

RECEIVED: February 9, 2021

REVISED: August 20, 2021

ACCEPTED: October 7, 2021

PUBLISHED: October 25, 2021

Learning to isolate muons

**Julian Collado,^a Kevin Bauer,^b Edmund Witkowski,^b Taylor Faucett,^b
Daniel Whiteson^b and Pierre Baldi^a**

^a*Department of Computer Science, University of California,
Irvine, CA, 92697, U.S.A.*

^b*Department of Physics and Astronomy, University of California,
Irvine, CA 92697, U.S.A.*

E-mail: julius506@gmail.com, kbauer1@uci.edu, witkows@uci.edu,
tfaucett@uci.edu, daniel@uci.edu, pfbaldi@uci.edu

ABSTRACT: Distinguishing between prompt muons produced in heavy boson decay and muons produced in association with heavy-flavor jet production is an important task in analysis of collider physics data. We explore whether there is information available in calorimeter deposits that is not captured by the standard approach of isolation cones. We find that convolutional networks and particle-flow networks accessing the calorimeter cells surpass the performance of isolation cones, suggesting that the radial energy distribution and the angular structure of the calorimeter deposits surrounding the muon contain unused discrimination power. We assemble a small set of high-level observables which summarize the calorimeter information and close the performance gap with networks which analyze the calorimeter cells directly. These observables are theoretically well-defined and can be studied with collider data.

KEYWORDS: Jet substructure, Hadron-Hadron scattering (experiments), Lepton production, QCD

ARXIV EPRINT: [2102.02278](https://arxiv.org/abs/2102.02278)

Contents

1	Introduction	1
2	Approach and dataset	2
2.1	Data generation	3
3	Networks and performance	4
4	Analysis	5
4.1	Search strategy	6
4.2	IRC safe observables	8
4.3	IRC-unsafe observables	9
5	Discussion	10
6	Conclusions	11
A	Neural network architectures	12
A.1	Muon image networks	12
A.2	Particle-flow networks	13
A.3	Isolation cone and EFP networks	13
B	ADO comparison	13

1 Introduction

Searches for new physics and precision tests of the Standard Model at hadron colliders have long relied on leptonic decays of heavy bosons, due to the relatively low background rates and excellent momentum resolution compared to hadronic final states. In the case of muons, the primary source of background to prompt muons (those from W, Z or other bosons) is production within a heavy-flavor jet. This non-prompt background is largest at lower values of muon transverse momentum, which has become important in searches for supersymmetry [1–3] as well as low-mass resonances [4].

The current state of the art strategy for distinguishing prompt and non-prompt muons in experimental searches involves techniques which integrate information from multiple detector systems [5, 6]. Critical to these strategies is the concept of isolation, which is sensitive to the presence of an associated jet that produces many tracks and calorimeter deposits. While the entire detector is worth studying [5], here we focus on the nature of the

information available in the calorimeter. There, the traditional approach is to use a robust and simple method, measuring:

$$I_\mu(R_0) = \sum_{i, R < R_0} \frac{p_T^{\text{cell } i}}{p_T^{\text{muon}}}$$

within a cone $R = \sqrt{\Delta\phi^2 + \Delta\eta^2} < R_0$ surrounding the muon [7]. Typically a single cone is used, with values of R_0 in the 0.1–0.45 range. This approach relies on identifying a typical characteristic of the signal, low calorimeter activity in the vicinity of the muon.

The traditional strategy, however, focuses on the simple nature of the signal and may overlook the rich set of characteristics offered by the background object, which can provide handles for additional rejection power. Related work, which approaches similar object classification tasks as a background jet rejection problem, has shown significant improvement in background discrimination when applied to photons [8, 9], pions [10] or electrons [11]. Other studies have shown that muons which fail the traditional isolation requirement can contain power to reveal new physics [12].

At the same time, there have been significant advances in machine learning techniques and their applications in physics [13, 14], specifically in the context of jet classification tasks, which take a fuller view of the object by directly analyzing the low-level calorimeter energy deposits, representing them either as a type of image [15, 16] or as a list [17].

It seems likely, therefore, that these machine learning strategies may identify the presence of significant additional calorimetric rejection power in the context of prompt muon identification. In this paper, we apply machine learning tools similar to those developed for jet calorimeter analysis to the task of distinguishing muons due to heavy boson decay from those produced within a heavy-flavor jet, analyze the nature of the information being used, and assemble a set of interpretable calorimeter features which capture that additional classification power.

2 Approach and dataset

The observable $I_\mu(R_0)$ is a powerful discriminator which reduces a large amount of information to a single high-level scalar. However, it is possible that it fails to capture the fullness of the calorimeter information available to distinguish prompt muons from those which are produced within a jet. To probe whether information has been lost, we compare the performance of deep neural networks which access the full calorimeter information to shallow networks which use one or more isolation cones.

Neural network decisions are notoriously difficult to reverse-engineer [13, 18–21], especially when the dimensionality of the data is large, as is the case for networks which directly use the low-level calorimeter cells. Understanding the nature of the decisions is particularly vital when the training is done with simulated samples, as it leads to valid concerns about the application of such complex strategies to collider data.

In this study, our goal is not to develop deep networks for use in collider data. Instead, we apply these deep networks as a probe, to measure a loose upper bound on the possible

classification performance, and provide insight into whether information has been lost in the reduction of the calorimeter cells to isolation cones.

Where information has been lost, we attempt to capture it, not by applying the deep network, but by assembling a small set of new high-level (HL) observables that bridge the performance gap and reproduce the classification decisions of the calorimeter cell networks [22]. These high-level observables are more compact, physically interpretable, can be validated in data, and allow for the straightforward assessment and propagation of systematic uncertainties.

2.1 Data generation

Samples of simulated prompt muons were generated via the process $pp \rightarrow Z' \rightarrow \mu^+ \mu^-$ with a Z' mass of 20 GeV. Non-prompt muons were generated via the process $pp \rightarrow b\bar{b}$. Both samples are generated at a center of mass energy $\sqrt{s} = 13$ TeV. Collisions and heavy boson decays are simulated with MADGRAPH5 v2.6.5 [23], showered and hadronized with PYTHIA v8.235 [24], and the detector response simulated with DELPHES v3.4.1 [25] using the standard ATLAS card and ROOT version 6.0800 [26]. The classification of these objects is sensitive to the presence of additional proton interactions, referred to as pile-up events. We overlay such interactions within the simulation with an average number of interactions per event of $\mu = 50$, as an estimate of LHC Run 2 experimental data.

Muons in the range $p_T \in [10, 15]$ GeV with $|\eta| < 2.53$ were considered; see figure 1. To avoid inducing biases from artifacts of the generation process, signal and background events are weighted such that the distributions in p_T and η are uniform, using 32 bins in each dimension. Only events where a muon is identified as a track in the muon spectrometer are used. In total, 499,970 events were used, where 249,991 were signal and 249,979 were background. Both the signal and background datasets are randomly split as: 83% training, 8.5% validation, and 8.5% testing sets.

Calorimeter deposits can be represented as images where each pixel value represents the E_T deposited by a particle [15]. Images are formed by considering cells in the calorimeter within a cone of radius up to $\Delta R = 0.45$ surrounding the muon location after propagating to the radius of the calorimeter.

We use a 32x32 grid, which approximately corresponds with the calorimeter granularity of ATLAS and CMS. Heat maps of the calorimeter energy deposits in $\eta - \phi$ space for both signal prompt muons and background non-prompt muons are shown in figure 2. The signal calorimeter deposits are uniform and can be attributed to pileup whereas the background deposits appear largely radially symmetric with a dense core from the jet.

We calculate the standard muon isolation observable $I_\mu(R_0)$ for a set of cones with $0.025 \leq R_0 \leq 0.45$ in 18 equally spaced steps.

Crucially, these isolation observables and all other calorimeter observables are calculated directly from the pixels of the muon images, ensuring that they contain a strict subset of the information available. This allows for direct and revealing comparisons of the performance between networks trained with the images and those trained with I_μ . Note that pixelization of the detector may incur some loss of information relative to the underlying segmentation

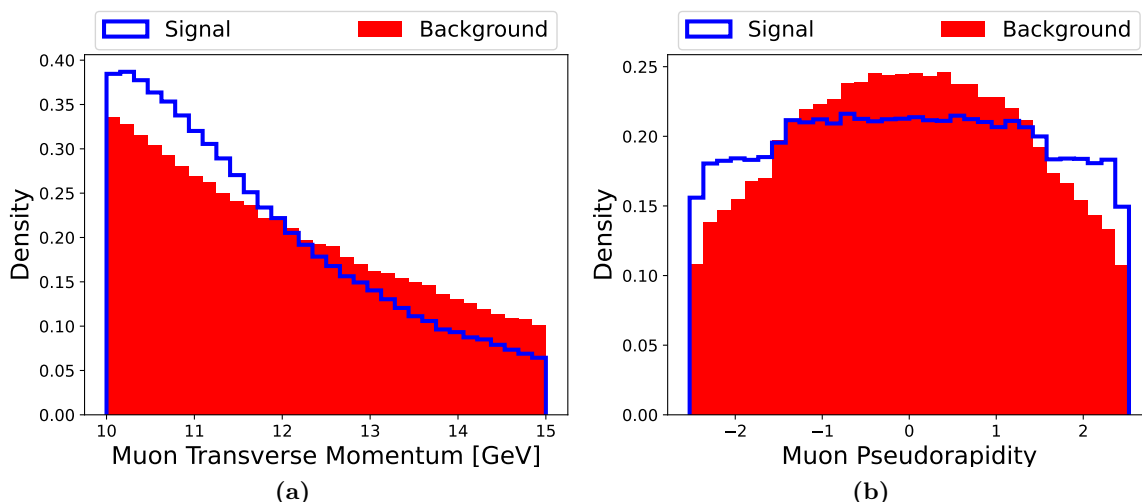


Figure 1. Distributions of muon transverse momentum (top) and pseudorapidity (bottom) for signal and background samples. Afterwards, the distributions are weighted to make both samples uniform.

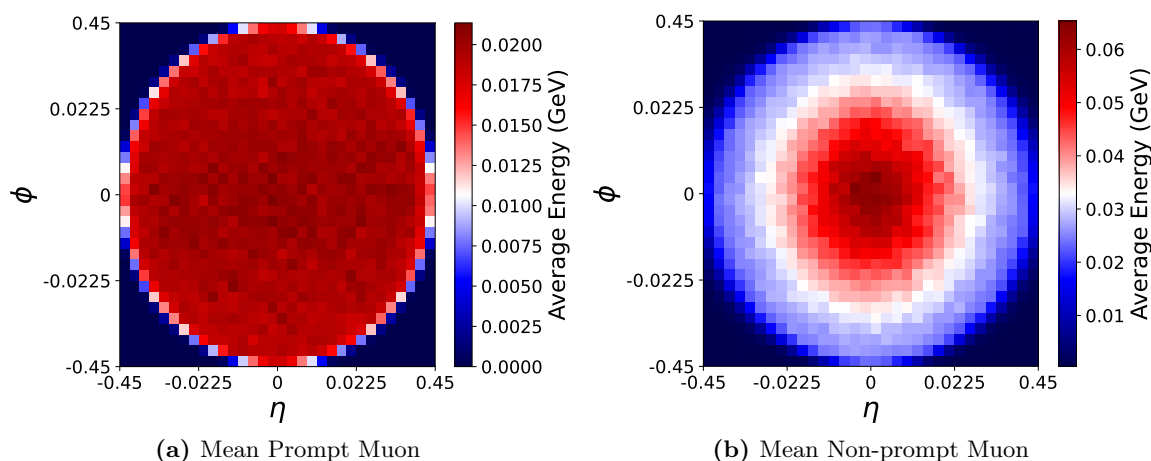


Figure 2. Mean calorimeter images for signal prompt muons (top) and muons produced within heavy-flavor jets (bottom), in the vicinity of reconstructed muons within a cone of $R = 0.45$. The color of each cell represents the sum of the E_T of the calorimeter deposits within the cell.

of the calorimeter. However, this work focuses on examining the relative power of different techniques, rather than identifying the best performance under the most realistic scenario.

3 Networks and performance

We apply several strategies to the task of classifying prompt and non-prompt muons, using both low-level calorimeter information and higher-level isolation quantities. We evaluate the performance of each approach by comparing the integral of the ROC (Receiver Operating Characteristic) curve, known as the AUC (Area Under the Curve). The uncertainty

for the AUC is calculated by training 100 randomly initialized models with the same hyperparameters on different bootstraps of the data. In this case, we seek to determine the statistical uncertainty due to the stochastic training method, rather than any systematic uncertainty due to the calorimeter resolution.

For the high-level quantities, the standard approach of using a single isolation cone yields an AUC of 0.787 for the optimal cone size, $R_0 = 0.425$.¹ We hypothesized that additional cones would provide useful information about the radial energy distribution. Including a second cone with a distinct R_0 value as input to a small neural network (see appendix A) slightly improves performance, with an AUC of 0.793. To estimate the full information available in the cones, we perform a greedy search through all 18 cones; we find that a set of 10 cones² yields another small boost in classification power up to an AUC of 0.803, as shown in figure 3. Performance was fairly insensitive to the specific choices of cone sizes, and does not grow significantly beyond 10 cones. Feed-forward dense networks are trained to use the information in one or more isolation cones (see the appendix for details on network architectures and training).

We next examine whether additional information is available by applying strategies which access the calorimeter information at the lowest-level and highest-dimensionality. Convolutional networks (CNN) are applied to the muon images [14–16]. As an alternative, we apply particle-flow networks (PFN) [17], which are mathematically structured as sums over inputs and thus are invariant to permutations of the inputs.

The muon image CNN achieves a significantly higher performance than the isolation-only networks, with an AUC of 0.841, and the particle flow network reaches 0.857, see figure 4 and table 1. This immediately suggests that there is significant additional information available to distinguish between the prompt and non-prompt muons beyond what is summarized in the isolation cones. A more restricted version of the PFN, an Energy-Flow Network [17] (EFN), which enforces infra-red and collinear (IRC) safety, achieves nearly the same performance, 0.849. This suggests that most of the additional information beyond the isolation cones is IRC-safe.

These results support the conventional wisdom that a significant fraction of the information relevant for classification is captured by a single, simple cone. However, they also indicate that there is additional information in the radial distribution of energy, which can be captured by using multiple cones. However, even many cones fail to match the performance of the networks which use the calorimeter cell information directly, suggesting that there is additional non-radial information relevant to the classification task not captured in the isolation cones. This is likely due to a difference between the muon axis, the center of the isolation cones, and the jet axis.

4 Analysis

The networks which use the calorimeter cells directly have the most powerful performance, but our aim is not simply to optimize classification performance in this particular simulated

¹Similar performance was seen for other cone sizes.

² $R_0 = [0.025, 0.05, 0.075, 0.125, 0.15, 0.225, 0.275, 0.325, 0.425, 0.45]$.

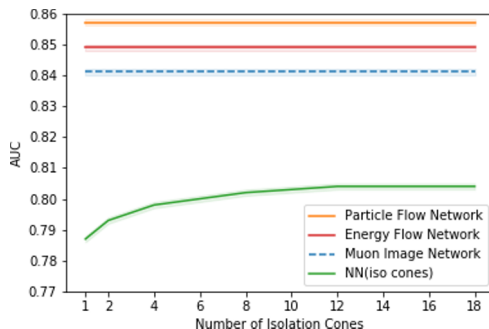


Figure 3. Comparison of classification performance using the metric AUC between Particle-Flow networks trained on lists of calorimeter deposits (orange, solid), Energy-Flow networks trained on lists of calorimeter deposits (red, solid), convolutional networks trained on muon images (blue, dashed) and networks which use increasing numbers of isolation cones (green, solid). For each number of cones, the optimal set is chosen.

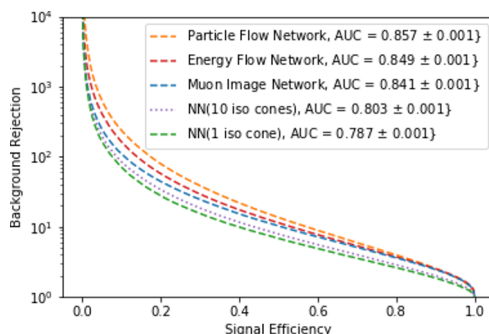


Figure 4. Background rejection versus signal efficiency for Particle-Flow networks trained on lists of calorimeter deposits (orange, dashed), Energy-Flow networks trained on lists of calorimeter deposits (red, dashed), convolutional networks trained on muon images (blue, dashed), networks trained on a set of isolation cones (purple, dotted) and the benchmark approach, a single isolation cone approach (green, dashed).

sample. Instead, we seek to understand the nature of the learned strategy in order to validate it and translate it into simpler, more easily interpretable high-level features which can be studied in other datasets, real or simulated. In addition, this understanding can reveal how well the strategy is likely to generalize to other kinds of jets that are not represented by this background sample, such as charm jets.

The CNN and PFN results indicate that the radially symmetric isolation cones are failing to utilize some information which is relevant to the classification task. In this section, we search for additional high-level observables which capture this information.

4.1 Search strategy

Interpreting the decisions of a deep network with a high-dimensional input vector is notoriously difficult. Instead, we attempt to translate its performance into a smaller set of interpretable observables [22]. This allows us to understand the nature of the information being used as well as to represent it more compactly.

One might imagine exploring a set of physically-motivated quantities, such as the relative p_T between the jet and the muon or the energy-weighted average distance between the jet and calorimeter cells. These particular quantities were considered and found to not contribute significant power in addition to the isolation cones.

Instead, we use a systematic approach and explore a formally complete set of observables. As the background non-prompt muons are due to jet production, we search within a set of observables originally intended for analysis of jets: the Energy Flow Polynomials (EFPs) [27], a formally infinite set of parameterized engineered functions, inspired by previous work on energy correlation functions [28], which sum over the contents of the cells scaled by relative angular distances. An EFP for a jet with M constituents which considers N correlators with angular connections k, l is written as:

$$\text{EFP} = \sum_{i_1=1}^M \dots \sum_{i_N=1}^M z_{i_1}^\kappa \dots z_{i_N}^\kappa \prod_{k,l} \theta_{i_k i_l}^\beta$$

where

$$(z_i)^\kappa = \left(\frac{p_{Ti}}{\sum_j p_{Tj}} \right)^\kappa, \tag{4.1}$$

$$\theta_{ij}^\beta = \left(\Delta\eta_{ij}^2 + \Delta\phi_{ij}^2 \right)^{\beta/2}. \tag{4.2}$$

Here, p_{Ti} is the transverse momentum of cell i , and $\Delta\eta_{ij}$ ($\Delta\phi_{ij}$) is the pseudorapidity (azimuth) difference between cells i and j . These parametric sums correspond to the set of all isomorphic multigraphs where:

$$\text{each node} \Rightarrow \sum_{i=1}^N z_i, \tag{4.3}$$

$$\text{each } k\text{-fold edge} \Rightarrow (\theta_{ij})^k. \tag{4.4}$$

As the EFPs are normalized, they capture only the relative information about the energy deposition. For this reason, in each network that includes EFP observables, we include as an additional input the sum of p_T over all cells, to indicate the overall scale of the energy deposition.

The original IRC-safe EFPs require $\kappa = 1$. To more broadly explore the space, we consider examples with $\kappa \neq 1$ to explore a broader space of observables.³

In principle, the space spanned by the EFPs is complete, such that any jet observable can be described by one or more EFPs of some degree. One might consider simply searching this space for all possible combinations of EFPs for a set which maximizes performance for this task. Such a search is computationally prohibitive; instead, we follow the black-box guided algorithm of ref. [22], which iteratively assembles a set of EFPs that mimic the decisions of another guiding network (the PFN in our case) by isolating the portion of the

³Also, note that $\kappa > 0$ generically corresponds to IR-safe but C-unsafe observables. For $\kappa < 0$, empty cells are omitted from the sum.

input space where the guiding network disagrees with the isolation network, and finding EFPs which mimic the guiding network’s decisions in that subspace.


Here, the agreement between networks $f(x)$ and $g(x)$ is evaluated over pairs of (x, x') by comparing their relative classification decisions, expressed mathematically as:

$$\text{DO}[f, g](x, x') = \Theta\left(\left(f(x) - f(x')\right)\left(g(x) - g(x')\right)\right), \quad (4.5)$$

and referred to as *decision ordering* (DO). A DO= 0 corresponds to inverted decisions over all input pairs and DO= 1 corresponds to the same decision ordering. As prescribed in ref. [22], we scan the space of EFPs to find the observable that has the highest average decision ordering (ADO) with the guiding network when averaged over disordered pairs. The selected EFP is then incorporated into the new network of HL features, HLN_{n+1} , and the process is repeated until the ADO plateaus.


4.2 IRC safe observables


As the elements of the EFP space are not orthogonal, there are potentially many combinations of EFP observables which capture the relevant information. As simpler EFPs may be more conducive to theoretical interpretation, we begin our search in a restricted subset of the EFP space. Specifically, we consider those which are IRC safe ($\kappa = 1$), have a simple angular weighting ($\beta \in [1, 2]$), and $n \leq 3$ fewer nodes with at most three edges between nodes. We also include $\sum p_T$, where the summation is over all calorimeter cells in the image, to set the scale accompanying the normalized EFPs. The first EFP observable identified is a simple three-point correlator:




$$(\kappa=1, \beta=1) = \sum_{a,b,c=1}^N z_a z_b z_c \theta_{ab} \theta_{bc} \theta_{ca}$$

which, when combined with the isolation cones and $\sum p_T$, yields an AUC of 0.838 and an ADO with the CNN of 0.891, a significant boost relative to just using the radial information of the isolation cones. The subsequent scans produce variants of this observable:



$$(\kappa=1, \beta=2) = \sum_{a,b,c=1}^N z_a z_b z_c \theta_{ab}^4 \theta_{bc}^6$$


$$(\kappa=1, \beta=1) = \sum_{a,b,c=1}^N z_a z_b z_c \theta_{ab}^2 \theta_{bc}^3$$


$$(\kappa=1, \beta=2) = \sum_{a,b=1}^N z_a z_b \theta_{ab}$$

with additional edges corresponding to higher powers of the angular information. Their power may come from their sensitivity to the collimated radiation pattern of the jet. Together with the isolation cones, these observables reach an AUC of 0.842 and an ADO with the PFN of 0.888, see table 1.

This set of observables partially closes the performance gap with the best calorimeter cell networks, indicating that angular information is relevant to the muon isolation classification

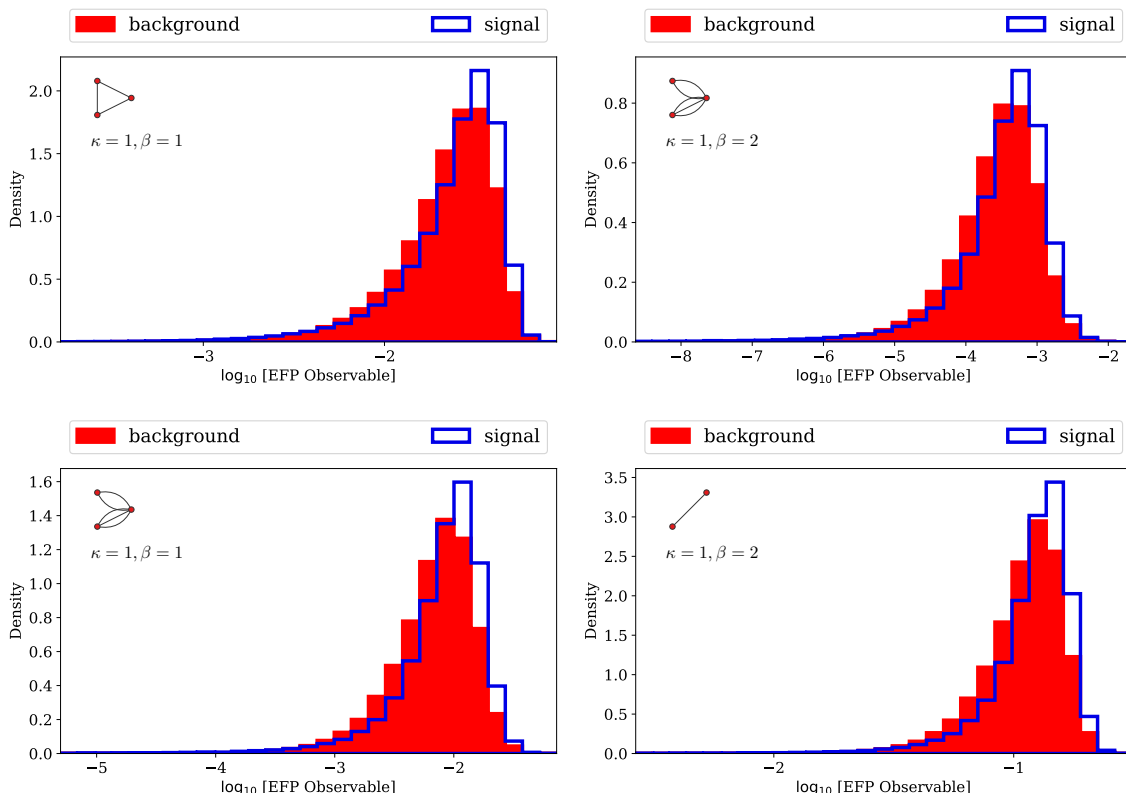


Figure 5. Distributions of the \log_{10} of the selected IRC-safe EFPs as chosen by the black-box guided strategy, for prompt (signal) muons and non-prompt (background) muons.

task, but fails to fully match its performance. Distributions of these EFPs for signal and background are shown in figure 5. Further scans in this limited space do not yield significant boost in AUC or ADO values. The strong result of the IRC-safe EFN indicates that it is possible to capture nearly all of the classification power using IRC-safe graphs, likely requiring graphs with complexity beyond what we have considered.

A scan guided by the CNN rather than the PFN yields very similar results, with identical choices for the first three EFPs.

4.3 IRC-unsafe observables

To understand the nature of the remaining information used by the PFN but not captured by the isolation cones and the IRC-safe observables, we expand the search space to include observables which are not IRC safe ($\kappa \in [-1, 0, \frac{1}{4}, \frac{1}{2}, 1, 2]$), with alternative angular powers ($\beta \in [\frac{1}{4}, \frac{1}{2}, 1, 2, 3, 4]$) and with up to $n = 7$ nodes and $d = 7$ edges.

A scan of these observables finds a set of 5 which, when combined with the isolation cones and $\sum p_T$ reach an AUC of 0.857. Figure 6 shows the EFP graphs as well as their distributions for prompt and non-prompt muons. They include single point-graphs, with no angular powers, as well as a two-point correlators with large angular power sensitive to high-angle effects, and more complex graphs with multiple nodes. We note that due to the overlapping nature of the large space of EFPs, there are several sets of EFPs which achieve

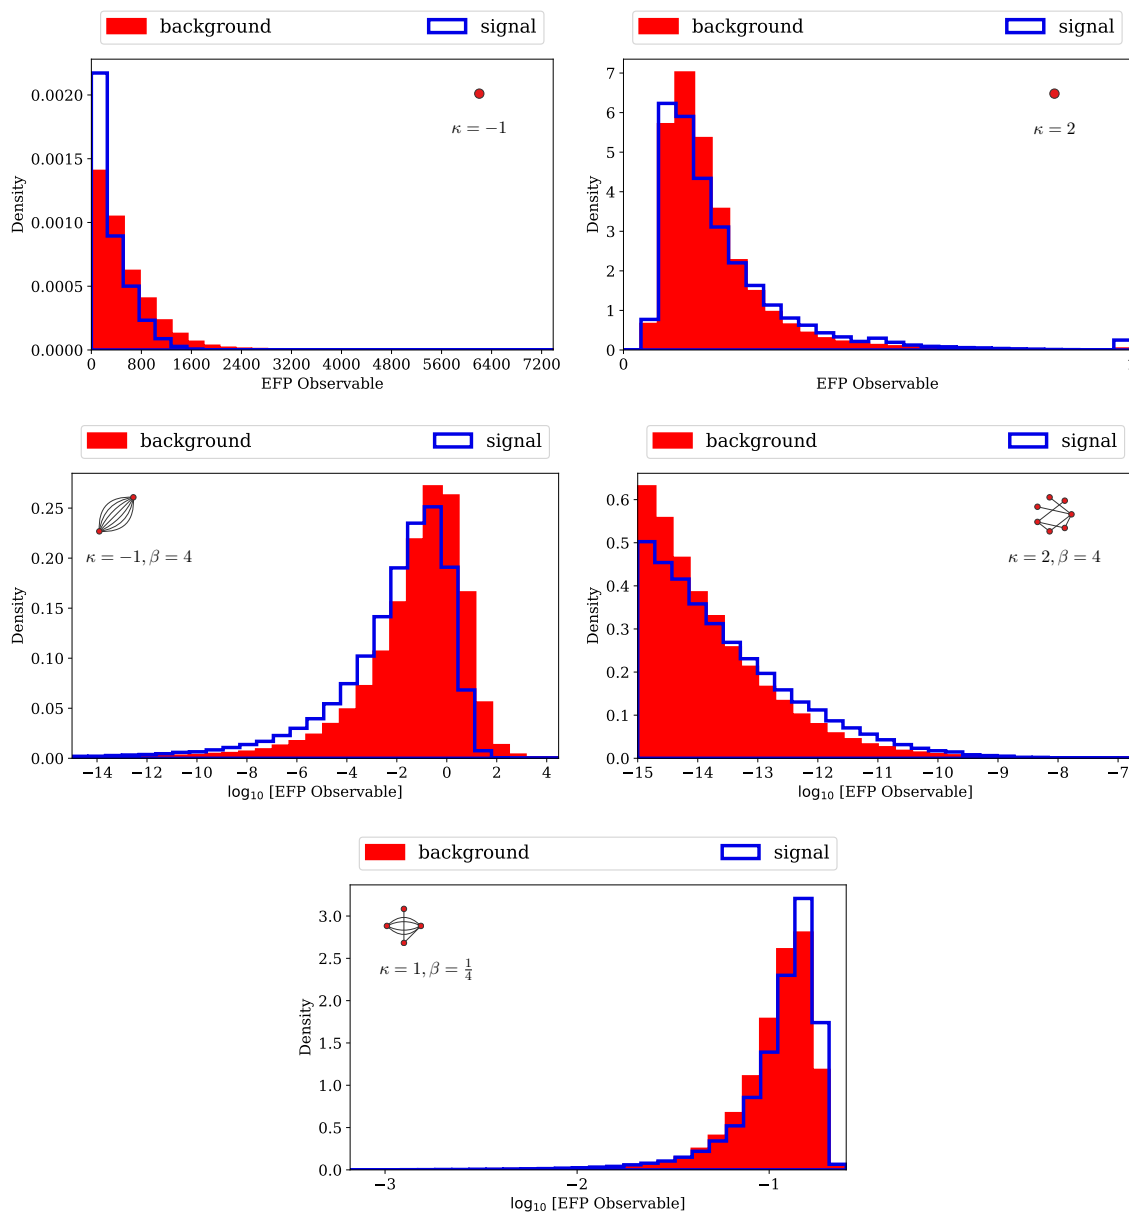


Figure 6. Distributions of the \log_{10} of the selected EFPs as chosen by the black-box guided strategy, for prompt (signal) muons and non-prompt (background) muons.

similar performance. Again, a similar scan guided by the CNN rather than the PFN yields very similar results.

5 Discussion

The performance of the networks which use the low-level calorimeter cells indicates that information exists in these cells which is not captured by the isolation cones, see table 1. A guided search through the space of IRC-safe EFPs closes most of the gap between these networks, giving us some insight as to the nature of the information. A broader search is

Method	N_{inputs}	AUC	ADO [PFN]	N_{Params}
Single Iso Cone	1	0.787	0.860	40k
10 Iso	10	0.803	0.877	41k
10 Iso, $\sum p_T$	11	0.807	0.884	42k
10 Iso, $\sum p_T$, 4 simple EFPs	15	0.842	0.888	42k
10 Iso, $\sum p_T$, 5 EFPs	16	0.857	0.900	43k
Calo image CNN	1024	0.841	0.950	167k
Calo cell Energy-Flow Net	102	0.849	0.951	453k
Calo cell Particle-Flow Net	102	0.857	1	453k

Table 1. Summary of performance (AUC) in the prompt muon classification task for various network architectures and input features. Statistical uncertainty in each case is ± 0.001 with 95% confidence, measured using bootstrapping over 100 models. Uncertainty due to the initial conditions of the network is found to be negligible. Also shown are the number of inputs to and parameters of each network.

able to complete the bridge, yielding the same performance as the low-level network, but employing IRC-unsafe EFPs. The multi-point correlators may be sensitive to the width of the jet, due to the momentum of the constituents relative to the jet axis, as a result of b - and c -quark decays.

A comparison of the network complexity for the various approaches is shown in table 1. The set of high-level features (isolation cones and EFP graphs) matches the PFN performance with 10 times fewer parameters, supporting the notion that the high-level features are effectively summarizing the relevant low-level information.

6 Conclusions

We have applied deep networks to low-level calorimeter deposits surrounding prompt and non-prompt muons in order to estimate the amount of classification power available and to probe whether the standard methods are fully capturing the relevant information.

The performance of the calorimeter cell networks significantly exceeds the benchmark approach, a single isolation cone. The use of several isolation cones provides some improvement, suggesting that there is additional useful information in the full radial energy distribution. However, a substantial gap remains, hinting that there is non-radial structure in the calorimeter cells which provides useful information for classification. We map the strategy of the calorimeter cell networks into a set of energy flow polynomials, finding four IRC-safe, simple three-point correlators which capture a significant amount of the missing information. As they are simple functions of the energy deposition, they can be physically interpreted, and the fidelity of their modeling can be studied in control regions in collider data. Any boost in the efficiency to identify prompt muons is extremely valuable to searches

at the LHC, especially those with multiple leptons, where event-level efficiencies depend sensitively on object-level efficiencies.

Additional, more complex EFPs provide a further modest boost in performance, closing the gap with the PFN. The strong performance of the IRC-safe EFN suggests that most of the additional information beyond the isolation cones is IRC-safe.

More broadly, the existence of a gap between the performance of state-of-the-art high-level features and networks using lower-level calorimeter information represents an opportunity to gather additional power in the battle to suppress lepton backgrounds. Rather than employing black-box deep networks directly, we have demonstrated the power of using them to identify the relevant observables from a large list of physically interpretable options. This allows the physicist to understand the nature of the information being used and to assess its systematic uncertainty. Here we have focused on two-dimensional projections of the calorimeter response, but longitudinal information expressed in three dimensions may offer additional power in future work. While these studies were performed with simulated samples, similar studies can be performed using unsupervised methods [29, 30] on samples of collider data, which we leave to future studies.

Acknowledgments

We would like to thank Michael Fenton, Dan Guest and Jesse Thaler for providing valuable feedback and insightful comments and Yuzo Kanomata for computing support. We also wish to acknowledge a hardware grant from NVIDIA. This material is based upon work supported by the National Science Foundation under grant number 1633631. DW is supported by the DOE Office of Science. The work of JC and PB in part supported by grants NSF 1839429 and NSF NRT 1633631 to PB.

A Neural network architectures

All networks were trained in Tensorflow [31] and Keras [32]. The networks were optimized with Adam [33] for up to 100 epochs with early stopping. For all networks except the PFNs, the weights were initialized using orthogonal weights [34]. Hyperparameters were optimized using Bayesian optimization with the Sherpa hyperparameter optimization library [35]. The variables and ranges for the hyperparameters are shown in tables 2 and 3.

Below are further details regarding the networks which use images and those which use isolation and EFP observables.

A.1 Muon image networks

The pixelated images were preprocessed to have zero mean and unit standard deviation. We tried rotating the images as in [16] but performance was considerably lowered by this preprocessing step. The best muon image network structure begins with three convolutional blocks. Each block contains three convolutional layers with 48 filters with rectified linear units [36], followed by a 2x2 pooling layer. Afterwards there are two fully connected layers with 74 rectified linear units and a final layer with a sigmoidal logistic activation function

Parameter	Range	Value
Num. of convolutional blocks	[1, 4]	3
Num. of filters	[16, 128]	48
Num. of fully connected layers	[2, 4]	2
Number of hidden units	[25, 200]	74
Learning rate	[0.0001, 0.01]	0.0003
Dropout	[0.0, 0.5]	0.2388

Table 2. Hyperparameter ranges for bayesian optimization of convolutional networks.

Parameter	Range	ISO Value
Num. of layers	[2, 8]	2
Num. of hidden units	[1, 200]	197
Learning rate	[0.0001, 0.01]	0.0003
Dropout	[0.0, 0.5]	0.0547

Table 3. Hyperparameter ranges for Bayesian optimization of fully connected networks.

to classify signal vs background. The model had dropout [37, 38] with value 0.2388 on the fully connected layers and an initial learning rate of 0.0003 and batch size of 128.

A.2 Particle-flow networks

The Particle Flow Network (PFN) is trained using the `energyflow` package [17]. Input features are taken from the muon image pixels and preprocessed by subtracting the mean and dividing by the variance. The PFN uses 3 dense layers in the per-particle frontend module and 3 dense layers in the backend module. Each layer uses 100 nodes, `relu` activation and `glorot_normal` initializer. The final output layer uses a sigmoidal logistic activation function to predict the probability of signal or background. The `Adam` optimizer is used with a learning rate of 0.0001 and trained with a batch size of 128.

A.3 Isolation cone and EFP networks

The isolation inputs and EFPs are preprocessed by subtracting the mean and dividing by the variance. We trained neural networks with two to eight fully connected hidden layers depending on the hyperparameter value and a final layer with a sigmoidal logistic activation function to predict the probability of signal or background.

For the minimal set of isolation inputs, the best model we found had 2 fully connected layers with 197 rectified linear hidden units [36] and a learning rate of 0.0003 and dropout rate of 0.0547.

B ADO comparison

In figure 7, the ADO between the various networks is shown.

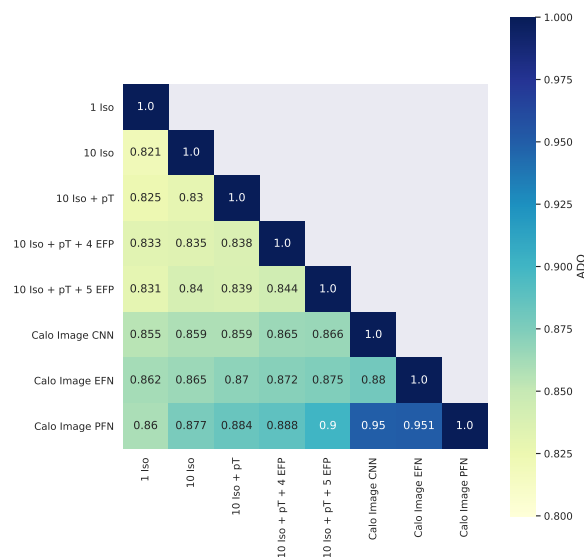


Figure 7. Comparison of the similarity of decisions made by pairs of networks, as quantified by the Average Decision Ordering (ADO) [22], defined in the text.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY 4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] ATLAS collaboration, *Search for electroweak production of supersymmetric states in scenarios with compressed mass spectra at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **97** (2018) 052010 [[arXiv:1712.08119](#)] [[INSPIRE](#)].
- [2] ATLAS and CMS collaborations, *Search for supersymmetry with extremely compressed spectra with the ATLAS and CMS detectors*, *Nucl. Part. Phys. Proc.* **273-275** (2016) 631 [[INSPIRE](#)].
- [3] CMS collaboration, *Search for supersymmetry in the vector-boson fusion topology in proton-proton collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **11** (2015) 189 [[arXiv:1508.07628](#)] [[INSPIRE](#)].
- [4] I. Hoenig, G. Samach and D. Tucker-Smith, *Searching for dilepton resonances below the Z mass at the LHC*, *Phys. Rev. D* **90** (2014) 075016 [[arXiv:1408.1075](#)] [[INSPIRE](#)].
- [5] CMS collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, *2017 JINST* **12** P10003 [[arXiv:1706.04965](#)] [[INSPIRE](#)].
- [6] J. Pata, J. Duarte, J.-R. Vlimant, M. Pierini and M. Spiropulu, *MLPF: Efficient machine-learned particle-flow reconstruction using graph neural networks*, *Eur. Phys. J. C* **81** (2021) 381 [[arXiv:2101.08578](#)] [[INSPIRE](#)].
- [7] ATLAS collaboration, *Muon reconstruction performance of the ATLAS detector in proton-proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **76** (2016) 292 [[arXiv:1603.05598](#)] [[INSPIRE](#)].
- [8] LHCb collaboration, *Search for Dark Photons Produced in 13 TeV pp Collisions*, *Phys. Rev. Lett.* **120** (2018) 061801 [[arXiv:1710.02867](#)] [[INSPIRE](#)].

- [9] E. Hall and J. Thaler, *Photon isolation and jet substructure*, *JHEP* **09** (2018) 164 [[arXiv:1805.11622](#)] [[INSPIRE](#)].
- [10] ATLAS collaboration, *Deep Learning for Pion Identification and Energy Calibration with the ATLAS Detector*, Tech. Rep. [ATL-PHYS-PUB-2020-018](#) (2020).
- [11] J. Collado, J.N. Howard, T. Faucett, T. Tong, P. Baldi and D. Whiteson, *Learning to identify electrons*, *Phys. Rev. D* **103** (2021) 116028 [[arXiv:2011.01984](#)] [[INSPIRE](#)].
- [12] C. Brust, P. Maksimovic, A. Sady, P. Saraswat, M.T. Walters and Y. Xin, *Identifying boosted new physics with non-isolated leptons*, *JHEP* **04** (2015) 079 [[arXiv:1410.0362](#)] [[INSPIRE](#)].
- [13] P. Baldi, P. Sadowski and D. Whiteson, *Searching for Exotic Particles in High-Energy Physics with Deep Learning*, *Nature Commun.* **5** (2014) 4308 [[arXiv:1402.4735](#)] [[INSPIRE](#)].
- [14] P. Baldi, *Deep Learning in Science*, Cambridge University Press, Cambridge, U.K. (2021).
- [15] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](#)] [[INSPIRE](#)].
- [16] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet Substructure Classification in High-Energy Physics with Deep Neural Networks*, *Phys. Rev. D* **93** (2016) 094034 [[arXiv:1603.09349](#)] [[INSPIRE](#)].
- [17] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy Flow Networks: Deep Sets for Particle Jets*, *JHEP* **01** (2019) 121 [[arXiv:1810.05165](#)] [[INSPIRE](#)].
- [18] S. Chang, T. Cohen and B. Ostdiek, *What is the Machine Learning?*, *Phys. Rev. D* **97** (2018) 056009 [[arXiv:1709.10106](#)] [[INSPIRE](#)].
- [19] T. Roxlo and M. Reece, *Opening the black box of neural nets: case studies in stop/top discrimination*, [arXiv:1804.09278](#) [[INSPIRE](#)].
- [20] S. Wunsch, R. Friese, R. Wolf and G. Quast, *Identifying the relevant dependencies of the neural network response on characteristics of the input space*, *Comput. Softw. Big Sci.* **2** (2018) 5 [[arXiv:1803.08782](#)] [[INSPIRE](#)].
- [21] G. Agarwal et al., *Explainable AI for ML jet taggers using expert variables and layerwise relevance propagation*, *JHEP* **05** (2021) 208 [[arXiv:2011.13466](#)] [[INSPIRE](#)].
- [22] T. Faucett, J. Thaler and D. Whiteson, *Mapping Machine-Learned Physics into a Human-Readable Space*, *Phys. Rev. D* **103** (2021) 036020 [[arXiv:2010.11998](#)] [[INSPIRE](#)].
- [23] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [24] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 Physics and Manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [25] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [26] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81 [[INSPIRE](#)].
- [27] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [[INSPIRE](#)].

- [28] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [29] L.M. Dery, B. Nachman, F. Rubbo and A. Schwartzman, *Weakly Supervised Classification in High Energy Physics*, *JHEP* **05** (2017) 145 [[arXiv:1702.00414](#)] [[INSPIRE](#)].
- [30] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: Learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](#)] [[INSPIRE](#)].
- [31] M. Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*, [arXiv:1603.04467](#) [[INSPIRE](#)].
- [32] F. Chollet et al., *Keras*, <https://keras.io> (2015).
- [33] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [34] A.M. Saxe, J.L. McClelland and S. Ganguli, *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*, [arXiv:1312.6120](#).
- [35] L. Hertel, J. Collado, P. Sadowski, J. Ott and P. Baldi, *Sherpa: Robust hyperparameter optimization for machine learning*, *SoftwareX* **12** (2020) 100591 [[arXiv:2005.04048](#)].
- [36] X. Glorot, A. Bordes and Y. Bengio, *Deep sparse rectifier neural networks*, in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, U.S.A., 11–13 April 2011, *Proc. Mach. Learn. Res.* **15** (2011) 315 [<http://proceedings.mlr.press/v15/glorot11a.html>].
- [37] N. Srivastava, G.E. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, *Dropout: a simple way to prevent neural networks from overfitting*, *J. Mach. Learn. Res.* **15** (2014) 1929.
- [38] P. Baldi and P. Sadowski, *The dropout learning algorithm*, *Artificial Intel.* **210** (2014) 78.