

RECEIVED: September 5, 2017

REVISED: October 25, 2017

ACCEPTED: November 24, 2017

PUBLISHED: December 12, 2017

Pileup Mitigation with Machine Learning (PUMML)

Patrick T. Komiske,^a Eric M. Metodiev,^a Benjamin Nachman^b and
Matthew D. Schwartz^c

^a*Center for Theoretical Physics, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge, MA, 02139 U.S.A.*

^b*Physics Division, Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA, 94720 U.S.A.*

^c*Department of Physics, Harvard University,
17 Oxford Street, Cambridge, MA, 02138 U.S.A.*

E-mail: pkomiske@mit.edu, metodiev@mit.edu, bpnachman@lbl.gov,
schwartz@physics.harvard.edu

ABSTRACT: Pileup involves the contamination of the energy distribution arising from the primary collision of interest (leading vertex) by radiation from soft collisions (pileup). We develop a new technique for removing this contamination using machine learning and convolutional neural networks. The network takes as input the energy distribution of charged leading vertex particles, charged pileup particles, and all neutral particles and outputs the energy distribution of particles coming from leading vertex alone. The PUMML algorithm performs remarkably well at eliminating pileup distortion on a wide range of simple and complex jet observables. We test the robustness of the algorithm in a number of ways and discuss how the network can be trained directly on data.

KEYWORDS: Jets

ARXIV EPRINT: [1707.08600](https://arxiv.org/abs/1707.08600)

Contents

1	Introduction	1
2	PUMML algorithm	3
3	Performance	6
4	Robustness	9
5	What is PUMML learning?	12
6	Conclusions	15

1 Introduction

The Large Hadron Collider (LHC) is operated at very high instantaneous luminosities to achieve the large statistics required to search for exotic Standard Model (SM) or beyond the SM processes as well as for precision SM measurements. At a hadron collider, protons are grouped together in bunches; as the luminosity increases for a fixed bunch spacing, the number of protons within each bunch that collide inelastically increases as well. Most of these inelastic collisions are soft, with the protons dissolving into mostly low-energy pions that disperse throughout the detector. A typical collision of this sort at the LHC will contribute about 0.6 GeV/rad^2 of energy [1, 2]. Occasionally, one pair of protons within a bunch crossing collides head-on, producing hard (high-energy) radiation of interest. At high luminosity, this hard collision, or leading vertex (LV), is always accompanied by soft proton-proton collisions called pileup. The data collected thus far by ATLAS and CMS have approximately 20 pileup collisions per bunch crossing on average ($\langle \text{NPU} \rangle \sim 20$); the data in Run 3 are expected to contain $\langle \text{NPU} \rangle \sim 80$; and the HL-LHC in Runs 4–5 will have $\langle \text{NPU} \rangle \sim 200$. Mitigating the impact of this extra energy on physical observables is one of the biggest challenges for data analysis at the LHC.

Using precision measurements, the charged particles coming from the pileup interactions can mostly be traced to collision points (primary vertices) different from that of the leading vertex. Indeed, due to the excellent vertex resolution at ATLAS and CMS [3–5] the charged particle tracks from pileup can almost completely be identified and removed.¹ This is the simplest pileup removal technique, called *charged-hadron subtraction*. The challenge with pileup removal is therefore how to distinguish neutral radiation associated with

¹Some detector systems have an integration time that is (much) longer than the bunch spacing of 25 ns, so there is also a contribution from pileup collisions happening before or after the collision of interest (out-of-time pileup). This contribution will not have charged particle tracks and can be at least partially mitigated with calorimeter timing information. Out-of-time pileup is not considered further in this analysis.

the hard collision from neutral pileup radiation.² Since radiation from pileup is fairly uniform,³ it can be removed on average, for example, using the jet areas technique [9]. The jet areas technique focuses on correcting the overall energy of collimated sprays of particles known as jets. Indeed, both the ATLAS and CMS experiments apply jet areas or similar techniques to calibrate the energy of their jets [1, 2, 10–13]. Unfortunately, for many measurements, such as those involving jet substructure or the full radiation patterns within the jet, removing the radiation on average is not enough.

Rather than calibrating only the energy or net 4-momentum of a jet, it is possible to correct the constituents of the jet. By removing the pileup contamination from each constituent, it should be possible to reconstruct more subtle jet observables. We can coarsely classify constituent pileup mitigation strategies into several categories: constituent preprocessing, jet/event grooming, subjet corrections, and constituent corrections. Grooming refers to algorithms that remove objects and corrections describe scale factors applied to individual objects. Both ATLAS and CMS perform preprocessing to all of their constituents before jet clustering. For ATLAS, pileup-dependent noise thresholds in topoclustering [14] suppresses low energy calorimeter deposits that are characteristic of pileup. In CMS, charged-hadron subtraction removes all of the pileup particle-flow candidates [15]. Jet grooming techniques are not necessarily designed to exclusively mitigate pileup but since they remove constituents or subjets in a jet (or event) that are soft and/or at wide angles to the jet axis, pileup particles are preferentially removed [6, 16–21]. Explicitly tagging and removing pileup subjets often performs comparably to algorithms without explicit pileup subjet removal [6]. A popular event-level grooming algorithm called SoftKiller [21] removes radiation below some cutoff on transverse momentum, p_T^{cut} chosen on an event-by-event basis so that half of a set of pileup-only patches are radiation free.

While grooming algorithms remove constituents and subjets, there are also techniques that try to reconstruct the exact energy distribution from the primary collision. One of the first such methods introduced was Jet Cleansing [22]. Cleansing works at the subjet level, clustering and declustering jets to correct each subjet separately based on its local energy information. Furthermore, Cleansing exploits the fact that the relative size of pileup fluctuations decreases as $\langle \text{NPU} \rangle \rightarrow \infty$ so that the neutral pileup-energy content of subjets can be estimated from the charged pileup-energy content. A series of related techniques operate on the constituents themselves [23–25]. One such technique called PUPPI also uses local charged track information but works at the particle level rather than subjet level. PUPPI computes a scale factor for each particle, using a local estimate inspired by the jets-without-jets paradigm [26]. In this paper, we will be comparing our method to PUPPI and SoftKiller.

In this paper, we present a new approach to pileup removal based on machine learning. The basic idea is to view the energy distribution of particles as the intensity of pixels

²Charged-hadron subtraction follows a particle-flow technique that removes calorimeter energy from pileup tracks. Due to the calorimeter energy resolution, there will be a residual contribution from charged-hadron pileup. This contribution is ignored but could in principle be added to the neutral pileup contribution.

³This work will not explicitly discuss identification of real high energy jets resulting from pileup collisions. The ATLAS and CMS pileup jet identification techniques are documented in refs. [6, 7] and [8], respectively.

in an image [27]. Convolutional neural networks applied to jet images [28] have found widespread applications in both classification [28–32] and generation [33, 34]. Previous jet-images applications have included boosted W -boson tagging [28–30], boosted top quark identification [31], and quark/gluon jet discrimination [32]. Most of these previous applications were all *classification* tasks: extracting a single binary classifier (quark or gluon, W jet or background jet, etc.) from a highly-correlated multidimensional input. The application to pileup removal is a more complicated *regression* task, as the output (a cleaned-up image) should be of similar dimensionality to the input. PUMML is among the first applications of modern machine learning tools to regression problems in high energy physics.

To apply the convolutional neural network paradigm to cleaning an image itself, we exploit the finer angular resolution of the tracking detectors relative to the calorimeters of ATLAS and CMS. Building on the use of multichannel inputs in [32], we give as input to our network three-channel jet images: one channel for the charged LV particles, one channel for the charged pileup particles, and one channel, at slightly lower resolution, for the total neutral particles. We then ask the network to reconstruct the unknown image for LV neutral particles. Thus our inputs are like those of Jet Cleansing but binned into a regular grid (as images) rather than single numbers for each subjet [22]. Further, the architecture is designed to be local (as with Cleansing or PUPPI), with the correction of a pixel only using information in a region around it. The details of our network architecture are described in section 2. Section 3 documents its performance in comparison to other state-of-the-art techniques. The remainder of the paper contains some robustness checks and a discussion in section 6 of the challenges and opportunities for this approach.

2 PUMML algorithm

The goal of the PUMML algorithm is to reconstruct the neutral leading vertex radiation from the charged leading vertex, charged pileup, and total neutral information. Since neutral particles do not have tracking information available, the challenge is to determine what fraction of the total neutral energy in each direction came from the leading vertex and what fraction came from pileup. To assist this discrimination, we take as inputs into our network the energy distribution of charged particles, separated into leading vertex and pileup contributions, in addition to the total neutral energy distribution.⁴ A natural way to combine these observables is using the multichannel images approach introduced in [32] based on color-image recognition technology.

We apply this machine learning technique to $R = 0.4$ anti- k_t jets. The jet image inputs are square grids in pseudorapidity-azimuth (η, ϕ) space of size 0.9×0.9 centered on the charged leading vertex transverse momentum (p_T)-weighted centroid of the jet. One could combine all layers to determine the jet axis, but in practice the axis determined from the charged leading vertex captures dominates because of its superior angular resolution and pileup robustness. To simulate the detector resolutions of charged and neutral calorimeters, charged images are discretized into $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ pixels and neutral images are

⁴Both ATLAS [35] and CMS [36, 37] are proposing precision timing detectors are part of their upgrades for the HL-LHC; such information could naturally be incorporated into another layer of the network.

discretized into $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ pixels.⁵ We use the following three input channels:

red = the transverse momenta of all neutral particles
green = the transverse momenta of charged pileup particles
blue = transverse momenta of charged leading vertex particles.

The output of our network is also an image:

output = the transverse momenta of neutral leading vertex particles.

Only charged particles with $p_T > 500$ MeV were included in the green or blue channels. Charged particles not passing this charged reconstruction cut were treated as if they were neutral particles. Otherwise, the separation into channels is assumed perfect. No image normalization or standardization was applied to the jet images, allowing the network to make use of the overall transverse momentum scale in each pixel. The different resolutions for charged and neutral particles initially present a challenge, since standard architectures assume identical resolution for each color channel. To avoid this issue, we perform a direct upsampling of each neutral pixel to 4×4 pixels of size $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ and divide each pixel value by 16 such that the total momentum in the image is unchanged.

In summary, the following processing was applied to produce the pileup images:

1. *Center*: center the jet image by translating in (η, ϕ) so that the total charged leading vertex p_T -weighted centroid pixel is at $(\eta, \phi) = (0, 0)$. This operation corresponds to rotating and boosting along the beam direction to center the jet.
2. *Pixelate*: crop to a 0.9×0.9 region centered at $(\eta, \phi) = (0, 0)$. Create jet images from the transverse momenta of all neutral particles, the charged leading vertex particles, the charged pileup particles, and the neutral leading vertex particles. Pixelizations of $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ and $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ are used for the charged and neutral jet images, respectively.
3. *Upsample*: upsample each neutral pixel to sixteen $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$ pixels, keeping the total transverse momentum in the image unchanged.

The convolutional neural net architecture used in this study took as input 36×36 pixel, three-channel pileup images. Two convolutional layers, each with 10 filters of size 6×6 with 2×2 strides, were used after zero-padding the input images and first convolutional layer with a 2-pixel buffer on all sides. The output of the second layer has size $9 \times 9 \times 10$, with the 9×9 part corresponding to the size of the target output and the 10 corresponding to the number of filters in the second layer. In order to project down to a $9 \times 9 \times 1$ output, a third convolution layer with filter size 1×1 is used. This last 1×1 convolutional layer is a standard scheme for dimensionality reduction. A rectified linear unit (ReLU) activation

⁵These dimensions are representative of typical tracking and calorimeter resolutions, but would be adapted to the particular detector in practice. We ignore other detector effects in this algorithm demonstration, as has also been done also for PUPPI and SoftKiller. In principle, additional complications due to the detector response can be naturally incorporated into the algorithm during training.

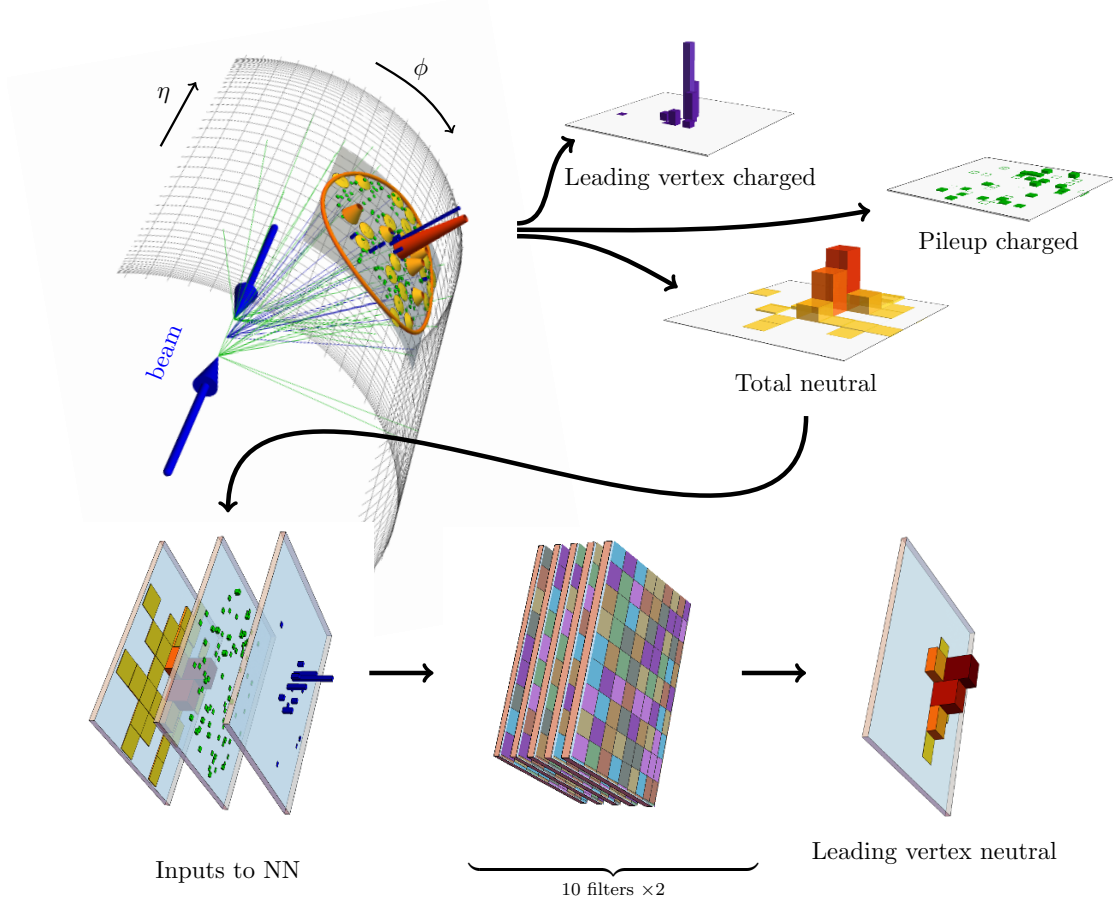


Figure 1. An illustration of the PUMML framework. The input is a three-channel image: blue/purple represents charged radiation from the leading vertex, green is charged pileup radiation, and yellow/orange/red is the total neutral radiation. The resolution of the charged images is higher than for the neutral one. These images are fed into a convolutional layer with several filters whose value at each pixel is a function of a patch around that pixel location in the input images. The output is an image combining the pixels of each filter to one output pixel.

function was applied at each stage. A schematic of the framework and architecture is shown in figure 1.

All neural network implementation and training was performed with the python deep learning libraries Keras [38] and Theano [39]. The dataset consisted of 56 k pileup images, with a 90%/10% train/test split. He-uniform initialization [40] was used to initialize the model weights. The neural network was trained using the Adam [41] algorithm with a batch size of 50 over 25 epochs with a learning rate of 0.001. The choice of loss function implicitly determines a preference for accuracy on harder pixels or softer pixels. To that end, the loss function used to train PUMML was a modified per-pixel logarithmic squared loss:

$$\ell = \left\langle \log \left(\frac{p_T^{(\text{pred})} + \bar{p}}{p_T^{(\text{true})} + \bar{p}} \right)^2 \right\rangle, \quad (2.1)$$

where \bar{p} is a hyperparameter that controls the choice between favoring all p_T equally ($\bar{p} \rightarrow$

∞) or favoring soft pixels ($\bar{p} \rightarrow 0$). After mild optimization, a value of $\bar{p} = 10 \text{ GeV}$ was chosen, though the performance of the model as measured by correlations between reconstructed and true observables is relatively robust to this choice. PUMML was found to give good performance even with a standard loss function such as the mean squared error, which favors all p_T equally.

The PUMML architecture is *local* in that the rescaling of a neutral pixel is a function solely of the information in a patch in (η, ϕ) -space around that pixel. The size of this patch can be controlled by tuning the filter sizes and number of layers in the architecture. Further, due to weight-sharing in convolutional layers, the same function is applied for all pixels. Building this locality and translation invariance into the architecture ensures that the algorithm learns a universal pileup mitigation technique, while carrying the benefit of drastically reducing the number of model parameters. Indeed, the PUMML architecture used in this study has only 4,711 parameters, which is small on the scale of deep learning architectures, but serves to highlight the effectiveness of using modern machine learning techniques (such as convolutional layers) in high energy physics without necessarily using large or deep networks.

While we considered jets and jet images in this study, the PUMML architecture using convolutional nets readily generalizes to event-level applications. The locality of the algorithm implies that the trained model can be applied to any desired region of the event using only the surrounding pixels. To train the model on the event level, either the existing PUMML architecture could be generalized to larger inputs and outputs or the event could be sliced into smaller images and the model trained as in the present study. The parameters of the PUMML architecture are the convolutional filter sizes, the number of filters per layer, and the number of convolutional layers, which may be optimized for a specific application. Here, we have presented an architecture optimized for simplicity and performance for jet-level pileup subtraction. PUMML is designed to be applicable at both jet- and event-level.

3 Performance

To test the PUMML algorithm, we consider $q\bar{q}$ light-quark-initiated jets coming from the decay of a scalar with mass $m_\phi = 500 \text{ GeV}$. Events were generated using Pythia 8.183 [42] with the default tune for pp collisions at $\sqrt{s} = 13 \text{ TeV}$. Pileup was generated by overlaying soft QCD processes onto each event. Final state particles except muons and neutrinos were kept. The events were clustered with FastJet 3.1.3 [43] using the anti- k_t algorithm [44] with a jet radius of $R = 0.4$. A parton-level p_T cut of 95 GeV was applied and up to two leading jets with $p_T > 100 \text{ GeV}$ and $\eta \in [-2.5, 2.5]$ were selected from each event. All particles were taken to be massless.

Samples were generated with the number of pileup vertices ranging from 0 to 180. Since the model must be trained to fix its parameters, the learned model depends on the pileup distribution used for training. For our pileup simulations, we trained on a Poisson distribution of NPUs with mean $\langle \text{NPU} \rangle = 140$. For robustness studies, we also tried training with $\text{NPU} = 140$ for each event or $\text{NPU} = 20$ for each event. The average jet image

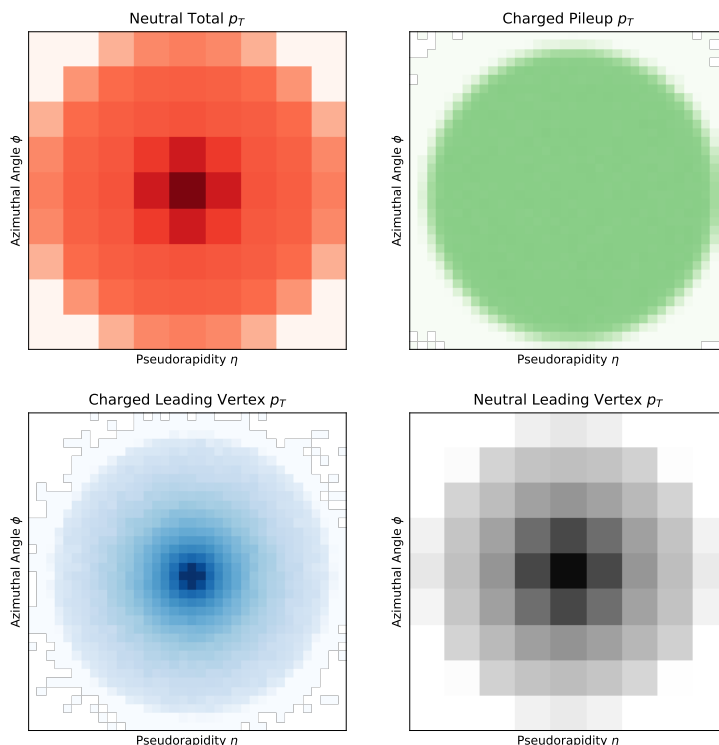


Figure 2. The average leading-jet images for a 500 GeV scalar decaying to light-quark jets with $\langle \text{NPU} \rangle = 140$ pileup, separated by all neutral particles (top left), charged pileup particles (top right), charged leading vertex particles (bottom left), and neutral leading vertex particles (bottom right). Different pixelizations are used for charged and neutral images to reflect the differences in calorimeter resolution. The charged and total neutral images comprise the three-channel input to the neural network, which is trained to predict the neutral leading vertex image.

inputs for this sample are shown in figure 2. For comparison, we show the performance of two powerful and widely used constituent-based pileup mitigation methods: PUPPI [23] and SoftKiller [21]. In both cases, default parameter values were used: $R_0 = 0.3$, $R_{\min} = 0.02$, $w_{\text{cut}} = 0.1$, $p_{\text{T}}^{\text{cut}}(\text{NPU}) = 0.1 + 0.007 \times \text{NPU}$ (PUPPI), grid size = 0.4 (SoftKiller). Variations in the PUPPI parameters did not yield a large difference in performance. Both PUPPI and SoftKiller were implemented at the particle level and then discretized for comparison with PUMML. We show the action of the various pileup mitigation methods on a random selection of events in figure 3. On these examples, PUMML more effectively removes moderately soft energy deposits that are retained by PUPPI and SoftKiller.

To evaluate the performance of different pileup mitigation techniques, we compute several observables and compare the true values to the corrected values of the observables. To facilitate a comparison with PUMML, which outputs corrected neutral calorimeter cells rather than lists of particles, a detector discretization is applied to the true and reconstructed events. Our comparisons focus on the following six jet observables:

- *Jet Mass*: invariant mass of the leading jet.
- *Dijet Mass*: invariant mass of the two leading jets.

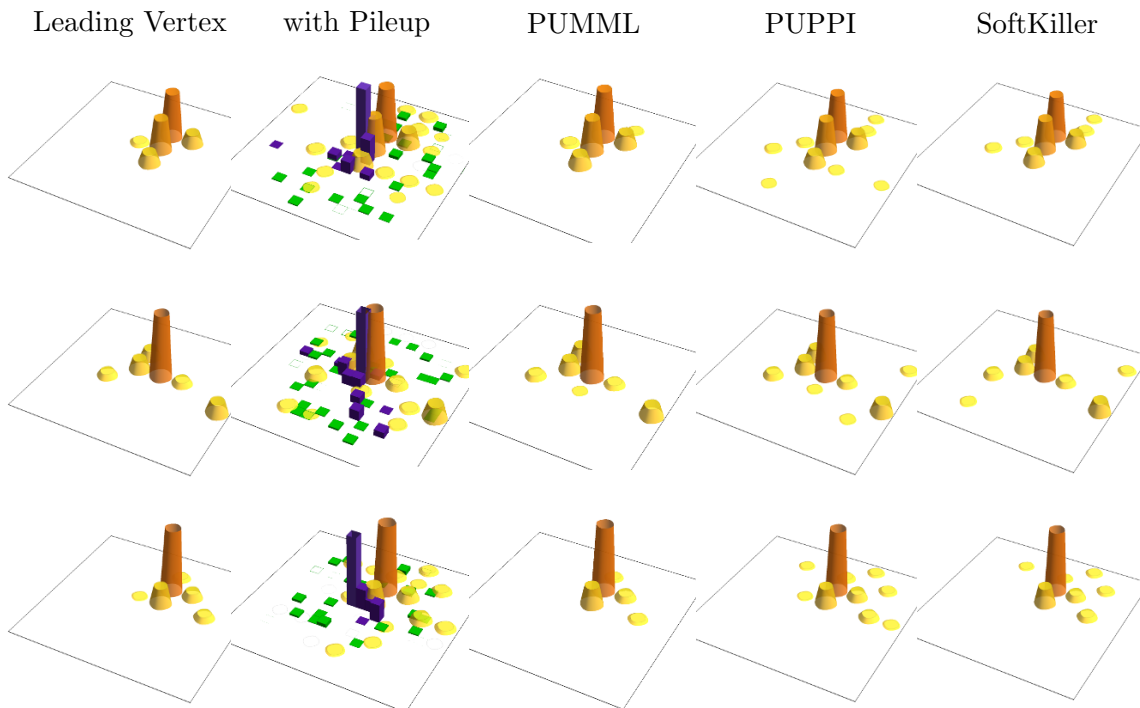


Figure 3. Depictions of three randomly chosen leading jets. Blue/purple represents charged radiation from the leading vertex, green is charged pileup radiation, and yellow/orange/red is the neutral radiation. Shown from left to right are the true neutral leading vertex particles, the event with pileup and charged leading vertex information, followed by the neutral leading vertex particles predicted by PUMML, PUPPI, and SoftKiller. From examining these events, it appears that PUMML has learned an effective pileup mitigation strategy.

- *Jet Transverse Momentum*: the total transverse momentum of the jet.
- *Neutral Image Activity*, N_{95} [45]: the number of neutral calorimeter cells which account for 95% of the total neutral transverse momentum.
- *Energy Correlation Functions*, $ECF_N^{(\beta)}$ [46]: specifically, we consider the logarithm of the two- and three-point ECFs with $\beta = 4$.

Figure 4 illustrates the distributions of several of these jet observables after applying the different pileup subtraction methods. While these plots are standard, they do not give a per-event indication of performance. A more useful comparison is to show the distributions of the per-event percent error in reconstructing the true values of the observables, which are shown in figure 5. To numerically explore the event-by-event effectiveness, we can look at the Pearson linear correlation coefficient between the true and corrected values or the interquartile range (IQR) of the percent errors. Table 1 summarizes the event-by-event correlation coefficients of the distributions shown in figure 4. Table 2 summarizes the IQRs of the distributions shown in figure 5. PUMML outperforms the other pileup mitigation techniques on both of these metrics, with improvements for jet substructure observables such as the jet mass and the energy correlation functions.

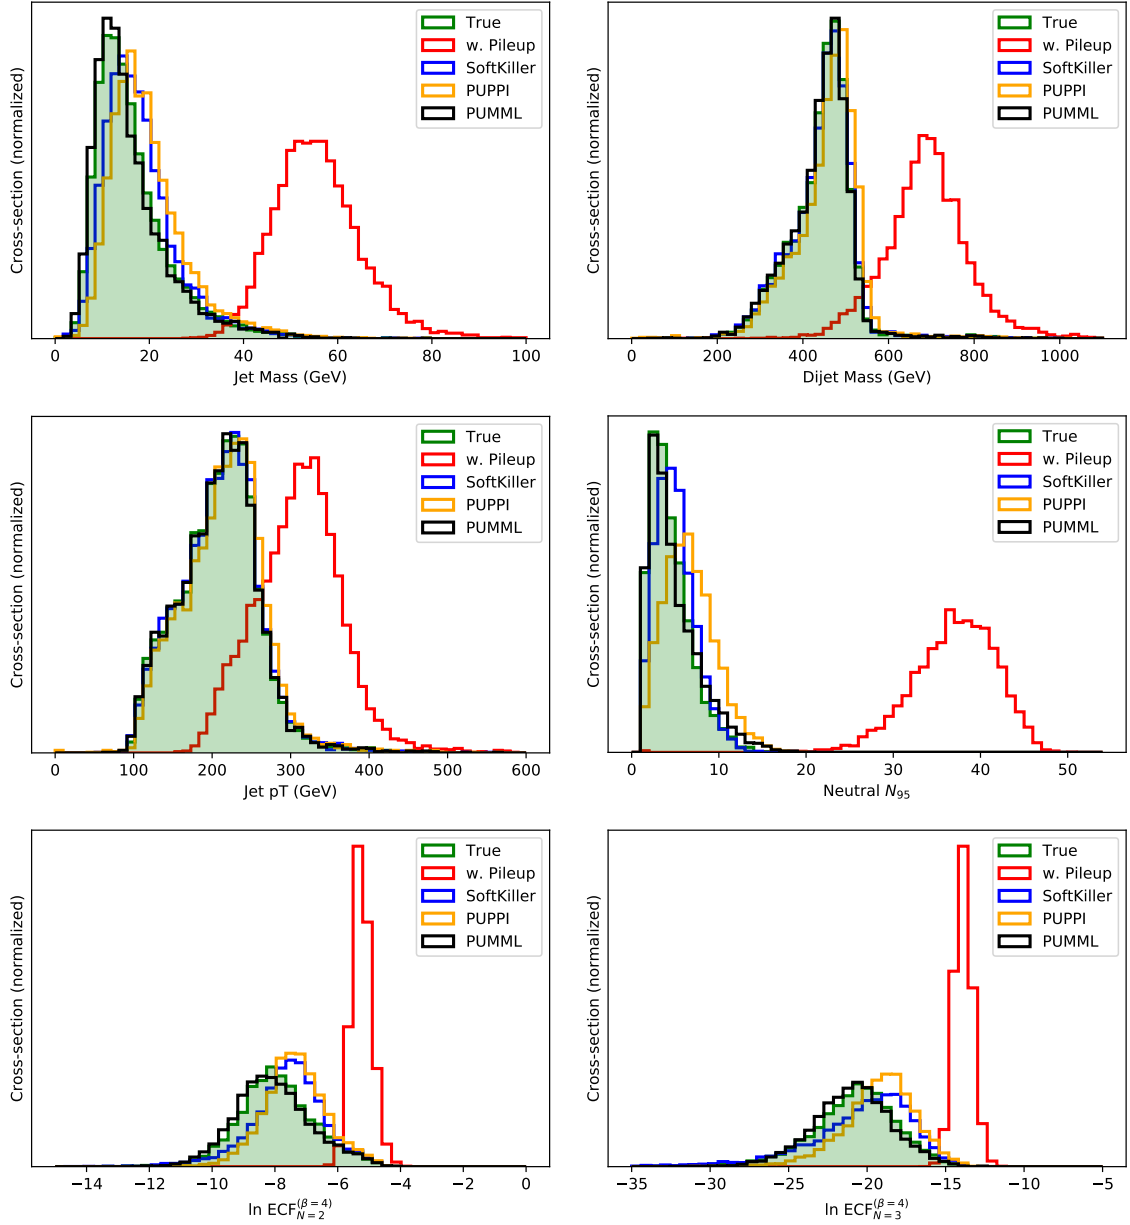


Figure 4. Distributions of leading jet mass (top left), dijet mass (top right), leading jet p_T (middle left), neutral N_{95} (middle right), $\ln \text{ECF}_{N=2}^{(\beta=4)}$ (bottom left), and $\ln \text{ECF}_{N=3}^{(\beta=4)}$ (bottom right) for the considered pileup subtraction methods with Poissonian $\langle \text{NPU} \rangle = 140$ pileup. While all of the pileup mitigation methods do well for observables such as the dijet mass and jet p_T , PUMML more closely matches the true distributions of more sensitive substructure observables like mass, neutral N_{95} , and the energy correlation functions.

4 Robustness

It is important to verify that PUMML learns a pileup mitigation function which is not overly sensitive to the NPU distribution of its training sample. Robustness to the NPU

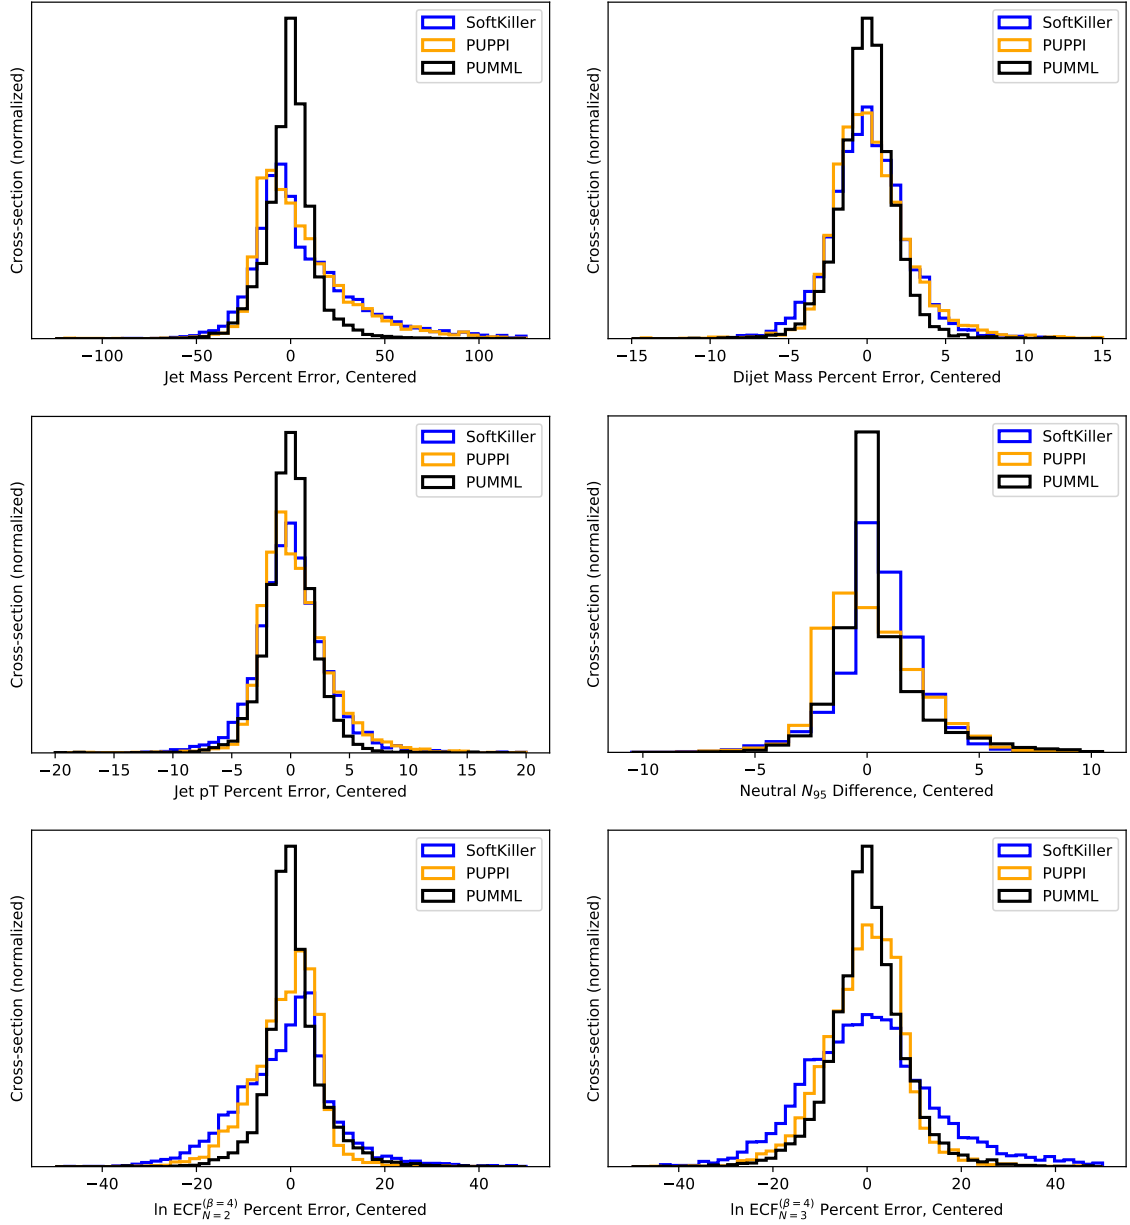


Figure 5. Distributions of the percent error between reconstructed and true values for leading jet mass (top left), dijet mass (top right), leading jet p_T (middle left), neutral N_{95} (middle right), $\ln \text{ECF}_{N=2}^{(\beta=4)}$ (bottom left), and $\ln \text{ECF}_{N=3}^{(\beta=4)}$ (bottom right) for the considered pileup subtraction methods with Poissonian $\langle \text{NPU} \rangle = 140$ pileup. For the discrete neutral N_{95} observable, only the difference is shown. All distributions are centered to have median at 0. The improved reconstruction performance of PUMML is highlighted by its taller and narrower peaks.

on which it is trained would indicate that PUMML is learning a universal subtraction strategy. To evaluate this robustness, PUMML was trained on 50k events with either $\text{NPU} = 20$ or $\text{NPU} = 140$ and then tested on samples with different NPUs. Figure 6 shows the jet mass correlation coefficients as a function of the test sample NPU. PUMML learns a

Correlation (%)	w. Pileup	PUMML	PUPPI	SoftKiller
Jet mass	65.5	97.4	94.0	91.3
Dijet mass	85.5	99.5	95.8	99.1
Jet p_T	94.4	99.7	98.0	99.4
Neutral N_{95}	36.2	75.3	70.4	67.7
$\ln \text{ECF}_{N=2}^{(\beta=4)}$	60.4	90.5	83.3	68.8
$\ln \text{ECF}_{N=3}^{(\beta=4)}$	41.6	77.2	69.1	45.7

Table 1. Correlation coefficients between the true and corrected values of different jet observables on an event-by-event level. The first column lists the correlation without any pileup mitigation applied to the event. Larger correlation coefficients are better.

IQR (%)	PUMML	PUPPI	SoftKiller
Jet mass	13.0	28.7	30.8
Dijet mass	2.02	2.95	2.97
Jet p_T	2.26	3.40	3.39
$\ln \text{ECF}_{N=2}^{(\beta=4)}$	5.63	8.82	11.9
$\ln \text{ECF}_{N=3}^{(\beta=4)}$	8.48	10.7	16.7

Table 2. The interquartile ranges (IQR) of the distributions in figure 5. Note that PUMML performs better than either PUPPI or SoftKiller. Lower IQR indicates better performance.

strategy that is surprisingly performant outside of the NPU range on which it was trained. Further, we see that by this measure of performance, PUMML consistently outperforms both PUPPI and SoftKiller.

A related robustness test is to probe how the performance of PUMML depends on the p_T spectrum of the training sample. To explore this, we generated two large training samples (50k events): one with a scalar mass of 200 GeV and one with a scalar mass of 2 TeV; we did not impose any parton-level p_T cuts on these samples. After training these two networks, we tested them on a set of samples generated from scalars with intermediate masses, from 300 GeV to 900 GeV. As can be seen in figure 7, the performance of PUMML is very robust to the p_T distribution of the jets in the training sample: the networks trained on the 200 GeV resonance and the 2 TeV resonance have identical performance. The figure also shows that the performance of PUMML is less sensitive to of the p_T of the testing sample than either PUPPI or Soft-Killer. This robustness test speaks to the PUMML algorithm’s ability to learn universal aspects of pileup mitigation.

A number of modifications of PUMML were also tried. Locally connected layers were tried instead of convolutional layers and were found to perform worse due to a large increase in the number of parameters of the model, while losing the translation invariance that makes PUMML powerful. We tried training without various combinations of the input channels;

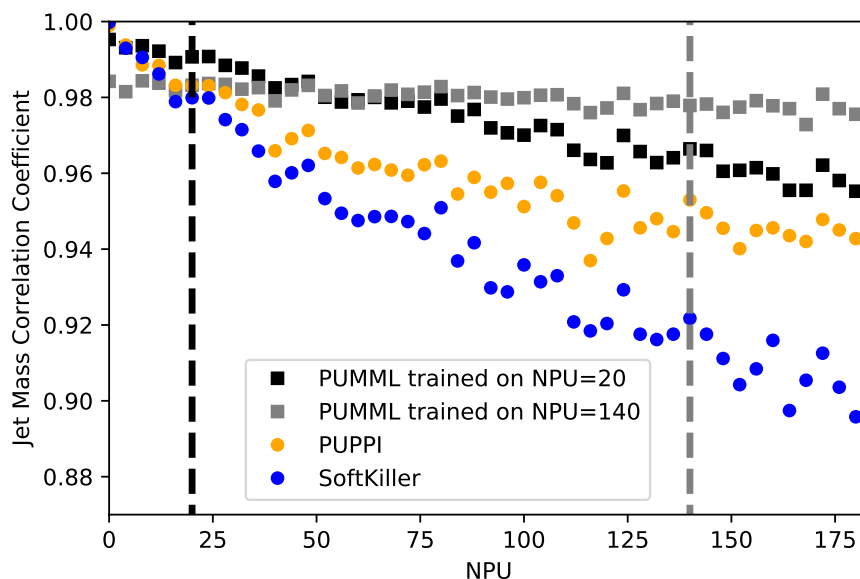


Figure 6. Correlation coefficients between reconstructed and true jet masses plotted as a function of NPU for the different pileup mitigation schemes. PUMML was trained on 50 k events with either $\text{NPU} = 20$ or $\text{NPU} = 140$ indicated by dashed vertical lines. The performance of PUMML with Poissonian $\langle \text{NPU} \rangle = 140$ is similar to the $\text{NPU} = 140$ curve. PUMML is surprisingly performant well outside the NPU range on which it was trained and consistently outperforms PUPPI and SoftKiller. Note that PUMML trained on the lower NPU sample better reconstructs the jet mass in the low pileup regime.

the model was found to perform moderately worse without either of the charged channels but suffered severe degradation without the total neutral channel. We tried using simpler models with only one layer or fewer filters per layer. Remarkably, even with only a single layer and a single 4×4 filter (a model that has just 49 parameters), PUMML performed only moderately worse than the version presented in this study, which was allowed to be more complicated in order to achieve even better performance.

5 What is PUMML learning?

While it is generally very difficult to determine what a network is learning, one possible probe is to examine the weights of the filter layers in the convolutional network. For our full network, these weights are complicated and the subtractor that the network learns is difficult to probe analytically. Instead, we trained a simplified PUMML network with a single 12×12 pixel filter, which spans 3×3 neutral pixels, with no bias term. The different channels of this filter are shown in figure 8. The neutral filter clearly selects the relevant neutral pixel for subtraction, while the charged pileup filter is approximately uniform (with the value dependent on the specific choice of loss function and activation function), and the charged leading vertex filter does not significantly contribute.

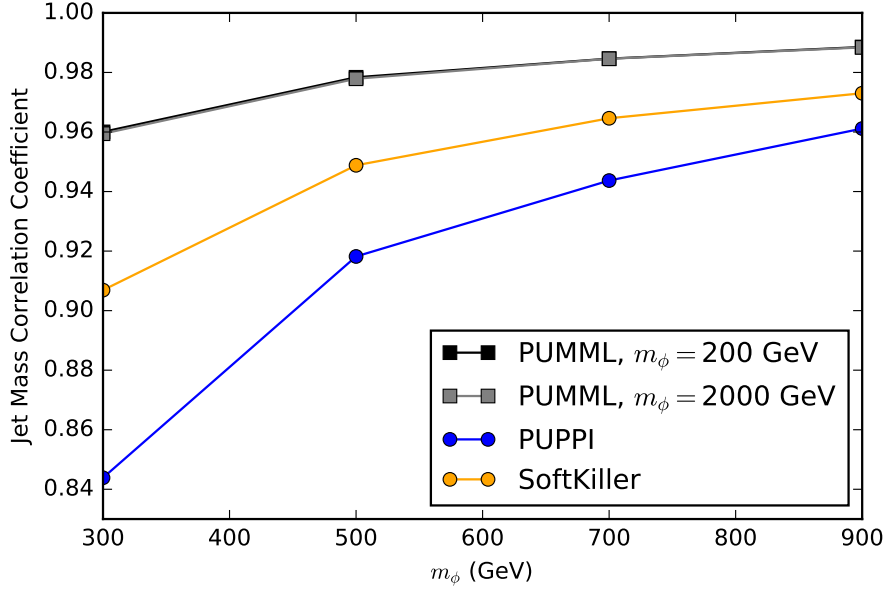


Figure 7. Correlation coefficients between reconstructed and true jet masses plotted as a function of the mass of the scalar resonance with $\text{NPU}=140$. A spread in scalar resonances is generated in order to produce a range in jet transverse momenta. In order to assess the impact of the p_T distribution used for training, one version of PUMML was trained with a scalar mass of 200 GeV (black) and one was trained with a mass of 2 TeV (gray). The two PUMML curves closely match one another.

The filter values motivate the following parameterization of what PUMML is learning:

$$p_T^{\text{N,LV}} = 1.0 p_T^{\text{N,total}} - \beta p_T^{\text{C,PU}} + 0.0 p_T^{\text{C,LV}}, \quad (5.1)$$

for some $\mathcal{O}(1)$ constant β , where $p_T^{\text{N,LV}}$, $p_T^{\text{N,total}}$, $p_T^{\text{C,PU}}$, and $p_T^{\text{C,LV}}$ are the neutral-pixel-level transverse momenta of the neutral leading-vertex particles, all neutral particles, charged pileup particles, and charged leading-vertex particles, respectively. The values 1.0 and 0.0 in eq. (5.1) are stable (to the 0.05 level) under variations in the loss and activation functions. This is reassuring as the learned subtractor is thereby robust in the $\text{NPU} \rightarrow 0$ limit despite being trained on $\langle \text{NPU} \rangle = 140$.

Eq. (5.1) is remarkably similar to the physically-motivated formula used in Jet Cleansing [22]. Cleansing is built on the observation that since pileup is the incoherent sum of many separate scattering events, its variance is smaller than the variance of the radiation from the leading-vertex. Thus, it is better to estimate $p_T^{\text{N,PU}}$ from $p_T^{\text{C,PU}}$ than to estimate $p_T^{\text{N,LV}}$ from $p_T^{\text{C,LV}}$. The simplest form of Cleansing (Linear Cleansing) gives the formula:

$$p_T^{\text{N,LV}} = p_T^{\text{N,tot}} - \left(\frac{1}{\overline{\gamma}_0} - 1 \right) p_T^{\text{C,PU}}, \quad (5.2)$$

where $\overline{\gamma}_0$ is the average ratio of charged p_T to total p_T in a subjet. Thus this simple one 12×12 filter PUMML network is learning a subtractor of precisely the same parametric form as Linear Cleansing!

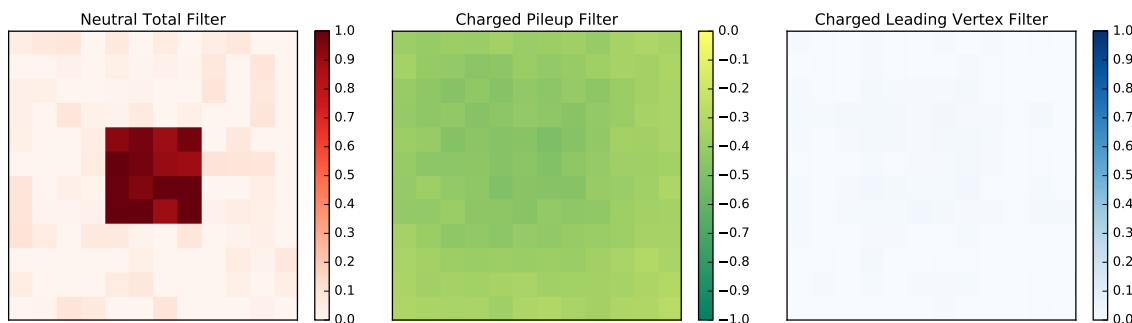


Figure 8. Filter weights for a simple PUMML network with a single 12×12 filter and a ReLU activation function trained with $\langle \text{NPU} \rangle = 140$. The network has selected the relevant neutral pixel, turned off the charged leading vertex contribution, and is using the charged pileup information uniformly.

The value of $\overline{\gamma_0}$ in Linear Cleansing and the value of β that is learned in eq. (5.1) depend on how soft radiation is handled. For example, if no reconstruction threshold is applied, $\overline{\gamma_0} \approx 2/3$ (since $2/3$ of pions are charged). In addition, the value of β depends on the loss function used. For example, if the loss function is minimized when the means of the true and predicted neutral transverse momenta are equal:

$$\ell = \left| \langle p_T^{(\text{true})} \rangle - \langle p_T^{(\text{pred})} \rangle \right| = \left| \langle p_T^{\text{N,LV}} \rangle - \langle p_T^{\text{N,total}} \rangle + \beta \langle p_T^{\text{C,PU}} \rangle \right|, \quad (5.3)$$

then we find that the optimal β is:

$$\beta = \frac{\langle p_T^{\text{N,PU}} \rangle}{\langle p_T^{\text{C,PU}} \rangle}. \quad (5.4)$$

Training the 12×12 PUMML filter without a ReLU or bias term, using the loss function of eq. (5.3) with the average taken pixel-wise over the batch, we find $\beta = 0.59$ with no charged reconstruction cut and $\beta = 1.18$ with the cut. These values are consistent with those predicted by eq. (5.4) of 0.62 and 1.26, respectively.

On the other hand, if we take a mean squared error loss function:

$$\ell = \left\langle \left(p_T^{(\text{true})} - p_T^{(\text{pred})} \right)^2 \right\rangle, \quad (5.5)$$

then the minimum occurs at:

$$\beta = \frac{\langle p_T^{\text{N,PU}} p_T^{\text{C,PU}} \rangle}{\langle p_T^{\text{C,PU}} p_T^{\text{C,PU}} \rangle}. \quad (5.6)$$

This still depends only on the pileup properties, as with Linear Cleansing, but also depends on correlations between neutral and charged radiation. For example, training the 12×12 PUMML filter without a ReLU or bias term using a mean squared error loss function, we find $\beta = 0.56$ with no charged reconstruction cut and $\beta = 0.97$ with the cut. These numbers are in general agreement (within 10–20%) with a direct evaluation of the right-hand side of eq. (5.6). In the limit that neutral and charged pileup radiation are constant, eq. (5.6) reduces to eq. (5.4).

Whether the loss function of eq. (5.4) or eq. (5.6) (or something else entirely) is better is not simple to establish. The inclusion of the ReLU activation function further complicates the analysis since the model is equally penalized for all non-positive predictions. We find with the single 12×12 filter, using the loss function of eq. (2.1) and including a ReLU and bias term, PUMML achieves a jet mass correlation coefficient of 90.4%. This is competitive with the values listed in table 1, as we might expect since Linear Cleansing has comparable performance to PUPPI and SoftKiller. The full network improves on Linear Cleansing by exploiting additional correlations that are hard to disentangle by looking at the filters.

6 Conclusions

In this paper, we have introduced the first application of machine learning to the critically important problem of pileup mitigation at hadron colliders. We have phrased the problem of pileup mitigation in the language of a machine learning regression problem. The method we introduced, PUMML, takes as input the transverse momentum distribution of charged leading-vertex, charged pileup, and all neutral particles, and outputs the corrected leading vertex neutral energy distribution. We demonstrated that PUMML works at least as well as, and often better than, the competing algorithms PUPPI and SoftKiller in their default implementations. It will be exciting to see these algorithms compared with a full detector simulation, where it will be possible to test the sensitivity to important experimental effects such as resolutions and inefficiencies.

There are several extensions and additional applications of the PUMML framework beyond the scope of this study. As mentioned in section 2, PUMML can very naturally be extended from jet images to entire events. Applying this event-level PUMML to the problem of missing transverse energy would be a natural next step. While the filter sizes can be the same for the event and jet images, the network training will likely require modification. Furthermore, the inhomogeneity of the detector response with $|\eta|$ will require attention. Another potentially useful modification to PUMML would be to train to predict the neutral pileup p_T rather than the neutral leading vertex p_T in order to increase out-of-sample robustness of the learned pileup mitigation algorithm. Additionally, using larger- R jets may be of interest, thereby necessitating a resizing of the local patch or other PUMML parameters, all of which is easily achieved.

An important consideration when using machine learning for particle physics applications is how the method can be used with data and whether or not the systematic uncertainties are under control. Unlike a purely physically-motivated algorithm, such as PUPPI or SoftKiller, machine learning runs the risk of being a “black-box” which can be difficult to understand. Nevertheless, machine learning is powerful, scaleable, and capable of complementing physical insight to solve complicated or otherwise intractable problems.

To prevent the model from learning simulation artifacts, it is preferable to train on actual data rather than simulation. In many machine learning applications in collider physics, obtaining truth-level training samples in data is a substantial challenge. To overcome this challenge in classification tasks, [47] introduces an approach to train from impure samples using class proportion information. For PUMML and pileup mitigation more broadly, a

more direct method to train on data is possible. To simulate pileup, we overlay soft QCD events on top of a hard scattering process, both generated with Pythia. Experimentally, there are large samples of minimum bias and zero-bias (i.e. randomly triggered) data. There are also samples of relatively pileup-free events from low luminosity runs. Thus we can construct high-pileup samples using purely data. This kind of *data overlay* approach, which has already been used by experimental groups in other contexts [48, 49], could be perfect for training PUMML with data. Therefore, an implementation of ML-based pileup mitigation in an actual experimental setting could avoid mis-modeling artifacts during training, thus adding more robustness and power to this new tool.

Acknowledgments

The authors would like to thank Philip Harris, Francesco Rubbo, Ariel Schwartzman and Nhan Tran for stimulating conversations, in particular for suggesting some of the extensions mentioned in the conclusions. We would also like to thank Jesse Thaler for helpful discussions. PTK and EMM would like to thank the MIT Physics Department for its support. Computations in this paper were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. This work was supported by the Office of Science of the U.S. Department of Energy (DOE) under contracts DE-AC02-05CH11231 and DE-SC0013607, the DOE Office of Nuclear Physics under contract DE-SC0011090, and the DOE Office of High Energy Physics under contract DE-SC0012567. Cloud computing resources were provided through a Microsoft Azure for Research award. Additional support was provided by the Harvard Data Science Initiative.

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

References

- [1] CMS collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, 2017 *JINST* **12** P02014 [[arXiv:1607.03663](https://arxiv.org/abs/1607.03663)] [[INSPIRE](#)].
- [2] ATLAS collaboration, *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **96** (2017) 072002 [[arXiv:1703.09665](https://arxiv.org/abs/1703.09665)] [[INSPIRE](#)].
- [3] CMS collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, 2014 *JINST* **9** P10009 [[arXiv:1405.6569](https://arxiv.org/abs/1405.6569)] [[INSPIRE](#)].
- [4] ATLAS collaboration, *Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC*, *ATLAS-CONF-2010-027* (2010) [[INSPIRE](#)].
- [5] ATLAS collaboration, *Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s} = 7$ TeV in the ATLAS experiment*, *ATLAS-CONF-2010-069* (2010) [[INSPIRE](#)].

- [6] ATLAS collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 581 [[arXiv:1510.03823](#)] [[INSPIRE](#)].
- [7] ATLAS collaboration, *Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector*, *Eur. Phys. J. C* **77** (2017) 580 [[arXiv:1705.02211](#)] [[INSPIRE](#)].
- [8] CMS collaboration, *Pileup Jet Identification*, CMS-PAS-JME-13-005 (2013) [[INSPIRE](#)].
- [9] M. Cacciari and G.P. Salam, *Pileup subtraction using jet areas*, *Phys. Lett. B* **659** (2008) 119 [[arXiv:0707.1378](#)] [[INSPIRE](#)].
- [10] CMS collaboration, *Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS*, 2011 JINST **6** P11002 [[arXiv:1107.4277](#)] [[INSPIRE](#)].
- [11] ATLAS collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **73** (2013) 2304 [[arXiv:1112.6426](#)] [[INSPIRE](#)].
- [12] ATLAS collaboration, *Monte Carlo Calibration and Combination of In-situ Measurements of Jet Energy Scale, Jet Energy Resolution and Jet Mass in ATLAS*, ATLAS-CONF-2015-037 (2015) [[INSPIRE](#)].
- [13] CMS collaboration, *Jet energy scale and resolution performances with 13 TeV data*, CMS-DP-2016-020 (2016).
- [14] ATLAS collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490 [[arXiv:1603.02934](#)] [[INSPIRE](#)].
- [15] CMS collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, 2017 JINST **12** P10003 [[arXiv:1706.04965](#)] [[INSPIRE](#)].
- [16] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](#)] [[INSPIRE](#)].
- [17] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [[INSPIRE](#)].
- [18] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, *Phys. Rev. D* **80** (2009) 051501 [[arXiv:0903.5081](#)] [[INSPIRE](#)].
- [19] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches*, *Phys. Rev. D* **81** (2010) 094023 [[arXiv:0912.0033](#)] [[INSPIRE](#)].
- [20] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [21] M. Cacciari, G.P. Salam and G. Soyez, *SoftKiller, a particle-level pileup removal method*, *Eur. Phys. J. C* **75** (2015) 59 [[arXiv:1407.0408](#)] [[INSPIRE](#)].
- [22] D. Krohn, M.D. Schwartz, M. Low and L.-T. Wang, *Jet Cleansing: Pileup Removal at High Luminosity*, *Phys. Rev. D* **90** (2014) 065020 [[arXiv:1309.4777](#)] [[INSPIRE](#)].
- [23] D. Bertolini, P. Harris, M. Low and N. Tran, *Pileup Per Particle Identification*, *JHEP* **10** (2014) 059 [[arXiv:1407.6013](#)] [[INSPIRE](#)].
- [24] P. Berta, M. Spousta, D.W. Miller and R. Leitner, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **06** (2014) 092 [[arXiv:1403.3108](#)] [[INSPIRE](#)].

- [25] ATLAS collaboration, *Constituent-level pile-up mitigation techniques in ATLAS*, [ATLAS-CONF-2017-065](#) (2017) [[INSPIRE](#)].
- [26] D. Bertolini, T. Chan and J. Thaler, *Jet Observables Without Jet Algorithms*, *JHEP* **04** (2014) 013 [[arXiv:1310.7584](#)] [[INSPIRE](#)].
- [27] 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](#)].
- [28] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](#)] [[INSPIRE](#)].
- [29] 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](#)].
- [30] J. Barnard, E.N. Dawe, M.J. Dolan and N. Rajcic, *Parton Shower Uncertainties in Jet Substructure Analyses with Deep Neural Networks*, *Phys. Rev. D* **95** (2017) 014018 [[arXiv:1609.00607](#)] [[INSPIRE](#)].
- [31] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning Top Taggers or The End of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [32] P.T. Komiske, E.M. Metodiev and M.D. Schwartz, *Deep learning in color: towards automated quark/gluon jet discrimination*, *JHEP* **01** (2017) 110 [[arXiv:1612.01551](#)] [[INSPIRE](#)].
- [33] L. de Oliveira, M. Paganini and B. Nachman, *Learning Particle Physics by Example: Location-Aware Generative Adversarial Networks for Physics Synthesis*, *Comput. Softw. Big Sci.* **1** (2017) 4 [[arXiv:1701.05927](#)] [[INSPIRE](#)].
- [34] M. Paganini, L. de Oliveira and B. Nachman, *CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calorimeters with Generative Adversarial Networks*, [arXiv:1705.02355](#) [[INSPIRE](#)].
- [35] ATLAS collaboration, *ATLAS Phase-II Upgrade Scoping Document*, [CERN-LHCC-2015-020](#) (2015).
- [36] CMS collaboration, *CMS Phase II Upgrade Scope Document*, [CERN-LHCC-2015-019](#) (2015).
- [37] CMS collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, [CERN-LHCC-2015-010](#) (2015) [[INSPIRE](#)].
- [38] F. Chollet, *Keras*, (2015) <https://github.com/fchollet/keras>.
- [39] J. Bergstra et al., *Theano: A cpu and gpu math compiler in python*, in proceedings of the *9th Python in Science Conference*, Austin, Texas, U.S.A., 28 June–3 July 2010, pp. 1–7.
- [40] K. He, X. Zhang, S. Ren and J. Sun, *Delving deep into rectifiers: Surpassing human-level performance on imagenet classification*, in proceedings of the *2015 IEEE International Conference on Computer Vision (ICCV)*, Santiago, Chile, 7–13 December 2015, pp. 1026–1034 [<https://doi.org/10.1109/ICCV.2015.123>] [[arXiv:1502.01852](#)].
- [41] D. Kingma and J. Ba, *Adam: A method for stochastic optimization*, [arXiv:1412.6980](#).
- [42] T. Sjöstrand et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [43] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].

- [44] M. Cacciari, G.P. Salam and G. Soyez, *The Anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [45] J. Pumplin, *How to tell quark jets from gluon jets*, *Phys. Rev. D* **44** (1991) 2025 [[INSPIRE](#)].
- [46] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [47] 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](#)].
- [48] ATLAS collaboration, *Simulation of Pile-up in the ATLAS Experiment*, *J. Phys. Conf. Ser.* **513** (2014) 022024 [[INSPIRE](#)].
- [49] ATLAS collaboration, *ATLAS Simulation using Real Data: Embedding and Overlay*, *J. Phys. Conf. Ser.* **898** (2017) 042004 [[INSPIRE](#)].