

Article

How to Use Machine Learning to Improve the Discrimination between Signal and Background at Particle Colliders

Xabier Cid Vidal ^{1,*} , Lorena Dieste Maroñas ^{1,2} and Álvaro Dosil Suárez ²

¹ Instituto Galego de Física de Altas Enerxías (IGFAE), Universidade de Santiago de Compostela, 15782 Santiago de Compostela, Spain; ldieste@triplealpha.in

² Triple Alpha Innovation, Cuns, 4, Outes, 15286 A Coruña, Spain; adosil@triplealpha.in

* Correspondence: xabier.cid.vidal@cern.ch

Abstract: The popularity of Machine Learning (ML) has been increasing in recent decades in almost every area, with the commercial and scientific fields being the most notorious ones. In particle physics, ML has been proven a useful resource to make the most of projects such as the Large Hadron Collider (LHC). The main advantage provided by ML is a reduction in the time and effort required for the measurements carried out by experiments, and improvements in the performance. With this work we aim to encourage scientists working with particle colliders to use ML and to try the different alternatives that are available, focusing on the separation of signal and background. We assess some of the most-used libraries in the field, such as Toolkit for Multivariate Data Analysis with ROOT, and also newer and more sophisticated options such as PyTorch and Keras. We also assess the suitability of some of the most common algorithms for signal-background discrimination, such as Boosted Decision Trees, and propose the use of others, namely Neural Networks. We compare the overall performance of different algorithms and libraries in simulated LHC data and produce some guidelines to help analysts deal with different situations. Examples include the use of low or high-level features from particle detectors or the amount of statistics that are available for training the algorithms. Our main conclusion is that the algorithms and libraries used more frequently at LHC collaborations might not always be those that provide the best results for the classification of signal candidates, and fully connected Neural Networks trained with Keras can improve the performance scores in most of the cases we formulate.

Keywords: high-energy physics; machine learning; LHC; LHCb



Citation: Cid Vida, X.; Dieste Maroñas, L.; Dosil Suárez, Á. How to Use Machine Learning to Improve the Discrimination between Signal and Background at Particle Colliders. *Appl. Sci.* **2021**, *11*, 11076. <https://doi.org/10.3390/app112211076>

Academic Editors: José D. Martín-Guerrero and Lucas Lamata

Received: 31 October 2021
Accepted: 18 November 2021
Published: 22 November 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Particle physics experiments, and especially those at particle colliders, have to deal with vast amounts of data where, very often, an elusive signal must be found against a much larger background. This has naturally paved the way for the usage of Machine Learning (ML) in the field. As in other problems where ML applies, the separation between signal and background relies on several variables (features) that behave differently in both categories. ML algorithms are able to statistically combine the separation power of different features, making the best of them all, and using the correlations between them in its favor. The resulting discrimination power is usually superior to anything a “manual” selection of requirements would achieve, and can be obtained in a more efficient way. The use of ML in particle physics is an emerging area, which is extending to more and more fields, as we shall see. A very complete (and live) review of this wide range of uses can be found in Reference [1].

To better understand the implementation of ML as a way of improving the discrimination between signal and background, we chose the LHCb experiment [2,3] at CERN as a benchmark. LHCb is one of the large collaborations in the Large Hadron Collider (LHC) project, where protons are accelerated to ultrarelativistic energies (~14 TeV) and then made

to collide against each other, allowing for the study of the smallest pieces of matter. These collisions happen at an incredibly high rate (~ 40 million times per second), giving rise to hundreds of other particles whose interaction is then recorded by detectors (such as LHCb), which act as gigantic photographic cameras that collect 90 petabytes of data per year [4]. From these vast datasets, scientists try to extract interesting and rare phenomena, such as undiscovered particles or very infrequent particle decays, whose characteristics help to understand the properties of the smallest pieces of matter.

The usual pipeline in analyses at a LHC experiment is as follows. Particles produced by collisions leave different types of electronic signals at the different components of the detectors (or subdetectors). These subdetectors are specialized in different aspects, from tracking to particle identification (PID). Their final goal is to provide as realistic a picture as possible of what happened after the collision. This picture involves knowing the energy, identity and position of the particles produced, and is usually referred to as “reconstruction”. To reconstruct each collision, different algorithms help to interpret the electronic signals grouping the particle “hits” into tracks or understanding which type of particle left those hits. Another type of object that is often reconstructed at particle colliders are jets, which are groups of particles flying very close to each other and usually originating from a single energetic particle of interest. Determining the nature of this particle from the properties of jets (known as tagging) is also a frequent task in the reconstruction of collisions. Although they are not the goal of this study, all these aspects have recently led to surge in the use of ML. Recent examples are the use of graph neural networks (aspects of some of these ML algorithms are discussed later in the paper) for PID [5], reinforcement learning for jet reconstruction [6] or deep learning for track reconstruction [7] or tagging of jets [8]. We refer one more time to Reference [1] for a more complete compendium of examples.

Once particle collisions are reconstructed, analysts are faced with a list of objects they must use to conduct measurements of interest (to perform what particle physicists denote “analyses”). Note that this reconstruction is not perfect, but instead limited by the resolution of the detectors. Although efforts can be made to improve this resolution, such as improving the calibration of detectors by means of convolutional neural networks [9], the reconstructed objects will always be different to their “original” counterparts, which limits what experiments can do with them. One of the main consequences of this is that, very often, it becomes very hard to distinguish the target signals that experiments look for from backgrounds composed of other particles with similar properties. For instance, if one looks for an exotic new particle decaying into a pair of hadrons, this could be faked by situations in which each hadron originates from a different standard particle, although they are located close to each other in space. This all fits very nicely with the use of ML algorithms, particularly with classifiers that are able to separate scarce signals hiding behind much larger amounts of background. Finding the best of these classifiers has been the subject of many studies in particle physics, using algorithms of very different kind, ranging from boosted decision trees [10] to deep learning [11]. Alternative involves unsupervised [12] or semi-supervised [13] algorithms, with the goal of optimizing searches for rare phenomena for which more traditional methods have failed to date. Despite the many efforts being conducted in this area of signal-background discrimination by means of ML, as we shall see, most analysts either still rely on more traditional approaches or do not use the latest, best-performing, tools available in terms of libraries and algorithms. Assessing this question in detail is one the main goals of this paper.

To some extent, the LHCb experiment has pioneered the application of ML at the LHC. This involves, for instance, the use of Multi-Variate Analysis (MVA) classification libraries in some of their first analyses [14] or the introduction of ML in the online trigger system of the detector [15,16], which runs automatically at every LHC proton collision. Even if LHCb is, in many ways, an example of good use of ML in particle physics, in this paper we first assess the extent to which ML is used in their analyses and then what kind of libraries and algorithms are currently being used in the experiment. We also compare these aspects to those of other experiments. We then contrast the performance of some of the most frequent

ML libraries and those used in particle physics with simulated LHCb data, to find out the extent to which there is room for improvement in the signal-background discrimination achieved in standard analyses. Note this is a novelty, since many other comparisons of this kind rely on existing High-Energy Physics (HEP) datasets designed for benchmarking, but corresponding to other topologies, kinematic ranges and experiments [17–19]. Next, we extend these tests to other libraries and algorithms. We quantify the overall margin of improvement and provide guidelines on how to use the latest libraries available in ML to improve the sensitivity of particle physics experiments. The conclusions of this study can mostly be expanded to experiments beyond LHCb. Therefore, instead of testing a specific algorithm to solve a specific problem, we try to provide a more global perspective, based on the actual popularity of algorithms and libraries in experiments and their potential performance. In summary, our main goals are:

- Finding out the popularity of different algorithms and libraries at LHC experiments.
- Determining whether the most popular methods are those that provide the best performance.
- Providing examples of potential alternatives, which might have a better performance.
- Describing how these alternatives might depend on the conditions of the analysis, such as the amount of statistics that are available for training.

This article is divided as follows. In Section 2 we present the main libraries and tools one can use for signal-background discrimination and check their use in different LHC experiments. Section 3 introduces the datasets we simulated to perform the comparisons in the rest of the paper, as well as the main features that we used to discriminate signal against the background. Section 4 presents a first comparison of the performance of different ML libraries for classification, out of those presented in Section 2, when facing simulated data. This comparison is extended in Section 5 to more libraries and algorithms. We then discuss our results in Section 6. Finally, we conclude in Section 7.

2. Main ML Algorithms for Signal and Background Discrimination and Their Use at Particle Colliders

When tackling signal-background discrimination by means of ML, most analysts in LHC experiments automatically limit themselves to one library, TMVA [20], and one algorithm, BDTs with AdaBoost. In this section, we describe both of these and discuss other convenient options on the market, in terms of libraries and methods.

The Toolkit for Multivariate Data Analysis (TMVA) used with ROOT [21], is an open-source data analysis framework that provides all the necessary functionalities for processing large volumes of data, statistic analysis, visualization and information storage. ROOT was created at CERN and, although it was designed for HEP, it is currently used in many other areas such as biology, chemistry, astronomy, etc. As ROOT is specifically designed for HEP, the library is very popular and well-known for LHC experiments. The same is true for TMVA. As both ROOT and TMVA were developed some time ago, some of the latest techniques and innovations concerning data analysis are not yet available using just these tools. In this regard, several different TMVA interfaces have been recently created to solve this issue, with many framework options developed to integrate TMVA with more sophisticated libraries, such as Keras and PyTorch (described below). Thanks to this expansion, the use of ML for particle physicists is becoming easier, wider and more common. TMVA can be convenient for Python users, providing easy interfaces to examine data correlations, overtraining checks and a simpler event weighting management. Even though these characteristics are available in other Python libraries, TMVA offers the possibility of integrating every step of the process without the need for additional ones. Moreover, it offers suitable pre-processing possibilities for the data before feeding them into any of the classifiers. These data must be given in form of either ROOT TTrees or ASCII text files. The first is the usual format in which scientists at CERN deal with their datasets. The analyses in this paper are based on ROOT v6.22/08.

Sklearn [22] is a native open-source library for Python, and is currently an essential tool for modern data analysis. It includes a large range of tools and algorithms, which allow for the appropriate statistical modeling of systems. It incorporates algorithms for classification, regression, clustering, and dimensionality reduction, and supports a wide variety of algorithms [23] such as KNNs, boosted decision trees, random forests and SVMs among others. Furthermore, it is compatible with other Python libraries such as matplotlib [24], pandas [25], SciPy [26] or NumPy [27]. Contrary to TMVA, Sklearn was originally designed for Python. This means that the library provides a variety of modules and algorithms that facilitate the learning and work of the data scientist in the early stages of its development. For TMVA users, Sklearn can provide a versatile and simple interface for ML, with very simple applications for the trained classifier in new datasets. In this paper, we use Sklearn version 0.20.4.

Other popular types of libraries were designed to better deal with Neural Networks (NNs), which are introduced below. PyTorch [28] created by Adam Paszke, is object-oriented, which allows for dealing with NNs in a natural and convenient way. An NN in PyTorch is an object, so each NN is an instance of the class through which the network is defined, all of which are inherited from the torch NNs Module. The torch NNs Module provides the main tools for creating, training, testing and deploying an NN. In our analyses, we use PyTorch 1.8.0. Keras [29] is an open-source software library that provides a Python interface for artificial NNs. It is capable of running on top of Tensorflow [30], Theano [31] and CNTK [32]. Keras has wide compatibility between platforms for the developed models and excellent support for multiple GPUs. For this paper, we use Keras Release 2.6.0.

Moving on to the methodologies, as explained above, BDTs are among the most popular at LHC experiments. BDT stands for “Boosted Decision Tree”, and is currently one of the most popular algorithms for classification in ML. BDTs are based on the combination of weak predictive models (decision trees or DTs) to create a stronger one. These DTs are generated sequentially, with each tree being created to correct the errors in the previous one. This process of sequencing is known as “Boosting”, and can be of different types. The usual trees in BDTs are typically “shallow”, which means they have just one, two, or three levels of deepness. BDTs were designed to improve the disadvantages of the regular DTs, such as a tendency to overfit and their sensitivity to unbalanced training data. When boosting DTs, a very common choice is Adaptive Boosting (AdaBoost). This boosting technique is one of the canonical examples of ensemble learning in ML. In this method, the weights that are assigned to the classified instances are re-assigned at each step of the process. AdaBoost gives higher values to the more incorrectly classified instances in order to reduce bias and the variance. With this, one can make weak learners, such as DTs, stronger. The main difference between other types of boosting and AdaBoost relates to the use of particular loss functions. BDTs boosted with AdaBoost are the most frequent algorithm used at the LHC for signal-background discrimination by means of ML. Due to this, in the first part of the paper, we focus on how the performance obtained by this algorithm depend on the library used (namely, TMVA vs. Sklearn).

A very popular alternative in ML is NNs. This is a type of model that can try to emulate the behaviour of the human brain, whose potential for HEP was first introduced decades ago [33]. In NNs, nodes, usually denominated “artificial neurons”, are connected to each other to transmit signals. The main goal is to transmit these signals from the input to the end in order to generate an output. The neurons of a network can be arranged in layers to form the NN. Each of the neurons in the network is given a weight, which modifies the received input. This weight changes the values that come through, to later continue on their way through the network. Once the end of the network has been reached, the final prediction calculated by the network is shown as the output. In general, it can be said that an NN is more complex the more layers and neurons it has.

Although they are not directly part of this study, for reference, we now briefly review some of the algorithms that data scientists have developed more recently. Even though some of these are slowly being incorporated into HEP, they are beyond the scope of our

analysis. One example of these new developments is that of Extreme Gradient Boosting (XGBoost) [34], which is yet another type of boosting. XGBoost modifies the more traditional BDT algorithms by using a different objective function, so that a convex loss function is combined with a penalty term that accounts for the model complexity. XGBoost can be used together with different types of techniques, such as independent component analysis, gray wolf optimization or the whale optimization algorithm, to improve the general BDT results [35,36]. Going beyond this, we should now mention “liquid learning”. Liquid Learning [37] is a new type of ML algorithm that continuously adjusts depending on new data inputs. This means that the algorithm is able to modify its internal equations so that it can always adapt to changes in the incoming data stream. The fluidity of this new type of learning makes the algorithms more robust against noise or unforeseen changes in the data. Finally, one of the most expected developments these days concerns quantum technologies, and, with these, quantum ML (QML) [38]. QML algorithms and models try to use the advantages of quantum technologies to improve classical ML, for instance, by developing efficient implementations of slow classical algorithms using quantum computing.

In the process of analyzing how to make the most of ML in analyses at particle colliders, as mentioned above, we first compare the two best-known libraries in particle physics, TMVA and Sklearn, to quantify which one is more effective for the same problem, i.e., using the same data with the same characteristics and the same algorithm: a BDT with AdaBoost. For the second part of the paper, we increase the range of methodologies beyond the usual BDTs with AdaBoost to include NNs, and enlarge the range of libraries to PyTorch and Keras. The goal is to determine how all of these methods and libraries perform in terms of signal-background discrimination.

To compare the performance of different classifiers, the usual methodology in ML involves the use of the so-called ROC curve. A Receiver Operating Characteristic (ROC) curve is a representation of a classifying model, built to show the discrimination power of the classifier at different thresholds. The two parameters represented in the curve are the rate of true positives vs. the rate of false positives. We built ROC curves for all the classifiers of interest in this paper. For each classifier, one can also integrate the area under the corresponding ROC curve (AUC). The AUC turns out to be a very useful benchmark when one wants to compare the performance of different classifiers, providing a metric to quantify the separation power between different classes. In general, the higher the AUC score, the better the model, i.e., the more often it correctly assigns each instance to the class to which it belongs. This parameter is taken into account to compare the different classifiers. Moreover, we also account for the learning time needed for the algorithms to process the data, as well as the correlations with some key independent or “spectator” variables, such as the invariant mass.

Before moving to the problem of creating the models, let us examine how different LHC experiments have been using ML in recent years. This is a counting exercise, with which we intend to evaluate the popularity of different ML algorithms and libraries, without directly comparing their performance, which will be performed in the following sections. After looking into all LHCb publications from 2010 to August 2021, we calculated the percentage of papers per year that use TMVA and the percentage of papers that use Sklearn or other popular libraries that deal with NNs, such as Keras and PyTorch. We also generated a label, which we named “Generic NN” for those papers in which the use of a Neural Network is implied but the library used is not mentioned. The result can be seen in Figure 1, which shows how the use of TMVA is higher than Sklearn. In Figures 2 and 3, we repeated this for ATLAS [39] and CMS [40], respectively. We observed a similar behaviour to LHCb, where the use of TMVA was preferred to the other libraries. In ATLAS, even though TMVA is still the most used tool, we can appreciate an increase in the use of Keras in the last two years. Note that, in all three experiments, the use of Keras, PyTorch or other types of general NNs usually was not focused on direct signal-background discrimination at the analysis level, but instead on other tasks related to the reconstruction of events at the LHC, such as jet reconstruction and tagging, PID or track reconstruction.

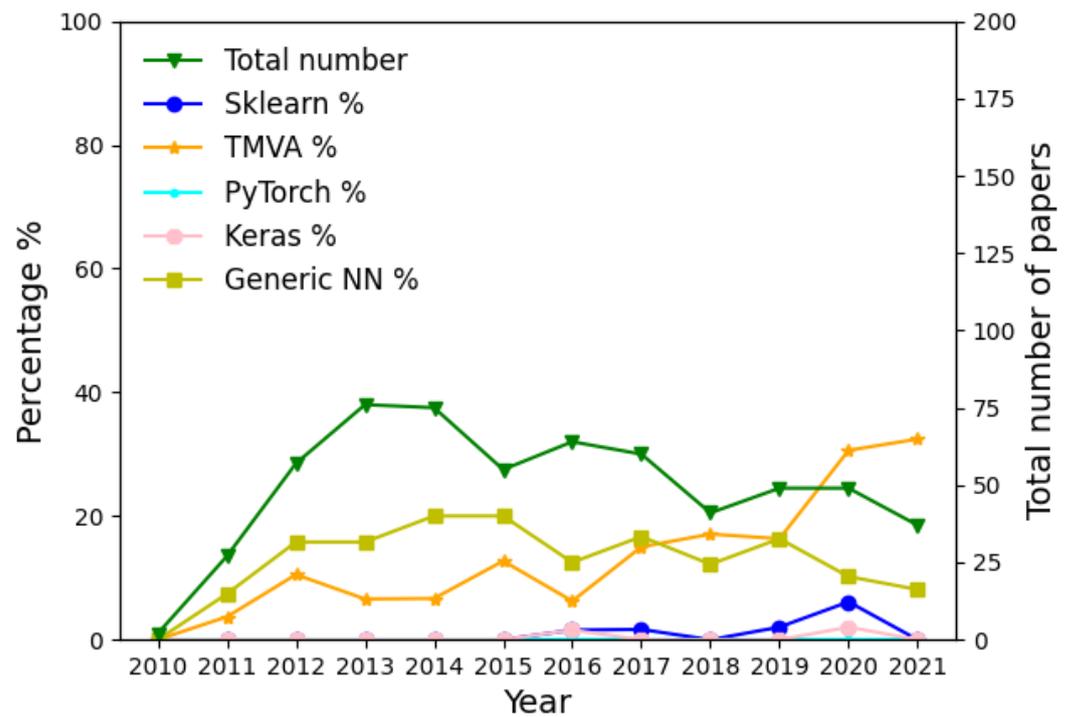


Figure 1. Usage of ML at LHCb across the years. We show the number of papers published every year (denoted as “Total number” and corresponding to the right axis), as well as the fraction of them reporting the use of TMVA, Sklearn, Keras, PyTorch and Generic NNs (in the left axis). The latter category corresponds to papers mentioning the use of NNs but never referencing any of the aforementioned libraries.

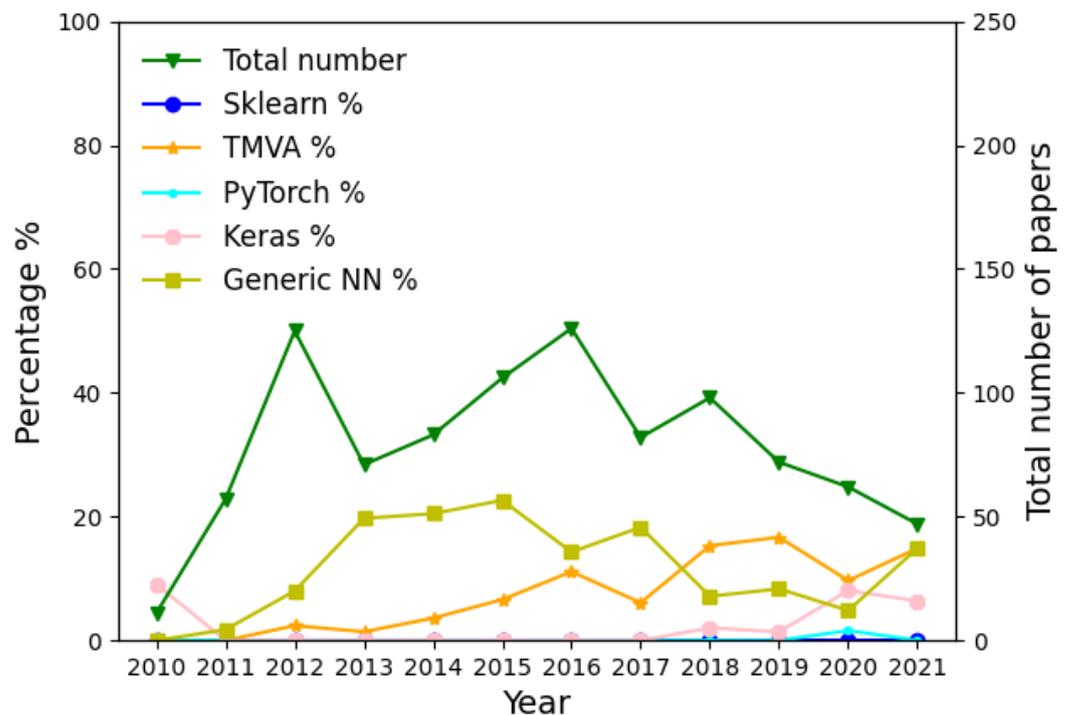


Figure 2. Usage of ML at ATLAS across the years. We show the number of papers published every year (denoted as “Total number” and corresponding to the right axis), as well as the fraction of them reporting the use of TMVA, Sklearn, Keras, PyTorch and Generic NNs (in the left axis). The latter category corresponds to papers mentioning the use of NNs but never referencing any of the aforementioned libraries.

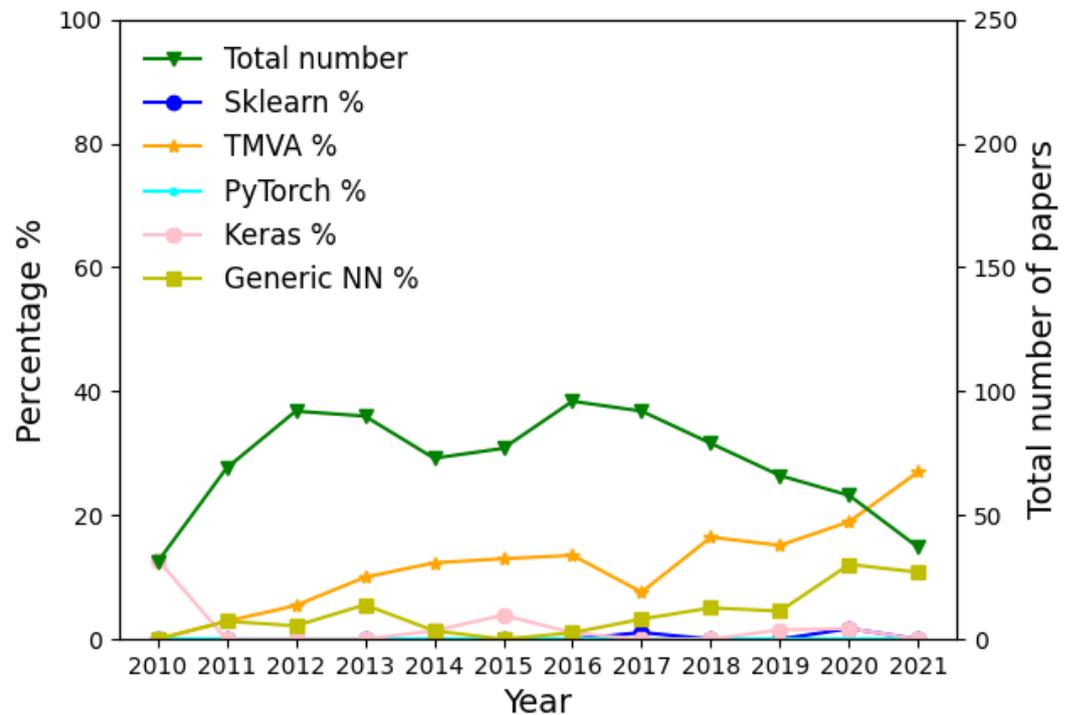


Figure 3. Usage of ML at CMS across the years. We show the number of papers published every year (denoted as “Total number” and corresponding to the right axis), as well as the fraction of them reporting the use of TMVA, Sklearn, Keras, PyTorch and Generic NNs (in the left axis). The latter category corresponds to papers mentioning the use of NNs but never referencing any of the aforementioned libraries.

3. The Data

As we mentioned before, the data we used to compare all the classifiers were the same. The main characteristics of the data are described in this section.

The data were generated using the simulation tool Pythia [41]. Pythia is a toolkit for the generation of high-energy physics events, simulating, for instance, the proton collisions that occur at the LHC. Given that we use LHCb as the benchmark for our studies, we focused on several B meson decay modes that were studied in the experiment. These correspond to different generic topologies and final-state particles. Note that we based our studies on the topology of the final states, and ignored any PID variable. LHCb has excellent libraries for PID, many of which rely on ML [42]. We generated the following decays:

- $B \rightarrow \mu^+ \mu^-$
- $B \rightarrow \pi^+ \pi^-$
- $B \rightarrow 3\pi$
- $B \rightarrow 4\pi$

The procedure to generate the signal samples is as follows. We enabled the `HardQCD:gg2bbbard` and `HardQCD:qqbar2bbbard` processes (which correspond to the production of a $b\bar{b}$ pair of quarks at the LHC) in Pythia at a collision energy of 14 TeV, and looked for B mesons. We then redecayed these to our desired final state using the `TGenPhaseSpace` tool from ROOT. Figure 4 characterizes the usual signal topology, with a B meson that flights a few mm and then decays to charged particles (such as pions or muons) whose trajectory and momentum can be fully reconstructed. We then applied the set of selection requirements included in Table 1. These are based on the quantities explained on Table 2 and Figure 4. On top of these quantities, we also applied selection requirements based on the pseudo-rapidity (η) of the final state particles. This relates to the angular acceptance of the detector, so that particles not falling in a specific η range fall outside the detector and therefore cannot be reconstructed.

Table 1. Values for the cuts used in Pythia to generate the signal and background samples for the analysis. The meaning of these variables can be found in Table 2 and Figure 4, as well as in the main text.

	η	p_{TB}	$p_{T\mu}, p_{T\pi}$	IP_B	IP_μ, IP_π	DOCA	DoF
μ	$2 < \eta < 5$	$>1000 \text{ MeV}/c$	$>500 \text{ MeV}/c$	$<0.1 \text{ mm}$	$>0.5 \text{ mm}$	$<0.1 \text{ mm}$	$>3 \text{ mm}$
π	$2 < \eta < 5$	$>1000 \text{ MeV}/c$	$>500 \text{ MeV}/c$	$<0.5 \text{ mm}$	$>0.5 \text{ mm}$	$<0.1 \text{ mm}$	$>3 \text{ mm}$

Table 2. List of features used to build the classifiers. For a mathematical definition of quantities such as the DOCA or IP see, e.g., Ref. [43]. More details are given in the text.

p_{TB}	Transverse momentum of the mother B meson
p_{Tdaug}	Transverse momentum of the daughter particles
p_x, p_y, p_z	Momentum components of the daughters
IP_B	Closest distance between the B mother trajectory and the proton–proton collision vertex
IP_{daug}	Closest distance between the daughter particle trajectory and the proton–proton collision vertex
DOCA	Distance of Closest Approach between daughter particles
DoF	Distance of Flight between the production and decay points of the mother B meson
Isolation μ_1	Minimum distance between μ_1 and any particle produced by the $b\bar{b}$ pair, excluding μ_2
Isolation μ_2	same but for μ_2 , excluding μ_1
Daughters _{pos}	Position of daughter particles
B _{pos}	Position of the B particle

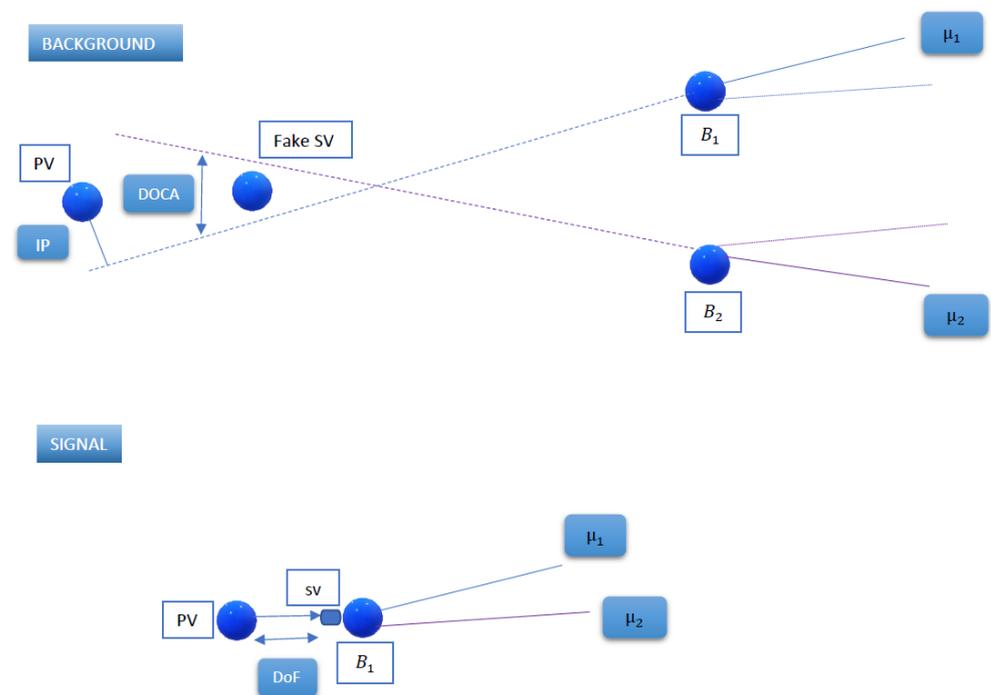


Figure 4. Characterization of the signal and background and definition of several variables to discriminate between them. For the $\mu\mu$ background, the usual candidates are formed by muons from different B mesons that are incorrectly matched together.

Figure 4 also explains the characterization of background events. We used the same type of Pythia processes to generate $b\bar{b}$ events, but there is no redecay at all this time, taking the event as it was originally fully simulated. We selected the charged pions and muons appearing in the event, grouped them, and applied the same cuts indicated in Table 1. For some of the cuts, we needed to generate a fake B meson decay vertex, which we used as the point in space minimizing the sum of distances to all the daughter particles.

In order to train the classifiers, we selected a list of variables (features) that was similar to those used to select the signal and background samples. These features are coherent between different types of classifiers to make the comparison fair. Note that some of the features chosen depend on the decay channel that was chosen. For instance, the $B \rightarrow 4\pi$ decay has four particles in the final state, which means the IP of two more particles is included compared to, e.g., $B \rightarrow \pi^+\pi^-$. As explained below, different sub-selections of features were also tried to provide as complete a picture as possible. The full list of features that were used can be found in Table A1. An additional aspect we have accounted for is the detector resolution. Our simulation provides the “true” value of features, but real-life detectors have some inherent inaccuracies due to their resolutions, as discussed in the introduction, so what they measure differs from this; for instance, the resolution damages the discrimination power of some of the features we work with. While, in principle, this would be only be a problem when trying to find a realistic performance of the classifiers, our goal here is only to compare their performance. In any case, to use a dataset that is as close as possible to that used by LHCb analysts, we conducted a Gaussian smear of the variables that are most affected by resolution effects. These are the momenta of the particles, the DOCAs and IP s. To determine the associated resolutions, we based ourselves in References [44,45].

Apart from the features introduced above, for the $B \rightarrow \mu^+\mu^-$ case we added an additional variable, the isolation, which is known to provide excellent signal-background discrimination in analyses of this kind [14]. The isolation exploits the fact that the usual $B \rightarrow \mu^+\mu^-$ decays have nothing “around” other than the muons, while the background contains additional objects produced in b -hadron decays (see Figure 4). To quantify this isolation, we calculated it as the minimum DOCA between our selected muons and every other charged particle in the event. These other particles were selected by applying the cuts in Table 3. Defined in this way, the isolation peaks at lower values for background and has larger peaks for the signal.

Table 3. List of cuts applied to charged particles entering the computation of the isolation. See text for details.

η	p_T	IP
$2 < \eta < 5$	$>250 \text{ MeV}/c$	$>0.1 \text{ mm}$

The simulated data consist of around 3000 samples of both signal and background events, generated as explained above. A normal LHC analysis may face difficulties in finding enough signal or background events to train a classifier, for instance, due to the very small selection efficiencies that require the simulation of an unreasonable amount of events. Additionally, analysts often face the choice between computing “low”- or “high”-level features. The former are those that were directly measured by the detector (for instance, the position or momentum or charged tracks), while the latter are combinations of those that are known to behave differently for the signal and backgrounds (such as the DOCA or IP , introduced above). While the usual classifiers typically rely on high-level features, a smart enough algorithm should be as effective only by using the low-level ones. To make all of this more concrete, we create six different training combinations between features and the number of samples. These are all features–high stats, all features–low stats, low-level features–high stats, low-level features–low stats, high-level features–high stats, high-level features–low stats. The meaning of these categories is as follows:

- High stats: In this category, we use all the 3000 samples for both signal and background.
- Low stats: In this category, we use 30% of the 3000 samples available for both signal and background.
- All features: In this category, we use all the features we have available for the data.
- High-level features: In this category, we only use high-level features.
- Low-level features: In this category, we only use low-level features.

With these categories, we aim to provide guidelines to analysts, helping them with different future analyses to choose the best tool to treat their data. All of the features used for each of the options mentioned above are shown in Appendix A.

One important last aspect to discuss, before moving on to the final analyses, is how the classifiers that we build must be uncorrelated with the invariant mass. In particle physics, one typically builds the invariant mass (the name “invariant” arises from special relativity, since this mass is the same for any reference frame) out of the properties of the final state particles. For instance, for the $B \rightarrow \mu^+ \mu^-$ decay, if we know the four-momenta of the muons, we can determine the invariant mass of the B mother. Since the value of the B meson mass is known a priori, the invariant mass is an excellent way to separate signal from background events, since, for signal, the distribution of this variable peaks at that known value. Another advantage of the invariant mass is that its distribution can be parameterized in an analytical way, which allows “counting” the amount of signal and background events by performing statistical fits to data. Since the invariant mass is a common and widely used feature in analyses in particle physics, as a rule of thumb, it is very important that any classifier we build is not correlated with the invariant mass, since the final count of signal events is performed based on this variable. This is so much the case that efforts have been made to develop classifiers that are explicitly trained so that no correlation is produced with the invariant mass (or other designed external variable) [46,47]. This is tricky, since a classifier might be able to learn the discrimination achieved by the invariant mass, provide an excellent AUC score and still not be useful for an actual analysis in particle physics. Accordingly, we followed a simple approach and ensured all our classifiers were not correlated with the invariant mass, which guarantees a fair comparison.

4. Scikit Learn vs. Tmva

As seen in Section 2, TMVA is the dominant reference tool for signal-background discrimination at LHC experiments. Therefore, in this section we check, using the generic data presented above, whether TMVA provides the optimal discrimination when compared to Sklearn, assuming that one is using the same datasets and the same algorithm. The exercises below are intended to show the same analysis being carried out with both TMVA and Sklearn libraries, with the goal of helping scientists to switch between the two according to their needs.

The main difference between Sklearn and TMVA is the way the data are processed. In Sklearn, we have to transform the ROOT TTrees to a numpy array or a pandas dataframe in order to work with them. At present, libraries such as uproot [48] or rootnumpy [49] allow us to make this change in the data format without extra effort. This option allows for the use of other libraries, instead of TMVA, even for the usual datasets in LHC experiments, which tend to be ROOT TTrees.

Both TMVA and Sklearn offer a simple way of handling algorithms, allowing for easy changes to their hyperparameters and providing useful information about the training time or even the ranking of features. To analyse and explain how the use of different libraries could be advantageous to the users, in Figure 5 and Table 4, we compare both libraries using BDT with AdaBoost for the $B \rightarrow \mu^+ \mu^-$ decay. In this case, we focus in the high-level features and high stats option for training. The comparison was made with the same hyperparameters and the same data for both libraries. This means that all the tunable parameters had the same value and the samples used for training and testing were the same for both cases. These hyperparameters were chosen as the best after performing a

grid search, i.e., running through a range of the best score was reached. The results of the classifiers can be seen in Figure 5, where Sklearn is shown to have better results. This is also shown in Table 4, where the AUC scores for $B \rightarrow \mu^+\mu^-$ and the rest of the decays can be found.

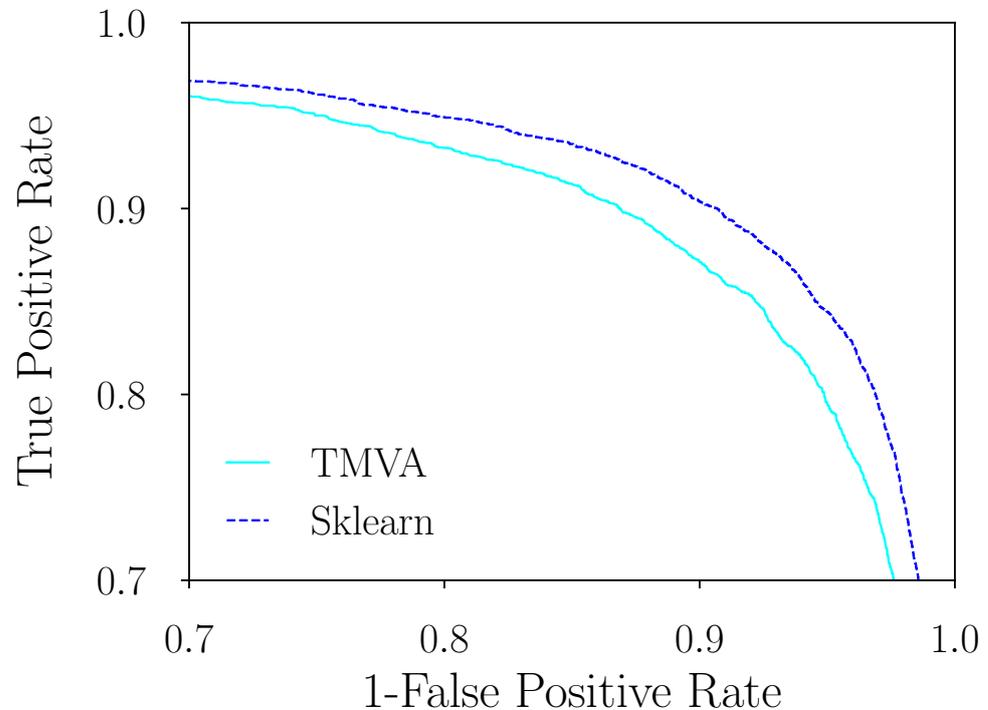


Figure 5. ROC Curve for the $B \rightarrow \mu^+\mu^-$ decay, corresponding to the high-level features and high stats option for training.

Table 4. TMVA vs. Sklearn results: AUC scores for all four decay channels, corresponding to the high-level features and high stats option for training. The z value corresponds to a score [50,51] comparing the AUCs of pairs of classifiers under the hypothesis that they are the same. The corresponding p -values are all ~ 0 .

	TMVA AUC	Sklearn AUC	z Value
$B \rightarrow \mu^+\mu^-$	0.954	0.963	88.884
$B \rightarrow \pi^+\pi^-$	0.952	0.961	90.047
$B \rightarrow 3\pi$	0.957	0.960	88.224
$B \rightarrow 4\pi$	0.957	0.961	88.508

In Table 4, we show all the results for the analyses that we defined in the previous section. For all these cases, Sklearn performs better than TMVA. This difference is small and not large enough to make a stronger statement, but it is still noticeable. To obtain a feeling of the statistical significance of the difference, we perform the DeLong test [50,51], which allows for the AUCs of pairs of classifiers to be compared. This test provides a z score that helps to quantify the probability that the two classifiers perform differently when looking at the same test data. The larger the z score, the more significant the separation between both AUCs. When comparing Sklearn and TMVA, we find that, for all decay channels, the AUC scores are statistically significantly different. Note that this statistical test is limited to pairs of classifiers, so we do not apply it further in the paper, given the large amount of comparisons being performed in Section 5.

One of the reasons that we think could explain the differences we observe between libraries relates to how these are designed and have evolved. Sklearn was released in 2007, and TMVA in 2009. Through the years, both libraries have changed and improved their performance until becoming those available at present. Still, as is shown in the tables and figures in this section, Sklearn always performs slightly better than TMVA. This means that even though both libraries a priori use the same algorithm, there are differences in the way these two frameworks operate. Different versions of these algorithms provide slightly different results, so one would be tempted to claim that Sklearn is able to provide a better optimization of the BDTs when compared to TMVA. This might be related to the application of the loss function or computational reasons.

5. Comparison of Popular ML Techniques beyond BDTs with Adaboost

In this section, we show the results achieved using a wider range of algorithms and libraries for each of the decays explained in Section 3. The main libraries we use are Sklearn, PyTorch and Keras. The main algorithms we use are BDT with AdaBoost and NNs, which were all introduced in Section 2. We limit ourselves to fully connected NNs with several layers.

Although we began this exercise using both types of algorithms for each library, once we obtained the results, we realized that the results for BDT with AdaBoost, when compared to NNs, were worse for every analysis for the PyTorch and Keras libraries and every classification of the data. This is the reason that we show only the results for the NNs for these two libraries. Furthermore, since, in Section 4, we found TMVA to have systematically lower AUC scores when compared to Sklearn, we did not include it in this second part. Finally, the results for Sklearn are shown using BDT with AdaBoost, since these were almost the same as this library when compared to NNs. Therefore, we analysed a total of 72 different classifiers (4 decay channels, 6 training options and 3 libraries). Overall, the use NNs was shown to generally be as the most competitive option for the analyses in this paper. Note that, in each case, we performed a grid search to choose the hyperparameters with the best AUC scores. These were used later for the actual comparison. The values of the hyperparameters are given in Appendix B.

We now show the results for all of the options we developed to select which is the best for each problem. Figures 6 and 7 show the results of all the previously mentioned analyses. For each column, we analysed a different decay and, for each row, we have a different training option for the analysis. For example, in row 0 column 0, we show the ROC Curves for the $B \rightarrow \mu^+ \mu^-$ decay and the low-level features and low stats option, referred to as “Low-Low”. Similar abbreviations appear for the other categories.

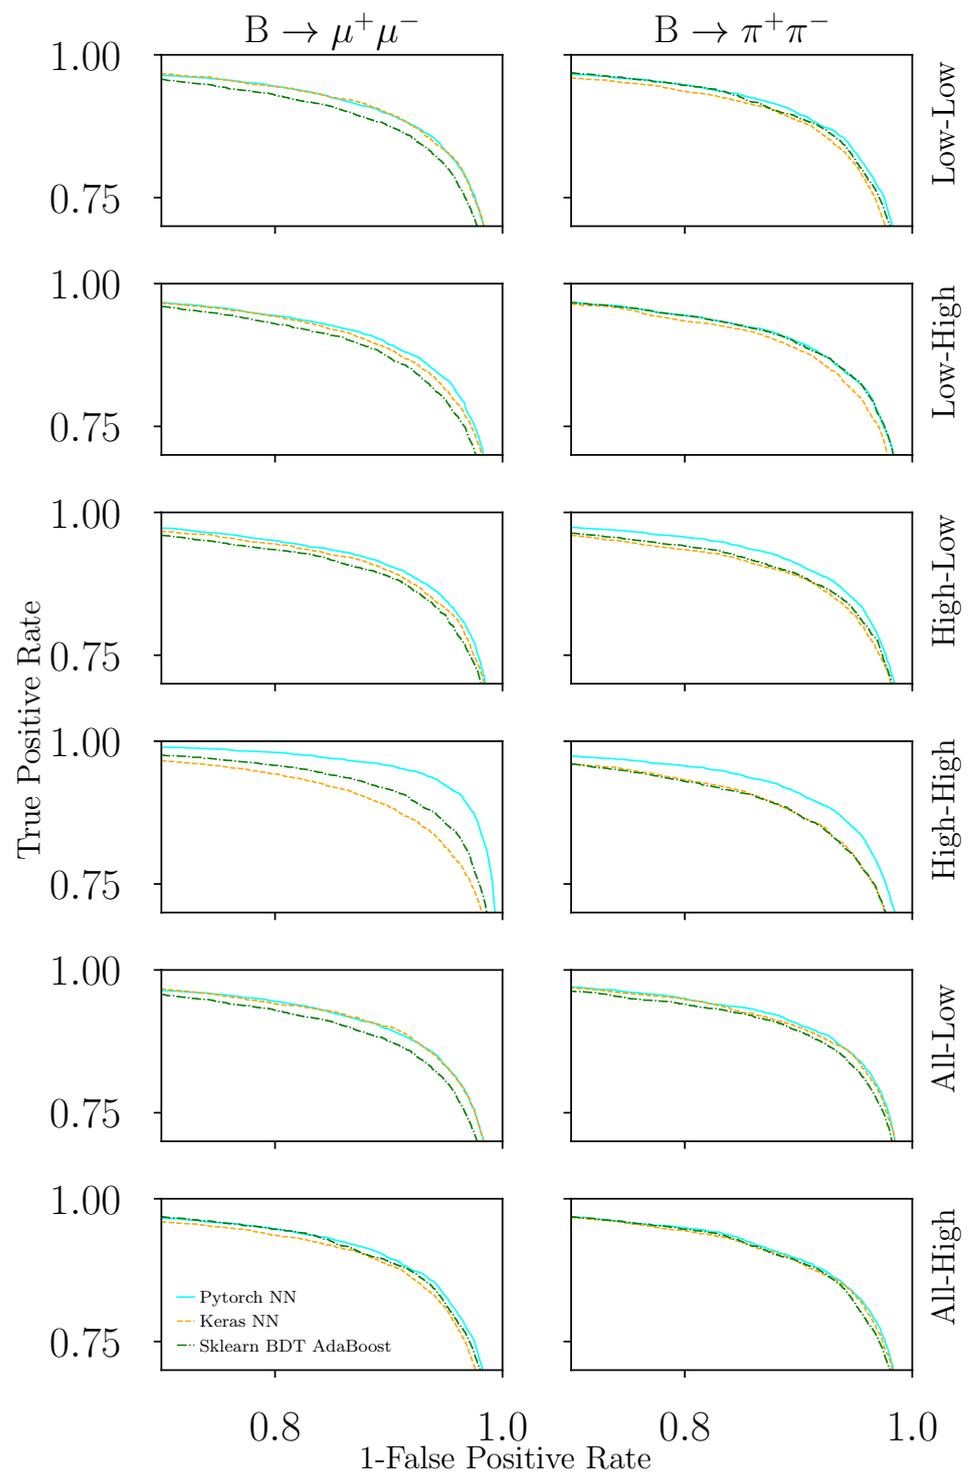


Figure 6. ROC Curve for the $B \rightarrow \mu^+\mu^-$ and $B \rightarrow \pi^+\pi^-$ decays and the four different options for training. The options are those listed in Section 3.

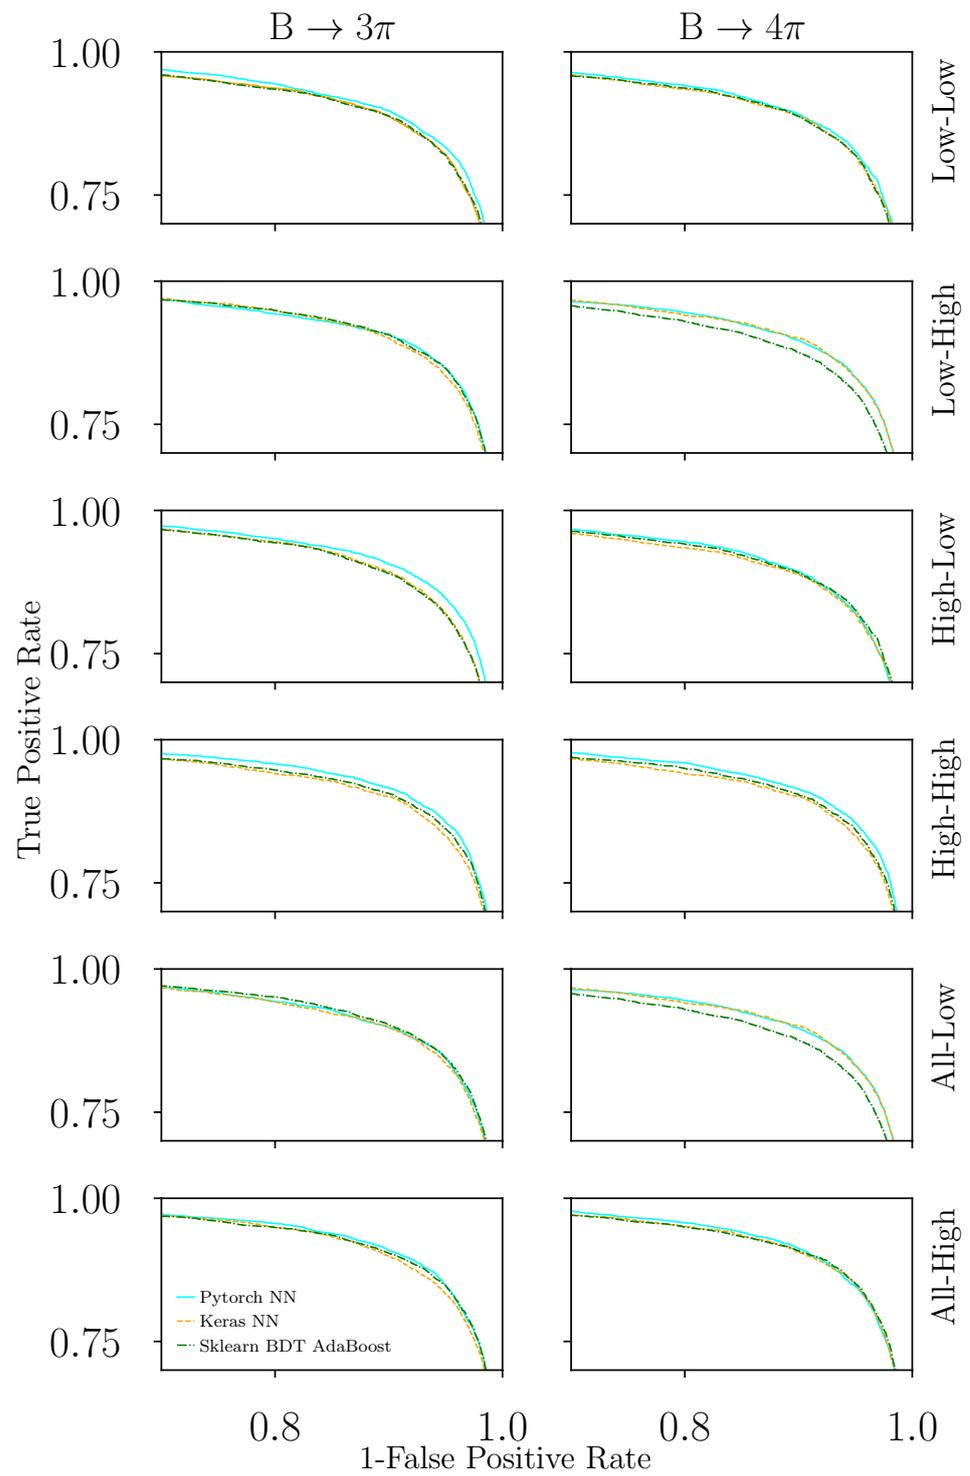


Figure 7. ROC Curve for the $B \rightarrow 3\pi$ and $B \rightarrow 4\pi$ decays and the four different options for training. The options are those listed in Section 3.

6. Results and Discussion

Once we have calculated the ROC curves and obtained the AUC scores for each of the algorithms, we proceed to compare them all for each channel. In the following tables, the results are shown for each tool and each training option. Tables 5–8 follow the results for each analysis. For reference, some of the training times, which are comparable in all cases, can be found in Appendix C. As shown in these tables, the AUC scores are similar for all of the options and most of the decays behave in a similar way, even when these scores are also subject to statistical fluctuations. The best results tend to be those obtained using PyTorch as the library, high-level features and all the available stats. This conclusion can be seen in Tables 9 and 10. Here, we can appreciate how PyTorch is the dominant library for most of the options, as it is the $B \rightarrow 4\pi$ decay, with all features and low stats and the $B \rightarrow 4\pi$ decay low-level features and high stats serving as the two exceptions. In these cases, Keras worked better. However, the result for high-level features and high stats is still the best option for all the decays, as shown in Table 10. We note again that most of the results are consistent across different decay channels, which correspond to different topologies, different types of background and even slightly different features (e.g., the use of isolation for the $B \rightarrow \mu^+\mu^-$ decay).

Table 5. $B \rightarrow \mu^+\mu^-$ channel: AUC scores for the different classifiers with all the libraries.

$B \rightarrow \mu^+\mu^-$	PyTorch	Keras	Sklearn
low-level features–low stats	0.961	0.959	0.952
low-level features–high stats	0.959	0.958	0.953
high-level features–low stats	0.964	0.959	0.956
high-level features–high stats	0.981	0.958	0.967
all features–low stats	0.961	0.960	0.953
all features–high stats	0.961	0.954	0.960

Table 6. $B \rightarrow \pi^+\pi^-$ channel: AUC scores for the different classifiers with all the libraries.

$B \rightarrow \pi^+\pi^-$	PyTorch	Keras	Sklearn
low-level features–low stats	0.960	0.954	0.959
low-level features–high stats	0.960	0.957	0.960
high-level features–low stats	0.965	0.956	0.958
high-level features–high stats	0.967	0.954	0.954
all features–low stats	0.963	0.962	0.959
all features–high stats	0.961	0.960	0.960

Table 7. $B \rightarrow 3\pi$ channel: AUC scores for the different classifiers with all the libraries.

$B \rightarrow 3\pi$	PyTorch	Keras	Sklearn
low-level features–low stats	0.961	0.954	0.956
low-level features–high stats	0.963	0.961	0.963
high-level features–low stats	0.965	0.960	0.953
high-level features–high stats	0.967	0.960	0.963
all features–low stats	0.961	0.960	0.963
all features–high stats	0.965	0.962	0.963

Table 8. $B \rightarrow 4\pi$ channel: AUC scores for the different classifiers with all the libraries.

$B \rightarrow 4\pi$	PyTorch	Keras	Sklearn
low-level features–low stats	0.958	0.956	0.955
low-level features–high stats	0.961	0.961	0.953
high-level features–low stats	0.960	0.958	0.958
high-level features–high stats	0.967	0.961	0.964
all features–low stats	0.960	0.961	0.954
all features–high stats	0.967	0.964	0.964

Table 9. Best library for each analysis and each training option for the data.

	$B \rightarrow \mu^+\mu^-$	$B \rightarrow \pi^+\pi^-$	$B \rightarrow 3\pi$	$B \rightarrow 4\pi$
low-level features–low stats	PyTorch	PyTorch	PyTorch	PyTorch
low-level features–high stats	PyTorch	PyTorch	PyTorch	Keras
high-level features–low stats	PyTorch	PyTorch	PyTorch	PyTorch
high-level features–high stats	PyTorch	PyTorch	PyTorch	PyTorch
all features–low stats	PyTorch	PyTorch	PyTorch	Keras
all features–high stats	PyTorch	PyTorch	PyTorch	PyTorch

Table 10. Best training option for the data and library for each analysis.

$B \rightarrow \mu^+\mu^-$	high-level features–high stats	PyTorch
$B \rightarrow \pi^+\pi^-$	high-level features–high stats	PyTorch
$B \rightarrow 3\pi$	high-level features–high stats	PyTorch
$B \rightarrow 4\pi$	high-level features–high stats	PyTorch

When using the same decay and the same classification features, we can compare how differently the low-stat and high-stat options behave. In some cases, as in $B \rightarrow 3\pi$, we see how the result improves for all options when using high stats. This is the type of behaviour we would generally expect to achieve when using ML. However, this is not something that happens in every case. For example, in the $B \rightarrow \pi^+\pi^-$ decay, we see that when using all features the AUC score does not improve, with more statistics, for Keras. The same occurs with Sklearn and high-level features and PyTorch and all features in the same channel. We observe that PyTorch and Sklearn tend to improve, with high stats, while Keras is more robust against a change. We also see a tendency to have a higher dependence on the statistics for the $B \rightarrow 3\pi$ and $B \rightarrow 4\pi$ decays, which present more final state particles and, therefore, have more features to account for.

To check how the selection of features affects the performance of these algorithms, we note there are no large differences between different decay channels and between the low- and high-stats categories. We then count how often the best AUC score is obtained for the different libraries depending on the feature selection. One would a priori expect that the addition of information from secondary features at the “all features” option would benefit the discrimination, especially for NNs. This must be balanced by the fact that adding too many or redundant variables can damage the overall performance in some cases. When looking at the global picture, we see that PyTorch performs better with the high-level features, while for Keras, the results tend to be better when choosing all features. The situation is slightly more balanced for Sklearn, although all features tend to provide a worse performance. Additionally, the use of low-level features does not provide the best performance in every library, but this still does not dramatically affect the power of discrimination, which indicates that the algorithms we use are robust. For reference, all the features used in this study can be found in Appendix A. Note that, as an additional element to judge the best choice, fewer features and smaller datasets can help to train the models faster and is more manageable in terms of, e.g., grid searches.

The convenience of using BDTs or NNs has long been a subject of discussion in the particle physics literature. For instance, Reference [52] discusses how BDTs can improve the performance of NNs for PID in neutrino detectors. Note that these results do not include some of the latest progress concerning the training of NNs, which provided essential improvements in the last decade [53]. It is remarkable how these improvements recover NNs as the best solution in the same dataset [54]. Other examples, which are more related to the type of analysis we present here, look into different algorithms in datasets that are often used by the ATLAS and CMS experiments, finding consistently better or similar results when comparing NNs to BDTs with different types of boosting, such as XGBoost [55,56]. One of the first attempts to apply QML in a data sample comparable to ours, composed of B meson decays [57], is also notable. The AUC achieved in this study manages to improve on those of other classical methods, although no NNs are included. Reference [58] compares QML algorithms to NNs in a different HEP dataset, achieving better or similar AUC scores, although this depends on the size of the training sample. Another interesting example concerns the use of these algorithms for the simulation of HEP data [59], where, again, NNs beat BDTs in terms of regression. Regarding the libraries used, the advantages of using PyTorch to deal with NNs are becoming well known, and several toolkits based on it have recently been developed for different HEP applications. Examples include Lumin [60,61] and Ariadne [7]. The first is a wrapper that includes some of the newest techniques that facilitate the training of NNs, while the second, as mentioned in the introduction, is used for tracking. As a final remark, the benefits NNs can bring over other methods, such as BDTs, not only concern physics, but also appear in fields such as medicine [62,63], finance [64], marketing [65], biology [66] or engineering [67]. As a general trend, NNs tend to perform better in these references, although the present developments in, e.g., BDTs, make these very competitive. We note the interest in looking at other sources, beyond particle physics, to find inspiration and discover better models for signal classification.

7. Conclusions

In this paper, we review the most frequent ML libraries used in HEP, checking their popularity and comparing their performance in terms of signal-background discrimination in an independent simulated dataset. These datasets are generated and selected using the LHCb detector at CERN as a benchmark.

In the first part of the paper, we observe that, even though TMVA is one of the most popular libraries, and by far the more frequently used for this type of analyses, there are other options that can work as well or even better. In fact, Sklearn performs better than TMVA in all the decays analysed in this paper when using BDTs boosted with AdaBoost. This can be a way to prompt scientists in HEP collaborations to try new alternatives and generate new content using the most modern libraries that are available according to Python and other languages. Thanks to the wide range of new libraries, the conversion between the ROOT files and Python-friendly data structures is growing easier over time, and this can result in the popularization of these modern ML libraries for particle physicists.

In the second part of the paper, we compared the obtained results with some of the most popular ML libraries in data science, namely, Keras, PyTorch and Sklearn, and showed how NNs can improve the results obtained by BDTs. Even if the results are similar between the three of them, PyTorch tends to provide the best scores. In any case, the final choice might depend on other aspects, such as the amount of statistics available, since, for instance, Keras appears to be more robust for lower training statistics. Regarding the selection of features, this depends on the library, but, as a rule of thumb, we recommend not enlarging the list unnecessarily, focusing on high-level libraries that are known to provide excellent signal-background discrimination. To our knowledge, this is the first time that a detailed study of the dependence of algorithms and libraries on the training statistics and number of features has been performed with an HEP dataset.

We should highlight that, even though we are able to see some differences among the libraries, algorithms and decays, all AUC scores are in the range of 0.95–0.96, and subject to statistical fluctuations. This means that the results we have obtained are indicative, but must still be made with caution. Therefore, even if our findings can be used as a guideline, we still advise analysts to check for the best classifier for their specific case, depending on aspects such as available statistics or the available CPU for training.

As a final note, we emphasize the need for particle physicists to enrich their perspective by looking at what is being carried out in other fields in terms of classification by means of ML. This not only concerns other areas of physics [68], but also fields that look far away, such as medicine [69], chemistry [70], industrial applications [71,72] or the interface with users [73].

Author Contributions: Conceptualization, X.C.V., L.D.M. and Á.D.S.; methodology, X.C.V., L.D.M. and Á.D.S.; software, X.C.V., L.D.M. and Á.D.S.; validation, X.C.V., L.D.M. and Á.D.S.; formal analysis, X.C.V., L.D.M. and Á.D.S.; investigation, X.C.V., L.D.M. and Á.D.S.; resources, X.C.V., L.D.M. and Á.D.S.; data curation, X.C.V., L.D.M. and Á.D.S.; writing—original draft preparation, X.C.V., L.D.M. and Á.D.S.; writing—review and editing, X.C.V., L.D.M. and Á.D.S.; visualization, X.C.V., L.D.M. and Á.D.S.; supervision, X.C.V. and Á.D.S.; project administration, X.C.V., L.D.M. and Á.D.S.; funding acquisition, X.C.V., L.D.M. and Á.D.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work has received financial support from Xunta de Galicia (Centro singular de investigación de Galicia accreditation 2019–2022), by European Union ERDF, and by the “María de Maeztu” Units of Excellence program MDM-2016-0692 and the Spanish Research State Agency. In particular: the work of X.C.V. is supported by MINECO (Spain) through the Ramón y Cajal program RYC-2016-20073 and by XuntaGAL under the ED431F 2018/01 project and the work of L.D.M. is supported by the Spanish Research State Agency (Spain) through the “Doctorados Industriales” program DIN2018-010092.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used for this analysis originates from a simple Pythia simulation, which should be trivial to reproduce.

Acknowledgments: We thank Alexandre Brea and Titus Mombacher for reading our manuscript and providing useful comments.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Description of Features

In this appendix, we show how the options of “high-level features”, “all features” and “low-level features” are designed. In Table A1, there is a tick for each feature that is used for the specific option. The option “all features” includes, as the name implies, all of the available features we simulate. In the other cases, we tried to use enough features to make the options distinguishable, while still allowing for the separation of different instances of the data. As explained in the main text, we create these options to prove how differently the models behave when giving redundant information versus how they work when giving just the necessary information.

Table A1. Features present in each training option for the data. Note that the isolation only applies to the $B \rightarrow \mu^+ \mu^-$ channel.

	High-Level Features	All Features	Low-Level Features
p_{TB}	✓	✓	
p_{Tdaug}	✓	✓	
IP_B	✓	✓	
$IP_{daughters}$	✓	✓	
Isolation μ	✓	✓	✓
Position of daughters		✓	✓
Position of mothers		✓	✓
Doca	✓	✓	
DoF	✓	✓	
p_x, p_y, p_z daughters		✓	✓

Appendix B. Hyperparameters for the Classifiers

In this appendix, we show all the hyperparameters needed by the libraries we selected after an exhaustive grid search. For each decay, training option and tool, the results are shown in the following tables, to allow the reader to replicate our results and estimate for which algorithms work better for the correspondent problem and decay.

Table A2. Neural Network parameters for PyTorch for the $B \rightarrow \mu^+ \mu^-$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	ReLU	9	0.0089	0.75	100
low-level features–high stats	ReLU	12	0.0093	0.8	180
high-level features–low stats	ReLU	11	0.009	0.9	80
high-level features–high stats	Tanh	9	0.0086	0.8	120
all features–low stats	ReLU	14	0.0075	0.85	110
all features–high stats	Tanh	12	0.0092	0.85	170

Table A3. Neural Network parameters for Keras for the $B \rightarrow \mu^+ \mu^-$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	ReLU	8	0.008	0.8	110
low-level features–high stats	Tanh	12	0.0085	0.85	180
high-level features–low stats	ReLU	11	0.0087	0.8	100
high-level features–high stats	Tanh	9	0.0082	0.8	130
all features–low stats	ReLU	14	0.0089	0.85	100
all features–high stats	ReLU	12	0.0082	0.9	190

Table A4. Neural Network parameters for PyTorch for the $B \rightarrow \pi^+ \pi^-$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	Tanh	14	0.0089	0.85	100
low-level features–high stats	ReLU	10	0.0091	0.8	140
high-level features–low stats	ReLU	10	0.0087	0.85	100
high-level features–high stats	Tanh	12	0.0088	0.8	120
all features–low stats	ReLU	10	0.0085	0.95	100
all features–high stats	Tanh	11	0.0089	0.85	130

Table A5. Neural Network parameters for Keras for the $B \rightarrow \pi^+ \pi^-$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	ReLU	10	0.0079	0.85	100
low-level features–high stats	Tanh	14	0.0081	0.85	130
high-level features–low stats	ReLU	10	0.0091	0.9	110
high-level features–high stats	Tanh	16	0.0093	0.85	140
all features–ow stats	ReLU	11	0.009	0.95	120
all features–high stats	ReLU	12	0.0086	0.85	120

Table A6. Neural Network parameters for PyTorch for the $B \rightarrow 3\pi$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	Tanh	11	0.009	0.9	110
low-level features–high stats	ReLU	13	0.009	0.85	100
high-level features–low stats	ReLU	12	0.0085	0.85	110
high-level features–high stats	Tanh	12	0.0087	0.75	130
all features–low stats	ReLU	12	0.0089	0.9	110
all features–high stats	Tanh	16	0.0086	0.8	130

Table A7. Neural Network parameters for Keras for the $B \rightarrow 3\pi$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	ReLU	9	0.0089	0.8	100
low-level features–high stats	Tanh	13	0.008	0.8	120
high-level features–low stats	ReLU	10	0.009	0.85	130
high-level features–high stats	Tanh	12	0.0087	0.85	150
all features–low stats	ReLU	11	0.0091	0.9	100
all features–high stats	ReLU	15	0.0085	0.8	140

Table A8. Neural Network parameters for PyTorch for the $B \rightarrow 4\pi$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	Tanh	11	0.009	0.81	100
low-level features–high stats	ReLU	13	0.0092	0.85	140
high-level features–low stats	ReLU	10	0.0089	0.85	110
high-level features–high stats	Tanh	10	0.0078	0.8	150
all features–low stats	ReLU	10	0.0075	0.9	130
all features–high stats	Tanh	14	0.0079	0.85	160

Table A9. Neural Network parameters for Keras for the $B \rightarrow 4\pi$ decay.

	A. Function	Hidden Layers	Learning Rate	Momentum	Epochs
low-level features–low stats	ReLU	10	0.008	0.8	100
low-level features–high stats	Tanh	16	0.0082	0.8	120
high-level features–low stats	ReLU	9	0.0081	0.7	100
high-level features–high stats	Tanh	12	0.0095	0.85	160
all features–low stats	ReLU	10	0.009	0.9	110
all features–high stats	ReLU	9	0.0086	0.9	140

Table A10. BDT with AdaBoost parameters for Sklearn for the $B \rightarrow \mu^+\mu^-$ decay.

	Base Estimator	N° Estimators	Learning Rate
low-level features–low stats	DecisionTreeClassifier	120	0.8
low-level features–high stats	DecisionTreeClassifier	100	0.8
high-level features–low stats	DecisionTreeClassifier	150	0.7
high-level features–high stats	DecisionTreeClassifier	170	0.9
all features–low stats	DecisionTreeClassifier	160	0.9
all features–high stats	DecisionTreeClassifier	200	0.8

Table A11. BDT with AdaBoost parameters for Sklearn for the $B \rightarrow \pi^+\pi^-$ decay.

	Base Estimator	N° Estimators	Learning Rate
low-level features–low stats	DecisionTreeClassifier	100	0.8
low-level features–high stats	DecisionTreeClassifier	130	0.8
high-level features–low stats	DecisionTreeClassifier	150	0.8
high-level features–high stats	DecisionTreeClassifier	200	0.9
all features–low stats	DecisionTreeClassifier	170	0.9
all features–high stats	DecisionTreeClassifier	250	0.9

Table A12. BDT with AdaBoost parameters for Sklearn for the $B \rightarrow 3\pi$ decay.

	Base Estimator	N° Estimators	Learning Rate
low-level features–low stats	DecisionTreeClassifier	110	0.75
low-level features–high stats	DecisionTreeClassifier	170	0.9
high-level features–low stats	DecisionTreeClassifier	190	0.75
high-level features–high stats	DecisionTreeClassifier	230	0.9
all features–low stats	DecisionTreeClassifier	220	0.85
all features–high stats	DecisionTreeClassifier	240	0.9

Table A13. BDT with AdaBoost parameters for Sklearn for the $B \rightarrow 4\pi$ decay.

	Base Estimator	N° Estimators	Learning Rate
low-level features–low stats	DecisionTreeClassifier	120	0.85
low-level features–high stats	DecisionTreeClassifier	120	0.85
high-level features–low stats	DecisionTreeClassifier	150	0.8
high-level features–high stats	DecisionTreeClassifier	190	0.9
all features–low stats	DecisionTreeClassifier	160	0.95
all features–high stats	DecisionTreeClassifier	230	0.9

Appendix C. Training Time for the Classifiers in a Personal Laptop for the $B \rightarrow \mu^+\mu^-$ Decay

In order to obtain another metric of how good an algorithm is, we present an example of how long it takes an algorithm to train with the samples we selected. We present all the different analyses for the decay $B \rightarrow \mu^+\mu^-$. Comparable times are obtained for other decays.

Table A14. Training time for each tool for the decay $B \rightarrow \mu^+\mu^-$

	PyTorch	Keras	Sklearn
high-level features–high stats	5 min 56 s	5 min 46 s	5 min 43 s
high-level features–low stats	4 min 35 s	4 min 43 s	4 min 30 s
low-level features–low stats	4 min 05 s	4 min 13 s	4 min 11 s
low-level features–high stats	5 min 34 s	5 min 51 s	6 min 36 s
all features–high stats	6 min 06 s	6 min 11 s	5 min 58 s
all features– low stats	6 min 46 s	6 min 03 s	5 min 20 s

References

- Feickert, M.; Nachman, B. A Living Review of Machine Learning for Particle Physics. 2021. Available online: <https://arxiv.org/abs/2102.02770> (accessed on 17 November 2021).
- Alves, A.A., Jr.; Filho, L.M.A.; Barbosa, A.F.; Bediaga, I.; Cernicchiaro, G.; Guerrer, G.; Lima, H.P.; Machado, A.A.; Magnin, J.; Marujo, F.; et al. The LHCb Detector at the LHC. *J. Instrum.* **2008**, *3*, S08005. [[CrossRef](#)]
- Aaij, R.; Adeva, B.; Adinolfi, M.; Affolder, A.; Ajaltouni, Z.; Akar, S.; Albrecht, J.; Alessio, F.; Alexander, M.; Ali, S.; et al. LHCb Detector Performance. *Int. J. Mod. Phys. A* **2015**, *30*, 1530022. [[CrossRef](#)]
- CERN Storage. Available online: <https://home.cern/science/computing/storage> (accessed on 1 October 2021).
- Qasim, S.R.; Kieseler, J.; Iiyama, Y.; Pierini, M. Learning representations of irregular particle-detector geometry with distance-weighted graph networks. *Eur. Phys. J. C* **2019**, *79*, 608. [[CrossRef](#)]
- Cranmer, K.; Drnevich, M.; Macaluso, S.; Pappadopulo, D. Reframing Jet Physics with New Computational Methods. *EPJ Web Conf.* **2021**, *251*, 03059. [[CrossRef](#)]
- Goncharov, P.; Schavelev, E.; Nikolskaya, A.; Ososkov, G. Ariadne: PyTorch Library for Particle Track Reconstruction Using Deep Learning. In *AIP Conference Proceedings*; AIP Publishing LLC: Melville, NY, USA, 2021; Volume 2377, p. 040004. [[CrossRef](#)]
- Andrews, M.; Burkle, B.; Chen, Y.; DiCroce, D.; Gleyzer, S.; Heintz, U.; Narain, M.; Paulini, M.; Pervan, N.; Shafi, Y.; et al. End-to-End Jet Classification of Boosted Top Quarks with CMS Open Data. *EPJ Web Conf.* **2021**, *251*, 04030. [[CrossRef](#)]
- Akchurin, N.; Cowden, C.; Damgov, J.; Hussain, A.; Kunori, S. Perspectives on the Calibration of CNN Energy Reconstruction in Highly Granular Calorimeters. 2021. Available online: <https://arxiv.org/abs/2108.10963> (accessed on 17 November 2021).

10. Cornell, A.S.; Doorsamy, W.; Fuks, B.; Harmsen, G.; Mason, L. Boosted Decision Trees in the Era of New Physics: A Smuon Analysis Case Study 2021. Available online: <https://arxiv.org/abs/2109.11815> (accessed on 17 November 2021).
11. Baldi, P.; Sadowski, P.; Whiteson, D. Searching for Exotic Particles in High-Energy Physics with Deep Learning. *Nat. Commun.* **2014**, *5*, 4308. [[CrossRef](#)] [[PubMed](#)]
12. Dillon, B.M.; Faroughy, D.A.; Kamenik, J.F.; Szewc, M. Learning the latent structure of collider events. *J. High Energy Phys.* **2020**, *10*, 206. [[CrossRef](#)]
13. Dahbi, S.E.; Choma, J.; Mellado, B.; Mokgatitwane, G.; Ruan, X.; Lieberman, B.; Celik, T. Machine Learning Approach for the Search of Resonances with Topological Features at the Large Hadron Collider. 2020. Available online: <https://arxiv.org/abs/2011.09863> (accessed on 17 November 2021).
14. Aaij, R.; Adeva, B.; Adinolfi, M.; Adrover, C.; Affolder, C.; Ajaltouni, Z.; Albrecht, J.; Alessio, F.; Alexander, M.; Alvarez Cartelle P.; et al. Search for the rare decays $B_s^0 \rightarrow \mu^+ \mu^-$ and $B^0 \rightarrow \mu^+ \mu^-$. *Phys. Lett. B* **2011**, *699*, 330–340. [[CrossRef](#)]
15. Williams, M.; Gligorov, V.V.; Thomas, C.; Dijkstra, H.; Nardulli, J.; Spradlin, P. *The HLT2 Topological Lines*; Technical Report; CERN: Geneva, Switzerland, 2011.
16. Likhomanenko, T.; Ilten, P.; Khairullin, E.; Rogozhnikov, A.; Ustyuzhanin, A.; Williams, M. LHCb Topological Trigger Reoptimization. *J. Phys. Conf. Ser.* **2015**, *664*, 082025. [[CrossRef](#)]
17. Adam-Bourdarios, C.; Cowan, G.; Germain, C.; Guyon, I.; Kégl, B.; Rousseau, D. The Higgs Boson Machine Learning Challenge. In Proceedings of the 2014 International Conference on High-Energy Physics and Machine Learning, JMLR.org, HEPML'14, Montreal, QC, Canada, 8–13 December 2014; Volume 42, pp. 19–55.
18. Aarrestad, T.; van Beekveld, M.; Bona, M.; Boveia, A.; Caron, S.; Davies, J.; De Simone, A.; Doglioni, C.; Duarte, J.M.; Farbin A.; et al. The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider. 2021. Available online: <https://arxiv.org/abs/2105.14027> (accessed on 17 November 2021).
19. Kasieczka, G.; Nachman, B.; Shih, D.; Amram, O.; Andreassen, A.; Benkendorfer, K.; Bortolato, B.; Brooijmans, G.; Canelli, F.; Collins, J.H.; et al. The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics. 2021. Available online: <https://arxiv.org/abs/2101.08320> (accessed on 17 November 2021).
20. Hocker, A.; Speckmayer, P.; Stelzer, J.; Therhaag, J.; von Toerne, E.; Voss, H.; Backes, M.; Carli, T.; Cohen O.; Christov A.; et al. *TMVA–Toolkit for Multivariate Data Analysis with ROOT: Users Guide*. *TMVA–Toolkit for Multivariate Data Analysis*; Technical Report; CERN: Geneva, Switzerland, 2007; *TMVA-v4 Users Guide*: 135 Pages, 19 Figures, Numerous Code Examples and References. Available online: <https://arxiv.org/abs/physics/0703039> (accessed on 17 November 2021).
21. Brun, R.; Rademakers, F.; Canal, P.; Naumann, A.; Couet, O.; Moneta, L.; Vassilev, V.; Linev, S.; Piparo, D.; Ganis, G.; et al. Root-Project/Root: v6.18/02. *Zenodo August* **2019**. [[CrossRef](#)]
22. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.
23. Géron, A. *Hands-on Machine Learning with Scikit-Learn and TensorFlow Concepts, Tools, and Techniques to Build Intelligentsystems*; O'Reilly Media: Sebastopol, CA, USA, 2017.
24. Hunter, J.D. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **2007**, *9*, 90–95. [[CrossRef](#)]
25. McKinney, W. Data structures for statistical computing in python. In Proceedings of the 9th Python in Science Conference, Austin, TX, USA, 28 June–3 July 2010; Volume 445, pp. 51–56.
26. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)] [[PubMed](#)]
27. Harris, C.R.; Millman, K.J.; van der Walt, S.J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N.J.; et al. Array programming with NumPy. *Nature* **2020**, *585*, 357–362. [[CrossRef](#)]
28. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*; Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2019; pp. 8024–8035.
29. Chollet, F. *Keras Software*. Available online: <https://keras.io> (accessed on 17 November 2021).
30. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. Software. Available online: <https://tensorflow.org> (accessed on 17 November 2021).
31. Al-Rfou, R.; Alain, G.; Almahairi, A.; Angermueller, C.; Bahdanau, D.; Ballas, N.; Bastien, F.; Bayer, J.; Belikov, A.; Belopolsky, A.; et al. Theano: A Python framework for fast computation of mathematical expressions. *arXiv* **2016**, arXiv:1605.02688.
32. Seide, F.; Agarwal, A. CNTK: Microsoft's Open-Source Deep-Learning Toolkit. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; Association for Computing Machinery: New York, NY, USA, 2016; p. 2135. [[CrossRef](#)]
33. Denby, B. Neural networks and cellular automata in experimental high energy physics. *Comput. Phys. Commun.* **1988**, *49*, 429–448. [[CrossRef](#)]
34. Chen, T.; Guestrin, C. XGBoost. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016.

35. Duan, J.; Asteris, P.; Nguyen, H.; Bui, X.N.; Moayedi, H. A Novel Artificial Intelligence Technique to Predict Compressive Strength of Recycled Aggregate Concrete Using ICA-XGBoost Model. *Eng. Comput.* **2021**, *37*, 3329–3346. [[CrossRef](#)]
36. Qiu, Y.; Zhou, J.; Khandelwal, M.; Yang, H.; Yang, P.; Li, C. Performance evaluation of hybrid WOA-XGBoost, GWO-XGBoost and BO-XGBoost models to predict blast-induced ground vibration. *Eng. Comput.* **2021**, 1–18. [[CrossRef](#)]
37. Hasani, R.; Lechner, M.; Amini, A.; Rus, D.; Grosu, R. Liquid Time-constant Networks. In Proceedings of the AAAI Conference on Artificial Intelligence (Virtual), 2–9 February 2021; Volume 35, pp. 7657–7666.
38. Schuld, M.; Sinayskiy, I.; Petruccione, F. An introduction to quantum machine learning. *Contemp. Phys.* **2014**, *56*, 172–185. [[CrossRef](#)]
39. Aad, G.; Bentvelsen, S.; Bobbink, G.J.; Bos, K.; Boterenbrood, H.; Brouwer, G.; Buis, E.J.; Buskop, J.J.F.; Colijn, A.P.; Dankers, R.; et al. The ATLAS Experiment at the CERN Large Hadron Collider. *J. Instrum.* **2008**, *3*, S08003. [[CrossRef](#)]
40. Chatrchyan, S.; Hmayakyan, G.; Khachatryan, V.; Sirunyan, A.M.; Adolphi, R.; Anagnostou, G.; Brauer, R.; Braunschweig, W.; Esser, H.; Feld, L.; et al. The CMS Experiment at the CERN LHC. *JINST* **2008**, *3*, S08004. [[CrossRef](#)]
41. Sjöstrand, T.; Ask, S.; Christiansen, J.R.; Corke, R.; Desai, N.; Ilten, P.; Mrenna, S.; Prestel, S.; Rasmussen, C.O.; Skands P.Z. An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **2015**, *191*, 159–177. [[CrossRef](#)]
42. Kazeev, N. Machine Learning for Particle Identification in the LHCb Detector. Ph.D. Thesis, Sapienza—University of Rome, Rome, Italy, 21 October 2020.
43. Buarque Franzosi, D.; Cacciapaglia, G.; Cid Vidal, X.; Ferretti, G.; Flacke, T.; Vázquez Sierra, C. Exploring New Possibilities to Discover a Light Pseudo-Scalar at LHCb. 2021. Available online: <https://arxiv.org/abs/2106.12615> (accessed on 17 November 2021).
44. Cid Vidal, X.; Ilten, P.; Plews, J.; Shuve, B.; Soreq, Y. Discovering True Muonium at LHCb. *Phys. Rev. D* **2019**, *100*, 053003. [[CrossRef](#)]
45. Hynds, D.P.M. Resolution Studies and Performance Evaluation of the LHCb VELO Upgrade. Ph.D. Thesis, University of Glasgow, Glasgow, UK, 27 November 2014.
46. Stevens, J.; Williams, M. uBoost: A boosting method for producing uniform selection efficiencies from multivariate classifiers. *J. Instrum.* **2013**, *8*, P12013. [[CrossRef](#)]
47. Rogozhnikov, A.; Bukva, A.; Gligorov, V.V.; Ustyuzhanin, A.; Williams, M. New approaches for boosting to uniformity. *J. Instrum.* **2015**, *10*, T03002. [[CrossRef](#)]
48. Pivarski, J.; Das, P.; Smirnov, D.; Burr, C.; Feickert, M.; Biederbeck, N.; Smith, N.; Rembser, J.; Schreiner, H.; Dembinski, H.; et al. Scikit-Hep/Uproot: 3.10.7. 2019. [[CrossRef](#)]
49. Dawe, N.; Ongmongkolkul, P.; Deil, C.; Stark, G.; Waller, P.; Howard, J.; Babuschkin, I. Root_NUMPY: 4.2.0. 2015. [[CrossRef](#)]
50. DeLong, E.R.; DeLong, D.M.; Clarke-Pearson, D.L. Comparing the Areas under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach. *Biometrics* **1988**, *44*, 837–845. [[CrossRef](#)]
51. Sun, X.; Xu, W. Fast Implementation of DeLong’s Algorithm for Comparing the Areas Under Correlated Receiver Operating Characteristic Curves. *IEEE Signal Process. Lett.* **2014**, *21*, 1389–1393. [[CrossRef](#)]
52. Roe, B.P.; Yang, H.J.; Zhu, J.; Liu, Y.; Stancu, I.; McGregor, G. Boosted decision trees as an alternative to artificial neural networks for particle identification. *Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrometers Detect. Assoc. Equip.* **2005**, *543*, 577–584. 10.1016/j.nima.2004.12.018. [[CrossRef](#)]
53. Wang, X.; Zhao, Y.; Pourpanah, F. Recent advances in deep learning. *Int. J. Mach. Learn. Cybern.* **2020**, *11*, 747–750. [[CrossRef](#)]
54. Stanev, D.; Riva, R.; Umassi, M. Deep Neural Network as an Alternative to Boosted Decision Trees for PID. 2021. Available online: <https://arxiv.org/abs/2104.14045> (accessed on 17 November 2021).
55. Alvestad, D.; Fomin, N.; Kersten, J.; Maeland, S.; Strümke, I. Beyond Cuts in Small Signal Scenarios—Enhanced Sneutrino Detectability Using Machine Learning. 2021. Available online: <https://arxiv.org/abs/2108.03125> (accessed on 17 November 2021).
56. Tannenwald, B.; Neu, C.; Li, A.; Buehlmann, G.; Cuddeback, A.; Hatfield, L.; Parvatam, R.; Thompson, C. Benchmarking Machine Learning Techniques with Di-Higgs Production at the LHC. 2020. Available online: <https://arxiv.org/abs/2009.06754> (accessed on 17 November 2021).
57. Heredge, J.; Hill, C.; Hollenberg, L.; Sevier, M. Quantum Support Vector Machines for Continuum Suppression in B Meson Decays. 2021. Available online: <https://arxiv.org/abs/2103.12257> (accessed on 17 November 2021).
58. Terashi, K.; Kaneda, M.; Kishimoto, T.; Saito, M.; Sawada, R.; Tanaka, J. Event Classification with Quantum Machine Learning in High-Energy Physics. *Comput. Softw. Big Sci.* **2021**, *5*, 2. [[CrossRef](#)]
59. Bendavid, J. Efficient Monte Carlo Integration Using Boosted Decision Trees and Generative Deep Neural Networks. 2017. Available online: <https://arxiv.org/abs/1707.00028> (accessed on 17 November 2021).
60. Strong, G.C. On the Impact of Selected Modern Deep-Learning Techniques to the Performance and Celerity of Classification Models in an Experimental High-Energy Physics Use Case. 2020. Available online: <https://arxiv.org/abs/2002.01427> (accessed on 17 November 2021). [[CrossRef](#)]
61. Strong, G.C. GilesStrong/Lumin: v0.8.0–Mistake Not. 2021. [[CrossRef](#)]

62. Hung, C.Y.; Chen, W.C.; Lai, P.T.; Lin, C.H.; Lee, C.C. Comparing deep neural network and other machine learning algorithms for stroke prediction in a large-scale population-based electronic medical claims database. In Proceedings of the 2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Jeju Island, Korea, 11–15 July 2017; pp. 3110–3113. [\[CrossRef\]](#)
63. Abdar, M.; Yen, N.Y.; Hung, J.C.S. Improving the Diagnosis of Liver Disease Using Multilayer Perceptron Neural Network and Boosted Decision Trees. *J. Med. Biol. Eng.* **2018**, *38*, 953–965. [\[CrossRef\]](#)
64. Botchkarev, A. Evaluating Hospital Case Cost Prediction Models Using Azure Machine Learning Studio. *CoRR* **2018**, abs/1804.01825. Available online: <https://arxiv.org/abs/1804.01825> (accessed on 17 November 2021).
65. Chen, C.; Liu, Z.; Zhou, J.; Li, X.; Qi, Y.; Jiao, Y.; Zhong, X. How Much Can A Retailer Sell? Sales Forecasting on Tmall. *CoRR* **2020**, abs/2002.11940. Available online: <https://arxiv.org/abs/2002.11940> (accessed on 17 November 2021).
66. Partin, A.; Brettin, T.; Evrard, Y.A.; Zhu, Y.; Yoo, H.; Xia, F.; Jiang, S.; Clyde, A.; Shukla, M.; Fonstein, M.; et al. Learning curves for drug response prediction in cancer cell lines. *BMC Bioinform.* **2021**, *22*, 252. [\[CrossRef\]](#)
67. Chen, Y.; Chen, W.; Pal, S.C.; Saha, A.; Chowdhuri, I.; Adeli, B.; Janizadeh, S.; Dineva, A.A.; Wang, X.; Mosavi, A. Evaluation efficiency of hybrid deep learning algorithms with neural network decision tree and boosting methods for predicting groundwater potential. *Geocarto Int.* **2021**, 1–21. [\[CrossRef\]](#)
68. Carleo, G.; Cirac, I.; Cranmer, K.; Daudet, L.; Schuld, M.; Tishby, N.; Vogt-Maranto, L.; Zdeborová, L. Machine learning and the physical sciences. *Rev. Mod. Phys.* **2019**, *91*, 045002. [\[CrossRef\]](#)
69. Rajkomar, A.; Dean, J.; Kohane, I. Machine learning in medicine. *N. Engl. J. Med.* **2019**, *380*, 1347–1358. [\[CrossRef\]](#) [\[PubMed\]](#)
70. Keith, J.A.; Vassilev-Galindo, V.; Cheng, B.; Chmiela, S.; Gastegger, M.; Müller, K.R.; Tkatchenko, A. Combining Machine Learning and Computational Chemistry for Predictive Insights Into Chemical Systems. *Chem. Rev.* **2021**, *121*, 9816–9872. [\[CrossRef\]](#) [\[PubMed\]](#)
71. Qian, X.; Yang, R. Machine learning for predicting thermal transport properties of solids. *Mater. Sci. Eng. R Rep.* **2021**, *146*, 100642. [\[CrossRef\]](#)
72. Nasiri, S.; Khosravani, M.R. Machine learning in predicting mechanical behavior of additively manufactured parts. *J. Mater. Res. Technol.* **2021**, *14*, 1137–1153. 10.1016/j.jmrt.2021.07.004. [\[CrossRef\]](#)
73. Dudley, J.J.; Kristensson, P.O. A Review of User Interface Design for Interactive Machine Learning. *ACM Trans. Interact. Intell. Syst. (TiiS)* **2018**, *8*, 1–37. [\[CrossRef\]](#)