

LHC physics dataset for unsupervised New Physics detection at 40 MHz

Ekaterina Govorkova¹, Ema Puljak^{1,*}, Thea Aarrestad¹, Maurizio Pierini¹, Kinga Anna Woźniak^{1,‡}, and Jennifer Ngadiuba^{2,†}

¹European Organization for Nuclear Research (CERN), CH-1211 Geneva 23, Switzerland

²Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

*Corresponding author: Ema Puljak (ema.puljak@cern.ch)

†Also at California Institute of Technology, USA

‡Also at University of Vienna, Austria

ABSTRACT

In particle detectors at the Large Hadron Collider, tens of terabytes of data are produced every second from proton-proton collisions occurring at a rate of 40 megahertz. This data rate is reduced to a sustainable level by a real-time event filter processing system which decides whether each collision event should be kept for further analysis or be discarded. We introduce a dataset of proton collision events that emulates a typical data stream collected by such a real-time processing system, pre-filtered by requiring the presence of at least one electron or muon. This dataset could be used to develop novel event selection strategies and assess their sensitivity to new phenomena. In particular, by publishing this dataset we intend to stimulate a community-based effort towards the design of novel algorithms for performing unsupervised New Physics detection, customized to fit the bandwidth, latency and computational resource constraints of the real-time event selection system of a typical particle detector.

Background & Summary

The proton beams of the CERN Large Hadron Collider (LHC) cross paths 40 million times per second in each of the experimental halls, each time generating a *collision event*. In each such event, multiple proton pairs collide, generating thousands of particles in the detectors located at the center of each hall. Detector sensors record the flow of emerging particles in the form of electronic signals, which globally amount to $\mathcal{O}(1 \text{ MB})$ of information. The resulting data throughput of about 40 TB/sec is too large to be recorded. This is why a typical detector, e.g., the two detectors (ATLAS¹ and CMS²) used to discover the Higgs boson^{3,4}, processes data in real-time to select a small fraction of events (about 1000/s), compatible with downstream computing resources. This filtering system, usually referred to as the *trigger*, consists of a two-stage selection, as illustrated in Fig. 1. The first selection stage, the Level-1 trigger (L1T), runs a set of algorithms deployed as logic circuits on custom electronic boards equipped with field-programmable gate arrays (FPGAs). This stage rejects more than 98% of the events, reducing the incoming data stream to 100K events/s. Due to the short time interval between two collisions (25 ns) and limited buffer capabilities, the entire pipeline of L1T algorithms has to be executed within $\mathcal{O}(1) \mu\text{s}$. The second stage, called the High-Level Trigger (HLT), consists of a computer farm processing events on commercial CPUs, running hundreds of complex selection algorithms within $\mathcal{O}(100) \text{ ms}$. The trigger selection algorithms are designed to guarantee a high acceptance rate for the physics processes under study. When designing searches for unobserved physics phenomena, one typically considers specific theory-motivated scenarios. This *supervised* strategy was proven successful when dealing with strongly theoretically motivated searches, e.g., the Higgs boson discovery^{3,4}. However, this approach might become a limiting factor in the absence of strong theoretical guidance. The ATLAS and CMS trigger systems could be discarding interesting events, jeopardizing the possibility of new physics discoveries.

Therefore, many recent works investigated *unsupervised* and *semi-supervised* approaches to data selection and analysis, focusing on Anomaly Detection (AD) strategies based on Deep Learning (DL) algorithms. These studies aim to learn a metric directly from the LHC data, with capability of ranking the events by typicality. One could identify the outliers of this typicality metric distribution as a subset of the collected data potentially enriched with rare (possibly unobserved) physics processes. Extensive reviews of several proposed methods are given in Refs.^{5,6} and references therein.

This effort, mainly targeting offline data analysis, should be paired with a similar effort to integrate AD algorithms in the trigger system of the LHC experiments, possibly already at the L1T. There, it would be possible to present an unbiased dataset to the AD algorithm, before discarding any event^{7,8}. One could then collect rare event topologies in a special data stream, similar to what was done at CMS with the *exotica hotline*^{9,10} during the first year of data taking at the LHC. By studying




Figure 1. The real-time data processing flow of the ATLAS and CMS experiments: 40M collisions are produced every second and processed by the hardware-based event selection system, consisting of algorithms implemented as logic circuits on custom electronic boards. Of these events, 100K are accepted and passed to the second selection stage, the HLT, which selects about 1000/sec for offline physics studies.

these events, one could formulate new theoretical models of new physics phenomena, that could be tested in future data-taking campaigns.

While the focus so far has been on the HLT, this strategy would be more effective if deployed in the LIT, before any selection bias is introduced. Each LIT event has to be processed within a few microseconds and therefore the trigger decision is taken by algorithms hard-coded in the hardware electronics as logic circuits. Algorithm complexity, and possibly accuracy, could be enhanced by deploying DL algorithms in the LIT FPGAs. To do so, the `hls4ml` library^{11–13} was introduced as a tool to translate a given DL model into an electronic circuit. By integrating the full network on the FPGA, `hls4ml` privileges inference speed and is ideal for small networks with $\mathcal{O}(100\text{ ns})$ latency.

Given this, all the ingredients are there to make it possible for the community to design an optimal AD strategy for the LIT system. In order to stimulate this effort as a community-based initiative, similarly to what was done with the LHC Olympics⁵ and the Dark Machine data challenge⁶, we present a new dataset designed to mimic a typical data stream from a LIT system.

Methods

Physics content of the dataset

The proton-proton collisions taking place at the LHC can lead to the production and observation of many different processes predicted by the Standard Model of particle physics^{14–16}. A brief summary of the SM particle content can be found in Refs.^{17,18}. The rate at which each of these processes occur can be calculated within the SM mathematical framework and then validated by the measurements performed by the experiments¹⁹. In this paper, we focus on events containing electrons (e) and muons (μ), light particles that, together with taus (τ) and their neutrino partners, form the three lepton families. In principle, we could have considered a dataset with no filter. While this would certainly be a more realistic representation of an unbiased LIT stream, generating such a dataset requires computing resources beyond our capabilities. Instead, we decided to use the lepton filter and make the dataset simulation tractable.

Within the limited size of a typical LHC detector, electrons and muons are stable particles, i.e., they do not decay in the detector and they are directly observed while crossing the detector material. Instead, τ leptons are much heavier than electrons and muons. They quickly decay into other particles. In a fraction of these decays, e and μ are produced. At the LHC, the most abundant source of high-energy leptons is the production of W and Z bosons²⁰ which are among the heaviest SM particles with a mass of ~ 80 and ~ 90 GeV, respectively. Once produced, they quickly decay into other particles, notably leptons. W and Z bosons are mainly produced directly in proton collisions. A sizable fraction of W bosons originate from the decay of top quarks (t) and anti-quarks (\bar{t}). Being the top quark heavy and highly unstable, it quickly decays into a W boson and a bottom quark, giving rise to signatures with only jets or with one e , μ , or τ , a neutrino and a jet. Leptons can originate from more rare W and Z production, such as from the decay of Higgs bosons or multi-boson production. Given the small production probability of these processes, we ignore them in this study.

An important source of leptons comes from the production of light quarks (up, down, charm, strange, and bottom) and gluons, predicted by Quantum Chromodynamics (QCD)²¹. As these quarks and gluons have a net colour charge and cannot exist freely due to colour-confinement, they are not directly observed. Instead, they come together to form colour-neutral hadrons, a process called hadronisation that leads to a collimated spray of hadrons called a *jet*. A jet is usually defined by the

algorithm used to cluster such spray of particles, e.g. the anti- k_T algorithm²². Rarely, leptons can be produced inside jets, typically from the decay of unstable hadrons. On the other hand, QCD multijet production is by far the most abundant process at the LHC, thereby this contribution is sizable and taken into account.

Dataset

The processes listed above are the main contributors to an e or μ data stream, i.e., the set of collision events selected for including an e or μ with energy above a defined threshold. One of the datasets presented in this paper consist of the simulation of such a stream. In addition, benchmark examples of new lepton-production processes are given. These processes consist of the production of postulated, but still unobserved particles. They serve as examples of data anomalies that could be used to validate the performance of an AD algorithm. Details on these processes can be found in Refs.^{7,8}.

Moreover, we published a *blackbox* dataset containing a mixture of SM processes and a *secret* signal process. Events in this dataset are uniquely labeled by an event number, which allows us to link each event to the corresponding *ground-truth* dataset, containing a set of 0 (for SM events) and 1 (for new physics events) bits. The ground truth dataset is stored on a private cloud storage area at CERN. Not publishing this dataset, we can assure that the distributed *blackbox* is unlabeled for external developers. It is intended to be used to independently validate the performance of AD algorithms developed based on this dataset. We consider the following Standard Model (SM) processes, with their relative contribution to the dataset listed in parenthesis:

- Inclusive W boson production, where the W boson decays to a charged lepton (ℓ) and a neutrino (ν), (59.2% of the dataset). The lepton could be a e , μ , or τ lepton.
- Inclusive Z boson production, with $Z \rightarrow \ell\ell$ ($\ell = e, \mu, \tau$) (6.7% of the dataset),
- $t\bar{t}$ production (0.3% of the dataset), and
- QCD multijet production (33.8% of the dataset).

These four samples are mixed to form a realistic data stream populated by known SM processes (collectively referred to as *background*), and is provided in Ref.²³. An AD algorithm can thus be trained on this sample to learn the underlying structure of the background to identify a new physics signature (the *signal*) as an outlier in the distribution of the learned metrics.

To study the performance of AD algorithms, four signal datasets are provided:

- A leptoquark (LQ) with mass 80 GeV, decaying to a b quark and a τ lepton²⁴,
- A neutral scalar boson (A) with mass 50 GeV, decaying to two off-shell Z bosons, each forced to decay to two leptons: $A \rightarrow 4\ell$ ²⁵,
- A scalar boson with mass 60 GeV, decaying to two tau leptons: $h^0 \rightarrow \tau\tau$ ²⁶,
- A charged scalar boson with mass 60 GeV, decaying to a tau lepton and a neutrino: $h^\pm \rightarrow \tau\nu$ ²⁷.

These samples are generated using the same code and workflow as for the SM events.

We use a Cartesian coordinate system with the z axis oriented along the beam axis, the x axis on the horizontal plane, and the y axis oriented upward as shown in Fig. 2. The x and y axes define the transverse plane, while the z axis identifies the longitudinal direction. Azimuth angle ϕ is computed with respect to the x axis. Its value is given in radians, in the $[-\pi, \pi]$ range. The polar angle θ is used to compute the pseudorapidity $\eta = -\log(\tan(\theta/2))$. The transverse momentum (p_T) is the projection of the particle momentum on the (x, y) plane. We fix units such that $c = h/2\pi = 1$, where c is the speed of light and h is the Planck constant. Charge conjugation is implicit.

Each event is represented by a list of four-momenta for high-level reconstructed objects: muons, electrons, and jets. In order to emulate the limited bandwidth of a typical LIT system, we consider only the first 4 muons, 4 electrons, and 10 jets in the event, selected after ordering the candidates by decreasing p_T . If an event contains less particles, the event is zero-padded to preserve the size of the input, as it is done in realistic LIT systems. Each particle is represented by its p_T , η , and ϕ values. In addition, we consider the absolute value and ϕ coordinates of the missing transverse energy, defined as the vector equal and opposite to the vectorial sum of the transverse momenta of all the reconstructed particles in the event.

Once generated, events are filtered using a custom selection algorithm, coded in Python. We require a reconstructed electron or a muon with $p_T > 23$ GeV within $|\eta| < 3$ and $|\eta| < 2.1$, respectively. Up to ten jets with $p_T > 15$ GeV within $|\eta| < 4$ are include in the event, together with up to four muons with $|\eta| < 2.1$ and $p_T > 3$ GeV, up to four electrons with $|\eta| < 3$ and $p_T > 3$ GeV, and the missing transverse energy (MET) defined as the vector equal and opposite to the vectorial




Figure 2. The reference system used to describe the momentum coordinates of the particles in the dataset.

sum of the momenta of all particles reconstructed in the event. Given these requirements, the four SM processes listed above provide a realistic approximation of a LIT data stream.

In total, the SM cocktail dataset consists of 8,209,492 events, of which 4 millions are used to define the training dataset²³. The rest is mixed to events from the secret new physics process, to generate the *blackbox* dataset²⁸. Together with the particle momenta, this dataset includes the event numbers needed to match each event to its ground-truth bit.

Data Records

The publication consists of six data records: one record containing the mixture of SM processes, four separate records for each of the BSM processes listed above and one record containing the *blackbox* data. These are listed in Table 1, together with the total number of events and whether the record is considered to be background, signal or a mixture of the two.

The datasets are revised versions of those utilized in Refs.^{7,8} and already published on Zenodo^{29–36}. They differ from each other on the data format, and the inclusion of a *blackbox* dataset containing secret new physics events. The benchmark new physics datasets have the same physics content as the original datasets, but the included events were generated specifically for this paper. The data records are published on Zenodo^{23–28}.

Sample name	Number of samples	Type
SM processes ²³	4,000,000	B
$LQ \rightarrow b\tau$ ²⁴	340,544	S
$A \rightarrow 4\ell$ ²⁵	55,969	S
$h^0 \rightarrow \tau\tau$ ²⁶	691,283	S
$h^\pm \rightarrow \tau\nu$ ²⁷	760,272	S
<i>blackbox</i> ²⁸	4,210,492	S+B

Table 1. The names and corresponding Zenodo reference for each data set, the total number of collision events and the dataset type (S for signal and B for background).

The data records are provided in Hierarchical Data Format version 5 (HDF5), and contains 3 datasets: "Particles", "Particles_Classes", "Particles_Names". The "Particles" dataset has a shape (N, 19, 4), where N is the number of events listed for each sample in Table 1. The second index runs over the different physics objects in the events: MET, 4 electrons, 4 muons, 10 jets. Its cardinality (19) is the maximum number of objects per event. If less objects are present, the event is zero padded in such a way that the 1st, 5th, and 9th positions correspond to the highest- p_T electron, muon, and jet, respectively. The last index (with cardinality 4) runs over the three features describing each physics object and a particle type index, which is equal to 1, 2, 3 and 4 for MET, electron, muon and jet correspondingly. Zero padding is done inclusively, e.g. for zero-padded particles particle type index is set to zero. The features are ordered as described in the "Particles_Names" dataset: p_T, η, ϕ . The *blackbox* sample includes an additional dataset ("EvtId") with dimension (N), containing an event ID which allows us to match each event to its ground truth (signal or background).

Technical Validation

The distributions of the features for the SM processes and for the chosen BSM models are shown in Fig. 3. All expected features are observed, e.g., the detector ϕ symmetry, the detection inefficiency in η in the transition regions between detector components, and the different p_T distributions for the different processes.

Usage Notes

The dataset publication is accompanied by a Python-based software package, providing examples on how to read the data, train an AD algorithm with them, apply it to the *blackbox* dataset and publish the result (the unique identifier of the 1000 most anomalous events according to some given AD metric)³⁷. We aim at challenging the interested data scientists within and outside the LHC community to design LIT-friendly AD algorithms, and to submit their list of 1000 most anomalous events to a specific GitHub repository³⁸. The final goal is to document these algorithms in a dedicated publication and to preserve the benchmark data for future studies.

Code availability

Data are generated using PYTHIA 8.240³⁹, setting the collision energy at 13 TeV. Unless otherwise specified, all parameters were fixed to their default values.

We set the beam parameters to produce proton-proton collisions at 13 TeV

```
Beams:idA = 2212           ! first beam, p = 2212, pbar = -2212 \\
Beams:idB = 2212           ! second beam, p = 2212, pbar = -2212 \\
Beams:eCM = 13000.         ! CM energy of collision \\
```

while the rest of the card is configured specifically for each process, as indicated in the PYTHIA manual³⁹. For example, $W \rightarrow \ell\nu$ decays are generating setting:

```
WeakSingleBoson:ffbar2W = on      ! switch on W production mode
24::onMode = off                  ! switch off any W decay
24:onIfAny = 11 13 15             ! switch on W-> lv decays.
```

The signal-specific parameters for the four benchmark signal models are set as follows:

- For $A \rightarrow 4\ell$: set the Higgs mass to 50 GeV, force the decay to Z^*Z^* final states, and force $Z^* \rightarrow \ell\ell$ decays ($\ell = e, \mu, \tau$).
- For $LQ \rightarrow b\tau$: set the LQ mass to 80 GeV and force its decays to a b quark and a τ lepton.
- For $h^0 \rightarrow \tau\tau$: set the Higgs boson mass to 60 GeV and switch off any decay mode other than $\tau\tau$.
- For $h^+ \rightarrow \tau\nu$: set the charged Higgs boson mass to 60 GeV and switch off any decay mode other than $\tau\nu$.

We emulate detector response with DELPHES 3.3.2⁴⁰, using the default Phase-II CMS detector card. For simplicity, we avoid degrading the detector resolution to account for the coarser nature of LIT event reconstruction. This simplification does not affect the aim of the study, which is not focused on assessing the absolute physics performance but instead on comparing different algorithms and their resource consumption. We include the effect of parasitic proton collisions, sampling the number of collisions according to a Poisson distribution centered at 20. The Delphes outcome is processed by a custom Python macro to store the aforementioned physics content on HDF5 files, which are then published.

References

1. Aad, G. *et al.* The ATLAS experiment at the CERN large hadron collider. *J. Instrumentation* **3**, S08003–S08003 (2008).
2. Chatrchyan, S. *et al.* The CMS Experiment at the CERN LHC. *JINST* **3**, S08004, [10.1088/1748-0221/3/08/S08004](https://arxiv.org/abs/10.1088/1748-0221/3/08/S08004) (2008).
3. Chatrchyan, S. *et al.* Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC. *Phys. Lett. B* **716**, 30–61, [10.1016/j.physletb.2012.08.021](https://arxiv.org/abs/10.1016/j.physletb.2012.08.021) (2012). [1207.7235](https://arxiv.org/abs/1207.7235).
4. Aad, G. *et al.* Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Phys. Lett. B* **716**, 1–29 (2012).
5. Kasieczka, G. *et al.* The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics. Tech. Rep. (2021). [2101.08320](https://arxiv.org/abs/2101.08320).

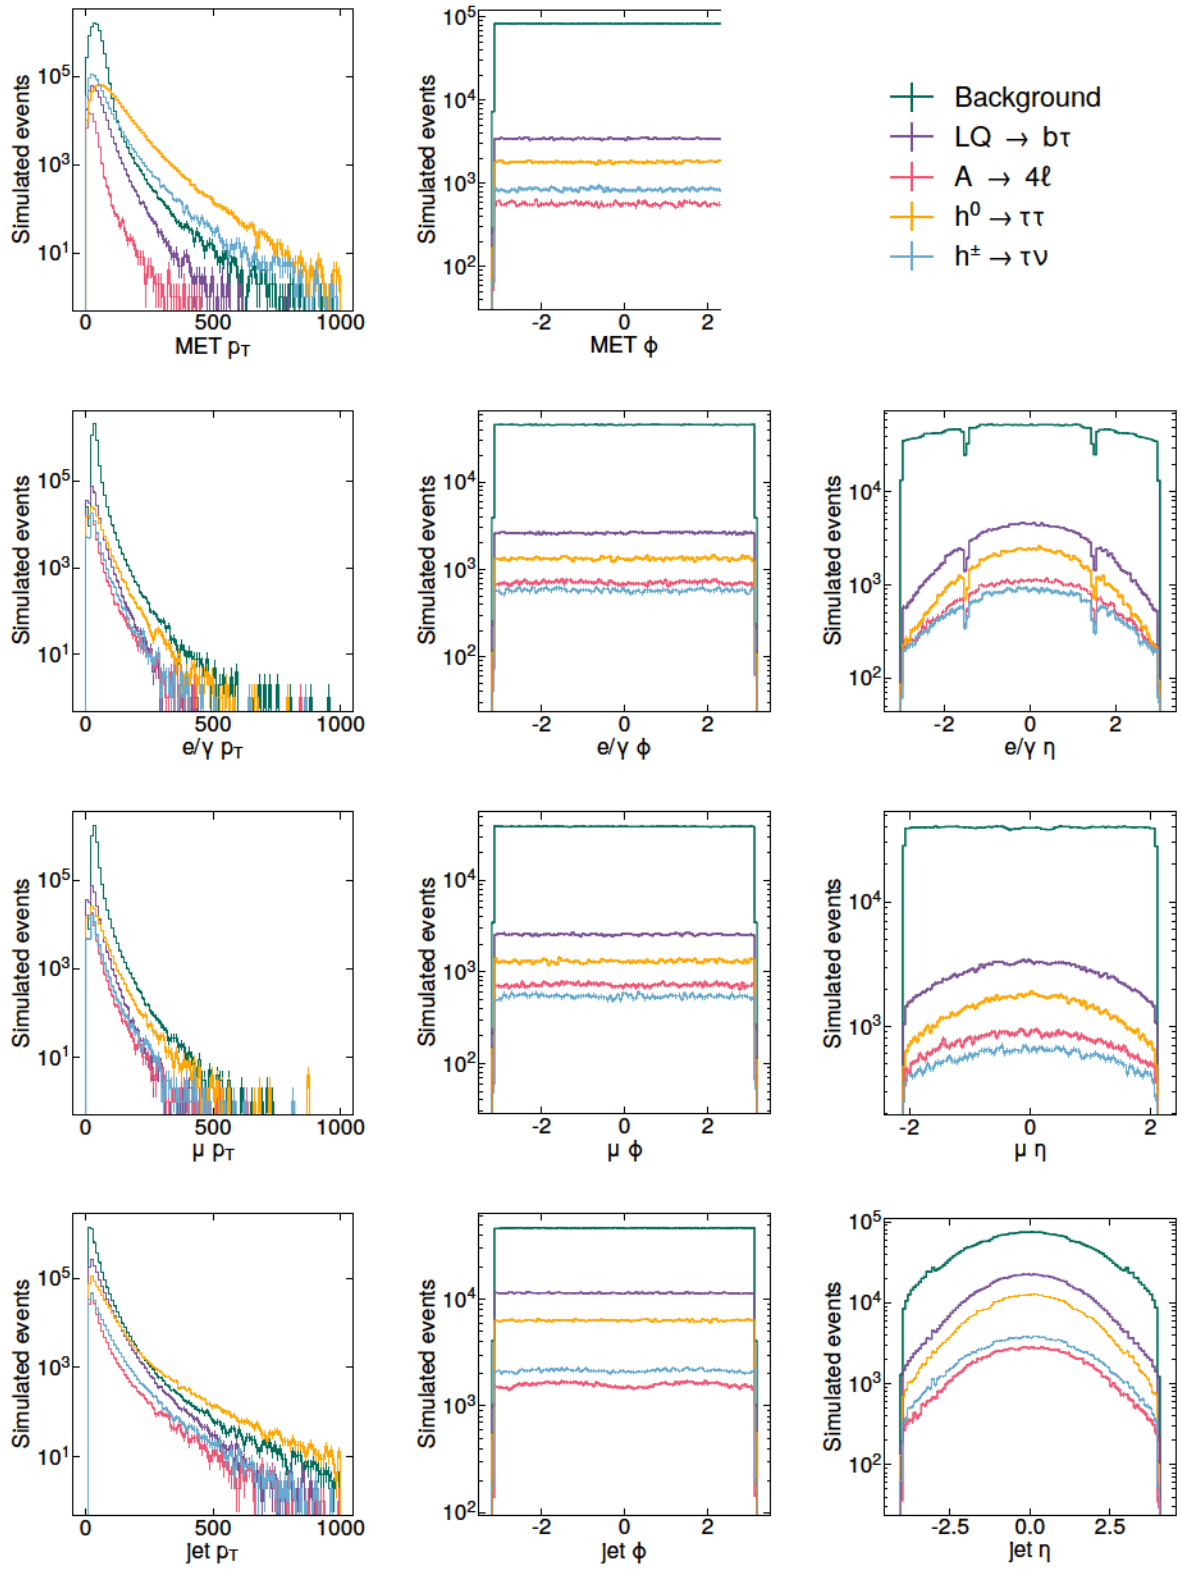


Figure 3. Distribution of the p_T (left), ϕ (center) and η (right) coordinates of the physics objects entering the dataset, for missing transverse energy, MET (top row), electrons (second row), muons (third row) and jets (bottom row).

6. Ostdiek, B. *et al.* The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider. Tech. Rep. (2021). [2105.14027](#).
7. Cerri, O., Nguyen, T. Q., Pierini, M., Spiropulu, M. & Vlimant, J.-R. Variational Autoencoders for New Physics Mining at the Large Hadron Collider. *JHEP* **05**, 036, [10.1007/JHEP05\(2019\)036](#) (2019). [1811.10276](#).
8. Knapp, O. *et al.* Adversarially Learned Anomaly Detection on CMS Open Data: re-discovering the top quark. *Eur. Phys. J. Plus* **136**, 236, [10.1140/epjp/s13360-021-01109-4](#) (2021). [2005.01598](#).
9. CMS Exotica hotline leads hunt for exotic particles. *Symmetry Magazine*, 2010.
10. Poppi, F. Is the bell ringing? *CERN Bull.* **BUL-NA-2010-317**, 14 (2010).
11. Duarte, J. *et al.* Fast inference of deep neural networks in FPGAs for particle physics. *JINST* **13**, P07027, [10.1088/1748-0221/13/07/P07027](#) (2018). [1804.06913](#).
12. Aarrestad, T. *et al.* Fast convolutional neural networks on fpgas with hls4ml. Tech. Rep. (2021). [2101.05108](#).
13. Coelho, C. N. *et al.* Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors. *Nat. Mach. Intell.* [10.1038/s42256-021-00356-5](#) (2021).
14. Glashow, S. L. Partial-symmetries of weak interactions. *Nucl. Phys.* **22**, 579–588, [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2) (1961).
15. Weinberg, S. A model of leptons. *Phys. Rev. Lett.* **19**, 1264–1266, [10.1103/PhysRevLett.19.1264](#) (1967).
16. Nobelstiftelsen. Elementary particle theory. : Relativistic groups and analyticity. proceedings of the eighth nobel symposium held may 19-25, 1968 at aspenäsgarden, lerum, in the county of Älvsborg, sweden / edited by nils svartholm. Tech. Rep., Stockholm : New York ; London (1968).
17. Wolchover, N. A new map of all the particles and forces. *Quanta Mag.* (2020).
18. Quigg, C. The Double simplex. In *GUSTAVOFEST: Symposium in Honor of Gustavo C. Branco: CP Violation and the Flavor Puzzle* (2005). [hep-ph/0509037](#).
19. Summary of the cross section measurements of standard model processes (2021).
20. Zyla, P. *et al.* Review of Particle Physics. *PTEP* **2020**, 083C01, [10.1093/ptep/ptaa104](#) (2020).
21. Ellis, R. K., Stirling, W. J. & Webber, B. R. *QCD and Collider Physics*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology (Cambridge University Press, 1996).
22. Cacciari, M., Salam, G. P. & Soyez, G. The anti- k_t jet clustering algorithm. *JHEP* **04**, 063, [10.1088/1126-6708/2008/04/063](#) (2008). [0802.1189](#).
23. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: Training dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5046389>, [10.5281/zenodo.5046389](#) (2021).
24. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: LQ \rightarrow b τ signal benchmark dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5055454>, [10.5281/zenodo.5055454](#) (2021).
25. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: A \rightarrow 4 leptons signal benchmark dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5046446>, [10.5281/zenodo.5046446](#) (2021).
26. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: $h^0 \rightarrow \tau\tau$ signal benchmark dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5061633>, [10.5281/zenodo.5061633](#) (2021).
27. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: $h^+ \rightarrow \tau\nu$ signal benchmark dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5061688>, [10.5281/zenodo.5061688](#) (2021).
28. Govorkova, E. *et al.* Unsupervised new physics detection at 40 mhz: Black box dataset. *Zenodo* <https://doi.org/10.5281/zenodo.5070455>, [10.5281/zenodo.5070455](#) (2021).
29. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $w \rightarrow \ell\nu$, [10.5281/zenodo.3675199](#) (2020).
30. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $z \rightarrow \ell\ell$, [10.5281/zenodo.3675203](#) (2020).
31. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: top pair production, [10.5281/zenodo.3675206](#) (2020).

32. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: Qcd multijet production, [10.5281/zenodo.3675210](https://doi.org/10.5281/zenodo.3675210) (2020).
33. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $Lq \rightarrow b\tau$, [10.5281/zenodo.3675196](https://doi.org/10.5281/zenodo.3675196) (2020).
34. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $a \rightarrow 4$ leptons, [10.5281/zenodo.3675159](https://doi.org/10.5281/zenodo.3675159) (2020).
35. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $h^0 \rightarrow \tau\tau$, [10.5281/zenodo.3675190](https://doi.org/10.5281/zenodo.3675190) (2020).
36. Cerri, O., Nguyen, T., Pierini, M. & Vlimant, J.-R. New physics mining at the large hadron collider: $h^\pm \rightarrow \tau\nu$, [10.5281/zenodo.3675178](https://doi.org/10.5281/zenodo.3675178) (2020).
37. Govorkova, E. *et al.* ADC2021 example code. *GitHub* <https://github.com/mpp-hep/ADC2021-examplecode> (2021).
38. Govorkova, E. *et al.* ADC2021 result repository. *GitHub* <https://github.com/mpp-hep/ADC2021-results> (2021).
39. Sjöstrand, T. *et al.* An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159–177, [10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024) (2015). [1410.3012](https://arxiv.org/abs/1410.3012).
40. de Favereau, J. *et al.* DELPHES 3, A modular framework for fast simulation of a generic collider experiment. *JHEP* **02**, 057, [10.1007/JHEP02\(2014\)057](https://doi.org/10.1007/JHEP02(2014)057) (2014). [1307.6346](https://arxiv.org/abs/1307.6346).

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 772369) and the ERC-POC programme (grant No. 996696).

Author contributions statement

J.N. conceived the idea of publishing the dataset and creating a data challenge on it; M.P. created the data in raw format; E.P. and E.G. applied the event selection and produced the dataset in its final format; T.A., J.N. and K.W. conceived the package with example code; E.P. designed the example autoencoder; all drafted the paper.

Competing interests

The authors declare no competing interests.