > 4.5 \text{ GeV}, \eta < 2, m_{\mu\mu} > 7 \text{ GeV}$ $2\mu, \text{OS}, p_T > 4 \text{ GeV}, \eta < 2.5, \Delta R < 1.2$ $2\mu, \text{OS}, p_T > 0 \text{ GeV}, \eta < 1.5, \Delta R < 1.4 \text{ (2023)}$ $3\mu, p_T > 5/3/3 \text{ GeV}$	$2\mu, p_T > 3 \text{ GeV}$
	$H_T > 280 \text{ (2023)}, 360 \text{ (2022)} \text{ GeV}$	
Jets/ H_T	$1 \text{ jet}, p_T > 180 \text{ GeV}$ $2 \text{ jets}, p_T > 30 \text{ GeV}, \eta < 2.5, \Delta\eta < 1.5,$ $m_{jj} > 250 \text{ (2023)}, 300 \text{ (2022)} \text{ GeV}$	

scenarios shown for the L1 input rates in Fig. 25.

Table 7: Comparison of the typical HLT trigger rates of the standard, parking, and scouting data streams during 2018 (Run 2), 2022, and 2023 (Run 3). The average $\mathcal{L}_{\text{inst}}$ value over one typical fill of a given data-taking year and the average pileup (PU) are also reported, coherently with the scenarios reported in Fig. 2.

Year	$\mathcal{L}_{\text{inst}} [\text{cm}^{-2}\text{s}^{-1}]$	PU	Standard rate [Hz]	Parking rate [Hz]	Scouting rate [Hz]
2018	1.2×10^{34}	38	1000	3000	5000
2022	1.5×10^{34}	46	1800	2440	22000
2023	1.7×10^{34}	48	1700	2660	17000

4.2 Event size and content

The event size reduction needed for the scouting strategy is achieved by applying selection criteria on the reconstructed physics objects and by storing high-multiplicity quantities with reduced numerical precision on the mantissa (10 bits). The amount of information in the data scouting output increased in Run 3 compared to Run 2 due to the added reconstruction of electrons and photons, and the storage of additional quantities. Nevertheless, the overall event size for the Run 3 scouting was kept smaller than 10 KB thanks to the selection and precision optimization.

Reconstructed PF candidates are stored together with their kinematic variables and their particle type as identified by the PF algorithm, if they have $p_T > 0.6 \text{ GeV}$ and $|\eta| < 3$ (similar to the Run 2 approach). If a PF candidate is a charged particle with an associated pixel-only track,

Table 8: HLT rates for the scouting paths seeded by different L1 algorithms during two reference runs with an average pileup of 60 recorded in 2022 and 2023.

Configuration	2022	2023
Scouting path	Rate per path [kHz]	
1 e / γ	—	9.1
≥ 2 e / γ	—	0.3
≥ 2 muons	—	3.4
Jets or H_T	—	11.0
e / γ , ≥ 2 muons, jets or H_T	31.3	—

basic track parameters and a reference to the vertex associated with that track are also stored. Track-related quantities of PF candidates were added in Run 3 that, for example, are used as input when training neural networks such as PARTICLENET [57] for tasks, such as jet flavor identification. Additionally, for reconstructed tracks with $p_T > 3$ GeV, the kinematic information, hit pattern, track parameters, and track fit quality, including the covariance matrix for the refitting of vertices, are stored. Finally, for vertices reconstructed from the pixel-only or muon tracks, their position and associated uncertainties as well as the χ^2 and number of degrees of freedom in the vertex fit, are stored.

Jet selection criteria require that jets have $p_T > 20$ GeV and $|\eta| < 3$ to be stored in the scouting stream. The same variables as in Run 2 are retained for each jet, which is detailed in Section 3.2. While the Run 3 event content stores only AK4 jets, the inclusion of PF candidates allows for subsequent offline clustering of PF jets with any distance parameter. The missing transverse momentum reconstructed from the PF candidates is stored on an event basis.

Muons with $|\eta| < 2.4$ are stored with their kinematic and isolation quantities, as previously discussed in Section 3.2 for the Run 2 scenario. Detailed track information is also available, enabling the refitting of dimuon vertices. All reconstructed dimuon vertices are stored separately with their positions, associated uncertainties, and fit qualities. Finally, electrons and photons are stored if they satisfy $p_T > 2$ GeV and $|\eta| < 2.5$. To reduce the processing workload, the time-consuming track reconstruction for electrons is initiated only when the energy in the HCAL directly behind the ECAL supercluster (within a cone $\Delta R < 0.15$) is less than 20% of the supercluster energy, and at least two hits are observed in the pixel tracker. The list of quantities stored in the scouting stream that relate to electrons and photons is reported in Table 9.

The reconstruction and identification performance of the jets, muons, electrons, and photons stored in the Run 3 scouting stream is discussed in the next sections. A special scouting monitoring data set that collects a randomly chosen reduced number of events (about 1% of the total) is used for most of the studies described in the following sections. Both offline and scouting objects are available in this data set, allowing for an easier comparison between the two strategies over the same set of events.

4.3 Jets

The scouting strategy enables lower hadronic trigger thresholds than the standard strategy relying on offline reconstructed data. As listed in Table 6, the data scouting trigger in 2022 included unprescaled L1 seeds targeting events based on the presence of at least one jet with p_T exceeding 180 GeV or H_T exceeding 360 GeV. In comparison, the lowest unprescaled triggers in the standard trigger strategy required jet p_T or H_T to exceed 500 GeV or 1050 GeV, respectively.

Table 9: List of observables related to the newly introduced scouting electrons (e) and photons (γ) stored in the scouting data set in Run 3. The calorimeter observables shared by e / γ objects are listed in the upper part of the table, while the track features specific to electrons are reported in the lower part. Tracker-based isolation for photons is to be computed offline from the stored PF candidates.

Observable	Definition
<i>Electron/photon common quantities (calorimeter-based)</i>	
(E, p_T, η, ϕ)	ECAL SC four-momentum
$\sigma_{i\eta i\eta}$	Spread of the ECAL shower from the central crystal
H/E	Ratio of energy deposit in HCAL to ECAL
I_E	ECAL isolation
I_H	HCAL isolation
r_9	Relative energy deposit in 3x3 η - ϕ matrix
seed ID	Crystal number of the central crystal
energy matrix	Energy deposit in each crystal of the SC
detector ID	Crystal number of each crystal of the SC
time matrix	Time stamp of each crystal of the SC
s_{minor} and s_{major}	Second moments of the SC energy matrix
rechitZeroSuppression	Flag indicating events with nonzero reconstructed hits
<i>Electron quantities only (tracker-based)</i>	
track (E, p_T, η, ϕ)	GSF track four-momentum
track d_0	Track d_0
track d_z	Track d_z
track χ^2/dof	Reduced- χ^2 of the track fit
track missing hits	Missing hits in the tracker inner pixel region
track q	Track charge
$\Delta\eta_{\text{in}}^{\text{seed}}$	Difference in η between central ECAL crystal and inner track
$\Delta\phi_{\text{in}}$	Difference in ϕ between SC and inner track
$1/E - 1/p$	Difference between the inverse of SC E and track p
I_{track}	Track isolation

The benefits of the lower trigger thresholds in scouting are quantified in Fig. 26, where the trigger efficiencies of the scouting and standard trigger selections are compared using collision data recorded in 2022. The trigger selection is a logical “OR” expression of all L1 seeds and HLT triggers. Events are selected based on the presence of at least one energetic jet or sufficiently energetic H_T . The efficiency is displayed as a function of the offline reconstructed AK4 PF jet p_T (left), AK8 PF jet p_T (center), and PF H_T (right).

The efficiency is measured using an unbiased sample of events, collected with a single-muon trigger and containing only one well-identified and isolated muon outside of the jet cone. Events with additional muons are excluded. At least one well-reconstructed PF jet is required in the event, and jets must further pass identification criteria that reject poorly reconstructed

jets or jets arising from detector noise. The AK4 PF jets are required to have $|\eta| < 2.5$ and $p_T > 30$ GeV, whereas the AK8 PF jets require $|\eta| < 2.5$ and $p_T > 170$ GeV. The efficiency is defined as the ratio of the number of events where an offline reconstructed PF jet is selected by the data scouting or standard triggers, relative to the total number of events with an offline reconstructed PF jet.

The low thresholds of the data scouting triggers are visible in the plot of each jet observable. The efficiency to select AK4 and AK8 scouting jets is about 100% for $p_T > 300$ GeV. In comparison, the standard trigger is fully efficient only from around 700–800 GeV. Similarly, data scouting is fully efficient for $H_T > 600$ GeV, compared to roughly 1300 GeV for the standard trigger. Therefore jet-based analyses relying on data scouting are able to probe regions of phase space inaccessible with the standard trigger strategy. By lowering the H_T threshold from 360 to 280 GeV in 2023, as discussed in Section 4.1, the scouting trigger improves even further the CMS acceptance to hadronic resonances.




Figure 26: Trigger efficiency as a function of AK4 jet p_T (left), AK8 jet p_T (center), and H_T (right). The efficiency is computed from collision data recorded in 2022 by the scouting (black points) and standard (red points) streams.

The quality of the Run 3 scouting jets is evaluated by computing the JES and JER with data recorded in 2022. The scouting jets in this study are reconstructed offline, using as input the scouting PF candidates that were reconstructed online during the data-taking. The offline reconstructed jets of the standard triggers are used as a reference. Before clustering, pileup is mitigated via the PUPPI technique for offline jets and via the CHS technique for scouting jets, as described in Section 2.3.5. Jets are corrected by applying detector response corrections computed from simulated samples. The corrections applied to the scouting jets are derived specifically for HLT jets, and differ from those applied to the offline jets. The same corrections are applied to simulated and observed jets. No in situ corrections are applied to observed jets either in the scouting or offline data sets.

The JES and JER derivations are performed with the same methods as described in Section 3.3.2. The measurement is performed in bins of jet η , where both leading jets are required to have $|\eta| < 1.3$ or $1.3 < |\eta| < 2.5$. Events where they are in different η regions are discarded. The final result of the JES measurement is presented in Fig. 27, where p_T^{scout} is the p_T of the scouting object, and p_T^{off} is the p_T of the offline object. The creation of mean scouting jet p_T ($\langle p_T^{\text{scout}} \rangle$) involves a mapping from p_T^{off} , resulting in varying bin widths across different η regions, as well as discrepancies between bin widths for simulation and recorded data. The uncertainties vary as a result of the trigger selection. The JES is similar between simulated and recorded events, at approximately 0.96–0.97 for $|\eta| < 1.3$ and 0.96–0.99 for $1.3 < |\eta| < 2.5$, when requiring jet $p_T > 200$ GeV.




Figure 27: The JES as a function of mean scouting jet p_T derived from simulated (left) and recorded (right) events. The red and black points correspond to events where the two leading jets have $|\eta| < 1.3$ and $1.3 < |\eta| < 2.5$, respectively.

The result of the JER measurement is presented in Figs. 28 and 29, as function of average p_T ($p_{T,\text{ave}} = (p_{T,1\text{st jet}} + p_{T,2\text{nd jet}})/2$) of the two highest- p_T jets in the event. The JER is stable from an average p_T value above 500 GeV, measuring approximately 5% (6%) in the barrel (endcap) region for both offline and scouting jets. For jet p_T below 500 GeV, scouting jets feature $\approx 10\%$ worse resolution compared to offline jets.

The jet performance discussed so far and the physics results presented earlier in Section 3.5 demonstrate that exploiting scouting jets is an effective strategy for resonance searches in hadronic final states. In addition, we now show that the scouting strategy provides a large data sample that is sensitive to a final-state topology featuring Higgs boson decays to bottom quark-antiquark pairs ($H \rightarrow b\bar{b}$). Recent searches by CMS for Higgs boson production in bottom [99, 100] or charm [101] quark decay channels improved signal sensitivity by targeting boosted final-state topologies that require reconstruction of the Higgs boson decay products within a single large-radius jet. The boosted object can then be identified with jet tagging techniques, e.g., by exploiting neural networks [57, 102]. As a preliminary study, we investigate the production of boosted Higgs bosons produced via gluon-gluon fusion (ggF) decaying to bottom quark-antiquark pairs.

The study is performed using a simulated sample of boosted Higgs boson events. Events are required to pass either a logical “OR” expression of jet-based scouting triggers or dedicated triggers targeting boosted Higgs boson topologies, deployed during Run 3 as part of the standard stream. Figures 30 and 31 show the performance of each trigger selection in terms of the trigger efficiency and the number of boosted Higgs boson events collected, respectively. The identification of such events is based on the requirement that the particle-level Higgs boson, together with its decay products, the bottom quark and antiquark, have a maximum angular distance $\Delta R < 0.8$ from the reconstructed AK8 jet with the highest p_T . The performance is shown as a function of both particle-level Higgs boson p_T and reconstructed AK8 jet p_T . The latter corresponds to the reconstructed jets originating from the boosted Higgs boson decays. The study indicates that data scouting increases the overall event selection efficiency for boosted $H \rightarrow b\bar{b}$ decays by approximately 20% compared to the standard triggers, particularly at low p_T . Since these results are preliminary and based only on simulated samples, this value is best




Figure 28: The JER as a function of average p_T . The JER is computed from simulated (upper) and recorded (lower) events, by requiring the two leading jets to have $|\eta| < 1.3$ (left) and $1.3 < |\eta| < 2.5$ (right). The red and black data points denote 2022 collision data of the scouting and offline reconstructions, respectively.




Figure 29: The ratio of the JER derived from scouting events to the JER derived from offline events as a function of average p_T . The ratio is computed from simulated (left) and recorded (right) events, by requiring the two leading jets to have $|\eta| < 1.3$ (red points) and $1.3 < |\eta| < 2.5$ (black points).

interpreted as an upper bound and can be affected by several factors, such as the background increase in scouting data or scouting jet-tagging efficiency, which require further investigation.

4.4 Muons

The dimuon scouting strategy, first developed in Run 2, offers numerous opportunities, as detailed in Section 3.5. In Run 3, the set of L1 algorithms used as input to the scouting stream in 2022 is the same as in Run 2, except for the temporary removal of the dimuon trigger with




Figure 30: Trigger efficiency for ggF boosted $H \rightarrow b\bar{b}$ events as a function of the highest particle-level Higgs boson p_T (left) and highest offline-reconstructed AK8 jet p_T (right), as determined from simulation. The black and red points correspond to the scouting and the standard trigger selection, respectively.




Figure 31: Number of ggF boosted $H \rightarrow b\bar{b}$ events as a function of the highest particle-level Higgs boson p_T (left) and highest offline-reconstructed AK8 jet p_T (right). The large-radius jet with highest p_T in each event is required to have a maximum angular distance $\Delta R < 0.8$ from the two final-state b quarks (blue). The events are then required to pass either the standard (red) or scouting (black) trigger selection. The number of events is computed from simulation with projected $\mathcal{L}_{\text{int}} = 100 \text{ fb}^{-1}$.

the loosest transverse momentum requirement, which is relevant for events in the very low dimuon mass window. This trigger was restored for the 2023 data-taking period. The dimuon mass spectrum obtained from opposite-sign muon pairs selected by requiring at least one of the 2022 L1 triggers to be satisfied is shown in Fig. 32. Considering data collected in 2022, corresponding to 17.6 fb^{-1} , all well-known dimuon resonances from meson or Z boson decays are visible. The breakdown of the individual L1 seed contributions is also shown.




Figure 32: Invariant mass distribution of opposite-sign muon pairs obtained with the scouting triggers, collected during 2022 with all Run 3 dimuon algorithms (blue curve), and with each individual algorithm (remaining colors).

As mentioned in Section 4.1, the removal of the requirement on the minimum number of hits in the pixel layers is a novelty with respect to Run 2. This update enables the collection of a larger number of events with high values of the dimuon transverse displacement, thus enhancing the ability to search for LLPs decaying to muons. The change occurred concurrently with the implementation of an updated version of the muon track finder algorithm in the barrel region at L1, which mainly improved the displaced-muon triggering performance [31]. Figure 33 shows the distribution of the dimuon vertex transverse displacement l_{xy} for events that contain at least one pair of OS muons associated with a selected secondary vertex. Events are collected with dimuon displacements up to ≈ 100 cm, corresponding to the end of the sensitive region of the tracker. At the positions of the pixel layers, with radii of 29, 68, 109, and 160 mm, photons undergoing conversion processes in the material lead to peaks in the l_{xy} distribution. These peaks are less pronounced in the Run 3 distribution because of the removal of the pixel-hit requirement, which leads to higher efficiency, but also lower purity – if no additional analysis-specific quality criteria are required, as is the case here.

The performance achieved by the online and offline reconstruction methods are compared using 2022 data collected with the scouting monitoring triggers. The resolution of the transverse momenta of muons reconstructed with the scouting algorithm is studied separately for muons reconstructed in the barrel and endcap regions. While the p_T range of interest for scouting muons is below 50 GeV, the study is performed for various p_T intervals between 3 and 100 GeV. Figure 34 shows the p_T resolution of scouting muons with respect to offline muons, computed as the standard deviation (σ) of the Gaussian fit to the following quantity:

$$\frac{p_T^{\text{scout}} - p_T^{\text{off}}}{p_T^{\text{scout}}}.$$

Leading and subleading p_T muons in events with exactly two muons are required to have $\Delta R > 0.2$ and to be geometrically matched to the corresponding offline muons. Differences in muon momentum resolution between the scouting and offline reconstruction algorithms are

found to be less than 1% for muons with $p_T < 60$ GeV and up to 1.5% for higher p_T values, in both barrel and endcap regions.

Figure 35 presents a comparison of the dimuon spectra obtained with the scouting and offline reconstruction, showing excellent agreement. The former is reconstructed with pairs of online muons associated with a common vertex and matched to the corresponding offline muons within $\Delta R < 0.1$, while the latter is composed of pairs of offline muons from selected events with exactly two muons. All dimuon resonances in the very low mass range below 11 GeV are reconstructed with excellent resolution compared to the offline algorithm. Differences between the mass resolution obtained with scouting and standard muons are observed to be less than 1.0–1.5%, both in the barrel and endcap regions.




Figure 33: Comparison between the l_{xy} distribution for Run 2 (orange) and Run 3 (blue) events in data that contain dimuon pairs with a common displaced vertex and a minimal selection on the vertex quality. The dashed vertical lines, placed at radii of 29, 68, 109, and 160 mm, correspond to the positions of the pixel layers where photons undergo conversion processes in the material, causing the observed peaks in the l_{xy} distribution.

4.5 Electrons and photons

The endeavor to reduce trigger thresholds for physics studies with electrons and photons represents a major and challenging novelty of the Run 3 scouting strategy. Background processes from soft hadronic interactions in pp collisions, electromagnetic activity within jets, and other low-energy deposits from pileup pose technical difficulties when lowering the trigger thresholds on the amount of energy in the ECAL. This complexity is further exacerbated by the time required to execute the GSF tracking algorithm for electrons. The inclusion of electrons and photons in the scouting stream was made possible by applying a background rejection strategy that focuses on the shower shape of the energy deposits in the ECAL, and by the reduced event content achieved.

The L1 and HLT requirements for the electron and photon scouting triggers are reported in Table 6. The efficiency of the scouting triggers to select events with a single electron or photon is calculated with an unbiased data set collected using a jet-based reference trigger. For the trigger efficiency measurements, both offline and scouting electrons and photons are used. An offline electron (photon) that passes a set of medium (loose) identification criteria is required to




Figure 34: Resolution on the transverse momentum of scouting muons compared to offline muons using data collected in 2022. Differences in muon momentum scale between the scouting and offline reconstruction algorithms are studied in bins of 1 GeV (10 GeV) for muon p_T smaller (larger) than 60 GeV. Values for the barrel (blue circles) and endcap (orange triangles) sections are shown separately.




Figure 35: Comparison of the dimuon spectra obtained with scouting (pink filled histogram) and offline (blue solid line) muons during the 2022 data-taking period. The ratio between the two distributions with a wider binning is also shown in the bottom panel as a gray band.

be within $\Delta R < 0.06$ of a scouting electron (photon). The medium and loose identification criteria are defined to achieve an isolated e/γ selection efficiency of 80% and 90%, respectively [39]. The trigger efficiency is computed as the ratio of the number of events that pass both the scouting single- e/γ trigger and the reference trigger, compared to the number of events that only pass the reference trigger.

Figure 36 shows the trigger efficiencies for electrons and photons in events triggered by the

single- e/γ trigger as a function of the offline object p_T . The trigger efficiency increases sharply for $p_T \approx 30$ GeV and reaches a plateau with $>90\%$ efficiency for $p_T > 45$ GeV. The low-energy reach for photon-triggered events in the scouting collection is therefore much improved compared to the 200 GeV (nonisolated) and 110 GeV (barrel only, isolated) thresholds in the standard trigger paths [103]. The trigger threshold for single-electron events in the scouting collection is at the minimum L1 threshold for triggering ECAL energy deposits. A minimal identification criteria is applied at the trigger level for both single-electron and single-photon paths. The typical offline criteria employed for physics analyses are expected to be much tighter than these selections. The scouting strategy thus maximizes the trigger efficiency for events with single electrons and photons in the CMS detector at the lowest energies.

The ability to perform physics studies with a combination of physics objects is new in Run 3 scouting. The reconstruction of scouting electrons and photons in paths seeded by L1 muons, jets and H_T becomes efficient at lower p_T thresholds compared to the thresholds required by trigger algorithms that exclusively target these objects, which must ensure that the trigger rates remain affordable. The scouting reconstruction efficiency of electrons and photons is identical to the one from the online HLT reconstruction.

The high quality of scouting electron objects is demonstrated by the ability to resolve decays of low mass mesons ($m_{ee} < 12$ GeV) to an electron-positron pair. Figure 37 shows the scouting dielectron mass spectrum where the J/ψ , $\psi(2S)$, and two of the resolved Υ meson peaks (1S and 2S) are visible. Events from the single- and double- e/γ scouting trigger paths collected during the 2023 data-taking period were combined for Fig. 37 including selections and corrections as follows. An electron-positron pair, each with $p_T > 12$ GeV, is required to pass an identification selection developed for mesons decaying to such pairs. To maximize the resonance signal over the background, both electrons are required to pass a tight identification selection. The reconstructed energy of the electron was corrected based on its position and shower shape in the ECAL using corrections derived from simulation. As a result, the peaks are shifted by $<2\%$ with respect to the actual mass of the corresponding SM mesons. The reconstructed dielectron mass resolution for the J/ψ peak is approximately 3%. Further calibration of the scouting electrons could make it possible to improve the resolution, and will depend on the physics analysis under consideration.




Figure 36: Trigger efficiencies for the scouting trigger paths seeded by L1 algorithms targeting either single-electron events (left) or single-photon events (right), as a function of the respective object p_T reconstructed offline. To be considered for scouting, the leading electron or photon must have $p_T > 30$ GeV. Results are only shown for electrons or photons detected in the barrel region.




Figure 37: Dielectron mass distribution observed with Run 3 scouting data collected during the 2023 data-taking period. The J/ψ and two of the Y meson peaks are visible.

The enhanced scouting program of CMS will play a pivotal role in carrying out low-mass searches and precision measurements during Run 3. The addition of pixel tracks will notably enhance the capability to conduct efficient searches and pioneer novel analysis strategies. This evolution will not only extend the results achieved during Run 2, but also broaden the data scouting physics program of CMS.

5 Data parking in Run 1 and Run 2

The data parking strategy addresses key computational challenges that impede prompt event reconstruction of all the data collected by the experiment. In Run 1, numerous physics processes were considered in the original parking approach, which is discussed in Section 5.1. In Run 2, the focus of the parking strategy shifted to the B physics program, apart from specialized backup parking data sets collected in 2016 and 2017 that remain unprocessed. Section 5.2 describes in detail the physics motivation and trigger strategy of the B parking campaign in Run 2. Finally, Section 5.3 presents an overview of the physics results obtained with this unique B parking data set.

5.1 Data parking in Run 1

An overview of the data parking triggers in 2012, which stored events at a total rate of about 300–350 Hz with $\mathcal{L}_{\text{inst}} \approx 4 \times 10^{33} \text{cm}^{-2}\text{s}^{-1}$, is shown in Table 10. The initial parking strategy focused on dijet and double- τ triggers for VBF topologies and Higgs boson measurements, multijet and single-photon plus $p_{\text{T}}^{\text{miss}}$ triggers to target compressed SUSY and DM models, and various dimuon triggers for B physics measurements.

A brief overview of the CMS results published with the Run 1 parking triggers is provided in the following paragraphs.

Table 10: Summary of the main physics targets of the Run 1 parking strategy and the corresponding parking triggers in 2012. The average HLT rates reserved for Higgs boson measurements, B physics measurements, and new-physics searches are also quoted for $\mathcal{L}_{\text{inst}} \approx 4 \times 10^{33} \text{cm}^{-2}\text{s}^{-1}$.

Physics motivation	Parking triggers	Average HLT rate [Hz]
VBF topologies	Dijet (large $\Delta\eta_{jj}$)	
Higgs boson measurements	Double τ_h	
B physics measurements	Double μ	95
Multijet searches	Four-jet	
DM and dark photons	$p_T^{\text{miss}} + \gamma$	
DM production	Razor variables	
SUSY hadronic searches	$H_T + \alpha_T$	
		$\left. \right\} 150$
		$\left. \right\} 105$

5.1.1 Higgs boson measurements

The distinctive pattern of VBF production, characterized by two jets with considerable angular separation in pseudorapidity and a high dijet invariant mass, was exploited for dedicated parking triggers that collected data in 2012, corresponding to \mathcal{L}_{int} of about 18.3 fb^{-1} . This inclusive trigger required two jets with p_T greater than 35 and 30 GeV, respectively, $|\Delta\eta_{jj}| > 3.5$, and $m_{jj} > 700 \text{ GeV}$. Two analyses were performed assuming a VBF production mode: a search for invisible Higgs boson decays [104], and a search for the SM Higgs boson decaying to a bottom quark-antiquark pair [105]. In both cases, the analysis sensitivity was significantly enhanced by the use of triggers recorded in the parked data stream. This trigger strategy has been resumed and refined in Run 3, as discussed in Section 6.2.1.

Similarly, new trigger paths with lower thresholds on the object transverse momentum were exploited to target both hadronic and leptonic decays of the tau lepton. In particular, the requirement of two isolated τ_h objects, each with $p_T > 35 \text{ GeV}$ and $|\eta| < 2.1$ enhanced the search for $H \rightarrow \tau\tau$ decays in both SM [106] and BSM [107] scenarios.

5.1.2 Dimuon final states for B physics

Final states with two muons provide a clean and distinctive experimental signature to identify interesting and rare processes involving for example b hadron, charmonium, or bottomonium decays. Inclusive low- p_T dimuon triggers without any mass constraint are ideal to cover a wide breadth of B physics analyses with maximum acceptance and simplicity, and were thus employed in CMS starting from 2011. In 2012, the instantaneous luminosity delivered by the LHC doubled relative to 2011, reaching $8 \times 10^{33} \text{cm}^{-2}\text{s}^{-1}$ and saturating the prompt event processing capacity of CMS. Dimuon triggers were nevertheless retained in 2012 but the reconstruction of collected events was postponed to 2013, during the LS1 of the LHC. This parked dimuon data set collected in 2012 allowed CMS to make important contributions to the accurate measurements of various b quark hadron lifetimes, including the rare and heavy B_c meson and the B_s meson in two final states corresponding to different mass eigenstate admixtures [108]. Moreover, CMS contributed to the measurement of angular parameters of the decays $B^0 \rightarrow K^*(892)^0 \mu^+ \mu^-$ [109, 110], $B^+ \rightarrow K \mu^+ \mu^-$ [111], and $B^+ \rightarrow K^*(892)^+ \mu^+ \mu^-$ [112], which are sensitive to new physics contributions in processes described by “penguin” diagrams and relevant to the ongoing puzzle surrounding flavor anomalies [24]. In Run 2, inclusive dimuon triggers could not be maintained anymore because both the instantaneous luminosity

and the center-of-mass energy nearly doubled relative to 2012. To keep the rate under control, these triggers were replaced by more complex and restrictive triggers, which were used to collect data in the prompt data sets. Additional requirements beyond the presence of two muons were applied, such as additional tracks to form three- or four-body vertices, dimuon mass requirements, higher dimuon p_T thresholds, and displacement conditions. Starting from 2022, thanks to a different resource allocation and to the enhanced capacity of the CMS DAQ and computing systems in Run 3, inclusive dimuon triggers were reinstated, as discussed in Section 6.1.3.

5.1.3 Searches for BSM physics

New trigger paths were introduced in the Run 1 data parking stream to recover sensitivity to new physics models in regions of phase space not covered by the standard trigger paths. The collection of parked data sets targeted three main physics cases: multijet searches, monophoton searches, and SUSY hadronic searches.

A four-jet trigger with loose thresholds on the transverse momentum of the jets, lowered to 45–50 GeV, was designed to target multijet searches and look for the top quark superpartner (top squark, or stop) predicted by natural SUSY models [113]. A subset of the 2012 data corresponding to 12.4 fb^{-1} was parked and used to extend the search in the mass region below 300 GeV, taking advantage of the lower jet p_T threshold.

A new parking trigger designed to extend the physics reach of monophoton searches in the low photon p_T and low- p_T^{miss} phase space was also introduced. It required a low- p_T photon with thresholds of 22 and 30 GeV at the L1 and at the HLT, respectively, and p_T^{miss} of at least 35 GeV, reducing the thresholds on these reconstructed objects by a factor of 2–3 compared to the standard triggers. Parking data collected at 8 TeV and corresponding to $\mathcal{L}_{\text{int}} = 7.3 \text{ fb}^{-1}$ were used to set limits on the exotic decays of the SM Higgs boson, and results were interpreted in the context of dark photon and dark matter pair production models [114].

For dark matter searches, new algorithms were designed to collect events with a large momentum imbalance, requiring at least two jets and no isolated leptons. Dedicated “razor” variables [115, 116] were computed from the momenta of the two leading jets and the p_T^{miss} in the event, in order to quantify the transverse momentum balance of the jet momenta. Trigger paths with loose requirements on these kinematic variables were introduced in the data parking stream. Parked data corresponding to $\mathcal{L}_{\text{int}} = 18.8 \text{ fb}^{-1}$ were collected at 8 TeV and enabled the exploration of events with moderate jet p_T , thereby improving the sensitivity to direct dark matter production [117].

Dedicated trigger algorithms that relied on the dimensionless variable α_T were used for SUSY searches in final states with jets and p_T^{miss} [118]. This kinematic variable is computed from the system of the two leading jets in the event. It is defined as the ratio between the transverse energy of the less energetic jet and the transverse mass of the dijet system, and it is used to discriminate between events with genuine p_T^{miss} associated with unobserved particles (e.g., neutralinos) and spurious values of p_T^{miss} arising from jet energy mismeasurements (e.g., QCD multijet background). The data sample, corresponding to $\mathcal{L}_{\text{int}} = 18.5 \text{ fb}^{-1}$, was used to search for evidence of SUSY models involving the pair production of top squarks. Parking data was recorded with a lower H_T threshold, extending the acceptance to a wide array of compressed-SUSY models, where the top squark and the lightest neutralino (a DM candidate) are nearly degenerate in mass.

5.2 Data parking for B physics in Run 2

This section details the main data parking strategy in Run 2, which focused on B physics. The physics motivation, experimental challenges, trigger strategy, and the performance of the parked triggers are all discussed in the following subsections. The physics results obtained with this approach are presented in Section 5.3.

5.2.1 Physics motivation

At the present time, several measurements of observables related to rare b hadron decays present some tension with respect to their predicted values from the SM. Collectively, these measurements are known as the “B flavor anomalies” and they are being interpreted by many in the physics community as potential evidence for BSM physics [24]. These anomalies have been observed throughout the last decade in both charged-current $b \rightarrow c\ell\nu$ and neutral-current $b \rightarrow s\ell\ell$ transitions by the BaBar [119], Belle [120], and LHCb [121] Collaborations.

The anomalous measurements can be divided into two categories of physics observables. First, there are those relating to (differential) branching fractions and the parametrization of four-body angular distributions for decays via the $b \rightarrow s\mu\mu$ transition. Second, there are observables constructed from ratios of branching fractions for semileptonic decays with final states that differ only by the lepton flavor. Several ratios R_X can be measured, where X represents the final-state hadron produced in the semileptonic decay.

The B flavor observables are particularly powerful probes of BSM physics because of the availability of both precise theoretical predictions and clean experimental signatures for processes involving semileptonic (and fully leptonic) decays of b hadrons. For instance, the R_X observables are sensitive to the violation of an accidental symmetry within the SM, known as lepton flavor universality (LFU), whereby the interactions between the gauge bosons and charged leptons are identical for all three lepton generations (beyond kinematical effects due to their differing masses). Confirmation of LFU violation would be a striking proof of the existence of BSM physics. In recent years, several key observables have received significant attention, some examples of which are given below.

The branching fraction $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ is an excellent probe to test the flavor sector of the SM, given its precise theoretical prediction and clean experimental signature. Furthermore, possible modifications of $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ relative to the SM expectations can be related to the same new physics operators responsible for LFU violation. The ATLAS [122], CMS [30], and LHCb Collaborations have reported several independent measurements [123–126], as well as combined measurements [127], in recent years. These measurements constitute one of the cleanest inputs to global fits aimed at providing a coherent global interpretation of the flavor anomalies [24].

In an effective field theory framework, the angular distributions of three- and four-body decays arising from $b \rightarrow s\ell\ell$ transitions offer sensitivity to new-physics operators. Multiple measurements of the P'_5 observable [128], constructed from the Wilson coefficients associated with these operators and optimized to mitigate the impact of QCD uncertainties, have been conducted using the $B^0 \rightarrow K^*(892)^0 \mu^+ \mu^-$ process. Most of these measurements have indicated tensions with the SM since 2013 [110, 129–131]. The theoretical predictions are affected by the limited knowledge of long-distance charm-loop contributions, which might enhance the apparent discrepancy with the SM.

The ratios of branching fractions $R_D = \mathcal{B}(B^- \rightarrow D^0 \tau^- \bar{\nu}_\tau) / \mathcal{B}(B^- \rightarrow D^0 \ell^- \bar{\nu}_\ell)$ and $R_{D^*} = \mathcal{B}(B^0 \rightarrow D^{*-} \tau^+ \nu_\tau) / \mathcal{B}(B^0 \rightarrow D^{*-} \ell^+ \nu_\ell)$ ($\ell = e, \mu$) involve tree-level $b \rightarrow c\ell\nu$ transitions. Two further observables are $R_K = \mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-) / \mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)$ and $R_{K^*} = \mathcal{B}(B^0 \rightarrow$

$K^*(892)^0 \mu^+ \mu^-) / \mathcal{B}(B^0 \rightarrow K^*(892)^0 e^+ e^-)$, which involve loop-level $b \rightarrow s \ell \ell$ transitions. Numerous measurements for $R_{D^{(*)}}$ [132–135] and $R_{K^{(*)}}$ [136–140] have been reported since 2012, culminating in a reported evidence of LFU violation for R_K in 2022 [141]. Most recently, the LHCb Collaboration has provided combined measurements for R_D and R_{D^*} [142], and for R_K and R_{K^*} [143, 144], which are now consistent with the SM at the level of 1.9 and 0.2σ , respectively, and the latter result supersedes the one reported in Ref. [141]. Regardless of this recent dilution of a pattern of anomalous behavior, there remains the potential for LFU violating processes and there is still substantial interest from the physics community for new results pertaining to the B flavor anomalies.

The CMS Collaboration recorded a unique data set of pp collisions at $\sqrt{s} = 13$ TeV in 2018 with the primary aim of extending its program of LFU tests. Dedicated trigger and data storage strategies were developed to record a large sample of events containing 10 billion unbiased b hadron decays. The reconstruction of physics events from the raw data sample was delayed until computing resources were available in 2019. The trigger and data storage strategies (known henceforth as “B parking”), the defining characteristics of the resulting data set, and some of the key physics results are described in the following sections.

5.2.2 Experimental challenges

Prior to 2018, measurements targeting the B flavor anomalies within the B physics program of CMS were restricted to observables involving dimuon final states, a consequence of the available dimuon trigger algorithms. Examples of measurements from CMS include $\mathcal{B}(B_s^0 \rightarrow \mu^+ \mu^-)$ [123] and several angular observables [110–112]. These successes can be contrasted with the absence of measurements of observables such as R_D and R_{D^*} , which rely on the reconstruction of single-muon final states resulting from the $b \rightarrow c \ell \nu$ transitions. Prior to 2018, the single-muon trigger algorithms were typically geared towards high- p_T physics processes, such as W boson production. The typical kinematical ($p_T > 20$ GeV) and topological (isolation from neighboring particles) requirements on the muon suppressed the acceptance to $b \rightarrow \mu X$ decays. Furthermore, prior to 2018, no trigger algorithms provided adequate fiducial acceptance to final states containing electron pairs that arise from rare $b \rightarrow se^+ e^-$ transitions; typically, the single- and dielectron trigger algorithms imposed p_T thresholds of 30 and 20 GeV, respectively, as well as isolation requirements. These constraints have thus far prohibited the measurement of observables such as R_K and R_{K^*} .

A new trigger strategy was devised and implemented in time for the LHC pp collision run of 2018. The trigger strategy relies on the accumulation of a very large sample of b quark-antiquark ($b\bar{b}$) pairs using a “tag-side” trigger logic that identifies the semileptonic decay of one of the b hadrons to a final state containing at least one displaced muon above an evolving p_T threshold in the range 7–12 GeV. Decays via the transition $b(\rightarrow c) \rightarrow \mu X$ have a total branching fraction of 18% [97] and thus approximately one in three $b\bar{b}$ events results in a final state containing at least one muon. The other “probe-side” b hadron is able to decay to all possible final states (including any flavor of lepton) with minimal kinematic bias from the tag-side trigger requirements. Thus, the study of b hadron decays that lead to final-state muons can focus on the tag-side muon candidate identified by the trigger system, while the study of processes involving other lepton flavors or nonleptonic final states can rely on the probe-side decays. Rare processes with branching fractions as small as $\mathcal{O}(10^{-7})$ are accessible if the sample of $b\bar{b}$ pairs is sufficiently large, e.g., $\mathcal{O}(10^{10})$.

Early feasibility studies, based on simulated data, demonstrated that a suitable data sample could be accumulated if high trigger rates and purities could be sustained throughout the

2018 data-taking period. The number of $b\bar{b}$ pairs $N_{b\bar{b}}$ produced at the LHC and subsequently recorded by the trigger logic was estimated with the following expression: $N_{b\bar{b}} = t_{\text{LHC}} \mathcal{R} P_{b\bar{b}}$, where t_{LHC} is the LHC operational running time (seconds), \mathcal{R} is the rate in Hz at which the trigger logic returns a positive decision, and $P_{b\bar{b}}$ is the purity of the resulting data stream, defined as the ratio of the numbers of genuine $b\bar{b}$ pairs and pp collision events recorded by the trigger system. The value of t_{LHC} is 6×10^6 s when assuming 140 days of LHC operations over a 6-month period and a beam duty cycle of 0.5. Hence, by maintaining an average trigger rate $\mathcal{R} \approx 2$ kHz throughout each LHC fill, and assuming a purity $P \approx 80\%$, it is possible to accumulate a data sample comprising 1.2×10^{10} events that contain 10^{10} $b\bar{b}$ pairs. Data parking is thus required to handle the high trigger rates, and the strategy was based on experience accumulated during Run 1, as detailed in Section 5.1.

Equivalently, $N_{b\bar{b}}$ can be estimated from the expression $N_{b\bar{b}} = \mathcal{L}_{\text{int}} \sigma_{b\bar{b}} \mathcal{B}' \varepsilon$, where the data recorded by the muon-based triggers correspond to $\mathcal{L}_{\text{int}} = 41.6 \text{ fb}^{-1}$ (described later in Section 5.2.9), the inclusive $b\bar{b}$ cross section is $\sigma_{b\bar{b}} = 4.7 \times 10^{11} \text{ fb}$ at $\sqrt{s} = 13 \text{ TeV}$ [145–148], the term $\mathcal{B}' = 1 - (1 - \mathcal{B})^2$ accounts for the fact that either b hadron can decay into a muon and the branching fraction for the $b(\rightarrow c) \rightarrow \mu X$ decay is $\mathcal{B} = 18\%$ [97], and the efficiency of the muon-based triggers to record an event containing at least one $b \rightarrow \mu X$ decay is $\varepsilon \approx 2 \times 10^{-3}$. These values yield $N_{b\bar{b}} \approx 10^{10}$.

A precise measurement of the R_{D^*} observable is feasible by reconstructing the $B^0 \rightarrow D^{*-} \mu^+ \nu_\mu$ and $B^0 \rightarrow D^{*-} \tau^+ (\rightarrow \mu^+ \nu_\mu \bar{\nu}_\mu) \nu_\tau$ decays using the large sample of 10^{10} tag-side muons: the sample is expected to contain $\mathcal{O}(10^6)$ and $\mathcal{O}(10^4)$ candidates for these decays, respectively. Similar arguments for precision can be made for other observables involving muon-based signatures, such as R_D and the numerator of the R_K observable, $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)$.

The limiting factor in the precision of a measurement of R_K is related to the number of reconstructed $B^+ \rightarrow K^+ e^+ e^-$ decays that can be identified unambiguously above background contributions. The number of $B^+ \rightarrow K^+ e^+ e^-$ decays found within the fiducial acceptance can be estimated from the expression $N_{B^+ \rightarrow K^+ e^+ e^-} \approx N_{b\bar{b}} f_{B^+} \mathcal{B} \mathcal{A}$, where: $N_{b\bar{b}} = 10^{10}$, the fragmentation fraction is $f_{B^+} = 0.4$ [97]; the branching fraction is $\mathcal{B}(B^+ \rightarrow K^+ e^+ e^-) = 4.5 \times 10^{-7}$ [97]; and the fiducial acceptance, defined by the fraction of probe-side B^+ decays with all daughter particles satisfying the requirements $p_T > 0.5 \text{ GeV}$ and $|\eta| < 2.5$, is $\mathcal{A} \approx 55\%$. Thus, approximately 1000 probe-side $B^+ \rightarrow K^+ e^+ e^-$ decays are expected within the fiducial volume. Subsequent reconstruction and selection requirements will further reduce the sample of identifiable $B^+ \rightarrow K^+ e^+ e^-$ decays.

5.2.3 Trigger strategy

The primary aim of the trigger strategy described here is to maximize the number of recorded $b\bar{b}$ events by maximizing the data stream purity and operating the trigger system close to its design limits. Crucially, this mode of operation must not compromise the availability of online resources for the core CMS physics program.

The new trigger strategy adopted a two-step optimization of the L1 and HLT trigger algorithms given the following key design constraints for 2018: the total L1 trigger rate was restricted to 90 kHz to maintain acceptable dead time from the subdetector readout systems; and the bandwidth of the B parking data stream could not substantially exceed an average of 2 GB/s because of the finite capacity of buffers at the CMS experimental site.

5.2.4 The L1 optimization

During an LHC fill, the $\mathcal{L}_{\text{inst}}$ slowly decreases with time. As a consequence, the number of pp interactions that occur within the same LHC bunch crossing, the L1 and HLT trigger rates, and the per-event HLT computational load are all observed to decrease with time. Hence, the availability of idle resources increases during the ongoing LHC fills, which can be leveraged by the trigger strategy described here.

The left panel of Fig. 38 shows the evolution of the total L1 trigger rate and pileup as a function of time during a typical LHC fill in the 2017 pp collisions run, prior to the implementation of the trigger strategy discussed here. Over a period of 14.5 hours, the pileup value decreases from 48 to 18 and, correspondingly, the total L1 trigger rate also decreases. At the beginning of an LHC fill, the L1 system typically operates at a total trigger rate of 90 kHz; towards the end of an LHC fill, there are up to several tens of kHz of spare-rate capacity available.




Figure 38: The L1 trigger rate and the amount of pileup as a function of time, shown for representative LHC fills during 2017 (left) and 2018 (right). Occasional lower rates are observed due to transient effects, such as the throttling of the trigger system in response to subdetector dead time [9]. Changes in the trigger configuration are indicated by vertical green dashed lines.

The new strategy for the L1 system repurposes existing algorithms to identify low- p_T muon candidates with high efficiency and purity. Simple kinematical requirements are used to identify interesting muon candidates. First, each muon candidate is required to be found centrally in the detector by satisfying $|\eta| < 1.5$, where the L1 muon identification and momentum-resolution performance is generally optimal; this requirement simultaneously improves the trigger purity and enhances the fiducial acceptance for the probe-side b hadron decays. Second, a variable p_T threshold is applied within the range of 12 GeV at $\mathcal{L}_{\text{inst}} = 1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ down to 7 GeV at $\mathcal{L}_{\text{inst}} = 0.9 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The threshold is progressively loosened within this range (via changes in the trigger configuration) as $\mathcal{L}_{\text{inst}}$ and the pileup decreases. Importantly, the threshold is tuned to ensure that the L1 system operates close to its design limit, i.e., at ≈ 90 kHz throughout the LHC fill, but not beyond so as to keep dead time from the subdetector readout systems to below $\approx 1\%$. Thus, there is negligible impact on the accumulated \mathcal{L}_{int} and the wider CMS physics program. While the evolution of the threshold improves the acceptance to $b\bar{b}$ candidates, it also degrades the purity of the data stream from the L1 system. The average purity (in terms of the fraction of selected events containing a $b\bar{b}$ candidate, based on studies of simulated data) is ≈ 0.3 . The remainder of the events contain muons from the direct production of charm mesons and their semileptonic decays, muons from kaon or pion decays, and misidentified muons.

The right panel of Fig. 38 shows the L1 trigger rate as a function of time during a typical LHC fill in the 2018 pp collisions run. While the average pileup during an LHC fill decreases similarly in

2017 and 2018, the trigger rates do not. The large instantaneous increases in the L1 trigger rate are coincident with changes in the trigger configuration. The total L1 trigger rate peaks at values close to 90 kHz throughout the LHC fill; this is because of an increasing rate contribution, as high as 40 kHz, from the L1 single-muon algorithm as the p_T threshold is reduced. At the start of each LHC fill, the trigger algorithms that serve the core physics program operate with a total L1 rate close to the 90 kHz ceiling; only once $\mathcal{L}_{\text{inst}}$ has dropped below $1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ are the dedicated triggers enabled, when sufficient online resources are available.

5.2.5 The HLT optimization

The purity of the data stream from the L1 system is substantially improved by the use of tailored muon algorithms at the HLT. The algorithms provide superior performance relative to the L1 logic, in terms of muon identification and momentum scale and resolution, because of the ability of the HLT software to reconstruct muons using tracking information from the silicon pixel and strip trackers.

The B parking data throughput, given by the product of the HLT trigger rate and the triggered event size, was limited in 2018 to an average of 2 GB/s for timescales longer than 24 hours because of the limited buffer capacity, as described in Section 5.2.7. The triggered event size has a linear dependence on $\mathcal{L}_{\text{inst}}$ and thus higher HLT trigger rates are accessible later during an LHC fill as both $\mathcal{L}_{\text{inst}}$ and the event size, with reduced pileup, decrease.

Various scenarios involving different assumptions on the LHC performance and load-balancing of the DAQ system, in terms of varying data throughput during an LHC fill, were investigated. The left panel of Fig. 39 shows an example scenario in which the data throughput is allowed to evolve during an LHC fill while ensuring that the average does not exceed 2 GB/s. The left panel also indicates the maximum HLT trigger rate permitted for each L1 trigger configuration, which changes as a function of $\mathcal{L}_{\text{inst}}$. The thresholds of the HLT trigger algorithms are therefore adjusted to operate close to these maximum values. This scheme, which allows for HLT trigger rates in the multi- kHz range, is optimized for long LHC fills, of duration 15 hours or more.

The HLT algorithms are tuned to maximize the number of recorded $b\bar{b}$ events and satisfy the trigger rate constraints. Two of the most discriminating variables are the muon p_T and the transverse impact parameter of the muon expressed in terms of its measurement significance, IP_{sig} , relative to the pp luminous region. By requiring a nonzero IP_{sig} value, the relatively long lifetime of the b hadron and the characteristic displacement (relative to the primary interaction point) of the muon from $b \rightarrow \mu X$ decays are leveraged to reduce prompt muon production from background processes such as D mesons and charmonium decay.

For each L1 trigger configuration, which changes as a function of $\mathcal{L}_{\text{inst}}$, a parameter scan is performed across all feasible combinations of p_T and IP_{sig} thresholds imposed by the HLT algorithm. The L1 configuration is determined by the procedure described above. The right panel of Fig. 39 shows the rate of $b\bar{b}$ candidates found within the experimental acceptance and the corresponding HLT trigger rate obtained for each unique $(p_T, \text{IP}_{\text{sig}})$ combination at the peak $\mathcal{L}_{\text{inst}} = 1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The optimal combination is one that lies along the upper boundary of the point set shown in the right panel of Fig. 39 and is as close as possible but does not exceed the maximum allowed HLT trigger rate for the given $\mathcal{L}_{\text{inst}}$ value. For the example shown, the procedure indicates that the HLT thresholds $p_T > 12 \text{ GeV}$ and $\text{IP}_{\text{sig}} > 6$ accumulate $b\bar{b}$ candidates at a rate of 1.3 kHz at a trigger rate of 1.5 kHz, which corresponds to a trigger purity of 85%. The procedure is repeated for each L1 trigger configuration, which varies with $\mathcal{L}_{\text{inst}}$ during an LHC fill.




Figure 39: Left: an example scenario in which the B parking data throughput per L1 trigger setting is adjusted to maintain an average of approximately 2GB/s throughout an LHC fill. The dotted red and dashed blue lines trace the B parking data throughput and maximum allowed HLT rate, respectively, determined for each trigger configuration. Changes in the trigger configuration are indicated by vertical green dashed lines. The trigger logic is adjusted to operate close to the permitted HLT rate. Right: rate of $b\bar{b}$ events in acceptance versus HLT rate for a parameter scan over p_T and IP_{sig} thresholds imposed in the HLT logic for the \mathcal{L}_{inst} and L1 requirements indicated in the legend; each point (blue circle) represents a unique pair of thresholds and the red star indicates the optimal pairing of $p_T > 12$ GeV and $IP_{sig} > 6$ at a peak $\mathcal{L}_{inst} = 1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and an HLT rate close to the maximum allowed value of ≈ 1.5 kHz.

As a result, the lower-bound thresholds on both p_T and IP_{sig} are relaxed, within the ranges 7–12 GeV and 3–6, respectively, as \mathcal{L}_{inst} decreases during an LHC fill. Table 11 summarizes the evolution of the trigger thresholds used to record pp collision data during the LHC fill shown in Figs. 38 and 40, as well as the resulting peak trigger rates and data stream purities. The settings are deployed sequentially via changes in the trigger configuration at different values of peak \mathcal{L}_{inst} ; each new setting corresponds to the enabling of new trigger logic with lower thresholds until the end of the LHC fill. No dedicated trigger logic is enabled for values of \mathcal{L}_{inst} above $1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. Minor adjustments were made to these settings during 2018 in response to the evolving LHC performance. This intra-fill evolution of trigger settings maximizes the estimated number of $b\bar{b}$ events recorded in the B parking data stream. Furthermore, the recording of $b\bar{b}$ events at higher rates towards the end of an LHC fill ensures that the pileup, and thus the additional activity, in these events is low.

Figure 40 shows the HLT trigger rates for both the promptly reconstructed data stream that serves the core CMS physics program, with a monotonically decreasing behavior, and the B parking data stream provided by the single-muon trigger algorithms. In the latter case, rates as high as 5.5 kHz are obtained late in the LHC fill because of the relaxed kinematical and topological thresholds.

5.2.6 Trigger purity

The purity of the data stream is estimated from simulated events to be in the range 60–90%, depending on the trigger thresholds, as indicated in Table 11, with an average of $\approx 80\%$. The estimates each have an associated uncertainty of 5%, arising from sources such as the kinematical modeling of the B meson decays.

The average purity is also determined from an analysis of the data sample itself, by estimating the total number of $b\bar{b}$ events contained in the sample by reconstructing D^{*+} candidates

Table 11: Single-muon trigger settings used during a typical LHC fill. The kinematical thresholds are changed when $\mathcal{L}_{\text{inst}}$ (and consequently pileup) fall below the listed values. Also listed are the lower-bound thresholds on the tag-side muon p_T (L1 and HLT) and IP_{sig} (HLT only), the corresponding L1 and HLT peak trigger rates, and the HLT data stream purity estimated from simulated events. No dedicated trigger is enabled at the start of each LHC fill when $\mathcal{L}_{\text{inst}}$ is typically above $1.7 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$.

$\mathcal{L}_{\text{inst}}$ [$10^{34} \text{cm}^{-2} \text{s}^{-1}$]	Pileup	L1 μ p_T [GeV]	HLT μ p_T [GeV]	HLT μ IP_{sig}	Peak L1 rate [kHz]	Peak HLT rate [kHz]	Purity [%]
2.0	54.0	—	—	—	—	—	—
1.7	45.9	12	12	6	20	1.5	92 ± 5
1.5	42.8	10	9	6	30	2.8	87 ± 4
1.3	35.1	9	9	5	32	3.0	86 ± 4
1.1	29.7	8	8	5	43	3.7	83 ± 4
0.9	24.3	7	7	4	53	5.4	59 ± 3




Figure 40: HLT trigger rates and the number of pileup events shown as a function of time during a representative LHC fill in 2018. The rates for the promptly reconstructed core physics (black solid markers) and B parking (blue open markers) data streams are shown separately. Occasional lower rates are observed due to transient effects, such as the throttling of the trigger system in response to subdetector dead time [9]. Changes in the trigger configuration are indicated by vertical green dashed lines.

from the production mode $B^0 \rightarrow D^{*+} \mu^- \bar{\nu}$ and the subsequent decay chain $D^{*+} \rightarrow D^0 \pi_{\text{soft}}^+ \rightarrow (K^- \pi^+) \pi_{\text{soft}}^+$, where π_{soft} indicates a low-momentum pion. The decay mode via the D^{*+} state is chosen for the purity measurement because of its large branching fraction and the fully reconstructable D^{*+} decay chain.

The method relies on extracting the number of D^{*+} candidates by exploiting the mass difference between the reconstructed D^{*+} and D^0 candidates, as shown in Fig. 41. The value of $\Delta m = m(K^- \pi^+ \pi_{\text{soft}}^+) - m(K^- \pi^+)$ is expected to peak at the mass difference between the D^{*+} and D^0 mesons if the kaon and muon candidates have same-sign charges, whereas a smooth combinatorial background shape is expected for opposite-sign charges.

In order to identify the $B^0 \rightarrow D^{*+} \mu^- \bar{\nu}$ decay, the muon responsible for the positive trigger de-




Figure 41: Mass difference between reconstructed D^{*+} and D^0 candidates from the production mode $B^0 \rightarrow D^{*+} \mu^- \bar{\nu}$ and the subsequent decay chain $D^{*+} \rightarrow D^0 \pi^+_{\text{soft}} \rightarrow (K^- \pi^+) \pi^+_{\text{soft}}$. Events containing kaon and muon candidates with same-sign (opposite-sign) charges are indicated by solid (open) markers.

selection and neighboring charged-particle tracks are considered. A candidate $D^0 \rightarrow K^- \pi^+$ decay is identified by considering pairs of oppositely charged tracks that form a vertex. The leading (subleading) track is required to satisfy $p_T > 5(3)$ GeV and quality criteria, and the vertex is subject to both displacement and quality criteria. The particle track with the same (opposite) electrical charge as the muon is assigned the mass of the kaon (pion). Only D^0 candidates with a reconstructed mass within a ± 30 MeV (corresponding to a $\pm 3\sigma$) window of 1.86 GeV [97] are selected for further consideration. Finally, a candidate $D^{*+} \rightarrow D^0 \pi^+_{\text{soft}}$ decay is identified by combining the selected D^0 candidate with a “soft” pion candidate ($p_T > 300$ MeV).

The number of $B^0 \rightarrow D^{*+} \mu^- \bar{\nu}$ decays is obtained by performing a binned maximum likelihood fit to the mass difference distribution obtained from data using Gaussian and second-order polynomial functions for the same-sign and opposite-sign charge hypotheses, respectively. The yield obtained from the fit is then corrected to determine the number of $b\bar{b}$ candidates $N_{b\bar{b}}$ contained in the data sample by accounting for differences in the fiducial acceptances and reconstruction efficiencies determined from simulation, and branching fractions for the $b \rightarrow \mu X$ decay and the $B^0 \rightarrow D^{*+} \mu^- \bar{\nu}$ decay chain. The procedure yields $N_{b\bar{b}} \approx 9 \times 10^9$ with an associated uncertainty of 5%. Given that the number of recorded LHC events is 11.9×10^9 , the purity obtained from data is $P_{b\bar{b}} = 0.75 \pm 0.04$, which agrees with the estimate obtained from simulation.

5.2.7 Data parking and processing

The B parking data are not transferred immediately to the permanent data processing center, but, instead, are temporarily stored in local buffers at the CMS experimental site and later transferred unprocessed to permanent tape storage. The buffers are capable of handling a total data throughput of 2 GB/s, limited by the maximum data transfer rate achievable from the buffer to the tape resources. The effective limit on data throughput is higher because of the LHC inter-fill downtime; for instance, a total of 3 GB/s when averaged over the timescale of a week. Hence, a throughput of 2 GB/s can be sustained for the B parking data stream, in addition to an allocation of 1 GB/s for the core CMS physics streams. Tape storage resources,

originally allocated to a data parking stream of 500 Hz throughout Run 2 and sufficient to allow for primary and backup copies of the data, were reallocated in full to the B parking campaign in 2018. Only a single copy of the B parking data is kept on tape.

The data parking strategy exploits the opportunistic use of computational resources for the delayed processing (i.e., global event reconstruction) of very large event samples that would otherwise not be possible during the LHC running periods. This includes the short end-of-year and long end-of-run shutdowns of the LHC complex, when the resource load from the core CMS physics program is reduced. Three processing campaigns of the B parking data set have been undertaken so far. An early “pilot” reconstruction campaign was performed in 2018 on a small fraction ($\approx 5\%$) of the data set to check the performance of the trigger strategy, via purity measurements (described in Section 5.2.3), and to validate new reconstruction algorithms, such as the one described in this section. The first full processing of the parked data was performed between May and December 2019, during the LHC LS2. Finally, the full data set was reprocessed in 2022 with the ultimate “legacy” reconstruction software and calibrations that provide the optimal physics performance for future data analysis.

5.2.8 Low- p_T electrons

The B parking data set provides access to a large sample of unbiased b hadron decays. The particles produced in these decays typically have very low p_T values. The left panel of Fig. 42 shows the p_T^{gen} distributions for the leading- and subleading- p_T^{gen} electrons from the $B^+ \rightarrow K^+ e^+ e^-$ decay, where p_T^{gen} is the generator-level p_T quantity obtained from simulation. The kinematical requirements of $p_T > 7 \text{ GeV}$ and $|\eta| < 1.5$ are imposed on the tag-side muon. The most probable p_T^{gen} values are below 2 GeV. The standard PF-based electron reconstruction algorithm [21, 39] relies primarily on information from both the tracking and ECAL subdetector systems and it is optimized to identify electrons with high efficiency for p_T^{gen} values above 10 GeV. The performance of the PF electron algorithm worsens significantly below 10 GeV, with efficiencies falling to zero at 2 GeV and below, as shown in the left panel of Fig. 42.

Given the limited low- p_T performance of the existing electron algorithm, a custom multistage electron reconstruction algorithm has been developed to improve the reconstruction and identification of genuine electrons with p_T values below 10 GeV. Given the superior physics performance of the tracker system relative to the calorimeter systems at very low p_T , in terms of momentum or energy scale and resolution, the new algorithm uses a tracker-seeded approach that first considers charged-particle tracks and attempts to match each one with a compatible pattern of calorimetric energy deposits. Primarily, the low- p_T electron algorithm targets electrons from b hadron decays, but it exhibits comparable performance for electrons, originating promptly or otherwise, from a broad range of physics processes.

Electrons undergo photon bremsstrahlung when traversing detector material. Thus, to accommodate these energy losses, the accurate determination of the track parameters for electron candidates relies on a technique that employs a GSF [60], as described in Section 2.3.7. The GSF approach delivers the optimal electron momentum scale and resolution, albeit at a high computational cost. Consequently, it is preceded by a more computationally efficient seeding algorithm.

The first stage of the reconstruction chain comprises seeding logic that exploits two BDT algorithms. The BDTs consider a range of kinematical and topological variables constructed from tracker- and calorimetry-based measurements that include: the shape of electromagnetic showers, the spatial compatibility between the track trajectory and energy deposits within the calorimetry systems, the compatibility of momentum and energy measurements, and the dif-




Figure 42: The left panel shows the p_T^{gen} spectra of the leading and subleading electrons (dashed and solid green histograms) from $B^+ \rightarrow K^+ e^+ e^-$ decays, and the efficiency to identify genuine electrons as a function of p_T^{gen} for PF (solid red squares) and low- p_T electron candidates (solid blue circles). Efficiencies for the low- p_T electron candidates that satisfy an ID score threshold, tuned to give the same misidentification probability as for PF electron candidates, are also shown (open markers). The right panel shows the performance of the PF (solid red square) and low- p_T (solid blue circle) reconstruction algorithms and their corresponding ID algorithms (curves). The efficiencies and misidentification probabilities are determined relative to charged-particle tracks obtained from simulation, for both $B^+ \rightarrow K^+ e^+ e^-$ decays and background processes, and satisfying $p_T > 2 \text{ GeV}$ and $|\eta| < 2.5$.

ferences in the momentum measurements determined at the innermost and outermost layers of the tracker systems. The set of variables considered is analogous to that used by the ID algorithms reported in Ref. [39]. A simplified version of the GSF track fit, with a reduced set of parameters, is used, which is then compared with the nominal tracking fit algorithm based on the Kalman filter [34]. The seeding logic is carefully tuned to balance signal-to-background discrimination performance against computational load. One BDT uses a kinematically agnostic approach that exploits only the aforementioned variables. The other BDT provides a kinematically aware model that also uses the p_T , η , and the transverse impact parameter significance $d_{xy}/\sigma_{d_{xy}}$ of the electron candidate to discriminate between genuine signal electrons from $B^+ \rightarrow K^+ e^+ e^-$ decays and misidentified electrons from background physics processes.

In order for the reconstruction to proceed to the next stage, the seeding logic requires the score produced by each BDT to satisfy an independent threshold value. The full GSF-based track fit is subsequently performed on each electron seed to determine the optimal track parameters. The resulting trajectory is used to identify a spatially compatible energy cluster in the ECAL that is assumed to be the electromagnetic shower from the incident electron. Additional clusters of energy that are spatially compatible with the expected position of bremsstrahlung photons within the ECAL are associated with the original cluster to form a super cluster. The logic is tuned to ensure that the electron seeds are promoted to electron candidates with high efficiency ($>95\%$). The associated cost of high-rate particle misidentification is mitigated in the next stage.

The final stage of the low- p_T algorithm aims to differentiate between genuine electrons and misidentified particles with the highest possible performance. A further BDT model takes information from all the preceding stages, namely the seeding, GSF-based tracking, and superclustering algorithms. An expanded set of input variables is used relative to the seeding BDTs, such as improved track parameter estimates from the full GSF-based tracking stage, and variables that test for consistency between the supercluster substructure and a bremsstrahlung energy-loss pattern. The BDTs are trained with a simulated sample of $B^+ \rightarrow K^+ e^+ e^-$ events

using the XGBOOST package [149].

Electron candidates are considered to be interesting for further analysis if they satisfy a threshold applied to the BDT discrimination score, known as an ID working point. Low- p_T electron candidates are reconstructed and identified for the kinematical regime $p_T > 0.5 \text{ GeV}$ and $|\eta| < 2.5$, whereas the PF electron algorithm is restricted to $p_T > 2 \text{ GeV}$. The right panel of Fig. 42 shows the performance of the two algorithms for electron candidates originating from both $B^+ \rightarrow K^+ e^+ e^-$ decays and randomly selected charged-particle tracks that satisfy $p_T > 2 \text{ GeV}$. The PF electron candidates are reconstructed with an efficiency and misidentification probability of 69% and 2%, respectively. The low- p_T reconstruction algorithm provides electron candidates with a higher efficiency of 92%, but also a substantially higher misidentification probability of 21%.

The ID performance for both the low- p_T and PF electron candidates is quantified in terms of the efficiency and misidentification probability as a function of the ID score threshold, as shown in the right panel of Fig. 42. The PF-based BDT model is based on the standard approach defined in Ref. [39], but it is retrained specifically for an extended kinematical regime down to 2 GeV. The performance of the low- p_T and PF IDs can be compared at working points that yield the same misidentification probability: for instance, efficiencies of 56% and 33% are obtained for low- p_T and PF electrons, respectively, for a misidentification probability of 0.1%. For the regime $0.5 < p_T < 2 \text{ GeV}$, the low- p_T ID yields an efficiency of 30% for a misidentification probability of 0.1%. Finally, a tag-and-probe technique [150] is utilized with a sample of $B^+ \rightarrow J/\psi (\rightarrow e^+ e^-) K^+$ decays in data to check the accuracy of the simulation modeling of the input variables to the BDTs and their output scores. The ID score distributions for both the low- p_T and PF electrons in data are consistent with those in simulation, within statistical uncertainties.

5.2.9 Characterization of the data set

The B parking data set comprises 1.2×10^{10} events and contains 1×10^{10} unbiased b hadron decays. The size of the single-copy unprocessed data sample is 7.6 PB. The reconstruction-level MINIAOD [22] data sample has a reduced footprint of 950 TB. Custom analysis-specific data samples, based on a common analysis-level data format (known as NANO AOD [23]), typically have a footprint of around 10 TB.

Table 12 summarizes all the unique combinations of thresholds used in the L1 and HLT algorithms. In total, the suite of algorithms recorded 41.9 fb^{-1} . The \mathcal{L}_{int} value, the mean number of pileup interactions, and the number of events recorded by each combination are provided. The values are determined from the periods for which each combination provided the lowest enabled L1 p_T threshold. The highest HLT rates are observed later in a fill and so a larger fraction of the data are recorded at lower pileup values than for the standard physics data streams. Figure 43 shows the pileup distribution for the B parking data set, along with the contributions from each of the individual trigger combinations.

Figure 44 shows the invariant mass distributions for pairs of oppositely charged muons originating from a common vertex, as obtained from the B parking data set. Both muons are required to satisfy a minimal set of kinematical and ID criteria, and one of the muons must be matched to the candidate responsible for the positive trigger decision. Peaks in the data resulting from the ρ/ω , ϕ , J/ψ , $\psi(2S)$, $\Upsilon(1S)$, and Z resonances are visible above the continuous background.

Table 12: Trigger configurations, defined by unique combinations of L1 p_T , HLT p_T , and IP_{sig} thresholds, used to record events containing $b \rightarrow \mu X$ decays. The \mathcal{L}_{int} value, the mean number of pileup interactions, and the number of events recorded by each combination are aggregated over the periods for which each combination provided the lowest enabled L1 p_T threshold.

L1 p_T [GeV]	HLT p_T [GeV]	HLT IP_{sig}	\mathcal{L}_{int} [fb $^{-1}$]	Mean PU	Events [$\times 10^9$]
12	12	6	8.1	37.7	0.72
10	9	6	8.4	32.9	1.67
10	9	5	1.6	33.9	0.37
9	9	6	1.6	28.2	0.34
9	9	5	5.2	28.3	1.30
9	8	5	1.6	29.2	0.52
8	9	6	1.8	24.2	0.40
8	9	5	3.8	23.9	1.00
8	8	5	1.7	24.2	0.60
8	7	4	1.5	24.5	0.84
7	8	3	0.8	19.1	0.45
7	7	4	5.5	18.6	3.56
Other combinations			0.3	—	0.12
Total			41.9	22.7	11.9




Figure 43: The pileup distribution obtained from the B parking data set. Contributions from each trigger combination are shown, with the histogram areas normalized to the number of events recorded by each trigger.

5.3 Physics results with the B parking data set

Various searches for LFU violation are being considered or are currently underway, e.g., measuring R_{D^*} using $B^0 \rightarrow D^{*-} \ell^+ \nu_\ell$ decays and the tag-side muon; searching for lepton flavor violation in (tag-side) $B_s^0 \rightarrow \mu^- e^+$ decays; and studying charge-parity violating processes in




Figure 44: The invariant mass distribution for pairs of oppositely charged muons originating from a common vertex, obtained from a subset of the B parking data.

fully reconstructed hadronic final states using probe-side b hadron decays, such as $D^0 \rightarrow K_S K_S$ and $B_s^0 \rightarrow \phi(\rightarrow K^+ K^-) \phi(\rightarrow K^+ K^-)$.

Beyond B physics, the data set provides a rich opportunity for the discovery of a broad range of BSM scenarios and will serve novel physics analyses for several years. It provides access to BSM models with low-mass states and/or very rare decays, a parameter space complementary to the one offered by data sets that serve the high- p_T searches typical at the LHC, and thus substantially extends the reach of the CMS physics program.

In addition to the $\approx 10^{10} b\bar{b}$ events, the B parking data set contains $\approx 3 \times 10^{11}$ pp pileup interactions, which may be of interest to searches for BSM processes that yield “untriggerable” signatures (i.e., for which there is no feasible trigger algorithm). Further, experimentally difficult signatures from the decay of low-mass BSM particles, which may be sensitive to combinatorial backgrounds from a high-pileup environment, can exploit the relatively low pileup observed for the B parking data set, particularly for data collected later in the LHC fills by the trigger algorithms with low p_T thresholds.

The following subsections highlight two key physics results obtained with the B parking data set: a measurement of R_K and of the differential branching fraction for the $B^+ \rightarrow K^+ \mu^+ \mu^-$ process, and a BSM search for heavy neutral leptons in b hadron decays.

5.3.1 Measurement of the R_K observable

Given the negligible masses of both electrons and muons in comparison to the B meson mass, the ratio of SM branching fractions, $\mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)$ and $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)$, is very close to unity. The presence of BSM physics could induce appreciable modifications to the branching fractions of the different lepton generations, such as via a leptoquark with flavor-dependent couplings. The observation of a nonunity ratio would provide compelling evidence for LFU violation through a BSM mechanism.

To enhance the sensitivity to LFU violation in $b \rightarrow s\ell\ell$ decays, it is advantageous to employ an observable that minimizes any associated theoretical uncertainties. The ratio of branching fractions $\mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)/\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)$, determined within a specific range of

the dilepton mass squared $q_{\min}^2 < q^2 < q_{\max}^2$, is robust against potentially substantial long-range corrections [151] and is known to be close to unity with a precision of approximately 1% [25–28]. Additionally, experimental systematic uncertainties (such as those related to lepton acceptances, momentum scale, and identification) can be reduced by measuring a double ratio, normalized to the corresponding $B^+ \rightarrow J/\psi(\rightarrow \ell^+ \ell^-)K^+$ decay channels:

$$R_K(q^2)[q_{\min}^2, q_{\max}^2] = \frac{\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)[q_{\min}^2, q_{\max}^2]}{\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow \mu^+ \mu^-)K^+)} \Big/ \frac{\mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)[q_{\min}^2, q_{\max}^2]}{\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow e^+ e^-)K^+)}$$

where the ratio $\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow \mu^+ \mu^-)K^+)/\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow e^+ e^-)K^+)$ has been experimentally determined to be nearly unity with a precision of 0.7% [97].

The R_K observable has been measured by the CMS Collaboration in the q^2 range spanning from 1.1 to 6.0 GeV² [152]. Within the same analysis, a differential branching fraction measurement is performed for the $B^+ \rightarrow K^+ \mu^+ \mu^-$ channel, divided into 15 bins of q^2 ranging from 0.98 to 22.9 GeV² and excluding the resonance regions of J/ψ and $\psi(2S)$.

To determine R_K , events are chosen in which either a $B^+ \rightarrow K^+ \mu^+ \mu^-$ candidate is identified on the tag side of the event (i.e., one of the muons is responsible for the positive trigger decision) or a $B^+ \rightarrow K^+ e^+ e^-$ candidate is identified from the sample of unbiased decays of the other b hadrons in the event.

The B^+ candidates are constructed by pairing two same-flavor leptons with opposite charges, whose invariant mass falls below 5 GeV, along with a positively charged track to which the kaon mass is assigned (in the absence of particle identification for pions and kaons). Rigorous quality criteria are imposed for each candidate, whether it be a muon, electron, or track, to minimize the occurrence of misreconstructed objects.

The tracks of the three particles constituting the B^+ candidate, including that of the muon responsible for the positive trigger decision, are required to have the same point of origin. The tracks are then used in a vertex fitting procedure, which relies on their measured momentum vectors along with associated uncertainties. This procedure enhances the accuracy of mass measurements for both the B^+ candidate and the lepton pair.

The kinematic fit algorithm [153], which is based on a least mean square minimization approach, constrains the tracks associated with each B^+ candidate, ensuring they originate from a common vertex. The particle trajectories are then refitted taking into account the common vertex as an additional constraint, and their momenta are recalculated.

In the electron channel, to enhance the sample purity, at least one of the electrons is required to be reconstructed by the PF algorithm. The second electron can be reconstructed using either the PF algorithm or the dedicated low- p_T (LP) algorithm described in Section 5.2.8.

The B^+ candidate is subjected to a minimal set of kinematic, topological, and quality criteria. A significant fraction of semileptonic decays of heavy-flavor hadrons, containing a $D^0 \rightarrow K^-\pi^+$ decay, remains within the sample, primarily due to the misidentification of the K^- meson as a lepton. To address this issue, a dedicated veto is implemented to remove such decays.

The final selection in each channel relies on a BDT, which combines several variables into a classifier constructed using XGBOOST. In the electron channel, independent BDTs are trained for the two exclusive PF-PF and PF-LP event categories, using the same input variables. Given the notably higher background levels in the electron channels relative to the muon channel, the corresponding BDTs utilize a greater number of input variables. The BDTs use a supervised training, which relies on simulated $B^+ \rightarrow K^+ \ell^+ \ell^-$ decays in the low- q^2 bin for signal, and

data sidebands to the low- q^2 bin for background. The optimal working point for each BDT is selected to maximize the expected significance of the $B^+ \rightarrow K^+ \ell^+ \ell^-$ signal within the low- q^2 region. Various checks are performed to ensure that the BDT performance is unbiased and robust for all three channels (muon, PF-PF, and PF-LP) using multiple data control regions.

After the final selection of B^+ candidates using the BDTs, the mass distributions obtained from the reconstructed B^+ candidates in each channel are fitted using analytical functions, as shown in Figs. 45 and 46. The fits for the muon, PF-PF, and PF-LP channels are statistically combined to yield the final R_K measurement. Systematic uncertainties are assessed independently for all channels. The uncertainty in the final measurement is statistically limited. The measurement of R_K is

$$R_K = 0.78^{+0.46}_{-0.23} \text{ (stat)}^{+0.09}_{-0.05} \text{ (syst)} = 0.78^{+0.47}_{-0.23},$$

which is consistent with the SM prediction within one standard deviation and can be compared with the corresponding LHCb measurement of $R_K = 0.949^{+0.048}_{-0.047}$ [143].

The sample of $B^+ \rightarrow K^+ \mu^+ \mu^-$ candidates is sufficiently large to enable both an inclusive and a differential measurement of $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)$. Figure 47 shows the measured differential branching fraction of the $B^+ \rightarrow K^+ \mu^+ \mu^-$ decay, as a function of q^2 , along with corresponding SM theoretical predictions as determined by various packages. The inclusive value for $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)$ within the q^2 region $1.1 < q^2 < 6.0 \text{ GeV}^2$ is $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-) = (12.42 \pm 0.68) \times 10^{-8}$, which is statistically limited and consistent with the world average [97].




Figure 45: Results of an unbinned likelihood fit to the invariant mass distributions for the the $B^+ \rightarrow J/\psi(\rightarrow \mu^+ \mu^-)K^+$ (left) and the $B^+ \rightarrow K^+ \mu^+ \mu^-$ (right) channels. Various functions are used to parametrize the contributions from the signal and various background processes. Taken from Ref. [152].

5.3.2 Search for long-lived heavy neutrinos in B meson decays

Heavy neutrinos (N) that feebly mix with the SM neutrinos constitute a possible appealing answer to several open questions in the SM, notably the evidence for nonzero neutrino masses in neutrino oscillations [154], the large amount of dark matter inferred from astrophysical and cosmological measurements [155], and the baryon asymmetry in the universe [156]. The small mixing amplitudes, V_{eN} , $V_{\mu N}$, and $V_{\tau N}$, between heavy neutrinos and their SM counterparts in the three flavor families imply that heavy neutrinos are long-lived. Indeed, the N particle proper lifetime scales as $c\tau_N \propto m_N^{-5}|V_N|^{-2}$, where m_N is the mass of the heavy neutrino and $|V_N|^2$ is defined as $|V_N|^2 = |V_{eN}|^2 + |V_{\mu N}|^2 + |V_{\tau N}|^2$. The large signal lifetimes give rise to displaced signatures. Moreover, thanks to mixing, flavor conservation may be violated.




Figure 46: Results of an unbinned likelihood fit to the invariant mass distributions for the $B^+ \rightarrow J/\psi(\rightarrow e^+e^-)K^+$ (left) and the $B^+ \rightarrow K^+e^+e^-$ (right) channels. The upper (lower) panels are for candidates using the PF-PF (PF-LP) category. Various functions are used to parametrize the contributions from the signal and various background processes. Taken from Ref. [152].

Numerous searches have been performed at the most diverse experiments covering a vast range of heavy neutrino masses, from the keV to the TeV scale, and mixing amplitudes, down to $|V_{N\ell}|^2 \approx 10^{-7}$. In CMS, virtually all searches for N with m_N up to around 100 GeV relied on the process $W \rightarrow \ell N$, where events are collected by high- p_T , single isolated lepton triggers [157]. However, electroweak processes with on-shell gauge boson production are not the sole or even the most abundant source of neutrinos, and hence also potentially of heavy neutrinos, at the LHC: leptonic and semileptonic B meson decays (shown in Fig. 48) are more than a thousand times more copious thanks to the much larger $pp \rightarrow b\bar{b}$ cross section. Furthermore, the final-state particles produced in B meson decays have lower momenta than those produced in W boson decays, because the B mesons have a lower mass than the W boson. For long-lived signatures, the softer momentum spectrum is an advantage as it leads to a higher fraction of N particles that decay within the CMS acceptance.

The 2018 B parking data set (described in Section 5.2) enabled CMS to perform a search for long-lived heavy neutrinos in B meson decays, using events with one lepton plus one displaced vertex comprising a lepton-pion pair and compatible with the decay of a long-lived heavy neutrino $N \rightarrow \ell^\pm \pi^\mp$ [158]. Both muons and electrons are considered as long as at least one muon matched to a B parking trigger algorithm is present. This search would have not been possible with standard triggers.

This analysis sets the most stringent upper limits, among those obtained at collider exper-




Figure 47: Comparison of the measured differential $B^+ \rightarrow K^+ \mu^+ \mu^-$ branching fraction with the theoretical predictions obtained using HEPFIT, SUPERISO, FLAVIO, and EOS packages. Reliable predictions are not available between the J/ψ and $\psi(2S)$ resonance regions. The HEPFIT predictions are available only for $q^2 < 8 \text{ GeV}^2$. Taken from Ref. [152].




Figure 48: Feynman diagram of a semileptonic decay of a B meson into the primary lepton (ℓ_p), a hadronic system (X), and a SM neutrino, which then mixes with a heavy neutrino (N). The N decays weakly into a charged lepton ℓ^\pm and a charged pion π^\mp , forming a vertex displaced from the pp interaction point. Taken from Ref. [158].

iments, at 95% CL on $|V_N|^2$ for heavy neutrinos with masses between 1.0 and 1.7 GeV and does so with unprecedented resolution as the N decay is fully reconstructed. Results are interpreted in various scenarios specified by different values of the mixing ratios $r_\ell \equiv |V_{\ell N}|^2/|V_N|^2$, $\ell = (\text{e}, \mu, \tau)$, as well as under either the Majorana (shown in Fig. 49) or Dirac hypotheses. For masses $m_N = 1.0, 1.5, 2.0 \text{ GeV}$, lower limits on $c\tau_N$ are provided for sixty-six possible flavor violating scenarios, depicted in Fig. 50.

6 Data parking in Run 3

The successful execution of the B parking program during Run 2 garnered significant interest, propelling the evolution of the parking technique into a comprehensive and diverse program. The enhancement of the B parking strategy in Run 3 is discussed in Section 6.1, which introduces novel inclusive triggers to collect data in the dimuon and dielectron final states.




Figure 49: Expected and observed 95% CL limits on $|V_N|^2$ as a function of m_N in the Majorana scenario, for the coupling hypotheses $(r_e, r_\mu, r_\tau) = (0, 1, 0)$ on the left and $(r_e, r_\mu, r_\tau) = (1/3, 1/3, 1/3)$ on the right. The mass range with no limits shown corresponds to the D^0 meson veto employed by the search. Taken from Ref. [158].




Figure 50: Observed limits on $c\tau_N$ as a function of the coupling ratios (r_e, r_μ, r_τ) for fixed N masses of 1 GeV (upper left), 1.5 GeV (upper right), and 2 GeV (lower center), in the Majorana scenario. A red cross indicates that no exclusion limit was set for that point. The tick orientation indicates the direction of reading. Taken from Ref. [158].

In Section 6.2, we provide an overview of other diverse opportunities arising from the parking strategy that extend beyond the B physics measurements, referred to as “alternative” data parking strategies.

6.1 Data parking for B physics in Run 3

The LHC Run 3 period started in 2022 and is currently expected to finish at the end of 2025. The LHC delivered $\mathcal{L}_{\text{int}} = 73.4 \text{ fb}^{-1}$ during 2022 and 2023, and the total may exceed 250 fb^{-1} by the end of the data-taking period in 2025. Thus, Run 3 represents an important opportunity to improve our understanding of b hadron production and decay modes. The LHC center-of-mass energy has also increased from 13 to 13.6 TeV between Runs 2 and 3, yielding a modest increase of 4% [145–148] in the $b\bar{b}$ inclusive production cross section.

A new B parking trigger strategy was designed and deployed in time for the beginning of pp collisions in 2022. The new strategy focuses on inclusive dilepton trigger algorithms, operating at high rates, that aim to provide gains in physics reach significantly beyond those expected from an increase in \mathcal{L}_{int} . The primary aim for Run 3 is to substantially expand the B physics program of CMS by using new or improved trigger algorithms and a data parking strategy to provide high acceptance for dimuon and dielectron final states from b hadron decays. Dimuon final states are central to the CMS B physics program and the majority of analyses leverage this clean experimental signature. An inclusive dimuon trigger strategy, which differs from the exclusive strategy deployed in Run 2, was identified as a key beneficiary of the data parking approach. Further, new triggers recording dielectron final states are specifically aimed at improving the statistical precision of the R_K measurement.

6.1.1 Trigger strategy

The B parking trigger and data parking strategies of 2018 have been adapted to accommodate changes in LHC operations and the evolving needs of the B physics program for Run 3. During 2022, the LHC machine parameters shifted to a new mode of operation in which $\mathcal{L}_{\text{inst}}$ is delivered to the experiments at a constant value of $\approx 2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (corresponding to $\text{PU} \approx 57$) for several hours. This luminosity-leveling mode, introduced in Section 1.3, allows experiments to accumulate large data samples as quickly as possible, under conditions suitable for data analysis—i.e., with a manageable amount of pileup—such that the quality of the event reconstruction and physics performance is not compromised. Following a leveling period of up to six hours in 2022, the $\mathcal{L}_{\text{inst}}$ value decreases during the remainder of the LHC fill, typically over a time period of 6–12 hours.

The change in the LHC machine parameters between Run 2 and Run 3 have important consequences for the trigger and data parking strategies described here. Since there are no idle resources available for the first few hours of an LHC fill, due to the aforementioned leveling, a fraction of the total rate budget for the L1 trigger has to be allocated to the dilepton trigger algorithms described here. This is in contrast to operations in 2018, during which the single-muon trigger algorithms were not enabled when the LHC operated with $\mathcal{L}_{\text{inst}}$ values above $1.7 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. In the case of the new dielectron trigger algorithms, the allocation is increased later in an LHC fill as $\mathcal{L}_{\text{inst}}$ falls and idle resources become increasingly available.

The standard dimuon trigger algorithms deployed in Run 2 to cater for B physics analyses typically required a reconstructed dimuon system plus additional constraints specific to a particular physics process to control trigger rates. For instance, in the HLT algorithms, the reconstructed invariant mass was restricted to a set of windows around the B_s meson and quarkonia masses (e.g., for the process $B_s^0 \rightarrow \mu^+ \mu^-$), or the presence of an additional particle track was

required (e.g., for $B^+ \rightarrow K^+ \mu^+ \mu^-$), or else the dimuon vertex position was required to be displaced with respect to the pp luminous region (i.e., for any $b \rightarrow \mu^+ \mu^- X$ process).

The new dimuon trigger algorithms do not impose any of the rate-reducing constraints highlighted above and are instead designed to provide broad coverage for a range of physics processes, by employing minimal kinematical and topological requirements. In contrast to the single-muon trigger strategy employed in 2018, where the thresholds were adjusted throughout the fill, the new dimuon triggers are always enabled and employ fixed thresholds. The dimuon trigger algorithms exploit improvements in both the L1 and HLT systems [31] implemented during LS2: the former now performs muon-track finding using the Kalman filter technique [159] to provide an improved muon p_T estimate, and the latter uses new algorithms implemented in a heterogeneous computing environment, comprising both CPU and GPU cores. The new algorithms both speed up track reconstruction and improve the track p_T resolution [36], which in turn can be exploited by the HLT-based muon algorithms. As a consequence, there is a substantial improvement in the experimental acceptance for a number of interesting production and decay modes that yield prompt and nonprompt dimuon final states, such as $\eta \rightarrow \mu^+ \mu^-$ and $B_s^0 \rightarrow \mu^+ \mu^-$. Furthermore, the new algorithms also improve the acceptance for b hadron decay chains that produce additional particles with nonnegligible lifetimes, such as $B^0 \rightarrow J/\psi K_S^0$.

The dielectron trigger algorithms target the measurement of the R_K observable. The measurement reported in Section 5.3.1 is limited by the finite number of reconstructed $B^+ \rightarrow K^+ e^+ e^-$ decays, because of the challenges associated with small branching fractions and low- p_T daughter particles for the probe-side B meson decays, as detailed in Section 5.2.2. Here, an alternative strategy is explored by using the trigger algorithms to directly identify the electron pairs produced in $b \rightarrow se^+ e^-$ transitions. This approach removes the aforementioned challenges but requires the use of electron reconstruction and identification algorithms, implemented in the L1 system, with limited performance. These L1 algorithms measure $E_T = E \sin \theta$, where θ is the polar angle [30]. Low E_T thresholds are required to provide adequate experimental acceptance for $B^+ \rightarrow K^+ e^+ e^-$ decays and, as a result, trigger rates in the multi-kHz range are required. The “dynamic threshold” strategy used in 2018, i.e., progressively reducing kinematical thresholds as $\mathcal{L}_{\text{inst}}$ falls during an LHC fill, is deployed again to access the lowest possible E_T thresholds and maximize the acceptance.

Additional details on the dimuon and dielectron trigger algorithms are provided in Sections 6.1.3 and 6.1.4, respectively.

6.1.2 Data parking strategy

In 2018, data parking referred to the delayed reconstruction of trigger data streams, on the order of several months (long-term parking). Since 2022, the definition of data parking has shifted somewhat to include prompt reconstruction (i.e., processing typically starting within 48 hours), contingent upon the availability of computing resources, and short-term parking, where the reconstruction is delayed until resources become available. Crucially, the ability to promptly reconstruct data streams serving the core CMS physics program must not be compromised by the data-parking strategy.

The initial budgeting of resources for data parking in 2022 was predicated on the assumption that 25% of the dimuon trigger data stream would undergo prompt reconstruction, while the remaining 75% would be (short-term) parked for delayed reconstruction. Additionally, it was foreseen that a small fraction of the dielectron data stream would also be promptly reconstructed, on an opportunistic basis and subject to resource availability. In practice, during

2022 and 2023, the data streams generated by both the dimuon and dielectron trigger algorithms were promptly reconstructed in full, since the availability of resources exceeded initial expectations.

6.1.3 Inclusive triggers for dimuon final states

The dimuon strategy for Run 3 is implemented in two L1 algorithms and three HLT algorithms. All algorithms were enabled at the beginning of the data-taking period in 2022.

Both L1 algorithms require two oppositely charged muons. The “central- η ” algorithm also imposes $|\eta| < 2$ and $p_T > 0$ GeV on each muon, even if an implicit threshold of $p_T \gtrsim 3$ GeV is required if the muons are to reach the muon detectors in the central region, and $|\Delta\eta(\mu_1, \mu_2)| < 1.6$. These thresholds were updated to $p_T > 3$ GeV and $\Delta R(\mu_1, \mu_2) < 1.4$ in 2023. The “higher p_T ” algorithm requires $p_T > 4$ GeV and $\Delta R(\mu_1, \mu_2) < 1.2$. The peak L1 trigger rate recorded by these two algorithms was 18 kHz at $\mathcal{L}_{\text{inst}} = 2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, which is approximately 20% of the total L1 trigger rate. This large allocation was possible as a result of an increase in the total L1 trigger rate from ≈ 90 kHz in 2018 to ≈ 100 kHz in 2022, stemming from operational improvements.

The first HLT algorithm, known henceforth as the “inclusive low-mass dimuon trigger”, imposes the following requirements: two opposite-charge muons, one muon satisfying $p_T > 4$ GeV and the other $p_T > 3$ GeV, a dimuon vertex fit probability of $P_{\text{vtx}}(\mu\mu) > 0.5\%$, and a reconstructed dimuon invariant mass satisfying $2m_\mu^{\text{PDG}} < m_{\mu\mu} < 8.5$ GeV. Here the superscript “PDG” refers to the global average of experimentally measured mass values as reported by the Particle Data Group (PDG) [97]. This algorithm now serves as the main trigger for all physics analyses of final states containing a low-mass dimuon system. The HLT algorithm results in a peak HLT rate of 1.6 kHz at $\mathcal{L}_{\text{inst}} = 2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

A second HLT algorithm, known as the “displaced low-mass dimuon trigger”, imposes tighter requirements: two opposite-charge muons with both muons satisfying $p_T > 4$ GeV; $P_{\text{vtx}}(\mu\mu) > 10\%$; a reconstructed dimuon vertex that is sufficiently displaced from the pp luminous region such that it satisfies $l_{xy}/\sigma_{l_{xy}} > 3$, where l_{xy} and $\sigma_{l_{xy}}$ are, respectively, the measured transverse decay length and its uncertainty; and $2m_\mu^{\text{PDG}} < m_{\mu\mu} < 8.5$ GeV. This trigger, with a lower trigger rate of 0.3 kHz at $\mathcal{L}_{\text{inst}} = 2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, acts as a backup to the inclusive trigger. Its event stream is a subset of that of the inclusive trigger. As a further fail-safe, the majority of all standard dimuon algorithms used during Run 2 are also maintained during Run 3, even if they also operate largely in the shadow of the new inclusive trigger.

The third algorithm records events that are likely to contain the production and decay of the $Y(nS)$ resonances, where $n = 1, 2, 3$. The requirements comprise: two oppositely charged muons, $P_{\text{vtx}}(\mu\mu) > 0.5\%$, a reconstructed dimuon system satisfying $p_T > 10$ GeV, $|y| < 1.4$, and $8.5 < m_{\mu\mu} < 11.5$ GeV.

Figure 51 shows the invariant mass distribution of dimuon final states for events from the data streams recorded by the trigger algorithms described above. The reconstructed dimuon candidates are required to satisfy loose kinematical, topological, and quality criteria. Several resonances are clearly visible over the continuum background, such as η , ω , ϕ , J/ψ , $\psi(2S)$, and $Y(nS)$ ($n = 1, 2, 3$).

Several physics analyses stand to benefit from the new inclusive approach. For example, the $B^0 \rightarrow J/\psi K_S^0$ decay is used to measure the $\sin(2\beta)$ angle in the CKM unitary triangle [160]. The new inclusive trigger provides a factor of 12 increase in the signal yield per fb^{-1} with




Figure 51: Dimuon mass spectra obtained from data recorded in 2022 during Run 3, corresponding to $\mathcal{L}_{\text{int}} = 3.2 \text{ fb}^{-1}$. In the range $2m_\mu^{\text{PDG}} < m_{\mu\mu} < 8.5 \text{ GeV}$, the light blue distribution represents the subset of dimuon events triggered by the inclusive low-mass trigger algorithm, while the dark blue distribution shows the subset of dimuon events triggered by the displaced low-mass trigger path. In the range $8.5 < m_{\mu\mu} < 11.5 \text{ GeV}$, dimuon events are instead triggered by the HLT paths targeting the $Y(nS)$ resonances, which are shown by the pink distribution.

respect to the standard dimuon triggers available in Run 2. The left panel of Fig. 52 shows the mass distribution of $B^0 \rightarrow J/\psi K_S^0$ candidates determined from events recorded by the dimuon triggers and selected according to the following requirements: $|m_{\pi^+\pi^-} - m_{K_S^0}^{\text{PDG}}| < 10 \text{ MeV}$, the reconstructed $J/\psi \rightarrow \mu^+\mu^-$ candidate satisfies $|m_{\mu\mu} - m_{J/\psi}^{\text{PDG}}| < 150 \text{ MeV}$ and $p_T > 8 \text{ GeV}$, the reconstructed B^0 candidate satisfies $p_T > 10 \text{ GeV}$, and the reconstructed J/ψ and K_S^0 candidates form a vertex with a fit probability greater than 10%. A one-dimensional unbinned extended maximum likelihood fit is performed to the mass distribution: two Gaussian functions with a common mean are used to model the signal distribution, an exponential function is used for the combinatorial background, and an error function is used to describe the $J/\psi K_S^0 + X$ background.




Figure 52: Left: invariant mass distribution for candidate $B^0 \rightarrow J/\psi K_S^0$ decays. Right: proper decay length (ct) distribution obtained from candidate $B^+ \rightarrow J/\psi (\rightarrow \mu^+\mu^-) K^+$ decays. Both types of candidates are reconstructed from events recorded using the dimuon triggers.

Another example of a substantial gain is in the decay $\Lambda_b \rightarrow J/\psi (\rightarrow \mu^+\mu^-) \Lambda$, which would permit a new angular analysis of this channel in CMS. In addition, significantly improved acceptance to low-mass resonances such as ω and ϕ allow for CPV measurements in decays such as $B_s^0 \rightarrow \phi (\rightarrow \mu^+\mu^-) \phi (\rightarrow K^+K^-)$.

The right panel of Fig. 52 shows the proper decay length (ct) distribution as obtained from a control sample of $B^+ \rightarrow J/\psi(\rightarrow \mu^+\mu^-)K^+$ decays in the events recorded by the dimuon triggers. The following selection criteria are also applied: the reconstructed kaon, J/ψ , and B^+ candidates are required to satisfy p_T thresholds of 2, 8, and 15 GeV, respectively, and the J/ψ and kaon candidates are required to form a vertex with a fit probability greater than 15%. The J/ψ candidate is also required to satisfy $|m_{\mu\mu} - m_{J/\psi}^{\text{PDG}}| < 150$ MeV. The ct distributions for both the signal and $J/\psi K^+ + X$ background processes are modeled with decaying exponential functions of the form $(e^{-ct/\lambda})$; the combinatorial background uses a decaying exponential function with both negative and positive terms. All exponential functions are convolved with a Gaussian function to reflect the finite time resolution. The absence of any displacement requirements in the inclusive trigger is expected to improve the accuracy of the lifetime measurements of the comparatively shorter-lived B_c^+ meson.

Figure 53 shows regions of the invariant mass distribution, as recorded by the inclusive low-mass trigger algorithm, that highlight the η and J/ψ resonances. Fits are performed to the data in these regions, as well as to other regions that cover the ϕ , $\psi(2S)$, and $Y(nS)$ ($n = 1, 2, 3$) resonances. These fits are used to extract the measured dimuon mass $m_{\mu\mu}$ and resolution $\sigma(m_{\mu\mu})$ parameters for each resonance. The measured dimuon mass scale $s_{\mu^+\mu^-}$ is defined in terms of the PDG mass value [97] for the resonance under consideration: $s_{\mu^+\mu^-} = |m_{\mu\mu} - m_{\mu^+\mu^-}^{\text{PDG}}| / m_{\mu^+\mu^-}^{\text{PDG}}$. For the J/ψ , $\psi(2S)$ and $Y(nS)$ resonances, the intrinsic width of the resonance is negligible with respect to the experimental resolution. These resonances are modeled by the CB and Gaussian functions that share the same $m_{\mu\mu}$ parameter. The latter function is used to correctly model effects due to the finite experimental resolution. The continuous background is described by an exponential function. For some resonances, such as the ϕ , the CB and Gaussian functions are convolved with a Breit–Wigner function that models the natural lineshape for the resonance. The dimuon mass scale is accurate to the per-mil level, and the relative dimuon mass resolution $\sigma(m_{\mu\mu})/m_{\mu\mu}$ is measured to be in the range 0.8–1.6% for the different resonances.




Figure 53: Dimuon invariant mass distributions in the η (left) and J/ψ (right) mass regions, as obtained from data recorded by the inclusive low-mass dimuon trigger algorithm.

6.1.4 Inclusive triggers for dielectron final states

The dielectron trigger algorithms require the presence of two electrons that satisfy loose kinematical and ID quality requirements. The L1 algorithm identifies electron/photon (e/γ) candidates from patterns of measured energy deposits in the ECAL system. Both e/γ candidates must satisfy $|\eta| < 1.2$ and a variable threshold applied to the E_T measurement of each object. The highest E_T threshold applied is 11 GeV, which results in an L1 trigger rate of ≈ 2 kHz for

$\mathcal{L}_{\text{inst}} = 2.0 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$. Copies of the algorithm exist with lower thresholds, available in 0.5 GeV steps down to 5 GeV, which are enabled later during an LHC fill when the $\mathcal{L}_{\text{inst}}$ and pileup are smaller and idle resources become available. The evolution of thresholds is tuned such that the L1 system operates close to its design limit of 100 kHz throughout each fill. The L1 trigger algorithms also impose an upper bound on the separation distance ΔR between the two e/γ candidates; the upper bound evolves between values from 0.6 to 0.9 with the decreasing E_T threshold. This requirement helps to control the trigger rate, while maintaining an efficiency of 95% for physics processes that produce two final-state electrons with four-momenta that yield an invariant mass below 6 GeV.

Positive trigger decisions from the L1 algorithms act as seeds for companion algorithms implemented in the HLT, which are able to exploit superior physics performance for reconstructed particle candidates by combining information from the tracker, ECAL, and HCAL subdetectors. Each version of the L1 algorithm, with its unique E_T threshold, is paired with a version of the HLT logic that imposes a corresponding threshold on the measured E_T values determined by the HLT algorithms. The HLT E_T thresholds, which range from 4 to 6.5 GeV, are lower than the L1 E_T thresholds so as to maintain high efficiency while controlling purity. Several quality-related criteria are also imposed by the HLT algorithms to help identify genuine electrons from interesting physics processes, while rejecting other particles misidentified as electrons. The ID criteria follow closely those used by the standard HLT electron reconstruction algorithms [39], with minor changes that improve the per-electron efficiency from 75 to 90%, relative to the L1 e/γ object. An upper bound of 6 GeV is imposed on the invariant mass of the dielectron system, which maps closely to the ΔR requirements imposed at L1. Overall, the HLT requirements are tuned with an emphasis on high signal efficiency, while substantially improving the purity of the L1 data stream.

Table 13 lists the unique combinations of E_T and ΔR thresholds used in the L1 and HLT algorithms to record events containing dielectron final states. Settings with lower E_T thresholds are enabled as the $\mathcal{L}_{\text{inst}}$ decreases during an LHC fill. At least one setting is enabled from the start of every LHC fill, i.e. throughout the period of luminosity leveling. In the L1 system, the highest E_T threshold is enabled for $\mathcal{L}_{\text{inst}}$ values of $2.0 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$ and above, which corresponds to as many as 57 pileup interactions. For each setting, the \mathcal{L}_{int} and the mean number of pileup interactions are aggregated over the periods for which the setting provides the lowest enabled L1 E_T threshold. The peak L1 and HLT trigger rates are listed for each setting. The efficiency for recording events containing $B^+ \rightarrow K^+ e^+ e^-$ decays changes significantly across the different settings, differing by a factor of ≈ 20 for the settings with the lowest and highest L1 E_T thresholds.

The left and right panels of Fig. 54 illustrate how the L1 and HLT trigger rates evolve during a typical LHC fill during 2022. Similar to 2018, instantaneous increases in rate are observed when changes in the trigger configuration occur, and the L1 system operates with a trigger rate close to 100 kHz throughout the LHC fill. The L1-based algorithms for the dielectron and dimuon triggers (the latter are described in Section 6.1.3) contribute up to 54 and 18 kHz to the total rate shown in the left panel of Fig. 54.

The right panel of Fig. 54 shows the total HLT trigger rates for both the promptly reconstructed data stream, which serves the core CMS physics program, and the B parking stream. The latter stream includes contributions from both the dielectron and dimuon algorithms. The HLT algorithms for the dielectron triggers typically operate at a rate of ≈ 100 Hz during the luminosity leveling period and as high as ≈ 1.3 kHz later in an LHC fill. The dimuon algorithms operate at a total rate of up to ≈ 1.2 kHz, which monotonically falls with $\mathcal{L}_{\text{inst}}$ during an LHC fill.

Table 13: Trigger configurations, defined by unique combinations of L1 and HLT E_T thresholds (applied to each electron candidate) and L1 ΔR , used to record events containing dielectron final states. The thresholds on the L1 and HLT E_T (L1 ΔR) values are lower (upper) bounds. The \mathcal{L}_{int} value and the mean number of pileup interactions recorded by each trigger combination are aggregated over periods for which each combination provided the lowest enabled L1 E_T threshold. Representative peak L1 and HLT trigger rates are given for each setting.

L1 E_T [GeV]	L1 ΔR	HLT E_T [GeV]	\mathcal{L}_{int} [fb $^{-1}$]	Mean PU	Peak L1 rate [kHz]	Peak HLT rate [kHz]
11.0	0.6	6.5	1.6	45.6	2.2	0.1
10.5	0.6	6.5	1.1	42.2	3.0	0.3
9.0	0.7	6.0	8.8	47.4	9.3	0.6
8.5	0.7	5.5	3.3	46.2	13	0.9
8.0	0.7	5.0	6.9	39.1	16	1.2
7.5	0.7	5.0	1.6	40.3	23	1.4
7.0	0.8	5.0	2.7	36.3	27	1.3
6.5	0.8	4.5	3.6	31.2	35	1.3
6.0	0.8	4.0	2.5	27.4	46	1.4
5.5	0.8	4.0	0.7	23.6	54	1.0
Other combinations		1.0	—	—	—	—
Total		33.9	34.8	—	—	—




Figure 54: Total L1 (left) and HLT (right) trigger rates, and the number of pileup interactions, shown as a function of time for a representative LHC fill during 2022. The rates for the promptly reconstructed core physics (black solid markers) and B parking (blue open markers) data streams are shown separately in the right panel. Occasional lower rates are observed due to transient effects, such as the throttling of the trigger system in response to subdetector dead time [9]. Changes in the trigger configuration are indicated by vertical green dashed lines.

Figure 55 shows the pileup distribution for the dielectron data set, along with the contributions from each of the individual trigger combinations. Figure 56 shows the invariant mass distributions for pairs of oppositely charged electrons originating from a common vertex, as obtained from the dielectron data set. Both electrons are required to satisfy a minimal set of kinematic and ID criteria, and both electrons must be matched to the electron candidates responsible for




Figure 55: Pileup distribution measured in the dielectron data set. Contributions from each trigger combination are shown, with the histogram areas normalized to the number of events recorded by each trigger.




Figure 56: The invariant mass distribution for pairs of oppositely charged electrons originating from a common vertex, reconstructed from the dielectron data set.

the positive trigger decision. Peaks in the data resulting from the $\rho/\omega/\phi$, J/ψ , $\psi(2S)$, and $Y(1S)$ resonances are clearly visible.

Figure 57 shows the invariant mass distributions for candidate $B^+ \rightarrow J/\psi(\rightarrow e^+e^-)K^+$ decays, as obtained from the dielectron data set. A pair of oppositely charged electrons originating from a common vertex, along with a charged-particle track assumed to be the kaon, are required to satisfy a minimal set of kinematic and ID criteria, and both electrons must be matched to the electron candidates responsible for the positive trigger decision. The invariant mass of the dielectron system is required to satisfy $2.9 < m_{ee} < 3.2 \text{ GeV}$ so as to be consistent with originating from the J/ψ meson decay.




Figure 57: The invariant mass distribution for candidate $B^+ \rightarrow J/\psi(\rightarrow e^+e^-)K^+$ decays, reconstructed from the dielectron data set. The histogram is normalized to unit area.

6.1.5 Physics potential

The new dilepton triggers will accumulate large data samples in Run 3, with substantially improved acceptances for many physics processes relative to Run 2. The dimuon triggers will provide substantial gains, beyond those expected from an increase in \mathcal{L}_{int} alone, for a range of interesting physics processes. In particular, processes that involve the associated production of $\mu^+\mu^-$ pairs and additional particles will benefit. Examples include $B^+ \rightarrow K^+\mu^+\mu^-$, $B_c^+ \rightarrow J/\psi\ell\nu$, $B^0 \rightarrow J/\psi K_S^0$, and $B_s^0 \rightarrow J/\psi\phi$ decays. The new dielectron triggers aim primarily to allow a precise measurement of R_K using the $B^+ \rightarrow K^+e^+e^-$ decay. Investigations are ongoing for the potential to perform an angular analysis or differential branching fraction measurements for decays such as $B^+ \rightarrow K^+e^+e^-$ and $B^0 \rightarrow K^*(892)^0e^+e^-$. Furthermore, both the dimuon and dielectron triggers provide ample scope for novel BSM searches. Both triggers also use L1 algorithms that are being adopted by the data scouting trigger streams to further improve acceptance to low-mass states.

6.2 Alternative data parking strategies in Run 3

After a break in Run 2 when the parking strategy focused predominantly on B physics analyses, the original idea of a more diversified parking strategy, meant to complement the existing standard triggers, was revived in Run 3. The main goal is to overcome the limited HLT bandwidth and to collect a sufficient number of events for specific physics goals, such as studying final states produced via VBF, exploring possible anomalous HH production, and improving the sensitivity to searches for exotic LLPs. These three novel parking approaches are detailed in the following sections.

6.2.1 VBF parking

Higgs boson production via the VBF channel is of paramount importance for the experimental study of the Higgs boson at the LHC. It is the second most common production mechanism, contributing about 10% of the total H production cross section. The sensitivity to several H decay modes, such as $H \rightarrow \tau^+\tau^-$ [161], $H \rightarrow \text{invisible}$ [162], and $H \rightarrow \mu^+\mu^-$ [163], is driven by the sensitivity to H VBF production. The VBF production is also important for a variety of

measurements, e.g., in effective field theory measurements that constrain dimension-6 operators [164], and in HH production, where it grants unique access to the VVHH coupling [165]. The VBF triggers constitute an interesting workaround to the low signal efficiencies obtained with triggers that must be sufficiently restrictive to keep rates under control. Instead of restricting the kinematic properties of the central physics objects from the signal of interest, VBF triggers place tighter constraints on the auxiliary jets. These requirements are often sufficient to significantly loosen or even entirely remove the selection on the central objects.

A dedicated inclusive VBF L1 seed was already introduced in 2017 [166]. This seed requires at least two jets with $p_T > 110$ and 35 GeV , respectively, and at least one pair of jets satisfying $m_{jj} > 650\text{ GeV}$ among all pairs of jets with $p_T > 35\text{ GeV}$ in the event.

In 2023, the VBF trigger strategy was extended with the introduction of a retuned inclusive VBF L1 seed, as well as a set of exclusive seeds, each requiring either one additional muon, τ_h , e/γ object, or p_T^{miss} , or two additional central jets, contributing an L1 rate of about 10 kHz with $\mathcal{L}_{\text{inst}} = 2 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$. These L1 triggers are described in Table 14.

Table 14: Definition and rates of the VBF algorithms at L1. The quoted rates are for $\mathcal{L}_{\text{inst}} = 2 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$ and do not account for overlaps with other seeds.

L1 trigger seed	VBF requirements ($p_T^{j1}, p_T^{j2}, m_{jj}$)	Requirements on additional objects	Rate [kHz]
VBF inclusive	(90, 30, 800)	—	5.0
VBF + 2 central jets	(60, 30, 500)	2 central jets, $p_T^j > 50\text{ GeV}$	3.0
VBF + p_T^{miss}	(65, 30, 500)	$p_T^{\text{miss}} > 65\text{ GeV}$	2.9
VBF + μ	(90, 30, 500)	1 muon, $p_T^\mu > 3\text{ GeV}$	2.3
VBF + τ_h	(35, 35, 450)	1 isolated τ_h , $p_T^\tau > 35\text{ GeV}$	3.1
VBF + e/γ	(40, 40, 450)	1 isolated e/γ object, $E_{e/\gamma} > 15\text{ GeV}$	1.0

The new L1 VBF triggers were used to define a set of loose HLT paths, described in Table 15, contributing a total HLT rate of about 1.2 kHz , directed to a dedicated VBF parking stream. The introduction of these new HLT paths is expected to significantly improve the acceptance to SM-like VBF signals. The low thresholds on central objects achieved with the exclusive paths will enable the efficient probing of final-state topologies sensitive to BSM models, such as dark photons [167], signatures with soft unclustered energy patterns (SUEPs) [168], and any currently unexplored experimental signatures that may become of interest in the future.

The efficiencies of the new VBF triggers were measured in data with a set of events passing a reference single-muon trigger available in the 2023 trigger menu, with muon p_T threshold of 27 GeV . Offline, events are required to contain exactly one isolated and well-identified muon with $p_T > 30\text{ GeV}$. Additional offline requirements on the kinematic variables of the VBF jets and of any central objects are applied to the events, depending on the trigger under consideration, in order to match the requirements made at the trigger level. Figure 58 shows the efficiencies of the inclusive VBF and VBF+ p_T^{miss} triggers, both at the L1 and HLT, as functions of m_{jj} . The HLT efficiencies plateau in both cases at around 80% because of suboptimal jet quality requirements in the online implementation. This has been fixed for the 2024 data-taking period.

Figure 59 compares the m_{jj} distribution of the VBF $H \rightarrow \text{invisible}$ events that pass only the Run 2 standard trigger path versus events that pass either the standard triggers or the VBF triggers (inclusive or VBF+ p_T^{miss}). The new algorithms show a significant acceptance improve-

Table 15: Definition and rates of the VBF paths at the HLT. The quoted rates are for $\mathcal{L}_{\text{inst}} = 2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$ and do not account for overlaps with other seeds.

HLT trigger path	VBF requirements ($p_T^{j1}, p_T^{j2}, m_{jj}, \Delta\eta_{jj}$)	Requirements on additional objects	Rate [Hz]
VBF inclusive	(105, 40, 1000, 3.5)	—	800
VBF + 2 central jets	(70, 40, 600, 2.5)	2 central jets, $p_T^j > 60 \text{ GeV}$	380
VBF + p_T^{miss}	(75, 40, 500, 2.5)	$p_T^{\text{miss}} > 85 \text{ GeV}$	110
VBF + μ	(90, 40, 600, 2.5)	1 muon, $p_T^\mu > 3 \text{ GeV}$	120
VBF + τ_h	(45, 45, 500, 2.5)	1 isolated τ_h , $p_T^\tau > 45 \text{ GeV}$	40
VBF + γ	(45, 45, 500, 2.5)	1 photon, $p_T^\gamma > 17 \text{ GeV}$	100
VBF + e	(45, 45, 500, 2.5)	1 isolated electron, $p_T^e > 17 \text{ GeV}$	2

ment, ranging from 20% at high m_{jj} to more than 300% at lower m_{jj} . The impact of the new VBF triggers on the overall acceptance to other benchmark signals is illustrated in Table 16.




Figure 58: The L1 (blue) and L1+HLT (red) efficiencies as a function of m_{jj} for the VBF inclusive (left) and VBF+ p_T^{miss} (right) parking triggers. In the right figure, p_T^{miss} (no- μ) refers to the event p_T^{miss} corrected for muons.

Table 16: Gains in signal acceptance from the new individual VBF parking paths for a selection of benchmark signals with respect to the relevant Run 2 triggers used to collect data in the standard data set.

Benchmark signal	Standard triggers	Parking triggers	Acceptance gain
VBF $H \rightarrow \text{SUEPs}$	H_T	VBF inclusive	+2500%
Fully hadronic VBS	Multijet, H_T	VBF + 2 central jets	+30%
VBF $H \rightarrow \text{invisible}$	p_T^{miss} , tight VBF+ p_T^{miss}	VBF + p_T^{miss}	+60%
VBF $H \rightarrow \tau\tau(\rightarrow \tau_h\tau_h)$	Di- τ_h , VBF+di- τ_h	VBF + τ_h	+50%
VBF $H \rightarrow \tau\tau(\rightarrow \mu\tau_h)$	Single- μ , $\mu + \tau_h$	VBF + μ	+30%
VBF $H \rightarrow \tau\tau(\rightarrow e\tau_h)$	Single- e , $e + \tau_h$	VBF + e	+40%
VBF $H \rightarrow \rho\gamma$	$\gamma + 2$ collimated tracks	VBF + γ	+25%




Figure 59: Distributions of m_{jj} for VBF $H \rightarrow$ invisible events passing the triggers used in the Run 2 analysis (blue), compared to events passing either one of the Run 2 triggers, the VBF+ p_T^{miss} parking trigger, or the VBF inclusive parking trigger implemented in Run 3 (black). The Run 2 trigger selection includes the VBF+ p_T^{miss} trigger algorithm introduced in Run 2, plus the standard trigger requiring $p_T^{\text{miss}} > 120 \text{ GeV}$ and $H_T^{\text{miss}} > 120 \text{ GeV}$. In all cases, loose, offline selections are applied to match the trigger-level requirements.

6.2.2 The HH parking

The Higgs boson self-coupling is a key parameter of the Higgs potential that remains unmeasured. It determines the strength of the double (HH) and triple (HHH) Higgs production at the LHC, which are sensitive probes of the electroweak symmetry breaking mechanism. The dominant decay mode of the Higgs boson is $H \rightarrow b\bar{b}$, leading to final states with multiple b jets in both HH and HHH searches. The most promising channels for HH observation are $\text{HH} \rightarrow 4b$, $\text{HH} \rightarrow 2b2\tau$ and $\text{HH} \rightarrow 2b2\gamma$. The photon triggers achieve high efficiency because of the calorimeter performance. The b jets and tau lepton triggers are subject to a larger contamination from jets faking their experimental signature in the detector. This effect is mitigated by applying an additional selection on the b tagging and τ lepton identification algorithms which results in lower efficiencies at the HLT.

In Run 2, the dedicated HLT path targeting the $\text{HH} \rightarrow 4b$ signal recorded events with multiple small-radius jets, of which at least three were identified as b jets with the online DEEPCSV [54] algorithm. The HLT path was seeded by L1 seeds that required H_T values of at least 280 GeV (2016) and 360 GeV (2018). At the HLT, the trigger required values of H_T above 340 GeV. The four most energetic small-radius jets were required to have p_T above 70, 50, 45, and 40 GeV, and at least three of these jets must pass a requirement on the DEEPCSV b jet discriminant.

In early 2022, a computationally lighter version of the PARTICLENET b tagging algorithm was deployed online [169]. The $\text{HH} \rightarrow 4b$ trigger was updated accordingly to benefit from the improved b tagging performance with respect to previous online algorithms. The switch to PARTICLENET allowed for more relaxed p_T thresholds for the jets (70, 50, 40, and 35 GeV) and a looser set of b tagging selection criteria, involving only the two most “b jet like” jets. At the L1, the H_T requirement remained 360 GeV. To increase the acceptance to the $\text{HH} \rightarrow 4b$ signal, in the low H_T region, the L1 H_T requirement was relaxed to 280 GeV in early 2023. Moreover, by exploiting the larger rate available via the data parking method, the trigger was updated to record events with at least four jets with p_T above 30 GeV and a looser threshold on the b tagging score, using the PARTICLENET online b tagging algorithm. In late 2023, this trigger achieved a rate of 180 Hz with $\mathcal{L}_{\text{inst}} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

The trigger efficiency is measured in simulated ggF $\text{HH} \rightarrow 4b$ and $\text{HH} \rightarrow 2b2\tau_h$ events and

defined as the number of events accepted by the signal triggers and an offline baseline selection, relative to the number of events satisfying the offline baseline selection alone. The offline selection requires the presence of at least four small-radius jets with p_T greater than 30 GeV and $|\eta| < 2.5$. The four most “b jet like” jets are used to reconstruct the HH pair. The trigger efficiency as a function of the reconstructed invariant mass of the two Higgs bosons is shown in Fig. 60.

For the $\text{HH} \rightarrow 4\text{b}$ events (left panel of Fig. 60), the 2023 HH trigger achieved an overall signal efficiency of 82%, corresponding to an improvement of about 60% (20%) with respect to the Run 2 (2022) trigger. Since the 2022 and 2023 di-Higgs triggers are designed to record events with at least two b jets in the final state, they are also suitable to record $\text{HH} \rightarrow 2\text{b}2\tau$ signal events (right panel of Fig. 60). For the given offline selection, the triggers requiring at least two τ_h candidates [170] achieve a signal efficiency of 34%, while the HH parking trigger results in an efficiency of 43%. By requiring events to satisfy either one of the two sets of triggers, the efficiency reaches 58%, demonstrating the complementarity of both sets in selecting signal events.




Figure 60: Trigger efficiency for selecting signal HH events, plotted as a function of the reconstructed invariant mass of the two Higgs bosons, as measured in simulated $\text{HH} \rightarrow 4\text{b}$ (left) and $\text{HH} \rightarrow 2\text{b}2\tau$ (right) samples corresponding to nominal Run 3 conditions.

6.2.3 The LLP parking

In 2023, new triggers were added to the parking stream to extend the physics reach in the search for exotic LLPs [171]. Particles with long lifetimes are often predicted in BSM theories and thus constitute interesting probes of new physics. Most conventional searches at the LHC target promptly decaying particles, and there are still vast regions of parameter space in the context of LLPs that remain unexplored. Thus, searches for new LLPs have a great potential for discovery.

The LLPs have distinct experimental signatures. They can decay far away from the pp interaction point, leaving decay products that are displaced from the PV. Specific examples of LLP signatures include displaced and delayed leptons [69, 172–174], photons [175], and jets [176–180]; disappearing tracks [181, 182]; and nonstandard tracks produced by monopoles or heavy stable charged particles [183]. Conventional trigger paths, object reconstruction algorithms, and background estimation strategies are usually inadequate for LLP searches because they are designed for promptly decaying particles, and custom techniques are needed for the online selection of interesting events and the offline analysis of the data.

The LLPs decaying to hadrons produce jets that contain tracks originating from a SV, spatially

displaced from the PV (displaced jets). In addition, since massive LLPs often have nonrelativistic velocities, the signal energy deposits in the calorimeters are expected to arrive late in time compared to those of relativistic background particles produced at the PV (delayed jets). In 2023, two sets of LLP triggers were added to the parking stream, targeting both displaced jet and delayed jet signatures.

A suite of displaced dijet paths, already available in the standard Run 2 data stream [176], underwent major improvements for the Run 3 data taking. The new displaced dijet paths in Run 3 require events passing two sets of selection criteria: either $H_T > 430\text{ GeV}$ plus two jets with zero prompt tracks, or $H_T > 240\text{ GeV}$ plus a soft muon with $p_T > 6\text{ GeV}$, plus two jets with zero prompt tracks and at least one displaced track. In all cases, the H_T is computed from calorimeter quantities. These new triggers used in Run 3 provide better efficiency in selecting low-mass LLPs, especially those that undergo heavy-flavor decays. In addition to the displaced dijet paths in the prompt reconstruction stream, several paths with lower H_T thresholds (down to 360 GeV) were added to the parking stream. The lower thresholds provide an increase in signal acceptance of 40–100% relative to prompt triggers for Higgs bosons that decay to long-lived scalars, with masses between 20 and 50 GeV and proper lifetimes $c\tau = 1\text{--}1000\text{ mm}$.

In addition to the displaced-jet triggers, new HLT paths were added in Run 3 that make use of ECAL timing. Two different types of delayed-jet triggers use the ECAL timing at the HLT: paths that are seeded by H_T , and paths that are seeded by L1 τ objects. Depending on the L1 seed, different requirements are placed at the HLT, e.g., one or two delayed jets, number of matched tracks in each jet, and the timing delay. In addition, for both types of seeds, some paths were added to the parking stream with a reduced timing delay. The latter increases the efficiency by a factor ranging from 30 to 800% for BSM delayed-jet signals, such as those produced by long-lived scalars that decay to four b jets or four τ leptons.

The combined trigger rate for the LLP parking triggers is 150 Hz with $\mathcal{L}_{\text{inst}} = 2 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$. The parked LLP triggers in Run 3 will play a crucial role in extending the sensitivity reach of displaced-jet and delayed-jets searches.

7 Summary

The extreme collision rate of the LHC, coupled with the data size needed to characterize complex interactions, poses a fundamental problem for collider experiments. Trigger, data acquisition, and downstream computing systems have finite resources that the experiments carefully allocate to the various parts of their physics programs. Searches for new physics at the energy frontier and measurements of electroweak-scale particles are one of the centerpieces of the CMS physics program, and therefore a significant portion of data acquisition resources is dedicated to triggers focusing on heavy-mass and high- p_T particles. There are, however, compelling reasons that new physics might manifest itself in other ways, for example in the existence of new light particles or, indirectly, as unexpected anomalies in precision measurements of benchmark standard model (SM) processes.

This review has highlighted two innovative techniques employed by CMS to extend the physics reach of the experiment beyond the one achieved with the original design of the detector and of the data processing pipeline: data parking and data scouting. After their inception in Run 1, both data scouting and data parking were expanded in Run 2, and became established and widely employed techniques in Run 3.

Data scouting records high-rate streams of data at the cost of a reduced event content utiliz-

ing physics objects such as jets, muons, and electrons reconstructed at the trigger level. The momentum and energy resolutions of these objects approach those achieved by the full offline event reconstruction, thereby facilitating searches for new physics in previously unexplored regions of phase space at the LHC. Notably, data scouting has enabled groundbreaking searches employing jet and muon objects that have already been published, while ongoing efforts involving electron and tau lepton objects further promise to extend the potential of this innovative technique.

The incorporation of jet objects obtained with data scouting has proven instrumental in pushing the boundaries of resonance searches in both dijet and multijet channels. These objects, which closely mirror nominal jets in energy resolution, have extended the reach of resonance searches to previously unexplored low-mass regions, overcoming limitations faced by standard analyses. Notably, the search for dijet resonances has been extended from the conventional lower limit of 1.5 TeV to 600 GeV. Moreover, multijet searches employing data scouting exhibit remarkable sensitivity down to masses as low as 70 GeV for both two- and three-parton decays, leveraging resolved and merged jets. The cross section sensitivity achieved in multijet analyses with data scouting surpasses previous searches by an order of magnitude, enabling probes of new physics sectors with electroweak couplings, such as Higgsinos, in fully hadronic final states. Furthermore, muon scouting has allowed the triggering of very low p_T dimuon events, including both prompt- and displaced-decay scenarios. These events have significantly extended the searches for new particles decaying to muon pairs, reaching the $2m_\mu$ kinematic limit. In the pursuit of long-lived dimuon resonances, the results are competitive with those from LHCb and B factories within certain mass and lifetime ranges. Additionally, decays to four muons have been successfully studied, culminating in the groundbreaking first-time observation of the rare decay of the η meson into four muons, $\eta \rightarrow \mu^+ \mu^- \mu^+ \mu^-$, with a measured branching fraction on the order of 10^{-9} . The result is in agreement with theoretical predictions and improves on the precision of previous upper-limit measurements by more than 5 orders of magnitude.

In the data-parking technique, data are recorded by high-rate triggers and temporarily stored (i.e., parked) in raw format until computing resources become available for full event reconstruction, for example during the end of the year or even during long shutdown periods. This contrasts with the typical reconstruction workflow, where data reconstruction begins within 48 hours after it is recorded. Throughout Run 1, different parking strategies were explored. Data collected with parking triggers during 2012 served the publication of impactful physics analyses including Higgs boson measurements, with a focus on vector boson fusion production, as well as precise measurements of various b quark hadron lifetimes in final states involving pairs of muons. Searches for beyond SM physics ranged from multijet searches to investigations into dark matter and supersymmetry in hadronic final states.

The 2018 data-parking analyses emphasized B physics. CMS collected a large data set comprising $\mathcal{O}(10^{10})$ $b\bar{b}$ events using a broad collection of single, displaced muon triggers. One distinctive feature of the parking data set is that different triggers with progressively lower thresholds were activated at different times during an LHC fill to maintain approximately constant (and high) L1 trigger rates, as the instantaneous luminosity decreased over the duration of the fill. With this method, we maximized the number of $b\bar{b}$ events recorded. At the start of Run 3, the parking strategy for B physics was expanded, notably including dedicated low- p_T triggers to collect dimuon and dielectron events. The single- and double-muon, and double-electron parked data sets enable CMS to perform a variety of B physics analyses of rare SM processes, as well as precision tests of lepton flavor universality. In addition, the B parking data enables several innovative searches beyond heavy flavor physics, such as for long-lived

heavy neutrinos.

Ongoing efforts in Run 3 have also seen the enhancement of alternative parking strategies inspired by the experience acquired in Runs 1 and 2. Dedicated parking triggers targeting vector boson fusion and double Higgs boson production have been designed to augment sensitivity to tests of the Higgs sector. Moreover, the implementation of distinct parking triggers dedicated to the exploration of long-lived particles decaying into jets, leptons, and photons offers the opportunity to substantially expand the CMS physics capabilities, extending its role at the forefront of high priority searches in the field.

These two techniques, scouting and parking, serve complementary purposes: data scouting is designed to accommodate searches akin to those conducted with standard triggers and data sets while overcoming the restrictions of stringent trigger thresholds, while data parking is particularly beneficial for analyses focused on optimal precision and accuracy but that also require a higher rate of data collection. Searches for new physics with low- p_T jets in the final state require scouting data streams. The parking data sets provide sensitivity to processes involving low- p_T single-muon and dilepton final states. Both scouting and parking data can be utilized to design comprehensive search and measurement strategies for low- p_T dimuon final states.

The comprehensive investigations conducted via scouting and parking analyses have significantly expanded the parameter space boundaries of CMS sensitivity, pushing beyond the anticipated limits of hadron colliders. The improvements achieved in trigger thresholds and data collection bandwidth, coupled with the implementation of pioneering methods for the reconstruction of essential entities such as jets, electrons, muons, tau leptons decaying hadronically, and photons, stand poised to elevate not only the quantity but also the caliber of data amassed with these sophisticated techniques. This breakthrough promises to deepen our understanding and unlock novel insights into the underlying physics, marking a significant advancement in the capabilities of the CMS experiment.

The importance of the data scouting and data parking approaches goes beyond enriching the current CMS physics program. The classic data acquisition model is not guaranteed to remain sustainable with the luminosity expected to be delivered by the LHC at the end of Run 3, and, crucially, during the high-luminosity LHC era. Moreover, computing and data acquisition resources will be designed to cope with peak luminosity demands, leaving significant spare computing power and data bandwidth at non-peak times. It is therefore essential to devise ways to develop real-time physics selection and analysis within the trigger and data acquisition systems themselves, and to optimize the utilization of idle computing resources during periods when the LHC is not operating at its maximum capacity. The CMS experience with data scouting and data parking in the last decade will prove decisive to tackling these challenges for particle physics experiments in the coming years.

Acknowledgments

We congratulate our colleagues in the CERN accelerator departments for the excellent performance of the LHC and thank the technical and administrative staffs at CERN and at other CMS institutes for their contributions to the success of the CMS effort. In addition, we gratefully acknowledge the computing centers and personnel of the Worldwide LHC Computing Grid and other centers for delivering so effectively the computing infrastructure essential to our analyses. Finally, we acknowledge the enduring support for the construction and operation of the LHC, the CMS detector, and the supporting computing infrastructure provided

by the following funding agencies: the Armenian Science Committee, project no. 22rl-037; the Austrian Federal Ministry of Education, Science and Research and the Austrian Science Fund; the Belgian Fonds de la Recherche Scientifique, and Fonds voor Wetenschappelijk Onderzoek; the Brazilian Funding Agencies (CNPq, CAPES, FAPERJ, FAPERGS, and FAPESP); the Bulgarian Ministry of Education and Science, and the Bulgarian National Science Fund; CERN; the Chinese Academy of Sciences, Ministry of Science and Technology, the National Natural Science Foundation of China, and Fundamental Research Funds for the Central Universities; the Ministerio de Ciencia Tecnología e Innovación (MINCIENCIAS), Colombia; the Croatian Ministry of Science, Education and Sport, and the Croatian Science Foundation; the Research and Innovation Foundation, Cyprus; the Secretariat for Higher Education, Science, Technology and Innovation, Ecuador; the Estonian Research Council via PRG780, PRG803, RVTT3 and the Ministry of Education and Research TK202; the Academy of Finland, Finnish Ministry of Education and Culture, and Helsinki Institute of Physics; the Institut National de Physique Nucléaire et de Physique des Particules / CNRS, and Commissariat à l'Énergie Atomique et aux Énergies Alternatives / CEA, France; the Shota Rustaveli National Science Foundation, Georgia; the Bundesministerium für Bildung und Forschung, the Deutsche Forschungsgemeinschaft (DFG), under Germany's Excellence Strategy – EXC 2121 "Quantum Universe" – 390833306, and under project number 400140256 - GRK2497, and Helmholtz-Gemeinschaft Deutscher Forschungszentren, Germany; the General Secretariat for Research and Innovation and the Hellenic Foundation for Research and Innovation (HFRI), Project Number 2288, Greece; the National Research, Development and Innovation Office (NKFIH), Hungary; the Department of Atomic Energy and the Department of Science and Technology, India; the Institute for Studies in Theoretical Physics and Mathematics, Iran; the Science Foundation, Ireland; the Istituto Nazionale di Fisica Nucleare, Italy; the Ministry of Science, ICT and Future Planning, and National Research Foundation (NRF), Republic of Korea; the Ministry of Education and Science of the Republic of Latvia; the Research Council of Lithuania, agreement No. VS-19 (LMTLT); the Ministry of Education, and University of Malaya (Malaysia); the Ministry of Science of Montenegro; the Mexican Funding Agencies (BUAP, CINVESTAV, CONACYT, LNS, SEP, and UASLP-FAI); the Ministry of Business, Innovation and Employment, New Zealand; the Pakistan Atomic Energy Commission; the Ministry of Education and Science and the National Science Center, Poland; the Fundação para a Ciência e a Tecnologia, grants CERN/FIS-PAR/0025/2019 and CERN/FIS-INS/0032/2019, Portugal; the Ministry of Education, Science and Technological Development of Serbia; MCIN/AEI/10.13039/501100011033, ERDF "a way of making Europe", Programa Estatal de Fomento de la Investigación Científica y Técnica de Excelencia María de Maeztu, grant MDM-2017-0765, projects PID2020-113705RB, PID2020-113304RB, PID2020-116262RB and PID2020-113341RB-I00, and Plan de Ciencia, Tecnología e Innovación de Asturias, Spain; the Ministry of Science, Technology and Research, Sri Lanka; the Swiss Funding Agencies (ETH Board, ETH Zurich, PSI, SNF, UniZH, Canton Zurich, and SER); the Ministry of Science and Technology, Taipei; the Ministry of Higher Education, Science, Research and Innovation, and the National Science and Technology Development Agency of Thailand; the Scientific and Technical Research Council of Turkey, and Turkish Energy, Nuclear and Mineral Research Agency; the National Academy of Sciences of Ukraine; the Science and Technology Facilities Council, UK; the US Department of Energy, and the US National Science Foundation.

Individuals have received support from the Marie-Curie programme and the European Research Council and Horizon 2020 Grant, contract Nos. 675440, 724704, 752730, 758316, 765710, 824093, 101115353, 101002207, and COST Action CA16108 (European Union) the Leventis Foundation; the Alfred P. Sloan Foundation; the Alexander von Humboldt Foundation; the Belgian Federal Science Policy Office; the Fonds pour la Formation à la Recherche dans l'Industrie et

dans l’Agriculture (FRIA-Belgium); the Agentschap voor Innovatie door Wetenschap en Technologie (IWT-Belgium); the F.R.S.-FNRS and FWO (Belgium) under the “Excellence of Science – EOS” – be.h project n. 30820817; the Beijing Municipal Science & Technology Commission, No. Z191100007219010; the Ministry of Education, Youth and Sports (MEYS) of the Czech Republic; the Shota Rustaveli National Science Foundation, grant FR-22-985 (Georgia); the Hungarian Academy of Sciences, the New National Excellence Program - ÚNKP, the NKFIH research grants K 131991, K 133046, K 138136, K 143460, K 143477, K 146913, K 146914, K 147048, 2020-2.2.1-ED-2021-00181, and TKP2021-NKTA-64 (Hungary); the Council of Scientific and Industrial Research, India; ICSC – National Research Center for High Performance Computing, Big Data and Quantum Computing, funded by the EU NexGeneration program, Italy; the Latvian Council of Science; the Ministry of Education and Science, project no. 2022/WK/14, and the National Science Center, contracts Opus 2021/41/B/ST2/01369 and 2021/43/B/ST2/01552 (Poland); the Fundação para a Ciência e a Tecnologia, grant FCT CEECIND/01334/2018; the National Priorities Research Program by Qatar National Research Fund; the Programa Estatal de Fomento de la Investigación Científica y Técnica de Excelencia María de Maeztu, grant MDM-2017-0765 and projects PID2020-113705RB, PID2020-113304RB, PID2020-116262RB and PID2020-113341RB-I00, and Programa Severo Ochoa del Principado de Asturias (Spain); the Chulalongkorn Academic into Its 2nd Century Project Advancement Project, and the National Science, Research and Innovation Fund via the Program Management Unit for Human Resources & Institutional Development, Research and Innovation, grant B37G660013 (Thailand); the Kavli Foundation; the Nvidia Corporation; the SuperMicro Corporation; the Welch Foundation, contract C-1845; and the Weston Havens Foundation (USA).

Glossary

AK	Anti- k_T
ALICE	A Large Ion Collider Experiment (experiment)
ATLAS	A Toroidal Lhc ApparatuS (experiment)
BDT	Boosted decision tree
BSM	Beyond standard model
CB	Crystal Ball
CHS	Charged hadron subtraction (algorithm)
CL	Confidence level
CMS	Compact Muon Solenoid (experiment)
CP	Charge-parity (symmetry)
CPU	Central processing unit
CR	Control region
DAQ	Data acquisition
DDT	Designing decorrelated taggers (procedure)
DNN	Deep neural network
DT	Deep tau (identification algorithm)
DY	Drell–Yan (process)
ECAL	Electromagnetic calorimeter
EFT	Effective field theory
EB	ECAL barrel section
EE	ECAL endcap section
EW	Electroweak
FPGA	Field programmable gate array
ggF	Gluon-gluon fusion (production process)
GPU	Graphics processing unit
GSF	Gaussian sum filter
GT	Global trigger
HCAL	Hadron calorimeter
HB	HCAL barrel section
HE	HCAL endcap section
HF	HCAL forward section
HO	HCAL outer section
HH	Higgs boson pair
HL-LHC	High-Luminosity Large Hadron Collider
HLT	High-level trigger
HPS	Hadrons-plus-strips (tau lepton reconstruction algorithm)
IO	Inside-out (method of matching muon tracks between the tracker and the muon detector)
IP	Interaction point
JER	Jet energy resolution
JES	Jet energy scale
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty (experiment)
L1	Level-1 trigger
L2	Level 2 muons (method of reconstructing muon candidates)
L3	Level 3 muons (method of reconstructing muon candidates)

LFU	Lepton flavor universality
LLP	Long-lived particle
LS1	Long shutdown 1, during the years 2013–2014
LS2	Long shutdown 2, during the years 2019–2021
LO	Leading order
LP	Low- p_T algorithm (for electron reconstruction)
MAHI	Minimization at HCAL, iteratively
MC	Monte Carlo (simulation)
ML	Machine learning
MSSM	Minimal supersymmetric standard model
MVA	Multi-variate analysis
NMSSM	Next-to-minimal supersymmetric standard model
NLO	Next-to-leading order
NN	Neural network
NWA	Narrow-width approximation
OI	Outside-in (method of matching muon tracks between the tracker and the muon detector)
PD	Primary data set
PDG	Particle Data Group
PF	Particle-flow (method of reconstructing particle candidates)
pp	Proton-proton
PU	Pileup
PUPPI	Pileup-per-particle identification (algorithm)
PV	Primary vertex
QCD	Quantum chromodynamics
QGD	Quark-gluon discriminator
R	distance in $(\Delta\eta, \Delta\phi)$ space
ROC	Receiver operating characteristic
RPV	R -parity violation (a supersymmetric model)
Run 1	The first run of the LHC, during the years 2010–2012
Run 2	The second run of the LHC, during the years 2015–2018
Run 3	The third run of the LHC, starting in 2022
SC	Supercluster (in the context of ECAL object reconstruction)
SD	Soft-drop (algorithm)
SM	Standard model
SR	Signal region
SUEP	Soft unclustered energy pattern
SUSY	Supersymmetry
SV	Secondary vertex
TP	Trigger primitive
TS	Time sample
VBF	Vector boson fusion (production process)
VH	Vector plus Higgs boson (production process or decay channel)
2HDM	Two-Higgs-doublet model
2HDM+S	Two-Higgs-doublet-plus-additional-singlet model

References

[1] CMS Collaboration, “A measurement of the Higgs boson mass in the diphoton decay channel”, *Phys. Lett. B* **805** (2020) 135425,

doi:10.1016/j.physletb.2020.135425, arXiv:2002.06398.

- [2] CMS Collaboration, “Measurements of production cross sections of the Higgs boson in the four-lepton final state in proton–proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **81** (2021) 488, doi:10.1140/epjc/s10052-021-09200-x, arXiv:2103.04956.
- [3] CMS Collaboration, “Measurement of the inclusive and differential Higgs boson production cross sections in the leptonic WW decay mode at $\sqrt{s} = 13$ TeV”, *JHEP* **03** (2021) 003, doi:10.1007/JHEP03(2021)003, arXiv:2007.01984.
- [4] CMS Collaboration, “A portrait of the Higgs boson by the CMS experiment ten years after the discovery.”, *Nature* **607** (2022) 60, doi:10.1038/s41586-022-04892-x, arXiv:2207.00043.
- [5] P. Agrawal et al., “Feebly-Interacting Particles: FIPs 2020 workshop report”, *Eur. Phys. J. C* **81** (2021) 1015, doi:10.1140/epjc/s10052-021-09703-7, arXiv:2102.12143.
- [6] C. Antel et al., “Feebly Interacting Particles: FIPs 2022 workshop report”, *Eur. Phys. J. C* **83** (2023) 1122, doi:10.1140/epjc/s10052-023-12168-5, arXiv:2305.01715.
- [7] O. S. Bruning et al., “LHC Design Report Vol.1: The LHC Main Ring”. CERN Yellow Reports: Monographs. CERN, Geneva, 2004. doi:10.5170/CERN-2004-003-v-1.
- [8] CMS Collaboration, “Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **15** (2020) P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165.
- [9] CMS Collaboration, “The CMS trigger system”, *JINST* **12** (2017) P01020, doi:10.1088/1748-0221/12/01/P01020, arXiv:1609.02366.
- [10] CMS Collaboration, “CMS Physics: Technical Design Report Volume 1: Detector Performance and Software”, Technical Report CERN-LHCC-2006-001, CMS-TDR-8-1, 2006.
- [11] CMS Collaboration, “CMS. The TriDAS project. Technical design report, vol. 1: The trigger systems”, Technical Report CERN-LHCC-2000-038, 2000.
- [12] CMS Collaboration, “CMS: The TriDAS project. Technical design report, Vol. 2: Data acquisition and high-level trigger”, Technical Report CERN-LHCC-2002-026, 2002.
- [13] G. Cerminara and B. van Besien, “Automated workflows for critical time-dependent calibrations at the CMS experiment”, *J. Phys. Conf. Ser.* **664** (2015) 072009, doi:10.1088/1742-6596/664/7/072009.
- [14] CMS Collaboration, “Data parking and data scouting at the CMS Experiment”, CMS Detector Performance Note CMS-DP-2012-022, 2012.
- [15] CMS Collaboration, “Search for narrow resonances in dijet final states at $\sqrt{s} = 8$ TeV with the novel CMS technique of data scouting”, *Phys. Rev. Lett.* **117** (2016) 031802, doi:10.1103/PhysRevLett.117.031802, arXiv:1604.08907.
- [16] S. Benson, V. V. Gligorov, M. A. Vesterinen, and M. Williams, “The LHCb Turbo Stream”, *J. Phys. Conf. Ser.* **664** (2015) 082004, doi:10.1088/1742-6596/664/8/082004.

- [17] ATLAS Collaboration, “Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **121** (2018) 081801, doi:10.1103/PhysRevLett.121.081801, arXiv:1804.03496.
- [18] ATLAS Collaboration, “Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector”, *Phys. Rev. D* **91** (2015) 052007, doi:10.1103/PhysRevD.91.052007, arXiv:1407.1376.
- [19] ALICE Collaboration, “Upgrade of the ALICE readout & trigger system”, Technical Design Report CERN-LHCC-2013-019, ALICE-TDR-015, 2013.
- [20] Kvapil, Jakub et al., “ALICE central trigger system for LHC Run 3”, *EPJ Web Conf.* **251** (2021) 04022, doi:10.1051/epjconf/202125104022, arXiv:2106.08353.
- [21] CMS Collaboration, “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965.
- [22] G. Petrucciani, A. Rizzi, and C. Vuosalo, “Mini-AOD: A new analysis data format for CMS”, *J. Phys. Conf. Ser.* **664** (2015) 072052, doi:10.1088/1742-6596/664/7/072052, arXiv:1702.04685.
- [23] M. Peruzzi, G. Petrucciani, and A. Rizzi, “The NanoAOD event data format in CMS”, *J. Phys. Conf. Ser.* **1525** (2020) 012038, doi:10.1088/1742-6596/1525/1/012038.
- [24] D. London and J. Matias, “B flavor anomalies: 2021 theoretical status report”, *Annu. Rev. Nucl. Part. Sci.* **72** (2022) 37, doi:10.1146/annurev-nucl-102020-090209, arXiv:2110.13270.
- [25] G. Hiller and F. Kruger, “More model-independent analysis of $b \rightarrow s$ processes”, *Phys. Rev. D* **69** (2004) 074020, doi:10.1103/PhysRevD.69.074020, arXiv:hep-ph/0310219.
- [26] M. Bordone, G. Isidori, and A. Patti, “On the standard model predictions for R_K and R_{K^*} ”, *Eur. Phys. J. C* **76** (2016) 440, doi:10.1140/epjc/s10052-016-4274-7, arXiv:1605.07633.
- [27] G. Isidori, S. Nabeboccus, and R. Zwicky, “QED corrections in $\bar{B} \rightarrow \bar{K}\ell^+\ell^-$ at the double-differential level”, *JHEP* **12** (2020) 104, doi:10.1007/JHEP12(2020)104, arXiv:2009.00929.
- [28] G. Isidori, D. Lancierini, S. Nabeboccus, and R. Zwicky, “QED in $\bar{B} \rightarrow \bar{K}\ell^+\ell^-$ LFU ratios: theory versus experiment, a Monte Carlo study”, *JHEP* **10** (2022) 146, doi:10.1007/JHEP10(2022)146, arXiv:2205.08635.
- [29] CMS Tracker Collaboration, “The CMS Phase-1 pixel detector upgrade”, *JINST* **16** (2021) P02027, doi:10.1088/1748-0221/16/02/P02027, arXiv:2012.14304.
- [30] CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [31] CMS Collaboration, “Development of the CMS detector for the CERN LHC Run 3”, 2023. arXiv:2309.05466. Submitted to JINST.

[32] G. Badaro et al., “The Phase-2 upgrade of the CMS data acquisition”, *EPJ Web Conf.* **251** (2021) 04023, doi:10.1051/epjconf/202125104023.

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

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

[35] CMS Collaboration, “Technical proposal for the Phase-II upgrade of the Compact Muon Solenoid”, CMS Technical Proposal CERN-LHCC-2015-010, CMS-TDR-15-02, 2015.

[36] A. Bocci et al., “Heterogeneous reconstruction of tracks and primary vertices with the CMS pixel tracker”, *Front. Big Data* **3** (2020) 601728, doi:10.3389/fdata.2020.601728, arXiv:2008.13461.

[37] CMS Collaboration, “Reconstruction of signal amplitudes in the CMS electromagnetic calorimeter in the presence of overlapping proton-proton interactions”, *JINST* **15** (2020) P10002, doi:10.1088/1748-0221/15/10/P10002, arXiv:2006.14359.

[38] CMS Collaboration, “CMS-EGM-18-002”, (2024). To be submitted for publication.

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

[40] CMS Collaboration, “Performance of the local reconstruction algorithms for the CMS hadron calorimeter with Run 2 data”, *JINST* **18** (2023) P11017, doi:10.1088/1748-0221/18/11/P11017, arXiv:2306.10355.

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

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

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

[44] CMS Collaboration, “Pileup mitigation at CMS in 13 TeV data”, *JINST* **15** (2020) P09018, doi:10.1088/1748-0221/15/09/p09018, arXiv:2003.00503.

[45] D. Bertolini, P. Harris, M. Low, and N. Tran, “Pileup per particle identification”, *JHEP* **10** (2014) 059, doi:10.1007/JHEP10(2014)059, arXiv:1407.6013.

[46] CMS Collaboration, “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017) P02014, doi:10.1088/1748-0221/12/02/P02014, arXiv:1607.03663.

[47] CMS Collaboration, “Jet algorithms performance in 13 TeV data”, CMS Physics Analysis Summary CMS-PAS-JME-16-003, 2017.

[48] CMS Collaboration, “Identification techniques for highly boosted W bosons that decay into hadrons”, *JHEP* **12** (2014) 017, doi:10.1007/JHEP12(2014)017, arXiv:1410.4227.

[49] D. Krohn, J. Thaler, and L.-T. Wang, “Jet trimming”, *JHEP* **02** (2010) 084, doi:10.1007/JHEP02(2010)084, arXiv:0912.1342.

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

[51] J. Thaler and K. Van Tilburg, “Identifying boosted objects with N -subjettiness”, *JHEP* **03** (2011) 015, doi:10.1007/JHEP03(2011)015, arXiv:1011.2268.

[52] I. Moult, L. Necib, and J. Thaler, “New angles on energy correlation functions”, *JHEP* **12** (2016) 153, doi:10.1007/JHEP12(2016)153, arXiv:1609.07483.

[53] CMS Collaboration, “Identification of b-quark jets with the CMS experiment”, *JINST* **8** (2013) P04013, doi:10.1088/1748-0221/8/04/P04013, arXiv:1211.4462.

[54] CMS Collaboration, “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”, *JINST* **13** (2018) P05011, doi:10.1088/1748-0221/13/05/p05011, arXiv:1712.07158.

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

[56] CMS Collaboration, “Performance of the DeepJet b tagging algorithm using 41.9/fb of data from proton-proton collisions at 13 TeV with Phase 1 CMS detector”, CMS Detector Performance Note CMS-DP-2018-058, 2018.

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

[58] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Methods Phys. Res. A* **262** (1987) 444, doi:10.1016/0168-9002(87)90887-4.

[59] CMS Collaboration, “Performance of the CMS muon trigger system in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JINST* **16** (2021) P07001, doi:10.1088/1748-0221/16/07/P07001, arXiv:2102.04790.

[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: Nucl. Part. Phys.* **31** (2005) N9, doi:10.1088/0954-3899/31/9/n01, arXiv:physics/0306087.

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

[62] 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** (2018) P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816.

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

[64] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas”, *Phys. Lett. B* **659** (2008) 119, doi:10.1016/j.physletb.2007.09.077, arXiv:0707.1378.

[65] CMS Collaboration, “Search for narrow and broad dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles”, *JHEP* **08** (2018) 130, doi:10.1007/jhep08(2018)130, arXiv:1806.00843.

[66] M. Zaheer et al., “Deep sets”, in *Advances in Neural Information Processing Systems*, I. Guyon et al., eds., volume 30, p. 3392. Curran Associates, Inc., 2017.

[67] CMS Collaboration, “Search for pair-produced multijet resonances using data scouting”, 2023. EXO-21-004. To be submitted for publication.

[68] CMS Collaboration, “Search for a narrow resonance lighter than 200 GeV decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **124** (2020) 131802, doi:10.1103/PhysRevLett.124.131802, arXiv:1912.04776.

[69] CMS Collaboration, “Search for long-lived particles decaying into muon pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV collected with a dedicated high-rate data stream”, *JHEP* **04** (2022) 062, doi:10.1007/JHEP04(2022)062, arXiv:2112.13769.

[70] B. P. Roe et al., “Boosted decision trees as an alternative to artificial neural networks for particle identification”, *Nucl. Instrum. Methods Phys. Res. A* **543** (2005) 577, doi:10.1016/j.nima.2004.12.018, arXiv:physics/0408124.

[71] CMS Collaboration, “Search for direct production of GeV-scale resonances decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **12** (2023) 070, doi:10.1007/JHEP12(2023)070, arXiv:2309.16003.

[72] M. J. Oreglia, “A study of the reactions $\psi' \rightarrow \gamma\gamma\psi$ ”. PhD thesis, Stanford University, 1980. SLAC Report SLAC-R-236, see Appendix D.

[73] J. E. Gaiser, “Charmonium spectroscopy from radiative decays of the J/ψ and ψ' ”. PhD thesis, Stanford University, 1982. SLAC Report SLAC-R-255.

[74] R. M. Harris and K. Kousouris, “Searches for dijet resonances at hadron colliders”, *Int. J. Mod. Phys. A* **26** (2011) 5005, doi:10.1142/S0217751X11054905, arXiv:1110.5302.

[75] B. A. Dobrescu and F. Yu, “Coupling-Mass Mapping of Dijet Peak Searches”, *Phys. Rev. D* **88** (2013) 035021, doi:10.1103/PhysRevD.88.035021, arXiv:1306.2629. [Erratum: Phys. Rev. D 90 (2014) 079901].

[76] 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** (2017) 520, doi:10.1016/j.physletb.2017.02.012, arXiv:1611.03568. [Erratum: doi:10.1016/j.physletb.2017.09.029].

[77] CMS Collaboration, “Search for dijet resonances using events with three jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **805** (2020) 135448, doi:10.1016/j.physletb.2020.135448, arXiv:1911.03761.

[78] CMS Collaboration, “Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **01** (2018) 097, doi:10.1007/JHEP01(2018)097, arXiv:1710.00159.

[79] CMS Collaboration, “Search for low-mass quark-antiquark resonances produced in association with a photon at $\sqrt{s} = 13$ TeV”, *Phys. Rev. Lett.* **123** (2019) 231803, doi:10.1103/PhysRevLett.123.231803, arXiv:1905.10331.

[80] R. Barbier et al., “R-parity violating supersymmetry”, *Phys. Rept.* **420** (2005) 1, doi:10.1016/j.physrep.2005.08.006, arXiv:hep-ph/0406039.

[81] J. Dolen et al., “Thinking outside the ROCs: Designing decorrelated taggers (DDT) for jet substructure”, *JHEP* **05** (2016) 156, doi:10.1007/jhep05(2016)156, arXiv:1603.00027.

[82] CDF Collaboration, “First search for multijet resonances in $\sqrt{s} = 1.96$ TeV $p\bar{p}$ collisions”, *Phys. Rev. Lett.* **107** (2011) 042001, doi:10.1103/PhysRevLett.107.042001, arXiv:1105.2815.

[83] ATLAS Collaboration, “Search for massive supersymmetric particles decaying to many jets using the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV”, *Phys. Rev. D* **91** (2015) 112016, doi:10.1103/PhysRevD.91.112016, arXiv:1502.05686. [Erratum: doi:10.1103/PhysRevD.93.039901].

[84] ATLAS Collaboration, “Search for R-parity-violating supersymmetric particles in multi-jet final states produced in p-p collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC”, *Phys. Lett. B* **785** (2018) 136, doi:10.1016/j.physletb.2018.08.021, arXiv:1804.03568.

[85] CMS Collaboration, “Search for three-jet resonances in pp collisions at $\sqrt{s} = 7$ TeV”, *Phys. Lett. B* **718** (2012) 329, doi:10.1016/j.physletb.2012.10.048, arXiv:1208.2931.

[86] CMS Collaboration, “Searches for light- and heavy-flavour three-jet resonances in pp collisions at $\sqrt{s} = 8$ TeV”, *Phys. Lett. B* **730** (2014) 193, doi:10.1016/j.physletb.2014.01.049, arXiv:1311.1799.

[87] CMS Collaboration, “Search for pair-produced three-jet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **99** (2019) 012010, doi:10.1103/PhysRevD.99.012010, arXiv:1810.10092.

[88] CMS Collaboration, “Search for pair-produced resonances decaying to quark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **98** (2018) 112014, doi:10.1103/PhysRevD.98.112014, arXiv:1808.03124.

[89] CMS Collaboration, “Search for resonant and nonresonant production of pairs of dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **07** (2023) 161, doi:10.1007/JHEP07(2023)161, arXiv:2206.09997.

[90] LHCb Collaboration, “Search for $A' \rightarrow \mu^+ \mu^-$ decays”, *Phys. Rev. Lett.* **124** (2020) 041801, doi:10.1103/PhysRevLett.124.041801, arXiv:1910.06926.

[91] D. Curtin, R. Essig, S. Gori, and J. Shelton, “Illuminating dark photons with high-energy colliders”, *JHEP* **02** (2015) 157, doi:10.1007/JHEP02(2015)157, arXiv:1412.0018.

[92] LHCb Collaboration, “Searches for low-mass dimuon resonances”, *JHEP* **10** (2020) 156, doi:10.1007/JHEP10(2020)156, arXiv:2007.03923.

[93] BaBar Collaboration, “Search for di-muon decays of a low-mass Higgs boson in radiative decays of the $\Upsilon(1S)$ ”, *Phys. Rev. D* **87** (2013) 031102, doi:10.1103/PhysRevD.87.031102, arXiv:1210.0287. [Erratum: doi:10.1103/PhysRevD.87.059903].

[94] F. Bezrukov and D. Gorbunov, “Light inflaton after LHC8 and WMAP9 results”, *JHEP* **07** (2013) 140, doi:10.1007/JHEP07(2013)140, arXiv:1303.4395.

[95] CMS Collaboration, “Observation of the rare decay of the η meson to four muons”, *Phys. Rev. Lett.* **131** (2023) 091903, doi:10.1103/PhysRevLett.131.091903, arXiv:2305.04904.

[96] R. Escribano and S. González-Solís, “A data-driven approach to π^0 , η and η' single and double Dalitz decays”, *Chinese Phys. C* **42** (2018) 023109, doi:10.1088/1674-1137/42/2/023109, arXiv:1511.04916.

[97] Particle Data Group, “Review of particle physics”, *Prog. Theor. Exp. Phys.* **2022** (2022) 083C01, doi:10.1093/ptep/ptac097.

[98] CMS Collaboration, “PF jet performances at high level trigger using Patatrack pixel tracks”, CMS Detector Performance Note CMS-DP-2021-005, 2021.

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

[100] CMS Collaboration, “CMS-HIG-21-020”, (2024). To be submitted for publication.

[101] 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** (2023) 041801, doi:10.1103/PhysRevLett.131.041801, arXiv:2211.14181.

[102] H. Qu, C. Li, and S. Qian, “Particle Transformer for Jet Tagging”, 2022. arXiv:2202.03772.

[103] CMS Collaboration, “Search for $W\gamma$ resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV using hadronic decays of Lorentz-boosted W bosons”, *Phys. Lett. B* **826** (2022) 136888, doi:10.1016/j.physletb.2022.136888, arXiv:2106.10509.

[104] CMS Collaboration, “Searches for invisible decays of the Higgs boson in pp collisions at $\sqrt{s} = 7, 8$, and 13 TeV”, *JHEP* **02** (2017) 135, doi:10.1007/JHEP02(2017)135, arXiv:1610.09218.

[105] CMS Collaboration, “Search for the standard model Higgs boson produced through vector boson fusion and decaying to $b\bar{b}$ ”, *Phys. Rev. D* **92** (2015) 032008, doi:10.1103/PhysRevD.92.032008, arXiv:1506.01010.

[106] CMS Collaboration, “Evidence for the 125 GeV Higgs boson decaying to a pair of τ leptons”, *JHEP* **05** (2014) 104, doi:10.1007/JHEP05(2014)104, arXiv:1401.5041.

[107] CMS Collaboration, “Search for neutral MSSM Higgs bosons decaying to a pair of tau leptons in pp collisions”, *JHEP* **10** (2014) 160, doi:10.1007/JHEP10(2014)160, arXiv:1408.3316.

[108] CMS Collaboration, “Measurement of b hadron lifetimes in pp collisions at $\sqrt{s} = 8$ TeV”, *Eur. Phys. J. C* **78** (2018) 457, doi:10.1140/epjc/s10052-018-5929-3, arXiv:1710.08949. [Erratum: doi:10.1140/epjc/s10052-018-6014-7].

[109] CMS Collaboration, “Angular analysis of the decay $B^0 \rightarrow K^{*0}\mu^+\mu^-$ from pp collisions at $\sqrt{s} = 8$ TeV”, *Phys. Lett. B* **753** (2016) 424, doi:10.1016/j.physletb.2015.12.020, arXiv:1507.08126.

[110] CMS Collaboration, “Measurement of angular parameters from the decay $B^0 \rightarrow K^{*0}\mu^+\mu^-$ in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *Phys. Lett. B* **781** (2018) 517, doi:10.1016/j.physletb.2018.04.030, arXiv:1710.02846.

[111] CMS Collaboration, “Angular analysis of the decay $B^+ \rightarrow K^+\mu^+\mu^-$ in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *Phys. Rev. D* **98** (2018) 112011, doi:10.1103/PhysRevD.98.112011, arXiv:1806.00636.

[112] CMS Collaboration, “Angular analysis of the decay $B^+ \rightarrow K^*(892)^+\mu^+\mu^-$ in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *JHEP* **04** (2021) 124, doi:10.1007/JHEP04(2021)124, arXiv:2010.13968.

[113] CMS Collaboration, “Search for pair-produced resonances decaying to jet pairs in proton-proton collisions at $\sqrt{s} = 8$ TeV”, *Phys. Lett. B* **747** (2015) 98, doi:10.1016/j.physletb.2015.04.045, arXiv:1412.7706.

[114] CMS Collaboration, “Search for exotic decays of a Higgs boson into undetectable particles and one or more photons”, *Phys. Lett. B* **753** (2016) 363, doi:10.1016/j.physletb.2015.12.017, arXiv:1507.00359.

[115] C. Rogan, “Kinematical variables towards new dynamics at the LHC”, 2010. arXiv:1006.2727.

[116] CMS Collaboration, “Inclusive search for squarks and gluinos in pp collisions at $\sqrt{s} = 7$ TeV”, *Phys. Rev. D* **85** (2012) 012004, doi:10.1103/PhysRevD.85.012004, arXiv:1107.1279.

[117] CMS Collaboration, “Search for dark matter particles in proton-proton collisions at $\sqrt{s} = 8$ TeV using the razor variables”, *JHEP* **12** (2016) 088, doi:10.1007/JHEP12(2016)088, arXiv:1603.08914.

[118] CMS Collaboration, “Search for top squark pair production in compressed-mass-spectrum scenarios in proton-proton collisions at $\sqrt{s} = 8$ TeV using the α_T variable”, *Phys. Lett. B* **767** (2017) 403, doi:10.1016/j.physletb.2017.02.007, arXiv:1605.08993.

[119] BaBar Collaboration, “The BaBar detector”, *Nucl. Instrum. Methods Phys. Res. A* **479** (2002) 1, doi:10.1016/S0168-9002(01)02012-5, arXiv:hep-ex/0105044.

[120] Belle Collaboration, “The Belle Detector”, *Nucl. Instrum. Methods Phys. Res. A* **479** (2002) 117, doi:10.1016/S0168-9002(01)02013-7.

[121] LHCb Collaboration, “The LHCb Detector at the LHC”, *JINST* **3** (2008) S08005, doi:10.1088/1748-0221/3/08/S08005.

[122] ATLAS Collaboration, “The ATLAS experiment at the CERN Large Hadron Collider”, *JINST* **3** (2008) S08003, doi:10.1088/1748-0221/3/08/S08003.

[123] CMS Collaboration, “Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ decay properties and search for the $B^0 \rightarrow \mu^+ \mu^-$ decay in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **842** (2023) 137955, doi:10.1016/j.physletb.2023.137955, arXiv:2212.10311.

[124] LHCb Collaboration, “Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ decay properties and search for the $B^0 \rightarrow \mu^+ \mu^-$ and $B_s^0 \rightarrow \mu^+ \mu^- \gamma$ decays”, *Phys. Rev. D* **105** (2022) 012010, doi:10.1103/PhysRevD.105.012010, arXiv:2108.09283.

[125] CMS Collaboration, “Measurement of properties of $B_s^0 \rightarrow \mu^+ \mu^-$ decays and search for $B^0 \rightarrow \mu^+ \mu^-$ with the CMS experiment”, *JHEP* **04** (2020) 188, doi:10.1007/JHEP04(2020)188, arXiv:1910.12127.

[126] ATLAS Collaboration, “Study of the rare decays of B_s^0 and B^0 mesons into muon pairs using data collected during 2015 and 2016 with the ATLAS detector”, *JHEP* **04** (2019) 098, doi:10.1007/JHEP04(2019)098, arXiv:1812.03017.

[127] CMS, LHCb Collaboration, “Observation of the rare $B_s^0 \rightarrow \mu^+ \mu^-$ decay from the combined analysis of CMS and LHCb data”, *Nature* **522** (2015) 68, doi:10.1038/nature14474, arXiv:1411.4413.

[128] S. Descotes-Genon, J. Matias, M. Ramon, and J. Virto, “Implications from clean observables for the binned analysis of $B \rightarrow K^* \mu^+ \mu^-$ at large recoil”, *JHEP* **01** (2013) 048, doi:10.1007/JHEP01(2013)048, arXiv:1207.2753.

[129] LHCb Collaboration, “Measurement of form-factor-independent observables in the decay $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ ”, *Phys. Rev. Lett.* **111** (2013) 191801, doi:10.1103/PhysRevLett.111.191801, arXiv:1308.1707.

[130] LHCb Collaboration, “Angular analysis of the $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decay using 3 fb^{-1} of integrated luminosity”, *JHEP* **02** (2016) 104, doi:10.1007/JHEP02(2016)104, arXiv:1512.04442.

[131] LHCb Collaboration, “Measurement of CP -averaged observables in the $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decay”, *Phys. Rev. Lett.* **125** (2020) 011802, doi:10.1103/PhysRevLett.125.011802, arXiv:2003.04831.

[132] BaBar Collaboration, “Evidence for an excess of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ decays”, *Phys. Rev. Lett.* **109** (2012) 101802, doi:10.1103/PhysRevLett.109.101802, arXiv:1205.5442.

[133] BaBar Collaboration, “Measurement of an excess of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ decays and implications for charged Higgs bosons”, *Phys. Rev. D* **88** (2013) 072012, doi:10.1103/PhysRevD.88.072012, arXiv:1303.0571.

[134] Belle Collaboration, “Measurement of the branching ratio of $\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau$ relative to $\bar{B} \rightarrow D^{(*)} \ell^- \bar{\nu}_\ell$ decays with hadronic tagging at Belle”, *Phys. Rev. D* **92** (2015) 072014, doi:10.1103/PhysRevD.92.072014, arXiv:1507.03233.

[135] LHCb Collaboration, “Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \tau^- \bar{\nu}_\tau) / \mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \mu^- \bar{\nu}_\mu)$ ”, *Phys. Rev. Lett.* **115** (2015) 111803, doi:10.1103/PhysRevLett.115.111803, arXiv:1506.08614. [Erratum: doi:10.1103/PhysRevLett.115.159901].

[136] LHCb Collaboration, “Test of lepton universality using $B^+ \rightarrow K^+ \ell^+ \ell^-$ decays”, *Phys. Rev. Lett.* **113** (2014) 151601, doi:10.1103/PhysRevLett.113.151601, arXiv:1406.6482.

[137] LHCb Collaboration, “Test of lepton universality with $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays”, *JHEP* **08** (2017) 055, doi:10.1007/JHEP08(2017)055, arXiv:1705.05802.

[138] LHCb Collaboration, “Search for lepton-universality violation in $B^+ \rightarrow K^+ \ell^+ \ell^-$ decays”, *Phys. Rev. Lett.* **122** (2019) 191801, doi:10.1103/PhysRevLett.122.191801, arXiv:1903.09252.

[139] BELLE Collaboration, “Test of lepton flavor universality and search for lepton flavor violation in $B \rightarrow K \ell \ell$ decays”, *JHEP* **03** (2021) 105, doi:10.1007/JHEP03(2021)105, arXiv:1908.01848.

[140] Belle Collaboration, “Test of lepton-flavor universality in $B \rightarrow K^* \ell^+ \ell^-$ decays at Belle”, *Phys. Rev. Lett.* (2021) 161801, doi:10.1103/PhysRevLett.126.161801, arXiv:1904.02440.

[141] LHCb Collaboration, “Test of lepton universality in beauty-quark decays”, *Nature Phys.* **18** (2022) 277, doi:10.1038/s41567-021-01478-8, arXiv:2103.11769. [Addendum: doi:10.1038/s41567-023-02095-3].

[142] LHCb Collaboration, “Measurement of the ratios of branching fractions $\mathcal{R}(D^*)$ and $\mathcal{R}(D^0)$ ”, *Phys. Rev. Lett.* **131** (2023) 111802, doi:10.1103/PhysRevLett.131.111802, arXiv:2302.02886.

[143] LHCb Collaboration, “Test of lepton universality in $b \rightarrow s \ell^+ \ell^-$ decays”, *Phys. Rev. Lett.* **131** (2023) 051803, doi:10.1103/PhysRevLett.131.051803, arXiv:2212.09152.

[144] LHCb Collaboration, “Measurement of lepton universality parameters in $B^+ \rightarrow K^+ \ell^+ \ell^-$ and $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays”, *Phys. Rev. D* **108** (2023) 032002, doi:10.1103/PhysRevD.108.032002, arXiv:2212.09153.

[145] M. Cacciari, M. Greco, and P. Nason, “The p_T spectrum in heavy-flavour hadroproduction.”, *JHEP* **05** (1998) 007, doi:10.1088/1126-6708/1998/05/007, arXiv:hep-ph/9803400.

[146] M. Cacciari, S. Frixione, and P. Nason, “The p_T spectrum in heavy flavor photoproduction”, *JHEP* **03** (2001) 006, doi:10.1088/1126-6708/2001/03/006, arXiv:hep-ph/0102134.

[147] M. Cacciari et al., “Theoretical predictions for charm and bottom production at the LHC”, *JHEP* **10** (2012) 137, doi:10.1007/JHEP10(2012)137, arXiv:1205.6344.

[148] M. Cacciari, M. L. Mangano, and P. Nason, “Gluon PDF constraints from the ratio of forward heavy-quark production at the LHC at $\sqrt{s} = 7$ and 13 TeV”, *Eur. Phys. J. C* **75** (2015) 610, doi:10.1140/epjc/s10052-015-3814-x, arXiv:1507.06197.

[149] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system", in *Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*. 2016.
`doi:10.1145/2939672.2939785`.

[150] CMS Collaboration, "Measurement of inclusive W and Z boson production cross sections in pp collisions at $\sqrt{s} = 8$ TeV", *Phys. Rev. Lett.* **112** (2014) 191802,
`doi:10.1103/PhysRevLett.112.191802`, [arXiv:1402.0923](https://arxiv.org/abs/1402.0923).

[151] M. Ciuchini et al., "Charming penguins and lepton universality violation in $b \rightarrow s\ell^+\ell^-$ decays", *Eur. Phys. J. C* **83** (2023) 64, `doi:10.1140/epjc/s10052-023-11191-w`,
[arXiv:2110.10126](https://arxiv.org/abs/2110.10126).

[152] CMS Collaboration, "Test of lepton flavor universality in $B^\pm \rightarrow K^\pm \mu^+ \mu^-$ and $B^\pm \rightarrow K^\pm e^+ e^-$ decays in proton-proton collisions at $\sqrt{s} = 13$ TeV", 2024. [arXiv:2401.07090](https://arxiv.org/abs/2401.07090). Submitted to *Rep. Prog. Phys.*

[153] K. Prokofiev and T. Speer, "A kinematic and a decay chain reconstruction library", in *14th International Conference on Computing in High-Energy and Nuclear Physics*, p. 411. 2005. `doi:10.5170/CERN-2005-002.411`.

[154] S. Bilenky, "Neutrino oscillations: From a historical perspective to the present status", *Nucl. Phys. B* **908** (2016) 2, `doi:10.1016/j.nuclphysb.2016.01.025`,
[arXiv:1602.00170](https://arxiv.org/abs/1602.00170).

[155] G. Bertone, ed., "Particle Dark Matter: Observations, Models and Searches". Cambridge Univ. Press, Cambridge, 2010. `doi:10.1017/CBO9780511770739`, ISBN 978-1-107-65392-4.

[156] G. R. Farrar and M. E. Shaposhnikov, "Baryon asymmetry of the universe in the standard model", *Phys. Rev. D* **50** (1994) 774, `doi:10.1103/PhysRevD.50.774`.

[157] CMS Collaboration, "CMS-EXO-23-006", (2024). To be submitted for publication.

[158] CMS Collaboration, "Search for long-lived heavy neutrinos in the decays of B mesons produced in proton-proton collisions at $\sqrt{s} = 13$ TeV", (2023). [arXiv:2403.04584](https://arxiv.org/abs/2403.04584). Submitted to *JHEP*.

[159] CMS Collaboration, "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* **343** (2019) 139, `doi:10.22323/1.343.0139`.

[160] LHCb Collaboration, "Measurement of the time-dependent CP asymmetries in $B_s^0 \rightarrow J/\psi K_S^0$ ", *JHEP* **06** (2015) 131, `doi:10.1007/JHEP06(2015)131`,
[arXiv:1503.07055](https://arxiv.org/abs/1503.07055).

[161] CMS Collaboration, "Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector", *Phys. Lett. B* **779** (2018) 283,
`doi:10.1016/j.physletb.2018.02.004`, [arXiv:1708.00373](https://arxiv.org/abs/1708.00373).

[162] CMS Collaboration, "Search for invisible decays of the Higgs boson produced via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV", *Phys. Rev. D* **105** (2022) 092007, `doi:10.1103/PhysRevD.105.092007`, [arXiv:2201.11585](https://arxiv.org/abs/2201.11585).

[163] CMS Collaboration, "Evidence for Higgs boson decay to a pair of muons", *JHEP* **01** (2021) 148, `doi:10.1007/JHEP01(2021)148`, [arXiv:2009.04363](https://arxiv.org/abs/2009.04363).

[164] SMEFiT Collaboration, “Combined SMEFT interpretation of Higgs, diboson, and top quark data from the LHC”, *JHEP* **11** (2021) 089, doi:10.1007/JHEP11(2021)089, arXiv:2105.00006.

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

[166] CMS Collaboration, “Level 1 tau trigger performance in 2016 data and VBF seeds at Level 1 trigger”, CMS Detector Performance Note CMS-DP-2017-022, 2017.

[167] CMS Collaboration, “Search for dark photons in Higgs boson production via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **03** (2021) 011, doi:10.1007/JHEP03(2021)011, arXiv:2009.14009.

[168] J. Barron et al., “Unsupervised hadronic SUEP at the LHC”, *JHEP* **12** (2021) 129, doi:10.1007/JHEP12(2021)129, arXiv:2107.12379.

[169] CMS Collaboration, “Performance of the ParticleNet tagger on small and large-radius jets at high level trigger in Run 3”, CMS Detector Performance Note CMS-DP-2023-021, 2023.

[170] CMS Collaboration, “Performance of tau lepton reconstruction at high level trigger using 2022 data from the CMS experiment at CERN”, CMS Detector Performance Note CMS-DP-2023-024, 2023.

[171] CMS Collaboration, “Performance of long lived particle triggers in Run 3”, CMS Detector Performance Note CMS-DP-2023-043, 2023.

[172] CMS Collaboration, “Search for long-lived particles decaying to leptons with large impact parameter in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **82** (2022) 153, doi:10.1140/epjc/s10052-022-10027-3, arXiv:2110.04809.

[173] CMS Collaboration, “Search for long-lived particles decaying to a pair of muons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **05** (2023) 228, doi:10.1007/JHEP05(2023)228, arXiv:2205.08582.

[174] CMS Collaboration, “Search for long-lived particles decaying to final states with a pair of muons in proton-proton collisions at $\sqrt{s} = 13.6$ TeV”, 2024. arXiv:2402.14491. Submitted to JHEP.

[175] CMS Collaboration, “Search for long-lived particles using delayed photons in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **100** (2019) 112003, doi:10.1103/PhysRevD.100.112003, arXiv:1909.06166.

[176] CMS Collaboration, “Search for long-lived particles using displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **104** (2021) 012015, doi:10.1103/PhysRevD.104.012015, arXiv:2012.01581.

[177] CMS Collaboration, “Search for long-lived particles decaying to jets with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **104** (2021) 052011, doi:10.1103/PhysRevD.104.052011, arXiv:2104.13474.

- [178] CMS Collaboration, “Search for long-lived particles using nonprompt jets and missing transverse momentum with proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **797** (2019) 134876, doi:10.1016/j.physletb.2019.134876, arXiv:1906.06441.
- [179] CMS Collaboration, “Search for long-lived particles using out-of-time trackless jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *JHEP* **07** (2023) 210, doi:10.1007/JHEP07(2023)210, arXiv:2212.06695.
- [180] CMS Collaboration, “Search for long-lived particles using displaced vertices and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”, 2024. arXiv:2402.15804. Submitted to PRD.
- [181] CMS Collaboration, “Search for supersymmetry in final states with disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV”, 2023. arXiv:2309.16823. Submitted to PRD.
- [182] CMS Collaboration, “Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Lett. B* **806** (2020) 135502, doi:10.1016/j.physletb.2020.135502, arXiv:2004.05153.
- [183] CMS Collaboration, “Search for long-lived charged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **94** (2016) 112004, doi:10.1103/PhysRevD.94.112004, arXiv:1609.08382.

8 The CMS Collaboration

Yerevan Physics Institute, Yerevan, Armenia

A. Hayrapetyan, A. Tumasyan¹ 

Institut für Hochenergiephysik, Vienna, Austria

W. Adam , J.W. Andrejkovic, T. Bergauer , S. Chatterjee , K. Damanakis , M. Dragicevic , P.S. Hussain , M. Jeitler² , N. Krammer , A. Li , D. Liko , I. Mikulec , J. Schieck² , R. Schöfbeck , D. Schwarz , M. Sonawane , S. Templ , W. Waltenberger , C.-E. Wulz²

Universiteit Antwerpen, Antwerpen, Belgium

M.R. Darwish³ , T. Janssen , P. Van Mechelen 

Vrije Universiteit Brussel, Brussel, Belgium

N. Breugelmans, J. D'Hondt , S. Dansana , A. De Moor , M. Delcourt , F. Heyen, S. Lowette , I. Makarenko , D. Müller , S. Tavernier , M. Tytgat⁴ , G.P. Van Onsem , S. Van Putte , D. Vannerom

Université Libre de Bruxelles, Bruxelles, Belgium

B. Clerbaux , A.K. Das, G. De Lentdecker , H. Evard , L. Favart , P. Gianneios , D. Hohov , J. Jaramillo , A. Khalilzadeh, F.A. Khan , K. Lee , M. Mahdavikhorrami , A. Malara , S. Paredes , M.A. Shahzad, L. Thomas , M. Vanden Bemden , C. Vander Velde , P. Vanlaer

Ghent University, Ghent, Belgium

M. De Coen , D. Dobur , G. Gokbulut , Y. Hong , J. Knolle , L. Lambrecht , D. Marckx , G. Mestdach, K. Mota Amarilo , A. Samalan, K. Skovpen , N. Van Den Bossche , J. van der Linden , L. Wezenbeek

Université Catholique de Louvain, Louvain-la-Neuve, Belgium

A. Benecke , A. Bethani , G. Bruno , C. Caputo , J. De Favereau De Jeneret , C. Delaere , I.S. Donertas , A. Giammanco , A.O. Guzel , Sa. Jain , V. Lemaitre, J. Lidrych , P. Mastrapasqua , T.T. Tran , S. Wertz

Centro Brasileiro de Pesquisas Fisicas, Rio de Janeiro, Brazil

G.A. Alves , M. Alves Gallo Pereira , E. Coelho , G. Correia Silva , C. Hensel , T. Menezes De Oliveira , A. Moraes , P. Rebello Teles , M. Soeiro, A. Vilela Pereira⁵ 

Universidade do Estado do Rio de Janeiro, Rio de Janeiro, Brazil

W.L. Aldá Júnior , M. Barroso Ferreira Filho , H. Brandao Malbouisson , W. Carvalho , J. Chinellato⁶, E.M. Da Costa , G.G. Da Silveira⁷ , D. De Jesus Damiao , S. Fonseca De Souza , R. Gomes De Souza, M. Macedo , J. Martins⁸ , C. Mora Herrera , L. Mundim , H. Nogima , J.P. Pinheiro , A. Santoro , A. Sznajder , M. Thiel

Universidade Estadual Paulista, Universidade Federal do ABC, São Paulo, Brazil

C.A. Bernardes⁷ , L. Calligaris , T.R. Fernandez Perez Tomei , E.M. Gregores , I. Maietto Silverio , P.G. Mercadante , S.F. Novaes , B. Orzari , Sandra S. Padula 

Institute for Nuclear Research and Nuclear Energy, Bulgarian Academy of Sciences, Sofia, Bulgaria

A. Aleksandrov , G. Antchev , R. Hadjiiska , P. Iaydjiev , M. Misheva , M. Shopova , G. Sultanov 

University of Sofia, Sofia, Bulgaria

A. Dimitrov , L. Litov , B. Pavlov , P. Petkov , A. Petrov , E. Shumka 

Instituto De Alta Investigación, Universidad de Tarapacá, Casilla 7 D, Arica, Chile
S. Keshri , S. Thakur 

Beihang University, Beijing, China

T. Cheng , T. Javaid , L. Yuan 

Department of Physics, Tsinghua University, Beijing, China

Z. Hu , Z. Liang, J. Liu, K. Yi^{9,10} 

Institute of High Energy Physics, Beijing, China

G.M. Chen¹¹ , H.S. Chen¹¹ , M. Chen¹¹ , F. Iemmi , C.H. Jiang, A. Kapoor¹² , H. Liao , Z.-A. Liu¹³ , R. Sharma¹⁴ , J.N. Song¹³, J. Tao , C. Wang¹¹, J. Wang , Z. Wang¹¹, H. Zhang , J. Zhao 

State Key Laboratory of Nuclear Physics and Technology, Peking University, Beijing, China

A. Agapitos , Y. Ban , S. Deng , B. Guo, C. Jiang , A. Levin , C. Li , Q. Li , Y. Mao, S. Qian, S.J. Qian , X. Qin, X. Sun , D. Wang , H. Yang, L. Zhang , Y. Zhao, C. Zhou 

Guangdong Provincial Key Laboratory of Nuclear Science and Guangdong-Hong Kong Joint Laboratory of Quantum Matter, South China Normal University, Guangzhou, China

S. Yang 

Sun Yat-Sen University, Guangzhou, China

Z. You 

University of Science and Technology of China, Hefei, China

K. Jaffel , N. Lu 

Nanjing Normal University, Nanjing, China

G. Bauer¹⁵, B. Li, J. Zhang 

Institute of Modern Physics and Key Laboratory of Nuclear Physics and Ion-beam Application (MOE) - Fudan University, Shanghai, China

X. Gao¹⁶ 

Zhejiang University, Hangzhou, Zhejiang, China

Z. Lin , C. Lu , M. Xiao 

Universidad de Los Andes, Bogota, Colombia

C. Avila , D.A. Barbosa Trujillo, A. Cabrera , C. Florez , J. Fraga , J.A. Reyes Vega

Universidad de Antioquia, Medellin, Colombia

F. Ramirez , C. Rendón, M. Rodriguez , A.A. Ruales Barbosa , J.D. Ruiz Alvarez 

University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, Split, Croatia

D. Giljanovic , N. Godinovic , D. Lelas , A. Sculac 

University of Split, Faculty of Science, Split, Croatia

M. Kovac , A. Petkovic, T. Sculac 

Institute Rudjer Boskovic, Zagreb, Croatia

P. Bargassa , V. Brigljevic , B.K. Chitroda , D. Ferencek , K. Jakovcic, S. Mishra , A. Starodumov¹⁷ , T. Susa 

University of Cyprus, Nicosia, Cyprus

A. Attikis , K. Christoforou , A. Hadjiagapiou, C. Leonidou, J. Mousa , C. Nicolaou, L. Paizanos, F. Ptochos , P.A. Razis , H. Rykaczewski, H. Saka , A. Stepennov 

Charles University, Prague, Czech Republic

M. Finger , M. Finger Jr. , A. Kveton 

Universidad San Francisco de Quito, Quito, Ecuador

E. Carrera Jarrin 

Academy of Scientific Research and Technology of the Arab Republic of Egypt, Egyptian Network of High Energy Physics, Cairo, Egypt

Y. Assran^{18,19}, B. El-mahdy, S. Elgammal¹⁹

Center for High Energy Physics (CHEP-FU), Fayoum University, El-Fayoum, Egypt

M. Abdullah Al-Mashad , M.A. Mahmoud 

National Institute of Chemical Physics and Biophysics, Tallinn, Estonia

K. Ehataht , M. Kadastik, T. Lange , S. Nandan , C. Nielsen , J. Pata , M. Raidal , L. Tani , C. Veelken 

Department of Physics, University of Helsinki, Helsinki, Finland

H. Kirschenmann , K. Osterberg , M. Voutilainen 

Helsinki Institute of Physics, Helsinki, Finland

S. Barthuar , N. Bin Norjoharuddeen , E. Brücke , F. Garcia , P. Inkaew , K.T.S. Kallonen , T. Lampén , K. Lassila-Perini , S. Lehti , T. Lindén , L. Martikainen , M. Myllymäki , M.m. Rantanen , H. Siikonen , J. Tuominiemi 

Lappeenranta-Lahti University of Technology, Lappeenranta, Finland

P. Luukka , H. Petrow 

IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France

M. Besancon , F. Couderc , M. Dejardin , D. Denegri, J.L. Faure, F. Ferri , S. Ganjour , P. Gras , G. Hamel de Monchenault , V. Lohezic , J. Malcles , F. Orlandi , L. Portales , A. Rosowsky , M.Ö. Sahin , A. Savoy-Navarro²⁰ , P. Simkina , M. Titov , M. Tornago 

Laboratoire Leprince-Ringuet, CNRS/IN2P3, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

F. Beaudette , P. Busson , A. Cappati , C. Charlot , M. Chiusi , F. Damas , O. Davignon , A. De Wit , I.T. Ehle , B.A. Fontana Santos Alves , S. Ghosh , A. Gilbert , R. Granier de Cassagnac , A. Hakimi , B. Harikrishnan , L. Kalipoliti , G. Liu , M. Nguyen , C. Ochando , R. Salerno , J.B. Sauvan , Y. Sirois , L. Urda Gómez , E. Vernazza , A. Zabi , A. Zghiche 

Université de Strasbourg, CNRS, IPHC UMR 7178, Strasbourg, France

J.-L. Agram²¹ , J. Andrea , D. Apparu , D. Bloch , J.-M. Brom , E.C. Chabert , C. Collard , S. Falke , U. Goerlach , R. Haeberle , A.-C. Le Bihan , M. Meena , O. Poncet , G. Saha , M.A. Sessini , P. Van Hove , P. Vauclerc 

Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules, CNRS/IN2P3, Villeurbanne, France

A. Di Florio 

Institut de Physique des 2 Infinis de Lyon (IP2I), Villeurbanne, France

D. Amram, S. Beauceron , B. Blançon , G. Boudoul , N. Chanon , D. Contardo , P. Depasse , C. Dozen²² , H. El Mamouni, J. Fay , S. Gascon , M. Gouzevitch 

C. Greenberg, G. Grenier , B. Ille , E. Jourd'huy, I.B. Laktineh, M. Lethuillier , L. Mirabito, S. Perries, A. Purohit , M. Vander Donckt , P. Verdier , J. Xiao 

Georgian Technical University, Tbilisi, Georgia

I. Lomidze , T. Toriashvili²³ , Z. Tsamalaidze¹⁷ 

RWTH Aachen University, I. Physikalisches Institut, Aachen, Germany

V. Botta , L. Feld , K. Klein , M. Lipinski , D. Meuser , A. Pauls , D. Pérez Adán , N. Röwert , M. Teroerde 

RWTH Aachen University, III. Physikalisches Institut A, Aachen, Germany

S. Diekmann , A. Dodonova , N. Eich , D. Eliseev , F. Engelke , J. Erdmann , M. Erdmann , P. Fackeldey , B. Fischer , T. Hebbeker , K. Hoepfner , F. Ivone , A. Jung , M.y. Lee , F. Mausolf , M. Merschmeyer , A. Meyer , S. Mukherjee , D. Noll , F. Nowotny, A. Pozdnyakov , Y. Rath, W. Redjeb , F. Rehm, H. Reithler , V. Sarkisovi , A. Schmidt , A. Sharma , J.L. Spah , A. Stein , F. Torres Da Silva De Araujo²⁴ , S. Wiedenbeck , S. Zaleski

RWTH Aachen University, III. Physikalisches Institut B, Aachen, Germany

C. Dzwikow , G. Flügge , T. Kress , A. Nowack , O. Pooth , A. Stahl , T. Ziemons , A. Zottz 

Deutsches Elektronen-Synchrotron, Hamburg, Germany

H. Aarup Petersen , M. Aldaya Martin , J. Alimena , S. Amoroso, Y. An , J. Bach , S. Baxter , M. Bayatmakou , H. Becerril Gonzalez , O. Behnke , A. Belvedere , S. Bhattacharya , F. Blekman²⁵ , K. Borras²⁶ , A. Campbell , A. Cardini , C. Cheng, F. Colombina , S. Consuegra Rodríguez , M. De Silva , G. Eckerlin, D. Eckstein , L.I. Estevez Banos , O. Filatov , E. Gallo²⁵ , A. Geiser , V. Guglielmi , M. Guthoff , A. Hinzmann , L. Jeppe , B. Kaech , M. Kasemann , C. Kleinwort , R. Kogler , M. Komm , D. Krücker , W. Lange, D. Leyva Pernia , K. Lipka²⁷ , W. Lohmann²⁸ , F. Lorkowski , R. Mankel , I.-A. Melzer-Pellmann , M. Mendizabal Morentin , A.B. Meyer , G. Milella , K. Moral Figueroa , A. Mussgiller , L.P. Nair , J. Niedziela , A. Nürnberg , Y. Otarid , J. Park , E. Ranken , A. Raspereza , D. Rastorguev , J. Rübenach, L. Rygaard, A. Saggio , M. Scham^{29,26} , S. Schnake²⁶ , P. Schütze , C. Schwanenberger²⁵ , D. Selivanova , K. Sharko , M. Shchedrolosiev , D. Stafford, F. Vazzoler , A. Ventura Barroso , R. Walsh , D. Wang , Q. Wang , Y. Wen , K. Wichmann, L. Wiens²⁶ , C. Wissing , Y. Yang , A. Zimermann Castro Santos 

University of Hamburg, Hamburg, Germany

A. Albrecht , S. Albrecht , M. Antonello , S. Bein , L. Benato , S. Bollweg, M. Bonanomi , P. Connor , K. El Morabit , Y. Fischer , E. Garutti , A. Grohsjean , J. Haller , H.R. Jabusch , G. Kasieczka , P. Keicher, R. Klanner , W. Korcari , T. Kramer , C.c. Kuo, V. Kutzner , F. Labe , J. Lange , A. Lobanov , C. Matthies , L. Moureaux , M. Mrowietz, A. Nigamova , Y. Nissan, A. Paasch , K.J. Pena Rodriguez , T. Quadfasel , B. Raciti , M. Rieger , D. Savoiu , J. Schindler , P. Schleper , M. Schröder , J. Schwandt , M. Sommerhalder , H. Stadie , G. Steinbrück , A. Tews, M. Wolf 

Karlsruher Institut fuer Technologie, Karlsruhe, Germany

S. Brommer , M. Burkart, E. Butz , T. Chwalek , A. Dierlamm , A. Droll, N. Faltermann , M. Giffels , A. Gottmann , F. Hartmann³⁰ , R. Hofsaess , M. Horzela , U. Husemann , J. Kieseler , M. Klute , R. Koppenhöfer , J.M. Lawhorn , M. Link, A. Lintuluoto , B. Maier , S. Maier , S. Mitra , M. Mormile , Th. Müller , M. Neukum,

M. Oh [id](#), E. Pfeffer [id](#), M. Presilla [id](#), G. Quast [id](#), K. Rabbertz [id](#), B. Regnery [id](#), N. Shadskiy [id](#), I. Shvetsov [id](#), H.J. Simonis [id](#), L. Sowa, L. Stockmeier, K. Tauqueer, M. Toms [id](#), N. Trevisani [id](#), R.F. Von Cube [id](#), M. Wassmer [id](#), S. Wieland [id](#), F. Wittig, R. Wolf [id](#), X. Zuo [id](#)

Institute of Nuclear and Particle Physics (INPP), NCSR Demokritos, Aghia Paraskevi, Greece

G. Anagnostou, G. Daskalakis [id](#), A. Kyriakis, A. Papadopoulos³⁰, A. Stakia [id](#)

National and Kapodistrian University of Athens, Athens, Greece

P. Kontaxakis [id](#), G. Melachroinos, Z. Painesis [id](#), I. Papavergou [id](#), I. Paraskevas [id](#), N. Saoulidou [id](#), K. Theofilatos [id](#), E. Tziaferi [id](#), K. Vellidis [id](#), I. Zisopoulos [id](#)

National Technical University of Athens, Athens, Greece

G. Bakas [id](#), T. Chatzistavrou, G. Karapostoli [id](#), K. Kousouris [id](#), I. Papakrivopoulos [id](#), E. Siamarkou, G. Tsipolitis, A. Zacharopoulou

University of Ioánnina, Ioánnina, Greece

K. Adamidis, I. Bestintzanos, I. Evangelou [id](#), C. Foudas, C. Kamtsikis, P. Katsoulis, P. Kokkas [id](#), P.G. Kosmoglou Kioseoglou [id](#), N. Manthos [id](#), I. Papadopoulos [id](#), J. Strologas [id](#)

HUN-REN Wigner Research Centre for Physics, Budapest, Hungary

C. Hajdu [id](#), D. Horvath^{31,32} [id](#), K. Márton, A.J. Rádl³³ [id](#), F. Sikler [id](#), V. Veszpremi [id](#)

MTA-ELTE Lendület CMS Particle and Nuclear Physics Group, Eötvös Loránd University, Budapest, Hungary

M. Csanad [id](#), K. Farkas [id](#), A. Feherkuti³⁴ [id](#), M.M.A. Gadallah³⁵ [id](#), . Kadlecsik [id](#), P. Major [id](#), G. Pasztor [id](#), G.I. Veres [id](#)

Faculty of Informatics, University of Debrecen, Debrecen, Hungary

B. Ujvari [id](#), G. Zilizi [id](#)

Institute of Nuclear Research ATOMKI, Debrecen, Hungary

G. Bencze, S. Czellar, J. Molnar, Z. Szillasi

Karoly Robert Campus, MATE Institute of Technology, Gyongyos, Hungary

T. Csorgo³⁴ [id](#), T. Novak [id](#)

Panjab University, Chandigarh, India

J. Babbar [id](#), S. Bansal [id](#), S.B. Beri, V. Bhatnagar [id](#), G. Chaudhary [id](#), S. Chauhan [id](#), N. Dhingra³⁶ [id](#), A. Kaur [id](#), A. Kaur [id](#), H. Kaur [id](#), M. Kaur [id](#), S. Kumar [id](#), K. Sandeep [id](#), T. Sheokand, J.B. Singh [id](#), A. Singla [id](#)

University of Delhi, Delhi, India

A. Ahmed [id](#), A. Bhardwaj [id](#), A. Chhetri [id](#), B.C. Choudhary [id](#), A. Kumar [id](#), A. Kumar [id](#), M. Naimuddin [id](#), K. Ranjan [id](#), M.K. Saini, S. Saumya [id](#)

Saha Institute of Nuclear Physics, HBNI, Kolkata, India

S. Baradia [id](#), S. Barman³⁷ [id](#), S. Bhattacharya [id](#), S. Das Gupta, S. Dutta [id](#), S. Dutta, S. Sarkar

Indian Institute of Technology Madras, Madras, India

M.M. Ameen [id](#), P.K. Behera [id](#), S.C. Behera [id](#), S. Chatterjee [id](#), G. Dash [id](#), P. Jana [id](#), P. Kalbhor [id](#), S. Kamble [id](#), J.R. Komaragiri³⁸ [id](#), D. Kumar³⁸ [id](#), P.R. Pujahari [id](#), N.R. Saha [id](#), A. Sharma [id](#), A.K. Sikdar [id](#), R.K. Singh, P. Verma, S. Verma [id](#), A. Vijay

Tata Institute of Fundamental Research-A, Mumbai, India

S. Dugad, M. Kumar [id](#), G.B. Mohanty [id](#), B. Parida [id](#), M. Shelake, P. Suryadevara

Tata Institute of Fundamental Research-B, Mumbai, India

A. Bala , S. Banerjee , R.M. Chatterjee, M. Guchait , Sh. Jain , A. Jaiswal, S. Kumar , G. Majumder , K. Mazumdar , S. Parolia , A. Thachayath

National Institute of Science Education and Research, An OCC of Homi Bhabha National Institute, Bhubaneswar, Odisha, India

S. Bahinipati³⁹ , C. Kar , D. Maity⁴⁰ , P. Mal , T. Mishra , V.K. Muraleedharan Nair Bindhu⁴⁰ , K. Naskar⁴⁰ , A. Nayak⁴⁰ , S. Nayak, K. Pal, P. Sadangi, S.K. Swain , S. Varghese⁴⁰ , D. Vats⁴⁰

Indian Institute of Science Education and Research (IISER), Pune, India

S. Acharya⁴¹ , A. Alpana , S. Dube , B. Gomber⁴¹ , P. Hazarika , B. Kansal , A. Laha , B. Sahu⁴¹ , S. Sharma , K.Y. Vaish

Isfahan University of Technology, Isfahan, Iran

H. Bakhshiansohi⁴² , A. Jafari⁴³ , M. Zeinali⁴⁴ 

Institute for Research in Fundamental Sciences (IPM), Tehran, Iran

S. Bashiri, S. Chenarani⁴⁵ , S.M. Etesami , Y. Hosseini , M. Khakzad , E. Khazaie⁴⁶ , M. Mohammadi Najafabadi , S. Tizchang 

University College Dublin, Dublin, Ireland

M. Felcini , M. Grunewald 

INFN Sezione di Bari^a, Università di Bari^b, Politecnico di Bari^c, Bari, Italy

M. Abbrescia^{a,b} , A. Colaleo^{a,b} , D. Creanza^{a,c} , B. D'Anzi^{a,b} , N. De Filippis^{a,c} , M. De Palma^{a,b} , L. Fiore^a , G. Iaselli^{a,c} , L. Longo^a , M. Louka^{a,b} , G. Maggi^{a,c} , M. Maggi^a , I. Margjeka^a , V. Mastrapasqua^{a,b} , S. My^{a,b} , S. Nuzzo^{a,b} , A. Pellecchia^{a,b} , A. Pompili^{a,b} , G. Pugliese^{a,c} , R. Radogna^{a,b} , D. Ramos^a , A. Ranieri^a , L. Silvestris^a , F.M. Simone^{a,c} , Ü. Sözbilir^a , A. Stamerra^{a,b} , D. Troiano^{a,b} , R. Venditti^{a,b} , P. Verwilligen^a , A. Zaza^{a,b}

INFN Sezione di Bologna^a, Università di Bologna^b, Bologna, Italy

G. Abbiendi^a , C. Battilana^{a,b} , D. Bonacorsi^{a,b} , L. Borgonovi^a , P. Capiluppi^{a,b} , A. Castro^{a,b} , F.R. Cavallo^a , M. Cuffiani^{a,b} , G.M. Dallavalle^a , T. Diotalevi^{a,b} , F. Fabbri^a , A. Fanfani^{a,b} , D. Fasanella^a , P. Giacomelli^a , L. Giommi^{a,b} , C. Grandi^a , L. Guiducci^{a,b} , S. Lo Meo^{a,47} , M. Lorusso^{a,b} , L. Lunerti^a , S. Marcellini^a , G. Masetti^a , F.L. Navarria^{a,b} , G. Paggi^{a,b} , A. Perrotta^a , F. Primavera^{a,b} , A.M. Rossi^{a,b} , S. Rossi Tisbeni^{a,b} , T. Rovelli^{a,b} , G.P. Siroli^{a,b}

INFN Sezione di Catania^a, Università di Catania^b, Catania, Italy

S. Costa^{a,b,48} , A. Di Mattia^a , A. Lapertosa^a , R. Potenza^{a,b} , A. Tricomi^{a,b,48} , C. Tuve^{a,b} 

INFN Sezione di Firenze^a, Università di Firenze^b, Firenze, Italy

P. Assiouras^a , G. Barbagli^a , G. Bardelli^{a,b} , B. Camaiani^{a,b} , A. Cassese^a , R. Ceccarelli^a , V. Ciulli^{a,b} , C. Civinini^a , R. D'Alessandro^{a,b} , E. Focardi^{a,b} , T. Kello^a , G. Latino^{a,b} , P. Lenzi^{a,b} , M. Lizzo^a , M. Meschini^a , S. Paoletti^a , A. Papanastassiou^{a,b} , G. Sguazzoni^a , L. Viliani^a

INFN Laboratori Nazionali di Frascati, Frascati, Italy

L. Benussi , S. Bianco , S. Meola⁴⁹ , D. Piccolo 

INFN Sezione di Genova^a, Università di Genova^b, Genova, Italy

P. Chatagnon^a , F. Ferro^a , E. Robutti^a , S. Tosi^{a,b} 

INFN Sezione di Milano-Bicocca^a, Università di Milano-Bicocca^b, Milano, Italy

A. Benaglia^a , G. Boldrini^{a,b} , F. Brivio^a , F. Cetorelli^{a,b} , F. De Guio^{a,b} , M.E. Dinardo^{a,b} , P. Dini^a , S. Gennai^a , R. Gerosa^{a,b} , A. Ghezzi^{a,b} , P. Govoni^{a,b} , L. Guzzi^a , M.T. Lucchini^{a,b} , M. Malberti^a , S. Malvezzi^a , A. Massironi^a , D. Menasce^a , L. Moroni^a , M. Paganoni^{a,b} , S. Palluotto^{a,b} , D. Pedrini^a , A. Perego^{a,b} , B.S. Pinolini^a, G. Pizzati^{a,b} , S. Ragazzi^{a,b} , T. Tabarelli de Fatis^{a,b} 

INFN Sezione di Napoli^a, Università di Napoli 'Federico II'^b, Napoli, Italy; Università della Basilicata^c, Potenza, Italy; Scuola Superiore Meridionale (SSM)^d, Napoli, Italy

S. Buontempo^a , A. Cagnotta^{a,b} , F. Carnevali^{a,b} , N. Cavallo^{a,c} , F. Fabozzi^{a,c} , A.O.M. Iorio^{a,b} , L. Lista^{a,b,50} , P. Paolucci^{a,30} , B. Rossi^a 

INFN Sezione di Padova^a, Università di Padova^b, Padova, Italy; Università di Trento^c, Trento, Italy

R. Ardino^a , P. Azzi^a , N. Bacchetta^{a,51} , D. Bisello^{a,b} , P. Bortignon^a , G. Bortolato^{a,b} , A. Bragagnolo^{a,b} , A.C.M. Bulla^a , R. Carlin^{a,b} , P. Checchia^a , T. Dorigo^a , F. Gasparini^{a,b} , U. Gasparini^{a,b} , E. Lusiani^a , M. Margoni^{a,b} , G. Maron^{a,52} , M. Migliorini^{a,b} , J. Pazzini^{a,b} , P. Ronchese^{a,b} , R. Rossin^{a,b} , F. Simonetto^{a,b} , G. Strong^a , M. Tosi^{a,b} , A. Triossi^{a,b} , S. Ventura^a , M. Zanetti^{a,b} , P. Zotto^{a,b} , A. Zucchetta^{a,b} , G. Zumerle^{a,b} 

INFN Sezione di Pavia^a, Università di Pavia^b, Pavia, Italy

C. Aimè^a , A. Braghieri^a , S. Calzaferri^a , D. Fiorina^a , P. Montagna^{a,b} , V. Re^a , C. Riccardi^{a,b} , P. Salvini^a , I. Vai^{a,b} , P. Vitulo^{a,b} 

INFN Sezione di Perugia^a, Università di Perugia^b, Perugia, Italy

S. Ajmal^{a,b} , M.E. Ascioti^{a,b} , G.M. Bilei^a , C. Carrivale^{a,b} , D. Ciangottini^{a,b} , L. Fanò^{a,b} , M. Magherini^{a,b} , V. Mariani^{a,b} , M. Menichelli^a , F. Moscatelli^{a,53} , A. Rossi^{a,b} , A. Santocchia^{a,b} , D. Spiga^a , T. Tedeschi^{a,b} 

INFN Sezione di Pisa^a, Università di Pisa^b, Scuola Normale Superiore di Pisa^c, Pisa, Italy; Università di Siena^d, Siena, Italy

C.A. Alexe^{a,c} , P. Asenov^{a,b} , P. Azzurri^a , G. Bagliesi^a , R. Bhattacharya^a , L. Bianchini^{a,b} , T. Boccali^a , E. Bossini^a , D. Bruschini^{a,c} , R. Castaldi^a , M.A. Ciocci^{a,b} , M. Cipriani^{a,b} , V. D'Amante^{a,d} , R. Dell'Orso^a , S. Donato^a , A. Giassi^a , F. Ligabue^{a,c} , D. Matos Figueiredo^a , A. Messineo^{a,b} , M. Musich^{a,b} , F. Palla^a , A. Rizzi^{a,b} , G. Rolandi^{a,c} , S. Roy Chowdhury^a , T. Sarkar^a , A. Scribano^a , P. Spagnolo^a , R. Tenchini^a , G. Tonelli^{a,b} , N. Turini^{a,d} , F. Vaselli^{a,c} , A. Venturi^a , P.G. Verdini^a 

INFN Sezione di Roma^a, Sapienza Università di Roma^b, Roma, Italy

C. Baldenegro Barrera^{a,b} , P. Barria^a , C. Basile^{a,b} , M. Campana^{a,b} , F. Cavallari^a , L. Cunqueiro Mendez^{a,b} , D. Del Re^{a,b} , E. Di Marco^a , M. Diemoz^a , F. Errico^{a,b} , E. Longo^{a,b} , J. Mijuskovic^{a,b} , G. Organtini^{a,b} , F. Pandolfi^a , R. Paramatti^{a,b} , C. Quaranta^{a,b} , S. Rahatlou^{a,b} , C. Rovelli^a , F. Santanastasio^{a,b} , L. Soffi^a 

INFN Sezione di Torino^a, Università di Torino^b, Torino, Italy; Università del Piemonte Orientale^c, Novara, Italy

N. Amapane^{a,b} , R. Arcidiacono^{a,c} , S. Argiro^{a,b} , M. Arneodo^{a,c} , N. Bartosik^a , R. Bellan^{a,b} , A. Bellora^{a,b} , C. Biino^a , C. Borca^{a,b} , N. Cartiglia^a , M. Costa^{a,b} , R. Covarelli^{a,b} , N. Demaria^a , L. Finco^a , M. Grippo^{a,b} , B. Kiani^{a,b} , F. Legger^a , F. Luongo^{a,b} , C. Mariotti^a , L. Markovic^{a,b} , S. Maselli^a , A. Mecca^{a,b} , L. Menzio^{a,b} , P. Meridiani^a , E. Migliore^{a,b} , M. Monteno^a , R. Mulargia^a , M.M. Obertino^{a,b} 

G. Ortona^a , L. Pacher^{a,b} , N. Pastrone^a , M. Pelliccioni^a , M. Ruspa^{a,c} , F. Siviero^{a,b} , V. Sola^{a,b} , A. Solano^{a,b} , A. Staiano^a , C. Tarricone^{a,b} , D. Trocino^a , G. Umoret^{a,b} , R. White^{a,b}

INFN Sezione di Trieste^a, Università di Trieste^b, Trieste, Italy

S. Belforte^a , V. Candelise^{a,b} , M. Casarsa^a , F. Cossutti^a , K. De Leo^a , G. Della Ricca^{a,b} 

Kyungpook National University, Daegu, Korea

S. Dogra , J. Hong , C. Huh , B. Kim , J. Kim, D. Lee, H. Lee, S.W. Lee , C.S. Moon , Y.D. Oh , M.S. Ryu , S. Sekmen , B. Tae, Y.C. Yang

Department of Mathematics and Physics - GWNU, Gangneung, Korea

M.S. Kim 

Chonnam National University, Institute for Universe and Elementary Particles, Kwangju, Korea

G. Bak , P. Gwak , H. Kim , D.H. Moon 

Hanyang University, Seoul, Korea

E. Asilar , J. Choi , D. Kim , T.J. Kim , J.A. Merlin, Y. Ryou

Korea University, Seoul, Korea

S. Choi , S. Han, B. Hong , K. Lee, K.S. Lee , S. Lee , J. Yoo 

Kyung Hee University, Department of Physics, Seoul, Korea

J. Goh , S. Yang 

Sejong University, Seoul, Korea

H. S. Kim , Y. Kim, S. Lee

Seoul National University, Seoul, Korea

J. Almond, J.H. Bhyun, J. Choi , J. Choi, W. Jun , J. Kim , S. Ko , H. Kwon , H. Lee , J. Lee , J. Lee , B.H. Oh , S.B. Oh , H. Seo , U.K. Yang, I. Yoon

University of Seoul, Seoul, Korea

W. Jang , D.Y. Kang, Y. Kang , S. Kim , B. Ko, J.S.H. Lee , Y. Lee , I.C. Park , Y. Roh, I.J. Watson 

Yonsei University, Department of Physics, Seoul, Korea

S. Ha , H.D. Yoo 

Sungkyunkwan University, Suwon, Korea

M. Choi , M.R. Kim , H. Lee, Y. Lee , I. Yu 

College of Engineering and Technology, American University of the Middle East (AUM), Dasman, Kuwait

T. Beyrouthy

Riga Technical University, Riga, Latvia

K. Dreimanis , A. Gaile , G. Pikurs, A. Potrebko , M. Seidel , D. Sidiropoulos Kontos

University of Latvia (LU), Riga, Latvia

N.R. Strautnieks 

Vilnius University, Vilnius, Lithuania

M. Ambrozas , A. Juodagalvis , A. Rinkevicius , G. Tamulaitis 

National Centre for Particle Physics, Universiti Malaya, Kuala Lumpur, Malaysia
 I. Yusuff⁵⁴ , Z. Zolkapli

Universidad de Sonora (UNISON), Hermosillo, Mexico

J.F. Benitez , A. Castaneda Hernandez , H.A. Encinas Acosta, L.G. Gallegos Maríñez, M. León Coello , J.A. Murillo Quijada , A. Sehrawat , L. Valencia Palomo 

Centro de Investigacion y de Estudios Avanzados del IPN, Mexico City, Mexico

G. Ayala , H. Castilla-Valdez , H. Crotte Ledesma, E. De La Cruz-Burelo , I. Heredia-De La Cruz⁵⁵ , R. Lopez-Fernandez , J. Mejia Guisao , C.A. Mondragon Herrera, A. Sánchez Hernández 

Universidad Iberoamericana, Mexico City, Mexico

C. Oropeza Barrera , D.L. Ramirez Guadarrama, M. Ramírez García 

Benemerita Universidad Autonoma de Puebla, Puebla, Mexico

I. Bautista , I. Pedraza , H.A. Salazar Ibarguen , C. Uribe Estrada 

University of Montenegro, Podgorica, Montenegro

I. Bubanja, N. Raicevic 

University of Canterbury, Christchurch, New Zealand

P.H. Butler 

National Centre for Physics, Quaid-I-Azam University, Islamabad, Pakistan

A. Ahmad , M.I. Asghar, A. Awais , M.I.M. Awan, H.R. Hoorani , W.A. Khan 

AGH University of Krakow, Faculty of Computer Science, Electronics and Telecommunications, Krakow, Poland

V. Avati, L. Grzanka , M. Malawski 

National Centre for Nuclear Research, Swierk, Poland

H. Bialkowska , M. Bluj , M. Górski , M. Kazana , M. Szleper , P. Zalewski 

Institute of Experimental Physics, Faculty of Physics, University of Warsaw, Warsaw, Poland

K. Bunkowski , K. Doroba , A. Kalinowski , M. Konecki , J. Krolkowski , A. Muhammad 

Warsaw University of Technology, Warsaw, Poland

K. Pozniak , W. Zabolotny 

Laboratório de Instrumentação e Física Experimental de Partículas, Lisboa, Portugal

M. Araujo , D. Bastos , C. Beirão Da Cruz E Silva , A. Boletti , M. Bozzo , T. Camporesi , G. Da Molin , P. Faccioli , M. Gallinaro , J. Hollar , N. Leonardo , G.B. Marozzo, T. Niknejad , A. Petrilli , M. Pisano , J. Seixas , J. Varela , J.W. Wulff

Faculty of Physics, University of Belgrade, Belgrade, Serbia

P. Adzic , P. Milenovic 

VINCA Institute of Nuclear Sciences, University of Belgrade, Belgrade, Serbia

M. Dordevic , J. Milosevic , L. Nadderd , V. Rekovic

Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

J. Alcaraz Maestre , Cristina F. Bedoya , Oliver M. Carretero , M. Cepeda , M. Cerrada , N. Colino , B. De La Cruz , A. Delgado Peris , A. Escalante Del Valle , D. Fernández Del Val , J.P. Fernández Ramos , J. Flix , M.C. Fouz , O. Gonzalez Lopez 

S. Goy Lopez [ID](#), J.M. Hernandez [ID](#), M.I. Josa [ID](#), E. Martin Viscasillas [ID](#), D. Moran [ID](#), C. M. Morcillo Perez [ID](#), Á. Navarro Tobar [ID](#), C. Perez Dengra [ID](#), A. Pérez-Calero Yzquierdo [ID](#), J. Puerta Pelayo [ID](#), I. Redondo [ID](#), S. Sánchez Navas [ID](#), J. Sastre [ID](#), J. Vazquez Escobar [ID](#)

Universidad Autónoma de Madrid, Madrid, Spain

J.F. de Trocóniz [ID](#)

Universidad de Oviedo, Instituto Universitario de Ciencias y Tecnologías Espaciales de Asturias (ICTEA), Oviedo, Spain

B. Alvarez Gonzalez [ID](#), J. Cuevas [ID](#), J. Fernandez Menendez [ID](#), S. Folgueras [ID](#), I. Gonzalez Caballero [ID](#), J.R. González Fernández [ID](#), P. Leguina [ID](#), E. Palencia Cortezon [ID](#), C. Ramón Álvarez [ID](#), V. Rodríguez Bouza [ID](#), A. Soto Rodríguez [ID](#), A. Trapote [ID](#), C. Vico Villalba [ID](#), P. Vischia [ID](#)

Instituto de Física de Cantabria (IFCA), CSIC-Universidad de Cantabria, Santander, Spain

S. Bhowmik [ID](#), S. Blanco Fernández [ID](#), J.A. Brochero Cifuentes [ID](#), I.J. Cabrillo [ID](#), A. Calderon [ID](#), J. Duarte Campderros [ID](#), M. Fernandez [ID](#), G. Gomez [ID](#), C. Lasosa García [ID](#), R. Lopez Ruiz [ID](#), C. Martinez Rivero [ID](#), P. Martinez Ruiz del Arbol [ID](#), F. Matorras [ID](#), P. Matorras Cuevas [ID](#), E. Navarrete Ramos [ID](#), J. Piedra Gomez [ID](#), L. Scodellaro [ID](#), I. Vila [ID](#), J.M. Vizan Garcia [ID](#)

University of Colombo, Colombo, Sri Lanka

B. Kailasapathy⁵⁶ [ID](#), D.D.C. Wickramarathna [ID](#)

University of Ruhuna, Department of Physics, Matara, Sri Lanka

W.G.D. Dharmaratna⁵⁷ [ID](#), K. Liyanage [ID](#), N. Perera [ID](#)

CERN, European Organization for Nuclear Research, Geneva, Switzerland

D. Abbaneo [ID](#), C. Amendola [ID](#), E. Auffray [ID](#), G. Auzinger [ID](#), J. Baechler, D. Barney [ID](#), A. Bermúdez Martínez [ID](#), M. Bianco [ID](#), B. Bilin [ID](#), A.A. Bin Anuar [ID](#), A. Bocci [ID](#), C. Botta [ID](#), E. Brondolin [ID](#), C. Caillol [ID](#), G. Cerminara [ID](#), N. Chernyavskaya [ID](#), D. d'Enterria [ID](#), A. Dabrowski [ID](#), A. David [ID](#), A. De Roeck [ID](#), M.M. Defranchis [ID](#), M. Deile [ID](#), M. Dobson [ID](#), G. Franzoni [ID](#), W. Funk [ID](#), S. Giani, D. Gigi, K. Gill [ID](#), F. Glege [ID](#), J. Hegeman [ID](#), J.K. Heikkilä [ID](#), B. Huber, V. Innocente [ID](#), T. James [ID](#), P. Janot [ID](#), O. Kaluzinska [ID](#), S. Laurila [ID](#), P. Lecoq [ID](#), E. Leutgeb [ID](#), C. Lourenço [ID](#), L. Malgeri [ID](#), M. Mannelli [ID](#), A.C. Marini [ID](#), M. Matthewman, A. Mehta [ID](#), F. Meijers [ID](#), S. Mersi [ID](#), E. Meschi [ID](#), V. Milosevic [ID](#), F. Monti [ID](#), F. Moortgat [ID](#), M. Mulders [ID](#), I. Neutelings [ID](#), S. Orfanelli, F. Pantaleo [ID](#), G. Petrucciani [ID](#), A. Pfeiffer [ID](#), M. Pierini [ID](#), H. Qu [ID](#), D. Rabady [ID](#), B. Ribeiro Lopes [ID](#), M. Rovere [ID](#), H. Sakulin [ID](#), S. Sanchez Cruz [ID](#), S. Scarfi [ID](#), C. Schwick, M. Selvaggi [ID](#), A. Sharma [ID](#), K. Shchelina [ID](#), P. Silva [ID](#), P. Sphicas⁵⁸ [ID](#), A.G. Stahl Leiton [ID](#), A. Steen [ID](#), S. Summers [ID](#), D. Treille [ID](#), P. Tropea [ID](#), D. Walter [ID](#), J. Wanczyk⁵⁹ [ID](#), J. Wang, S. Wuchterl [ID](#), P. Zehetner [ID](#), P. Zejdl [ID](#), W.D. Zeuner

Paul Scherrer Institut, Villigen, Switzerland

T. Bevilacqua⁶⁰ [ID](#), L. Caminada⁶⁰ [ID](#), A. Ebrahimi [ID](#), W. Erdmann [ID](#), R. Horisberger [ID](#), Q. Ingram [ID](#), H.C. Kaestli [ID](#), D. Kotlinski [ID](#), C. Lange [ID](#), M. Missiroli⁶⁰ [ID](#), L. Noehte⁶⁰ [ID](#), T. Rohe [ID](#)

ETH Zurich - Institute for Particle Physics and Astrophysics (IPA), Zurich, Switzerland

T.K. Arrestad [ID](#), K. Androsov⁵⁹ [ID](#), M. Backhaus [ID](#), G. Bonomelli, A. Calandri [ID](#), C. Cazzaniga [ID](#), K. Datta [ID](#), P. De Bryas Dexmiers D'archiac⁵⁹ [ID](#), A. De Cosa [ID](#), G. Dissertori [ID](#), M. Dittmar, M. Donegà [ID](#), F. Eble [ID](#), M. Galli [ID](#), K. Gedia [ID](#), F. Glessgen [ID](#), C. Grab [ID](#), N. Härringer [ID](#), T.G. Harte, D. Hits [ID](#), W. Lustermann [ID](#), A.-M. Lyon [ID](#), R.A. Manzoni [ID](#),

M. Marchegiani , L. Marchese , C. Martin Perez , A. Mascellani⁵⁹ , F. Nessi-Tedaldi , F. Pauss , V. Perovic , S. Pigazzini , C. Reissel , T. Reitenspiess , B. Ristic , F. Riti , R. Seidita , J. Steggemann⁵⁹ , A. Tarabini , D. Valsecchi , R. Wallny

Universität Zürich, Zurich, Switzerland

C. Amsler⁶¹ , P. Bärtschi , M.F. Canelli , K. Cormier , M. Huwiler , W. Jin , A. Jofrehei , B. Kilminster , S. Leontsinis , S.P. Liechti , A. Macchiolo , P. Meiring , F. Meng , U. Molinatti , J. Motta , A. Reimers , P. Robmann, M. Senger , E. Shokr, F. Stäger , R. Tramontano

National Central University, Chung-Li, Taiwan

C. Adloff⁶², D. Bhowmik, C.M. Kuo, W. Lin, P.K. Rout , P.C. Tiwari³⁸ , S.S. Yu 

National Taiwan University (NTU), Taipei, Taiwan

L. Ceard, K.F. Chen , P.s. Chen, Z.g. Chen, A. De Iorio , W.-S. Hou , T.h. Hsu, Y.w. Kao, S. Karmakar , G. Kole , Y.y. Li , R.-S. Lu , E. Paganis , X.f. Su , J. Thomas-Wilsker , L.s. Tsai, H.y. Wu, E. Yazgan

High Energy Physics Research Unit, Department of Physics, Faculty of Science, Chulalongkorn University, Bangkok, Thailand

C. Asawatangtrakuldee , N. Srimanobhas , V. Wachirapusanand 

Çukurova University, Physics Department, Science and Art Faculty, Adana, Turkey

D. Agyel , F. Boran , F. Dolek , I. Dumanoglu⁶³ , E. Eskut , Y. Guler⁶⁴ , E. Gurpinar Guler⁶⁴ , C. Isik , O. Kara, A. Kayis Topaksu , U. Kiminsu , G. Onengut , K. Ozdemir⁶⁵ , A. Polatoz , B. Tali⁶⁶ , U.G. Tok , S. Turkcapar , E. Uslan , I.S. Zorbakir

Middle East Technical University, Physics Department, Ankara, Turkey

G. Sokmen, M. Yalvac⁶⁷ 

Bogazici University, Istanbul, Turkey

B. Akgun , I.O. Atakisi , E. Gürmez , M. Kaya⁶⁸ , O. Kaya⁶⁹ , S. Tekten⁷⁰ 

Istanbul Technical University, Istanbul, Turkey

A. Cakir , K. Cankocak^{63,71} , G.G. Dincer⁶³ , Y. Komurcu , S. Sen⁷² 

Istanbul University, Istanbul, Turkey

O. Aydilek⁷³ , B. Hacisahinoglu , I. Hos⁷⁴ , B. Kaynak , S. Ozkorucuklu , O. Potok , H. Sert , C. Simsek , C. Zorbilmez

Yildiz Technical University, Istanbul, Turkey

S. Cerci⁶⁶ , B. Isildak⁷⁵ , D. Sunar Cerci , T. Yetkin 

Institute for Scintillation Materials of National Academy of Science of Ukraine, Kharkiv, Ukraine

A. Boyaryntsev , B. Grynyov 

National Science Centre, Kharkiv Institute of Physics and Technology, Kharkiv, Ukraine

L. Levchuk 

University of Bristol, Bristol, United Kingdom

D. Anthony , J.J. Brooke , A. Bundock , F. Bury , E. Clement , D. Cussans , H. Flacher , M. Glowacki, J. Goldstein , H.F. Heath , M.-L. Holmberg , L. Kreczko , S. Paramesvaran , L. Robertshaw, S. Seif El Nasr-Storey, V.J. Smith , N. Stylianou⁷⁶ , K. Walkingshaw Pass

Rutherford Appleton Laboratory, Didcot, United Kingdom

A.H. Ball, K.W. Bell [ID](#), A. Belyaev⁷⁷ [ID](#), C. Brew [ID](#), R.M. Brown [ID](#), D.J.A. Cockerill [ID](#), C. Cooke [ID](#), A. Elliot [ID](#), K.V. Ellis, K. Harder [ID](#), S. Harper [ID](#), J. Linacre [ID](#), K. Manolopoulos, D.M. Newbold [ID](#), E. Olaiya, D. Petyt [ID](#), T. Reis [ID](#), A.R. Sahasransu [ID](#), G. Salvi [ID](#), T. Schuh, C.H. Shepherd-Themistocleous [ID](#), I.R. Tomalin [ID](#), K.C. Whalen [ID](#), T. Williams [ID](#)

Imperial College, London, United Kingdom

I. Andreou [ID](#), R. Bainbridge [ID](#), P. Bloch [ID](#), C.E. Brown [ID](#), O. Buchmuller, V. Cacchio, C.A. Carrillo Montoya [ID](#), G.S. Chahal⁷⁸ [ID](#), D. Colling [ID](#), J.S. Dancu, I. Das [ID](#), P. Dauncey [ID](#), G. Davies [ID](#), J. Davies, M. Della Negra [ID](#), S. Fayer, G. Fedi [ID](#), G. Hall [ID](#), M.H. Hassanshahi [ID](#), A. Howard, G. Iles [ID](#), M. Knight [ID](#), J. Langford [ID](#), J. León Holgado [ID](#), L. Lyons [ID](#), A.-M. Magnan [ID](#), S. Mallios, M. Mieskolainen [ID](#), J. Nash⁷⁹ [ID](#), M. Pesaresi [ID](#), P.B. Pradeep, B.C. Radburn-Smith [ID](#), A. Richards, A. Rose [ID](#), K. Savva [ID](#), C. Seez [ID](#), R. Shukla [ID](#), A. Tapper [ID](#), K. Uchida [ID](#), G.P. Uttley [ID](#), L.H. Vage, T. Virdee³⁰ [ID](#), M. Vojinovic [ID](#), N. Wardle [ID](#), D. Winterbottom [ID](#)

Brunel University, Uxbridge, United Kingdom

K. Coldham, J.E. Cole [ID](#), A. Khan, P. Kyberd [ID](#), I.D. Reid [ID](#)

Baylor University, Waco, Texas, USA

S. Abdullin [ID](#), A. Brinkerhoff [ID](#), B. Caraway [ID](#), E. Collins [ID](#), J. Dittmann [ID](#), K. Hatakeyama [ID](#), J. Hiltbrand [ID](#), B. McMaster [ID](#), J. Samudio [ID](#), S. Sawant [ID](#), C. Sutantawibul [ID](#), J. Wilson [ID](#)

Catholic University of America, Washington, DC, USA

R. Bartek [ID](#), A. Dominguez [ID](#), C. Huerta Escamilla, A.E. Simsek [ID](#), R. Uniyal [ID](#), A.M. Vargas Hernandez [ID](#)

The University of Alabama, Tuscaloosa, Alabama, USA

B. Bam [ID](#), A. Buchot Perraguin [ID](#), R. Chudasama [ID](#), S.I. Cooper [ID](#), C. Crovella [ID](#), S.V. Gleyzer [ID](#), E. Pearson, C.U. Perez [ID](#), P. Rumerio⁸⁰ [ID](#), E. Usai [ID](#), R. Yi [ID](#)

Boston University, Boston, Massachusetts, USA

A. Akpinar [ID](#), C. Cosby [ID](#), G. De Castro, Z. Demiragli [ID](#), C. Erice [ID](#), C. Fangmeier [ID](#), C. Fernandez Madrazo [ID](#), E. Fontanesi [ID](#), D. Gastler [ID](#), F. Golf [ID](#), S. Jeon [ID](#), J. O'cain, I. Reed [ID](#), J. Rohlf [ID](#), K. Salyer [ID](#), D. Sperka [ID](#), D. Spitzbart [ID](#), I. Suarez [ID](#), A. Tsatsos [ID](#), A.G. Zecchinelli [ID](#)

Brown University, Providence, Rhode Island, USA

G. Benelli [ID](#), X. Coubez²⁶, D. Cutts [ID](#), L. Gouskos [ID](#), M. Hadley [ID](#), U. Heintz [ID](#), J.M. Hogan⁸¹ [ID](#), T. Kwon [ID](#), G. Landsberg [ID](#), K.T. Lau [ID](#), D. Li [ID](#), J. Luo [ID](#), S. Mondal [ID](#), M. Narain[†] [ID](#), N. Pervan [ID](#), T. Russell, S. Sagir⁸² [ID](#), F. Simpson [ID](#), M. Stamenkovic [ID](#), N. Venkatasubramanian, X. Yan [ID](#), W. Zhang

University of California, Davis, Davis, California, USA

S. Abbott [ID](#), C. Brainerd [ID](#), R. Breedon [ID](#), H. Cai [ID](#), M. Calderon De La Barca Sanchez [ID](#), M. Chertok [ID](#), M. Citron [ID](#), J. Conway [ID](#), P.T. Cox [ID](#), R. Erbacher [ID](#), F. Jensen [ID](#), O. Kukral [ID](#), G. Mocellin [ID](#), M. Mulhearn [ID](#), S. Ostrom [ID](#), W. Wei [ID](#), Y. Yao [ID](#), S. Yoo [ID](#), F. Zhang [ID](#)

University of California, Los Angeles, California, USA

M. Bachtis [ID](#), R. Cousins [ID](#), A. Datta [ID](#), G. Flores Avila, J. Hauser [ID](#), M. Ignatenko [ID](#), M.A. Iqbal [ID](#), T. Lam [ID](#), E. Manca [ID](#), A. Nunez Del Prado, D. Saltzberg [ID](#), V. Valuev [ID](#)

University of California, Riverside, Riverside, California, USA

R. Clare [ID](#), J.W. Gary [ID](#), M. Gordon, G. Hanson [ID](#), W. Si [ID](#), S. Wimpenny[†] [ID](#)

University of California, San Diego, La Jolla, California, USA

A. Aportela, A. Arora [ID](#), J.G. Branson [ID](#), S. Cittolin [ID](#), S. Cooperstein [ID](#), D. Diaz [ID](#), J. Duarte [ID](#), L. Giannini [ID](#), Y. Gu, J. Guiang [ID](#), R. Kansal [ID](#), V. Krutelyov [ID](#), R. Lee [ID](#), J. Letts [ID](#), M. Masciovecchio [ID](#), F. Mokhtar [ID](#), S. Mukherjee [ID](#), M. Pieri [ID](#), M. Quinnan [ID](#), B.V. Sathia Narayanan [ID](#), V. Sharma [ID](#), M. Tadel [ID](#), E. Vourliotis [ID](#), F. Würthwein [ID](#), Y. Xiang [ID](#), A. Yagil [ID](#)

University of California, Santa Barbara - Department of Physics, Santa Barbara, California, USA

A. Barzdukas [ID](#), L. Brennan [ID](#), C. Campagnari [ID](#), K. Downham [ID](#), C. Grieco [ID](#), J. Incandela [ID](#), J. Kim [ID](#), A.J. Li [ID](#), P. Masterson [ID](#), H. Mei [ID](#), J. Richman [ID](#), S.N. Santpur [ID](#), U. Sarica [ID](#), R. Schmitz [ID](#), F. Setti [ID](#), J. Sheplock [ID](#), D. Stuart [ID](#), T.Á. Vámi [ID](#), S. Wang [ID](#), D. Zhang

California Institute of Technology, Pasadena, California, USA

A. Bornheim [ID](#), O. Cerri, A. Latorre, J. Mao [ID](#), H.B. Newman [ID](#), G. Reales Gutiérrez, M. Spiropulu [ID](#), J.R. Vlimant [ID](#), C. Wang [ID](#), S. Xie [ID](#), R.Y. Zhu [ID](#)

Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

J. Alison [ID](#), S. An [ID](#), M.B. Andrews [ID](#), P. Bryant [ID](#), M. Cremonesi, V. Dutta [ID](#), T. Ferguson [ID](#), A. Harilal [ID](#), A. Kallil Tharayil, C. Liu [ID](#), T. Mudholkar [ID](#), S. Murthy [ID](#), P. Palit [ID](#), K. Park, M. Paulini [ID](#), A. Roberts [ID](#), A. Sanchez [ID](#), W. Terrill [ID](#)

University of Colorado Boulder, Boulder, Colorado, USA

J.P. Cumalat [ID](#), W.T. Ford [ID](#), A. Hart [ID](#), A. Hassani [ID](#), G. Karathanasis [ID](#), N. Manganelli [ID](#), A. Perloff [ID](#), C. Savard [ID](#), N. Schonbeck [ID](#), K. Stenson [ID](#), K.A. Ulmer [ID](#), S.R. Wagner [ID](#), N. Zipper [ID](#), D. Zuolo [ID](#)

Cornell University, Ithaca, New York, USA

J. Alexander [ID](#), S. Bright-Thonney [ID](#), X. Chen [ID](#), D.J. Cranshaw [ID](#), J. Fan [ID](#), X. Fan [ID](#), S. Hogan [ID](#), P. Kotamnives, J. Monroy [ID](#), M. Oshiro [ID](#), J.R. Patterson [ID](#), M. Reid [ID](#), A. Ryd [ID](#), J. Thom [ID](#), P. Wittich [ID](#), R. Zou [ID](#)

Fermi National Accelerator Laboratory, Batavia, Illinois, USA

M. Albrow [ID](#), M. Alyari [ID](#), O. Amram [ID](#), G. Apollinari [ID](#), A. Apresyan [ID](#), L.A.T. Bauerick [ID](#), D. Berry [ID](#), J. Berryhill [ID](#), P.C. Bhat [ID](#), K. Burkett [ID](#), J.N. Butler [ID](#), A. Canepa [ID](#), G.B. Cerati [ID](#), H.W.K. Cheung [ID](#), F. Chlebana [ID](#), G. Cummings [ID](#), J. Dickinson [ID](#), I. Dutta [ID](#), V.D. Elvira [ID](#), Y. Feng [ID](#), J. Freeman [ID](#), A. Gandrakota [ID](#), Z. Gecse [ID](#), L. Gray [ID](#), D. Green, A. Grummer [ID](#), S. Grünendahl [ID](#), D. Guerrero [ID](#), O. Gutsche [ID](#), R.M. Harris [ID](#), R. Heller [ID](#), T.C. Herwig [ID](#), J. Hirschauer [ID](#), B. Jayatilaka [ID](#), S. Jindariani [ID](#), M. Johnson [ID](#), U. Joshi [ID](#), T. Klijnsma [ID](#), B. Klima [ID](#), K.H.M. Kwok [ID](#), S. Lammel [ID](#), D. Lincoln [ID](#), R. Lipton [ID](#), T. Liu [ID](#), C. Madrid [ID](#), K. Maeshima [ID](#), C. Mantilla [ID](#), D. Mason [ID](#), P. McBride [ID](#), P. Merkel [ID](#), S. Mrenna [ID](#), S. Nahn [ID](#), J. Ngadiuba [ID](#), D. Noonan [ID](#), S. Norberg, V. Papadimitriou [ID](#), N. Pastika [ID](#), K. Pedro [ID](#), C. Pena⁸³ [ID](#), F. Raveras [ID](#), A. Reinsvold Hall⁸⁴ [ID](#), L. Ristori [ID](#), M. Safdari [ID](#), E. Sexton-Kennedy [ID](#), N. Smith [ID](#), A. Soha [ID](#), L. Spiegel [ID](#), S. Stoynev [ID](#), J. Strait [ID](#), L. Taylor [ID](#), S. Tkaczyk [ID](#), N.V. Tran [ID](#), L. Uplegger [ID](#), E.W. Vaandering [ID](#), I. Zoi [ID](#)

University of Florida, Gainesville, Florida, USA

C. Aruta [ID](#), P. Avery [ID](#), D. Bourilkov [ID](#), P. Chang [ID](#), V. Cherepanov [ID](#), R.D. Field, E. Koenig [ID](#), M. Kolosova [ID](#), J. Konigsberg [ID](#), A. Korytov [ID](#), K. Matchev [ID](#), N. Menendez [ID](#), G. Mitselmakher [ID](#), K. Mohrman [ID](#), A. Muthirakalayil Madhu [ID](#), N. Rawal [ID](#), S. Rosenzweig [ID](#), Y. Takahashi [ID](#), J. Wang [ID](#)

Florida State University, Tallahassee, Florida, USA

T. Adams , A. Al Kadhim , A. Askew , S. Bower , R. Habibullah , V. Hagopian , R. Hashmi , R.S. Kim , S. Kim , T. Kolberg , G. Martinez, H. Prosper , P.R. Prova, M. Wulansatiti , R. Yohay , J. Zhang

Florida Institute of Technology, Melbourne, Florida, USA

B. Alsufyani, M.M. Baarmand , S. Butalla , S. Das , T. Elkafrawy⁸⁵ , M. Hohlmann , M. Rahmani, E. Yanes

University of Illinois Chicago, Chicago, USA, Chicago, USA

M.R. Adams , A. Baty , C. Bennett, R. Cavanaugh , R. Escobar Franco , O. Evdokimov , C.E. Gerber , M. Hawksworth, A. Hingrajiya, D.J. Hofman , J.h. Lee , D. S. Lemos , A.H. Merrit , C. Mills , S. Nanda , G. Oh , B. Ozek , D. Pilipovic , R. Pradhan , E. Prifti, T. Roy , S. Rudrabhatla , M.B. Tonjes , N. Varelas , M.A. Wadud , Z. Ye , J. Yoo

The University of Iowa, Iowa City, Iowa, USA

M. Alhusseini , D. Blend, K. Dilsiz⁸⁶ , L. Emediato , G. Karaman , O.K. Köseyan , J.-P. Merlo, A. Mestvirishvili⁸⁷ , O. Neogi, H. O gul⁸⁸ , Y. Onel , A. Penzo , C. Snyder, E. Tiras⁸⁹

Johns Hopkins University, Baltimore, Maryland, USA

B. Blumenfeld , L. Corcodilos , J. Davis , A.V. Gritsan , L. Kang , S. Kyriacou , P. Maksimovic , M. Roguljic , J. Roskes , S. Sekhar , M. Swartz

The University of Kansas, Lawrence, Kansas, USA

A. Abreu , L.F. Alcerro Alcerro , J. Anguiano , S. Arteaga Escatel , P. Baringer , A. Bean , Z. Flowers , D. Grove , J. King , G. Krintiras , M. Lazarovits , C. Le Mahieu , J. Marquez , N. Minafra , M. Murray , M. Nickel , M. Pitt , S. Popescu⁹⁰ , C. Rogan , C. Royon , R. Salvatico , S. Sanders , C. Smith , G. Wilson

Kansas State University, Manhattan, Kansas, USA

B. Allmond , R. Guju Gurunadha , A. Ivanov , K. Kaadze , Y. Maravin , J. Natoli , D. Roy , G. Sorrentino 

University of Maryland, College Park, Maryland, USA

A. Baden , A. Belloni , J. Bistany-riebman, Y.M. Chen , S.C. Eno , N.J. Hadley , S. Jabeen , R.G. Kellogg , T. Koeth , B. Kronheim, Y. Lai , S. Lascio , A.C. Mignerey , S. Nabili , C. Palmer , C. Papageorgakis , M.M. Paranjpe, L. Wang

Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

J. Bendavid , I.A. Cali , P.c. Chou , M. D'Alfonso , J. Eysermans , C. Freer , G. Gomez-Ceballos , M. Goncharov, G. Gross, P. Harris, D. Hoang, D. Kovalskyi , J. Krupa , L. Lavezzi , Y.-J. Lee , K. Long , C. Mcginn, A. Novak , C. Paus , D. Rankin , C. Roland , G. Roland , S. Rothman , G.S.F. Stephans , Z. Wang , B. Wyslouch , T. J. Yang

University of Minnesota, Minneapolis, Minnesota, USA

B. Crossman , B.M. Joshi , C. Kapsiak , M. Krohn , D. Mahon , J. Mans , B. Marzocchi , M. Revering , R. Rusack , R. Saradhy , N. Strobbe

University of Nebraska-Lincoln, Lincoln, Nebraska, USA

K. Bloom , D.R. Claes , G. Haza , J. Hossain , C. Joo , I. Kravchenko , J.E. Siado , W. Tabb , A. Vagnerini , A. Wightman , F. Yan , D. Yu

State University of New York at Buffalo, Buffalo, New York, USA

H. Bandyopadhyay [ID](#), L. Hay [ID](#), H.w. Hsia, I. Iashvili [ID](#), A. Kalogeropoulos [ID](#), A. Kharchilava [ID](#), M. Morris [ID](#), D. Nguyen [ID](#), S. Rappoccio [ID](#), H. Rejeb Sfar, A. Williams [ID](#), P. Young [ID](#)

Northeastern University, Boston, Massachusetts, USA

G. Alverson [ID](#), E. Barberis [ID](#), J. Bonilla [ID](#), J. Dervan, Y. Haddad [ID](#), Y. Han [ID](#), A. Krishna [ID](#), J. Li [ID](#), M. Lu [ID](#), G. Madigan [ID](#), R. Mccarthy [ID](#), D.M. Morse [ID](#), V. Nguyen [ID](#), T. Orimoto [ID](#), A. Parker [ID](#), L. Skinnari [ID](#), D. Wood [ID](#)

Northwestern University, Evanston, Illinois, USA

J. Bueghly, S. Dittmer [ID](#), K.A. Hahn [ID](#), Y. Liu [ID](#), Y. Miao [ID](#), D.G. Monk [ID](#), M.H. Schmitt [ID](#), A. Taliercio [ID](#), M. Velasco

University of Notre Dame, Notre Dame, Indiana, USA

G. Agarwal [ID](#), R. Band [ID](#), R. Bucci, S. Castells [ID](#), A. Das [ID](#), R. Goldouzian [ID](#), M. Hildreth [ID](#), K.W. Ho [ID](#), K. Hurtado Anampa [ID](#), T. Ivanov [ID](#), C. Jessop [ID](#), K. Lannon [ID](#), J. Lawrence [ID](#), N. Loukas [ID](#), L. Lutton [ID](#), J. Mariano, N. Marinelli, I. Mcalister, T. McCauley [ID](#), C. Mcgrady [ID](#), C. Moore [ID](#), Y. Musienko¹⁷ [ID](#), H. Nelson [ID](#), M. Osherson [ID](#), A. Piccinelli [ID](#), R. Ruchti [ID](#), A. Townsend [ID](#), Y. Wan, M. Wayne [ID](#), H. Yockey, M. Zarucki [ID](#), L. Zygala [ID](#)

The Ohio State University, Columbus, Ohio, USA

A. Basnet [ID](#), B. Bylsma, M. Carrigan [ID](#), L.S. Durkin [ID](#), C. Hill [ID](#), M. Joyce [ID](#), M. Nunez Ornelas [ID](#), K. Wei, B.L. Winer [ID](#), B. R. Yates [ID](#)

Princeton University, Princeton, New Jersey, USA

H. Bouchamaoui [ID](#), P. Das [ID](#), G. Dezoort [ID](#), P. Elmer [ID](#), A. Frankenthal [ID](#), B. Greenberg [ID](#), N. Haubrich [ID](#), K. Kennedy, G. Kopp [ID](#), S. Kwan [ID](#), D. Lange [ID](#), A. Loeliger [ID](#), D. Marlow [ID](#), I. Ojalvo [ID](#), J. Olsen [ID](#), A. Shevelev [ID](#), D. Stickland [ID](#), C. Tully [ID](#)

University of Puerto Rico, Mayaguez, Puerto Rico, USA

S. Malik [ID](#)

Purdue University, West Lafayette, Indiana, USA

A.S. Bakshi [ID](#), S. Chandra [ID](#), R. Chawla [ID](#), A. Gu [ID](#), L. Gutay, M. Jones [ID](#), A.W. Jung [ID](#), A.M. Koshy, M. Liu [ID](#), G. Negro [ID](#), N. Neumeister [ID](#), G. Paspalaki [ID](#), S. Piperov [ID](#), V. Scheurer, J.F. Schulte [ID](#), M. Stojanovic [ID](#), J. Thieman [ID](#), A. K. Virdi [ID](#), F. Wang [ID](#), W. Xie [ID](#)

Purdue University Northwest, Hammond, Indiana, USA

J. Dolen [ID](#), N. Parashar [ID](#), A. Pathak [ID](#)

Rice University, Houston, Texas, USA

D. Acosta [ID](#), T. Carnahan [ID](#), K.M. Ecklund [ID](#), P.J. Fernández Manteca [ID](#), S. Freed, P. Gardner, F.J.M. Geurts [ID](#), W. Li [ID](#), J. Lin [ID](#), O. Miguel Colin [ID](#), B.P. Padley [ID](#), R. Redjimi, J. Rotter [ID](#), E. Yigitbasi [ID](#), Y. Zhang [ID](#)

University of Rochester, Rochester, New York, USA

A. Bodek [ID](#), P. de Barbaro [ID](#), R. Demina [ID](#), J.L. Dulemba [ID](#), A. Garcia-Bellido [ID](#), O. Hindrichs [ID](#), A. Khukhunaishvili [ID](#), N. Parmar, P. Parygin⁹¹ [ID](#), E. Popova⁹¹ [ID](#), R. Taus [ID](#)

Rutgers, The State University of New Jersey, Piscataway, New Jersey, USA

B. Chiarito, J.P. Chou [ID](#), S.V. Clark [ID](#), D. Gadkari [ID](#), Y. Gershtein [ID](#), E. Halkiadakis [ID](#), M. Heindl [ID](#), C. Houghton [ID](#), D. Jaroslawski [ID](#), O. Karacheban²⁸ [ID](#), S. Konstantinou [ID](#), I. Laflotte [ID](#), A. Lath [ID](#), R. Montalvo, K. Nash, J. Reichert [ID](#), H. Routray [ID](#), P. Saha [ID](#), S. Salur [ID](#), S. Schnetzer, S. Somalwar [ID](#), R. Stone [ID](#), S.A. Thayil [ID](#), S. Thomas, J. Vora [ID](#), H. Wang [ID](#)

University of Tennessee, Knoxville, Tennessee, USA

H. Acharya, D. Ally , A.G. Delannoy , S. Fiorendi , S. Higginbotham , T. Holmes , A.R. Kanuganti , N. Karunaratna , L. Lee , E. Nibigira , S. Spanier 

Texas A&M University, College Station, Texas, USA

D. Aebi , M. Ahmad , T. Akhter , O. Bouhali⁹² , R. Eusebi , J. Gilmore , T. Huang , T. Kamon⁹³ , H. Kim , S. Luo , R. Mueller , D. Overton , D. Rathjens , A. Safonov 

Texas Tech University, Lubbock, Texas, USA

N. Akchurin , J. Damgov , N. Gogate , V. Hegde , A. Hussain , Y. Kazhykarim, K. Lamichhane , S.W. Lee , A. Mankel , T. Peltola , I. Volobouev 

Vanderbilt University, Nashville, Tennessee, USA

E. Appelt , Y. Chen , S. Greene, A. Gurrola , W. Johns , R. Kunnawalkam Elayavalli , A. Melo , F. Romeo , P. Sheldon , S. Tuo , J. Velkovska , J. Viinikainen 

University of Virginia, Charlottesville, Virginia, USA

B. Cardwell , B. Cox , J. Hakala , R. Hirosky , A. Ledovskoy , C. Neu 

Wayne State University, Detroit, Michigan, USA

S. Bhattacharya , P.E. Karchin 

University of Wisconsin - Madison, Madison, Wisconsin, USA

A. Aravind, S. Banerjee , K. Black , T. Bose , S. Dasu , I. De Bruyn , P. Everaerts , C. Galloni, H. He , M. Herndon , A. Herve , C.K. Koraka , A. Lanaro, R. Loveless , J. Madhusudanan Sreekala , A. Mallampalli , A. Mohammadi , S. Mondal, G. Parida , L. Pétré , D. Pinna, A. Savin, V. Shang , V. Sharma , W.H. Smith , D. Teague, H.F. Tsoi , W. Vetens , A. Warden 

Authors affiliated with an institute or an international laboratory covered by a cooperation agreement with CERN

S. Afanasiev , V. Alexakhin , V. Andreev , Yu. Andreev , T. Aushev , M. Azarkin , A. Babaev , V. Blinov⁹⁴, E. Boos , V. Borshch , D. Budkouski , V. Bunichev , M. Chadeeva⁹⁴ , V. Chekhovsky, R. Chistov⁹⁴ , A. Dermenev , T. Dimova⁹⁴ , D. Druzhkin⁹⁵ , M. Dubinin⁸³ , L. Dudko , A. Ershov , G. Gavrilov , V. Gavrilov , S. Gninenko , V. Golovtcov , N. Golubev , I. Golutvin , I. Gorbunov , A. Gribushin , Y. Ivanov , V. Kachanov , V. Karjavin , A. Karneyeu , V. Kim⁹⁴ , M. Kirakosyan, D. Kirpichnikov , M. Kirsanov , V. Klyukhin , O. Kodolova⁹⁶ , D. Konstantinov , V. Korenkov , A. Kozyrev⁹⁴ , N. Krasnikov , A. Lanev , P. Levchenko⁹⁷ , N. Lychkovskaya , V. Makarenko , A. Malakhov , V. Matveev⁹⁴ , V. Murzin , A. Nikitenko^{98,96} , S. Obraztsov , V. Oreshkin , V. Palichik , V. Perelygin , S. Petrushanko , S. Polikarpov⁹⁴ , V. Popov , O. Radchenko⁹⁴ , M. Savina , V. Savrin , V. Shalaev , S. Shmatov , S. Shulha , Y. Skovpen⁹⁴ , S. Slabospitskii , V. Smirnov , A. Snigirev , D. Sosnov , V. Sulimov , E. Tcherniaev , A. Terkulov , O. Teryaev , I. Tlisova , A. Toropin , L. Uvarov , A. Uzunian , A. Vorobyev[†], N. Voytishin , B.S. Yuldashev⁹⁹, A. Zarubin , I. Zhizhin , A. Zhokin 

[†]: Deceased

¹Also at Yerevan State University, Yerevan, Armenia

²Also at TU Wien, Vienna, Austria

³Also at Institute of Basic and Applied Sciences, Faculty of Engineering, Arab Academy for Science, Technology and Maritime Transport, Alexandria, Egypt

⁴Also at Ghent University, Ghent, Belgium

⁵Also at Universidade do Estado do Rio de Janeiro, Rio de Janeiro, Brazil

⁶Also at Universidade Estadual de Campinas, Campinas, Brazil

⁷Also at Federal University of Rio Grande do Sul, Porto Alegre, Brazil

⁸Also at UFMS, Nova Andradina, Brazil

⁹Also at Nanjing Normal University, Nanjing, China

¹⁰Now at The University of Iowa, Iowa City, Iowa, USA

¹¹Also at University of Chinese Academy of Sciences, Beijing, China

¹²Also at China Center of Advanced Science and Technology, Beijing, China

¹³Also at University of Chinese Academy of Sciences, Beijing, China

¹⁴Also at China Spallation Neutron Source, Guangdong, China

¹⁵Now at Henan Normal University, Xinxiang, China

¹⁶Also at Université Libre de Bruxelles, Bruxelles, Belgium

¹⁷Also at an institute or an international laboratory covered by a cooperation agreement with CERN

¹⁸Also at Suez University, Suez, Egypt

¹⁹Now at British University in Egypt, Cairo, Egypt

²⁰Also at Purdue University, West Lafayette, Indiana, USA

²¹Also at Université de Haute Alsace, Mulhouse, France

²²Also at Department of Physics, Tsinghua University, Beijing, China

²³Also at Tbilisi State University, Tbilisi, Georgia

²⁴Also at The University of the State of Amazonas, Manaus, Brazil

²⁵Also at University of Hamburg, Hamburg, Germany

²⁶Also at RWTH Aachen University, III. Physikalisches Institut A, Aachen, Germany

²⁷Also at Bergische University Wuppertal (BUW), Wuppertal, Germany

²⁸Also at Brandenburg University of Technology, Cottbus, Germany

²⁹Also at Forschungszentrum Jülich, Juelich, Germany

³⁰Also at CERN, European Organization for Nuclear Research, Geneva, Switzerland

³¹Also at Institute of Nuclear Research ATOMKI, Debrecen, Hungary

³²Now at Universitatea Babes-Bolyai - Facultatea de Fizica, Cluj-Napoca, Romania

³³Also at MTA-ELTE Lendület CMS Particle and Nuclear Physics Group, Eötvös Loránd University, Budapest, Hungary

³⁴Also at HUN-REN Wigner Research Centre for Physics, Budapest, Hungary

³⁵Also at Physics Department, Faculty of Science, Assiut University, Assiut, Egypt

³⁶Also at Punjab Agricultural University, Ludhiana, India

³⁷Also at University of Visva-Bharati, Santiniketan, India

³⁸Also at Indian Institute of Science (IISc), Bangalore, India

³⁹Also at IIT Bhubaneswar, Bhubaneswar, India

⁴⁰Also at Institute of Physics, Bhubaneswar, India

⁴¹Also at University of Hyderabad, Hyderabad, India

⁴²Also at Deutsches Elektronen-Synchrotron, Hamburg, Germany

⁴³Also at Isfahan University of Technology, Isfahan, Iran

⁴⁴Also at Sharif University of Technology, Tehran, Iran

⁴⁵Also at Department of Physics, University of Science and Technology of Mazandaran, Behshahr, Iran

⁴⁶Also at Department of Physics, Isfahan University of Technology, Isfahan, Iran

⁴⁷Also at Italian National Agency for New Technologies, Energy and Sustainable Economic Development, Bologna, Italy

⁴⁸Also at Centro Siciliano di Fisica Nucleare e di Struttura Della Materia, Catania, Italy

⁴⁹Also at Università degli Studi Guglielmo Marconi, Roma, Italy

⁵⁰Also at Scuola Superiore Meridionale, Università di Napoli 'Federico II', Napoli, Italy

⁵¹Also at Fermi National Accelerator Laboratory, Batavia, Illinois, USA

⁵²Also at Laboratori Nazionali di Legnaro dell'INFN, Legnaro, Italy

⁵³Also at Consiglio Nazionale delle Ricerche - Istituto Officina dei Materiali, Perugia, Italy

⁵⁴Also at Department of Applied Physics, Faculty of Science and Technology, Universiti Kebangsaan Malaysia, Bangi, Malaysia

⁵⁵Also at Consejo Nacional de Ciencia y Tecnología, Mexico City, Mexico

⁵⁶Also at Trincomalee Campus, Eastern University, Sri Lanka, Nilaveli, Sri Lanka

⁵⁷Also at Saegis Campus, Nugegoda, Sri Lanka

⁵⁸Also at National and Kapodistrian University of Athens, Athens, Greece

⁵⁹Also at Ecole Polytechnique Fédérale Lausanne, Lausanne, Switzerland

⁶⁰Also at Universität Zürich, Zurich, Switzerland

⁶¹Also at Stefan Meyer Institute for Subatomic Physics, Vienna, Austria

⁶²Also at Laboratoire d'Annecy-le-Vieux de Physique des Particules, IN2P3-CNRS, Annecy-le-Vieux, France

⁶³Also at Near East University, Research Center of Experimental Health Science, Mersin, Turkey

⁶⁴Also at Konya Technical University, Konya, Turkey

⁶⁵Also at Izmir Bakircay University, Izmir, Turkey

⁶⁶Also at Adiyaman University, Adiyaman, Turkey

⁶⁷Also at Bozok Universitetesi Rektörlüğü, Yozgat, Turkey

⁶⁸Also at Marmara University, Istanbul, Turkey

⁶⁹Also at Milli Savunma University, Istanbul, Turkey

⁷⁰Also at Kafkas University, Kars, Turkey

⁷¹Now at stanbul Okan University, Istanbul, Turkey

⁷²Also at Hacettepe University, Ankara, Turkey

⁷³Also at Erzincan Binali Yıldırım University, Erzincan, Turkey

⁷⁴Also at Istanbul University - Cerrahpasa, Faculty of Engineering, Istanbul, Turkey

⁷⁵Also at Yıldız Technical University, Istanbul, Turkey

⁷⁶Also at Vrije Universiteit Brussel, Brussel, Belgium

⁷⁷Also at School of Physics and Astronomy, University of Southampton, Southampton, United Kingdom

⁷⁸Also at IPPP Durham University, Durham, United Kingdom

⁷⁹Also at Monash University, Faculty of Science, Clayton, Australia

⁸⁰Also at Università di Torino, Torino, Italy

⁸¹Also at Bethel University, St. Paul, Minnesota, USA

⁸²Also at Karamanoğlu Mehmetbey University, Karaman, Turkey

⁸³Also at California Institute of Technology, Pasadena, California, USA

⁸⁴Also at United States Naval Academy, Annapolis, Maryland, USA

⁸⁵Also at Ain Shams University, Cairo, Egypt

⁸⁶Also at Bingol University, Bingol, Turkey

⁸⁷Also at Georgian Technical University, Tbilisi, Georgia

⁸⁸Also at Sinop University, Sinop, Turkey

⁸⁹Also at Erciyes University, Kayseri, Turkey

⁹⁰Also at Horia Hulubei National Institute of Physics and Nuclear Engineering (IFIN-HH), Bucharest, Romania

⁹¹Now at an institute or an international laboratory covered by a cooperation agreement with CERN

⁹²Also at Texas A&M University at Qatar, Doha, Qatar

⁹³Also at Kyungpook National University, Daegu, Korea

⁹⁴Also at another institute or international laboratory covered by a cooperation agreement with CERN

⁹⁵Also at Universiteit Antwerpen, Antwerpen, Belgium

⁹⁶Also at Yerevan Physics Institute, Yerevan, Armenia

⁹⁷Also at Northeastern University, Boston, Massachusetts, USA

⁹⁸Also at Imperial College, London, United Kingdom

⁹⁹Also at Institute of Nuclear Physics of the Uzbekistan Academy of Sciences, Tashkent, Uzbekistan