

---

# Towards an Interpretable Data-driven Trigger System for High-throughput Physics Facilities

---

Chinmaya Mahesh<sup>†</sup> Kristin Dona\* David W. Miller\* Yuxin Chen\*

<sup>†</sup> University of Illinois Urbana-Champaign, cmahesh2@illinois.edu

\*University of Chicago, {kdona, davemilr, chenYuxin}@uchicago.edu

## Abstract

Data-intensive science is increasingly reliant on real-time processing capabilities and machine learning workflows, in order to filter and analyze the extreme volumes of data being collected. This is especially true at the energy and intensity frontiers of particle physics where bandwidths of raw data can exceed 100 Tb/s of heterogeneous, high-dimensional data sourced from hundreds of millions of individual sensors. In this paper, we introduce a new data-driven approach for designing and optimizing high-throughput data filtering and trigger systems such as those in use at physics facilities like the Large Hadron Collider (LHC). Concretely, our goal is to design a data-driven filtering system with a minimal run-time cost for determining which data event to keep, while preserving (and potentially improving upon) the distribution of the output as generated by the hand-designed trigger system. We introduce key insights from interpretable predictive modeling and cost-sensitive learning in order to account for non-local inefficiencies in the current paradigm and construct a cost-effective data filtering and trigger model that does not compromise physics coverage.

## 1 Introduction and Background

Data filtering algorithms — or *trigger algorithms* [1, 2, 3] — targeted at discovery science (e.g. identification of a data event containing evidence of dark matter produced at the Large Hadron Collider), must operate at the level of 1 part in  $10^5$  due to numerous bandwidth, computation, and storage-related constraints. Once executed, these algorithms often drive the data curation process, funneling event records with certain features into categories that are predefined based on the labels extracted by those algorithms. Although discovery science has been based on this approach for decades, the design of these systems relies heavily on excellent prior knowledge of the feature space being probed, detailed simulations of hand-designed feature extraction algorithms, and heavily modularized approaches to optimization. Consequently, significant inefficiencies exist in these systems, including both redundant feature labeling schemes and cost-ineffective algorithm execution.

In this paper, we describe the design and initial results of the construction of a data-driven approach to refine the trigger and data filtering algorithms at future physics facilities. To accomplish such a daunting task, it is essential that the *explainability* of such an approach first be established. Whereas in other contexts, the explainability of a classifier decision, for example, can be related to expert variables that describe the physics of a process [4, 5, 6, 7], we use the relationship to a hand-designed set of trigger decisions in the approach described here. As illustrated in Figure 1, we may therefore construct a cost model by which the data filtering and curation process can be quantifiably assessed. This approach allows for *pruning* of costly but ineffective individual algorithm labels. In addition, progress towards a fully functional automated data filtering and processing system must start by constructing a learning-based filter system that is able to reproduce the current, hand-designed, trigger “menus.” More importantly, though, in order to be useful, this new system must be able to *explain*

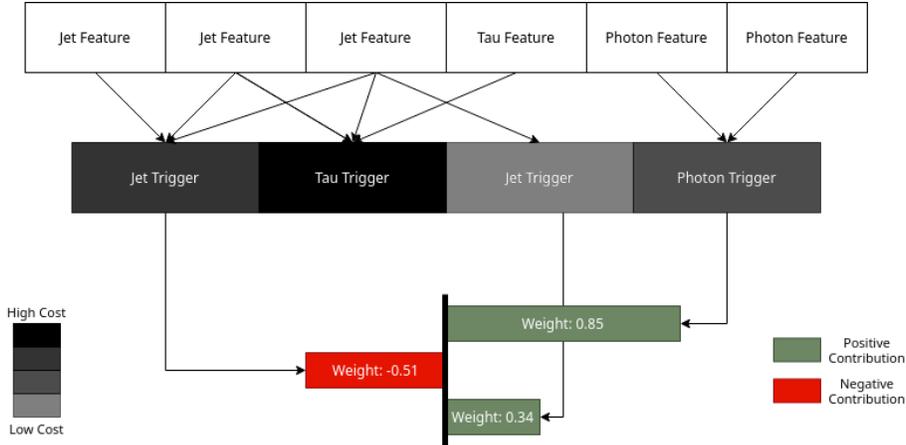


Figure 1: An example cost-effective explanation of an event. In this case, the Tau Trigger has the highest cost and thus the weight associated with the Tau Trigger was driven to 0. The remaining weights result in the final cost-effective explanation of the event, with the weights with the highest absolute value being the most important.

its selections in the context of both low-level features (e.g. sensor signals) and high-level features (e.g. physics object multiplicities and kinematics). This is the first step towards the ultimate goal of constructing an active continuous learning model that is able to update itself and provide *explanations* for those updates. A distinguishing aspect of this research effort is its focus on model *interpretability*: Instead of a closed-box model that is capable of recovering the original data distribution, we aim to design an *open-box* predictive model, which, for any given input, not only outputs a decision (e.g., “keep this data point”), but also explains why we should do so, by associating the decision with the existing rules in the hand-designed trigger menu.

In addition to the development of an open-box predictive model, we also aim to minimize the latency of the trigger system. Given an incoming data event, each trigger algorithm incurs a latency at runtime – assuming that algorithms are run in parallel, the latency of the trigger system depends on the worst-case running time of all trigger algorithms. Thus, for each data event, finding the most efficient set of trigger algorithms at run time is crucial for a real-time trigger system.

To address the real-time data processing challenge, we investigate the following combinatorial optimization problem: given a ground set of candidate trigger algorithms from the existing trigger menu and a latency cost for each trigger algorithm, we seek an optimal subset of trigger algorithms for each incoming data event, such that the selected algorithms can jointly make the correct filtering decision with the *minimal latency cost*. We then use the solution of this optimization problem as the explanation of our open-box predictive model, which in turn will be used to optimize the latency of the existing trigger system. Here, we highlight a significant challenge during this process: Since the effectiveness of any subset of trigger algorithm is often unknown *a priori* and needs to be learned from data, this task essentially amounts to building a data-driven, cost-sensitive explainable model—an emerging topic at the frontier of machine learning and operations research.

## 2 Related Work

**Interpretable models** Interpretable predictive modeling is useful for very wide application domains where black-box models are not preferred (e.g., medical / biological research, emergency response planning, etc). In these scenarios, decisions are critical and often have serious consequences, and domain experts hence would like to know how these predictions are made. Depending on different measures of *interpretability*, existing methods can be categorized into three types: *sparse models*, which use a small number of features / parameters, such as sparse linear classifiers [8, 9], (low-rank) latent factor models, etc; *discretization methods*, which split up a problem into several (independent) sub-problems, such as decision trees and decision lists [10]; and *prototype-based classifiers*, which use examples and basic features instead of formal models and constructed features, such as nearest

neighbors [11]. Recently developments in interpretable models, e.g., Local interpretable model-agnostic explanations (LIME) [12], have proposed to use local surrogate models to explain individual predictions of black box machine learning models, where each local surrogate model is interpretable by itself. Our work extends the LIME framework and can be viewed as a sparsity-based locally interpretable model, where we seek a minimal-cost explanation for the LHC trigger outputs.

**Resource-efficient machine learning** In many data-driven decision making tasks such as active learning [13, 14, 15, 16], recommender systems [17, 18], experimental design [19] and surveillance [20, 21, 22, 23], learning and prediction are often under stringent resource constraints—As an example, features are associated by costs (e.g. monetary cost for conducting a medical tests, computational cost for extracting a feature) both in training [24, 25] and test time [26, 27]. In connection to our work, a data-driven trigger system must maintain a throughput constraint to keep pace with arriving traffic of the collision event. A common challenge among these applications is the need to properly trade-off prediction accuracy against the prediction cost [28, 29, 30]. Our work aligns with contemporary models that demand cost-effectiveness at run time.

### 3 Automated Trigger Menu Design

We are given a dataset  $X \in \mathbb{R}^{n \times p}$  of physical collision data (where each of the  $n$  records are known as *events*), a set of labels  $T$  (known as *triggers*), and an outcome matrix  $y = \{0, 1\}^{n \times |T|}$  which denotes the triggers each event satisfies. Let each event in  $X$  be described by numerical values of a set of  $p$  features  $F = \{f_1, f_2, \dots, f_p\}$ . We are also given a mapping  $c : F \rightarrow \mathbb{R}^+$ , denoting the *cost function*, that is,  $c(f_i)$  is the cost of using feature  $f_i$  to predict the outcome of an event. The goal is to identify the most *cost-efficient* subset  $F' \subset F$  that enables us to maximize coverage of  $X$  in the trained model while using only the features in  $F'$  to make predictions.

#### 3.1 LIME with Elastic Net

We utilize the LIME framework [12], which creates an interpretable prediction for an event  $x$  by training a sparse model with a dataset of perturbations of  $x$ . Using the trained weight vector  $\hat{\beta}$  of this sparse model we are able to understand the importance of each feature in the final prediction, as a higher value of  $|\hat{\beta}_i|$  indicates higher influence of feature  $f_i$  in the final prediction. In order to obtain an interpretable yet accurate prediction of the outcome of an event, we consider to apply a sparsity regularizer to the local linear model being trained. As discussed in §2, interpretability is achieved because the non-zero entries of a sparse weight vector of a linear model can be inspected manually. Note that sparse models might be far less accurate, and how to strike a good balance between performance and interpretability remains an open problem in many practical applications. In this work, we adopt *elastic net* [31] as a general formulation that trades off model interpretability (sparsity) and accuracy. Formally, given a dataset  $X \in \mathbb{R}^{n \times p}$ , an outcome vector  $y \in \mathbb{R}^n$ , and a tunable parameter  $\lambda > 0$ , the final  $\hat{\beta}$  obtained can be defined as

$$\hat{\beta} = \arg \min_{\beta} \left( |y - X\beta|^2 + (1 - \alpha)\lambda \sum_{i=1}^p |\beta_i| + \alpha\lambda \sum_{i=1}^p |\beta_i|^2 \right).$$

The LASSO and ridge regressions are special cases of the elastic net where  $\alpha = 0$  and  $\alpha = 1$  respectively.

#### 3.2 Cost-Effective Elastic Net

To obtain a value of  $\hat{\beta}$  which is both sparse and cost efficient, we propose adding a coefficient of  $c(f_i)$ , which is the cost of feature  $i$ , to each respective term  $|\beta_i|$  and  $|\beta_i|^2$  in the elastic net penalty:

$$\hat{\beta} = \arg \min_{\beta} \left( |y - X\beta|^2 + (1 - \alpha)\lambda \sum_{i=1}^p |\beta_i| \cdot c(f_i) + \alpha\lambda \sum_{i=1}^p |\beta_i|^2 \cdot c(f_i) \right)$$

As we have the restriction  $\forall f_i \in F, c(f_i) \in \mathbb{R}^+$ , the above modification to the elastic net (henceforth referred to as the *cost-effective elastic net*), maintains all the convexity properties of the elastic net. Thus by using LIME with the cost-effective elastic net as the sparse model, we have modified LIME to create interpretable explanations which additionally have the property of being cost-effective.

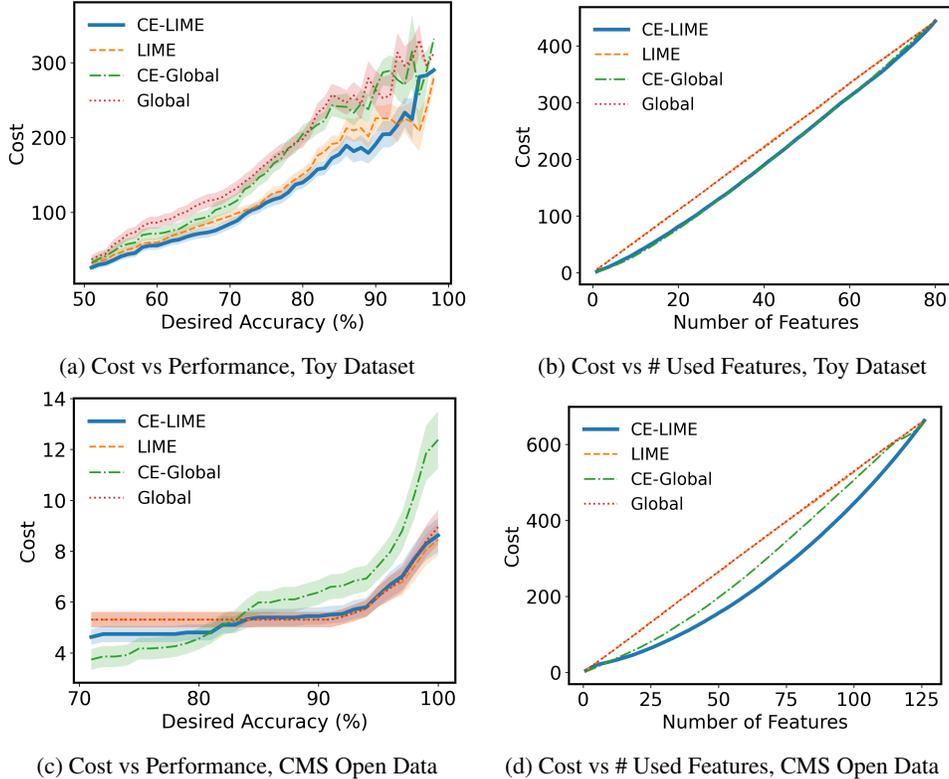


Figure 2: Experimental results on synthetic data (top) and CMS Open Data (bottom). We compare our proposed algorithm cost-effective LIME (CE-LIME) with three baseline approaches, namely LIME with submodular pick (LIME), a sparse regressor via elastic net regression (Global), and a cost-effective elastic net (CE-Global). Error bars represent standard error over all trials.

### 3.3 LIME with Submodular Pick

With the above approach, we are able to obtain information about the features which are the most important while being cost-effective. However, a larger model-wide explanation to analyse the scenarios in which certain features are determined be both important to the prediction and cost-effective is desired. To achieve this, a modified version of SP-LIME [12] is used. While calculating the set of representative explanations for the whole model, SP-LIME creates an importance vector  $I$  such that if  $I_i > I_j$  for any  $i$  and  $j$ , then feature  $i$  is considered to be more important than feature  $j$ . This gives us a total ordering of  $F$  which enables us to select an optimal subset of  $F$  containing the most important and cost-effective features for any arbitrary cutoff. We call this method of using a modified SP-LIME with a cost-effective elastic net CE-LIME.

## 4 Experiments

**Toy Dataset in a Controlled Environment** In order to empirically test the effectiveness of the cost-effective elastic net, 100 trials were run on randomly generated classification problems created using the `make_classification` feature of scikit-learn [32]. In each trial, a matrix of predictor variables  $X \in \mathbb{R}^{1000 \times 80}$  and an outcome vector  $y \in \{0, 1\}^{1000}$  were created such that 20 of the features were indicative, 20 were linear combinations of the indicative features, and 20 were duplicated features sampled from the previous two groups of features, in order to create the usual scenario of a high dimensional dataset with several redundant or correlated features. Furthermore, a cost function  $c$  was created with a uniformly randomized distribution in the interval  $[0, 10]$ . For any event  $x_{\text{new}}$  to be predicted, initially the value of every feature  $f_i$  to be used by the classifier is set to the mean of the feature's value in  $X$ . In this way, a higher accuracy level can be obtained by iteratively adding the value of more features in accordance to the total feature order  $I$ .

In Fig. 2a, we see that CE-LIME can achieve the same accuracy level as LIME with submodular pick (LIME), a sparse regressor via elastic net regression (Global), and a cost-effective elastic net (CE-Global) with an overall lower cost. In Fig. 2b, we observe a dip of the required cost at certain accuracy levels due to the decrease in probability of classifiers obtaining high accuracy levels. This leads to only the surviving models being counted, which tend to have lower overall cost requirements due to the simpler nature of their respective classification problems.



Figure 3: Trigger label overlap amongst (a) all labels and (b) groups of related triggers separated into categories.

**Case Study on CMS Open Data** To demonstrate the potential value of these optimization methods at the LHC, we use publicly available datasets from the CMS Experiment at the LHC, available in the public CERN Open Data Portal [33]. The Portal represents an exciting opportunity for the general public to access LHC data for studies such as the trigger optimization presented in this paper. In particular, we use a simulated dataset of  $t\bar{t}$  events [34], generated for collisions at a center-of-mass energy of 8 TeV and reconstructed to provide information on the physics objects such as electrons, photons, muons, taus, and jets, as well as their respective masses, transverse momenta, and pseudorapidity. The events also include Boolean values for over 400 different triggers that we have analyzed using the physics information and trigger decisions.

Similar to the experimental setup of the toy dataset, 10 trials were run on 9 different triggers with randomized cost of features in every trial, with the costs being uniformly distributed in  $[0, 10]$ . Figure 3 is generated with the CMS Open Data and shows the fractional overlap between features which share trigger labels and trigger label categories. The large fractional overlap emphasises the potential for these algorithms to be optimized. Fig. 2c shows that the required cost to achieve any accuracy level is the lowest for CE-LIME for most cases, and Figure 2d shows that the cost of using a certain number of features in the order given by CE-LIME is competitive with the globally trained cost-effective elastic net (CE-GLOBAL).

## Acknowledgment

The authors thank Dr. Cecilia Toscani for the helpful discussion. The authors also thank the CMS Open Data Workshop for Theorists for assistance with accessing the CMS public data<sup>1</sup>. This work was done when Chinmaya Mahesh was at the Center for Data and Computing (CDAC) Data & Computing Summer Lab at UChicago. The project was supported by a CDAC Discovery Grant.

<sup>1</sup><https://indico.cern.ch/event/882586/>

## Broader Impact

The broader impact of this work derives both from the nexus of novel ideas and the application to which we've applied those ideas. Interpretable predictive modeling combined with the approaches to cost-sensitive learning represent a new direction that aims to introduce a quantifiable metric by which to optimize black-box systems, and then to provide a mechanism for "opening" that box. In addition, the novel application of these ideas to the challenging real-time data filtering and processing systems that are in use at large-scale physics facilities offer the potential for a paradigm shift in the operational models, and perhaps scientific capabilities, of those facilities. Future work aims to explore the use of active learning models that capitalize on the interpretability and cost-sensitive optimization discussed in this work to provide for a continuous learning model that is able to "explain" what it has learned in terms of both cost and performance.

## References

- [1] ATLAS Collaboration. ATLAS high-level trigger, data-acquisition and controls: Technical Design Report. CERN-LHCC-2003-022, 2003.
- [2] CMS Collaboration. CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems. CERN-LHCC-2000-038, 2000.
- [3] CMS Collaboration. CMS The TriDAS Project: Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger. CMS trigger and data-acquisition project. CERN-LHCC-2002-026, 2002.
- [4] Daniel Guest, Julian Collado, Pierre Baldi, Shih-Chieh Hsu, Gregor Urban, and Daniel Whiteson. Jet flavor classification in high-energy physics with deep neural networks. *Physical Review D*, 94(11), Dec 2016.
- [5] Taylor Faucett, Jesse Thaler, and Daniel Whiteson. Mapping machine-learned physics into a human-readable space. 2020.
- [6] Julian Collado, Jessica N. Howard, Taylor Faucett, Tony Tong, Pierre Baldi, and Daniel Whiteson. Learning to identify electrons. 2020.
- [7] Garvita Agarwal, Lauren Hay, Ia Iashvili, Benjamin Mannix, Christine McLean, Margaret Morris, Salvatore Rappoccio, and Ulrich Schubert. Explainable ai for ml jet taggers using expert variables and layerwise relevance propagation. 2020.
- [8] Jonathan Chang, Sean Gerrish, Chong Wang, Jordan L Boyd-Graber, and David M Blei. Reading tea leaves: How humans interpret topic models. In *Advances in neural information processing systems*, pages 288–296, 2009.
- [9] Berk Ustun and Cynthia Rudin. Methods and models for interpretable linear classification. *arXiv preprint arXiv:1405.4047*, 2014.
- [10] Benjamin Letham, Cynthia Rudin, Tyler H McCormick, David Madigan, et al. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics*, 9(3):1350–1371, 2015.
- [11] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- [12] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016.
- [13] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. *Journal of machine learning research*, 2(Nov):45–66, 2001.
- [14] Simon Tong and Edward Chang. Support vector machine active learning for image retrieval. In *Proceedings of the ninth ACM international conference on Multimedia*, pages 107–118. ACM, 2001.

- [15] B. Settles. *Active Learning*. Synthesis digital library of engineering and computer science. Morgan & Claypool Publishers, 2012.
- [16] Burr Settles. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1467–1478. Association for Computational Linguistics, 2011.
- [17] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [18] Yuxin Chen, Jean-Michel Renders, Morteza Haghiri Chehreghani, and Andreas Krause. Efficient online learning for optimizing value of information: Theory and application to interactive troubleshooting. In *Proc. Conference on Uncertainty in AI (UAI)*, August 2017.
- [19] Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.
- [20] Chen Change Loy, Timothy M Hospedales, Tao Xiang, and Shaogang Gong. Stream-based joint exploration-exploitation active learning. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1560–1567, 2012.
- [21] Wei Chu, Martin Zinkevich, Lihong Li, Achint Thomas, and Belle Tseng. Unbiased online active learning in data streams. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 195–203. ACM, 2011.
- [22] Mohamed-Rafik Bouguelia, Yolande Belaid, and Abdel Belaid. A stream-based semi-supervised active learning approach for document classification. In *2013 12th International Conference on Document Analysis and Recognition*, pages 611–615. IEEE, 2013.
- [23] Yuxin Chen, Hiroaki Shioi, Cesar Fuentes Montesinos, Lian Pin Koh, Serge Wich, and Andreas Krause. Active detection via adaptive submodularity. In *Proc. International Conference on Machine Learning (ICML)*, 2014.
- [24] Huan Liu, Hiroshi Motoda, and Lei Yu. A selective sampling approach to active feature selection. *Artificial Intelligence*, 159(1-2):49–74, 2004.
- [25] Hema Raghavan, Omid Madani, and Rosie Jones. Active learning with feedback on features and instances. *Journal of Machine Learning Research*, 7(Aug):1655–1686, 2006.
- [26] Jacqueline Gottlieb. Attention, learning, and the value of information. *Neuron*, 76(2):281–295, 2012.
- [27] Kirstin Early, Stephen E Fienberg, and Jennifer Mankoff. Cost-effective feature selection and ordering for personalized energy estimates. In *AAAI Workshop: AI for Smart Grids and Smart Buildings*, 2016.
- [28] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- [29] Koby Crammer, Jaz Kandola, and Yoram Singer. Online classification on a budget. In *Advances in neural information processing systems*, pages 225–232, 2004.
- [30] Ashish Kumar, Saurabh Goyal, and Manik Varma. Resource-efficient machine learning in 2 kb ram for the internet of things. In *International Conference on Machine Learning*, pages 1935–1944, 2017.
- [31] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320, 2005.
- [32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

- [33] Aashish Tripathee, Wei Xue, Andrew Larkoski, Simone Marzani, and Jesse Thaler. Jet substructure studies with cms open data. *Phys. Rev. D*, 96:074003, Oct 2017.
- [34] CMS Collaboration. Simulated Dataset: TTbar\_8TeV-Madspin\_aMCatNLO-herwig in AOD-SIM Format for 2012 Collision Data. CERN Open Data Portal. 2017.