Pileup mitigation at the Large Hadron Collider with Graph Neural Networks

J. Arjona Martínez¹, O. Cerri², M. Spiropulu², JR. Vlimant², and M. Pierini³

¹ University of Cambridge, Trinity Ln, Cambridge CB2 1TN, UK

² California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125

³ CERN, CH-1211 Geneva, Switzerland
e-mail: maurizio.pierini@cern.ch

Abstract. At the Large Hadron Collider, the high transverse-momentum events studied by experimental collaborations occur in coincidence with parasitic low transverse-momentum collisions, usually referred to as pileup. Pileup mitigation is a key ingredient of the online and offline event reconstruction as pileup affects the reconstruction accuracy of many physics observables. We present a classifier based on Graph Neural Networks, trained to retain particles coming from high-transverse-momentum collisions, while rejecting those coming from pileup collisions. This model is designed as a refinement of the PUPPI algorithm [1], employed in many LHC data analyses since 2015. Thanks to an extended basis of input information and the learning capabilities of the considered network architecture, we show an improvement in pileup-rejection performances with respect to state-of-the-art solutions.

1 Introduction

In order to deliver large datasets to experimental collaborations, the CERN Large Hadron Collider (LHC) operates with proton bunches, each containing as many as $\mathcal{O}(10^{11})$ protons. These protons are densely packed to increase the beam luminosity and, with it, the collision rate. This luminosity increase comes at the cost of an increasing number of parasitic collisions (pileup), typically consisting of soft-QCD events at small transverse momentum (p_T). At the end of LHC Run II, an average of ~ 40 (and a maximum of ~ 80) collisions happened simultaneously to each interesting high- p_T event.

Due to the shape of the luminous region, the (x, y) coordinates¹ of the pileup collisions in an event are aligned to those of the high- p_T interesting collision, referred to as the leading vertex (LV). On the other hand, pileup vertices can be displaced by several cm from the LV in the z direction. The impact of pileup collisions is reduced by applying targeted algorithms designed to estimate and mitigate this effect.

Being detected in the inner layers of a typical multi-purpose detector such as ATLAS [2] or CMS [3], charged particles can be tracked back to their point of origin, while this is not typically the case for neutral particles. Thanks to the CMS [4] and ATLAS [5] vertex resolution, charged particles from pileup can then be associated to vertices other than the LV and consequently removed. This technique, referred to as charged hadron subtraction (CHS), greatly simplifies the problem, as can be seen in Figure 1. The main challenge becomes correcting for the neutral pileup contribution, for which sufficient vertex information is typically unavailable. Early approaches, such as the area-subtraction method [6, 7, 8, 9] employed in LHC Run I (2009-2012) analyses, correct the event based only on the characteristic per-event pileup energy density. While these approaches help in obtaining unbiased estimates of the jets four-momenta, they are affected by a serious resolution loss for a large number of pileup interactions, even when extended to jet shapes [10]. This motivated the introduction of new algorithms for the LHC Run II (2015-2018).

The currently adopted pileup mitigation techniques consist of rule-based algorithms and usually operate on a per-particle basis, tailored to suppress particles believed to originate from pileup interactions, or to weight them proportionally to their probability of originating from the hard interactions. Examples of the former category include

¹ We use a Cartesian coordinate system with the z axis oriented along the beam axis, the x on the horizontal plane, and the y axis oriented upward. The x and y axes define the transverse plane, while the z axis identifies the longitudinal direction. The azimuth angle ϕ is computed with respect to the x axis. The polar angle θ is used to compute the pseudorapidity $\eta = -\log(\tan(\theta/2))$. We fix units such that $c = \hbar = 1$.

Fig. 1. Depiction of the effect of CHS. The full event (left), the event after CHS is applied (middle) and the Ground Truth (right) are shown. A $Z \rightarrow \nu\bar{\nu} + \text{jets}$ event is superimposed to 80 pileup events. Particles from the LV are shown in orange (dark) and those from pileup in blue (light).

the *SoftKiller* algorithm [11]. The PileUp Per Particle Identification (PUPPI) algorithm [1] employed by the CMS collaboration in the LHC Run II is an example of the latter category. These two algorithms fairly represent the state-of-the-art performance for what concerns pileup mitigation algorithms, having being successfully tested on real LHC collision data [12, 13]. For this reason, we take them as a baseline in this paper.

Our study takes as starting point the PUPPI algorithm. Unlike many other pileup-removal algorithms, PUPPI is designed to assign a weight to each particle, the so-called PUPPI weights. These weights quantify how likely it is that each particle originated from the LV. As described in Ref. [1], the weight computation is based on the per-event distribution of the quantity:

$$\alpha_i^\beta = \log \sum_{j \in \text{event}} \xi_{ij} \times \Theta(\Delta R_{ij} < R_{\min}) \times \Theta(\Delta R_{ij} < R_0). \quad (1)$$

In Eq.(1), i is the label of the considered particle in the event and $\xi_{ij} = p_{Tj}/\Delta R_{ij}^\beta$; $\Delta R_{ij} = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ is the distance between the i -th and j -th particle in the plane identified by the pseudorapidity η and the azimuth angle ϕ ; $R_0 = 0.3$ defines a cone around the i -th particle, and $R_{\min} = 0.02$ removes the region surrounding the i -th particle. In Ref. [1], α^1 is found to be the optimal metric to quantify the PUPPI weight. When CHS is applied upstream to PUPPI, the sum in Eq.(1) is performed over the charged particles from the LV, as opposed to the full event. By definition, the PUPPI algorithm requires to be tuned on the specific dataset it is applied to. For a fair comparison, the PUPPI results presented in this work are obtained applying the tuning procedure outlined in Ref. [1] to the datasets described in Section 3.

Our work aims to extend the traditional PUPPI setup using Deep Learning. In particular, we present a classifier, based on Graph Networks, trained to identify particles originating from the high- p_T events and discard the others. We consider an architecture based on Gated Graph Neural Networks (GGNN) [14] as our final model, called PUPPIML. In addition, a model based on fully-connected Neural Networks [15] and one based on Gated Recurrent Units (GRU) [16] are presented for comparison.

This paper is structured as follows: Section 2 describes related literature. Section 3 and 4 describe the utilized data and models, respectively. Results are presented in Section 5, while the robustness of the proposed approach is discussed in Section 5.2. Conclusions are given in Section 6.

2 Related work

Machine Learning (ML) traditionally played a prominent role in High Energy Physics (HEP), as discussed for instance in Ref. [17]. Among the many existing proof-of-principle studies, few applications have already been deployed in the central data processing of major HEP experiments. For instance, recurrent architectures proved fruitful in bottom-jet identification [18, 19] and Convolutional Neural Networks (CNNs) in neutrino physics [20]. Graph Networks have very recently been used for jet tagging, matching the performances of other deep learning approaches [21], and to identify interesting typologies at the LHC [22], and in IceCube [23].

A first application of machine learning to the problem of pileup subtraction is presented in Ref. [24], describing the PUMML algorithm. This work clearly demonstrates that the use of Deep Learning techniques to pileup removal results in a performance improvement. The PUMML algorithm is based on CNNs. To apply it, one needs to represent the collision event as an image. This is usually done binning the portion of the (η, ϕ) plane covered by detector acceptance

and creating a map of the deposited p_T in each bin. Such a practice poses some problem when confronting the main features of a typical collider detector. For instance::

1. PUMML is applied to images of cropped regions around the reconstructed jets, derived by binning the (η, ϕ) plane. The binning choice is intended to be “representative of typical tracking and calorimeter resolutions” [24]. This approach is very suitable for the central barrel of a typical detector, but it overlooks the complexity of a typical detector geometry. In particular, it neglects the existence of complicated overlap regions between the barrel and the endcap regions, and the irregular geometry of the endcaps (e.g., the variable segmentation of the CMS ECAL and HCAL endcaps). On the other hand, pileup subtraction is particularly important in the endcap regions, where the solenoid magnetic field of a typical cylindrical detector pushes the abundant low- p_T particles produced in pileup interactions.
2. The use of a calorimeters-inspired (η, ϕ) fixed-size grid neglects the fact that charged particles are mainly detected through the inner tracker. A typical-tracker (η, ϕ) resolution can hardly be represented as a fixed number, since it depends on the track p_T (as shown for instance in Ref. [25] for the CMS detector). This fact might not be crucial for a calorimeter-centric jet reconstruction. On the other hand, such an assumption is extremely sub-optimal for local or global reconstruction algorithms based on particle flow [26, 27]. In these cases, the enhanced tracker resolution plays a crucial role in driving the measurement of particle momenta.
3. The assumed (η, ϕ) resolution (0.025x0.025) is realistic for the barrel region of a typical electromagnetic calorimeter, but substantially underestimates a realistic tracker (η, ϕ) resolution. As an example, the (η, ϕ) resolution of the CMS tracker for a 1 GeV (10 GeV) track is ~ 10 (~ 100) times better than what assumed in Ref. [24]. Even if a fixed-size regular (η, ϕ) grid could faithfully represent the (η, ϕ) resolution of a silicon-based tracker device, one would be confronted with a very large number of cells and, consequently, many technical complications in training the algorithm (e.g., GPU memory consumption and image sparsity).

While these might be marginal problems for traditional approaches to event reconstruction, we see them as extremely sub-optimal for the CMS global event description, which is what we are more familiar with. This fact motivated us to extend the work of Ref. [24], investigating network architectures that would not rely on an image-based representation of the event or, more generally, on specific assumptions on the detector geometry or granularity. We consider three different network architectures: a fully-connected deep neural network (DNN), a network with GRU layers [16], and a GGNN [14]. By utilizing feature- or particle-based representations of the event, these three architectures don’t rely on specific aspects of the underlying detector geometry and can be easily integrated in a global or local particle-flow reconstruction like those employed at the LHC [26, 27].

3 Dataset

The dataset employed in this work consists of simulations of LHC proton-proton collisions at a center-of-mass energy of 13 TeV, created using the PYTHIA 8.223 [28] event generator, tune 4C [29]. The generated events correspond to a sample of Z bosons decaying to a pair of neutrinos and produced in association with at least one quark or gluon, resulting in at least one jet. Following Refs. [1, 11], the generation of the underlying event is turned off, so that the reconstructed particles in the event could be divided into two well defined categories: particles from the high- p_T LV or pileup particles. This technical aspect facilitates the sample preparation and training while having a negligible impact on performances. The list of particles originating from the hard collision and the consequent shower of quarks and gluons are given as input to the DELPHES 3.3.2 [30] detector-simulation software. As typically done in previous studies of this kind [1, 11, 24], no detector reconstruction effect is applied at this stage. DELPHES is mainly used as a convenient tool to read the HEPMC files generated by PYTHIA, overlay the pileup events, and store the event content in a ROOT [31] file, preserving the provenance information for each particle (LV or pileup events). We set the average number of pileup interactions $\bar{n}_{PU}$ to 20, 80 or 140, and randomly generate the per-event pileup multiplicity according to a Poisson distribution. A corresponding number of events is then sampled from a so-called minimum-bias library, consisting of an inclusive sample of PYTHIA-generated QCD events. The Z +jets events are generated in batches of 100 events each. For each batch, pileup events are samples from a batch-specific library of 10000 events. This allows to reduce the probability of event re-usage to a negligible level, preventing the network to learn specific patterns of specific events in the training datasets.

All final-state particles except neutrinos are kept and assumed to be massless, as normally done for high- p_T physics studies. For each particle, the following information is stored:

- The momentum coordinates (p_T, η, ϕ) .
- The electric charge of the particle.
- The local shape α_i^β for $\beta = 1, 2, 3$, as defined in Ref. [1] (see Eq. 1), both summing over the full event ($\alpha_i^{\beta,F}$) or only over the charged particles from the LV ($\alpha_i^{\beta,C}$).

Fig. 2. Conceptual depiction of the GGNN model architecture. The event is pre-processed by linking local particles together, after which it is fed to 3 GGNN layers with time-steps [2, 1, 1] and including a residual connection from the first to the third layer. This is then passed, individually per graph node, to a fully-connected network that outputs a [0,1] pileup classification score. Adam is used with a learning rate of 0.004 to minimize the binary cross-entropy. The output of the network is checked to be a well-calibrated probability.

- The PUPPI weight associated to the particle: a number between 0 and 1 that can be related to the probability of the particle being pileup.
- A flag set to 1 for charged particles from the LV, -1 for charged pileup particles and 0 for all neutral particles. This flag provides a simple encoding for when CHS is used.
- A *pileup flag* indicating whether the particle belongs to the hard scattering or to any of the pileup vertices. This information is used as the ground truth later on.

We assume units such that $\hbar = c = 1$. Furthermore, we store for each event the median p_T per unit area in (η, ϕ) for all particles (ρ), for all charged particles alone (ρ_C), and for all neutral particles alone (ρ_N).

4 Network architectures

PUPPI can be straightforwardly interpreted as a per particle classification algorithm. Under this point of view, traditional metrics such as the Receiver Operating Characteristic (ROC) curve (true positive rate against false positive rate) or the accuracy (fraction of correctly classified particles) can be used. The choice of the shape variable α^1 is then driven by its discriminating power, with the underlying assumption that a better classification performance should correlate with a better reconstruction of physics-motivated quantities which are relevant to study these data. For all the investigated network architectures, we generalize this approach to multiple shape variables, indicated from now on as features. We feed as input to our networks all the particles, with all the features discussed in Sec. 3 except for the pileup flag, which we use as the training ground truth. The global features are concatenated to each particle's individual features. An generalization of PUPPI by mean of ML techniques is already discussed in Ref. [1], where it is asserted that training a Boosted Decision Tree modestly improves performance when compared to the use of α^1 as discriminating quantity.

Our most straightforward model makes use of two stacked fully-connected hidden layers and a final single-neuron layer with a sigmoid activation function. This network is trained, as all the other models, to minimize a binary cross-entropy loss function using the Adam optimizer [32]. This model stands out for its simplicity, as it operates on each particle completely independently of the others, but suffers from a clear issue: while the input includes global (ρ, ρ_C, ρ_N) and local (α_i^C, α_i^F) features, the network has no mechanism by which it could learn these or similar features.

Extending the network architecture beyond a simple per-particle processing, one could overcome this limitation. To this purpose, different network architectures can be chosen. Reference [24] describes an approach based on CNNs. Motivated by the arguments described in Sec. 2, we complement the results of Ref. [24] by studying GRUs and GGNNs. Both these architectures take as input the full list of particles in the event, outputting a per-particle label.

The GRU is a recurrent neural network architecture that sequentially processes each item of an input list, based on the outcome of previous-item processing. While making no assumption on the underlying detector geometry, the GRU architecture implies the use of a ranking principle to order the items in the input list. In our study, the input list contains the particles in the event, which are ordered by their p_T value. This is one of the many arbitrary choices that one could make. In the network, we make use of a bidirectional GRU layer, i.e., we consider both increasing- and decreasing- p_T ordering. The output of this layer is concatenated to each particle's features. We show that this approach does not improve the classification performance with respect to DNNs and traditional methods. This is mainly due by the fact that GRUs require a global ordering criterion, while the information determining if a given particles belongs to the LV or originates from pileuphas mainly to due with the particle's local neighborhood.

The architecture of a GGNN (and in general any graph network) provides the capability of learning specific features of an input item from the nearby neighbors, delivering state-of-the-art performance (see Section 5) without making any assumption on the underlying geometry or input ordering. A GGNN can be considered a special type of Message Passing Neural Network [33]. Each example in the input dataset is represented as a set of vertices $v \in \mathcal{V}$. The vertices are connected by directional edges $e_t \in \mathcal{E}$. The combination of the vertices and the connection edges forms the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. At input, each vertex of the graph is represented by a feature vector $\mathbf{h}_v^0$ of dimension h^0 . During a graph-layer processing, the graph evolves in a set of time steps. Each time step $i > 0$ consists of two operations:

1. For each node v , we generate an incoming message from each connected neighbour vertex v_j . The message $\mathbf{a}_{v,v_j}^i$ is computed multiplying a learnable $h^{i-1} \times h^{i-1}$ matrix $\mathbf{A}_t$ by the feature vector v_j from the previous time interaction:

$$\mathbf{a}_{v,v_j}^i = \mathbf{A}_t \mathbf{h}_{v_j}^{i-1}, \forall v_j : v_j \xrightarrow{t} v. \quad (2)$$

When different kinds of edges are present in the graph, edge-type specific A matrices are used.

2. All messages coming to a node v are averaged. The average is used to create a new representation of the node. Repeating this operation for each node, a new graph is generated, isomorphic to the original one but characterized by a new node representation, learned through the message exchange. In the specific case of a GGNN, the node update is performed utilizing one step of a GRU, initializing the GRU memory cell with the node representation from the previous time-step $\mathbf{h}_v^{i-1}$ and giving the incoming message in input to the GRU cell:

$$\mathbf{h}_v^i = \text{GRU}(\mathbf{h}_v^{i-1}, \overline{\mathbf{a}_{v,v_0}^i, \dots, \mathbf{a}_{v,v_N}^i}). \quad (3)$$

Our final PUPPIML model is pictorially represented in Fig. 2. The GGNN processing procedure is iterated stacking three GGNN layers one after the other. Each layer is characterized by its own node representation (100 for all of them), a set of $\mathbf{A}_t$ matrices and its own number of time steps (two for the first GGNN layer, one for the others). A given layer processes the output representation of the previous layer and creates a graph which is isomorphic to the input graph and characterized by a new vertex representation. In our case, the graph vertices are the particles in the event, and their input representation consists of the (p_T, η, ϕ) vector. Following Ref. [14], the initial input is added to the output of the next-to-last GGNN layer, with a residual connection analogous to what is done with ResNet [34], as shown in Fig. 2. In the last step of the network, the per-vertex (i.e., per-particle) output of the last GGNN layer is given as input to a fully connected DNN, which produces a pileup score, estimating the probability of a given particle to come from any of the pileup collisions. The hyperparameters defining the PUPPIML architecture are indicated in Fig. 2. We include as part of the graph the charged particles even when CHS is applied, as they aid in the classification of neighbouring particles. We train a different network for each mean pileup level (20, 80 and 140). Section 5.2 discusses the generalization to a range of pileup configurations.

For our graph representation, we choose to connect all pairs of particles separated by a distance $\Delta R < R_1$ in the (η, ϕ) plane, uniformly binning the distance ΔR into N_0 discrete edge types. We make all graphs undirected by introducing edges in both the forward and backward direction for each pair of connected particles. Moreover, as opposed to PUPPI, we don't rescale the particles' four-momenta and simply discard all particles for which the network predicts a probability lower than a threshold p_{cut} . The choice of R_1 , N_0 , and p_{cut} is made in order to maximise the network classification performance, as discussed in Sec. 5.

All our models are trained in Keras v2.1.2 [35] with a Tensorflow v1.2.1 [36] backend, or in Tensorflow directly. The implementation of the GGNN is based on publicly available code [37] published by Microsoft under the MIT license. We train the model using the Adam optimizer [32] with a learning rate of 0.004 and early stopping, and apply no pre-processing to our inputs.

5 Results

We make a general comparison of algorithm performances by comparing the ROC curves and the corresponding areas under the curve (auc) for the network models introduced in Section 4 and the baseline rule-based approaches listed in Section 1. For PUPPIML, we tested different architectures, varying the number of GGNN layers, the number of time steps, and the dimensionality of the node representation. We trained 20 randomly-picked configurations for these parameters and considered the one with best performance, measured in terms of the auc on the validation data. We observe a very weak dependence of the auc value on the network architecture.

Using a sample of test data, we derive the auc values for different pileup configurations: a mean pileup $\bar{n}_{\text{PU}} = 20, 80$ or 140 when CHS is applied, and $\bar{n}_{\text{PU}} = 80$ when CHS is not used.

As an example, Figure 3 shows the ROC curves corresponding to $\bar{n}_{\text{PU}} = 80$ with CHS. Table 1 reports the auc for this and the other pileup configurations. Since *SoftKiller* and PUPPI base the decision of whether a particle comes from the LV on its transverse momentum or on the PUPPI weight calculated from α^1 respectively, the first two rows are meant

Fig. 3. Receiver Operating Characteristic (ROC) curve for our proposed features and models. Classifiers based on PUPPI weight and p_T are included as an indicator of the expected performance of PUPPI and *SoftKiller* respectively. GGNN outperform other proposed architectures.

	$\bar{n}_{PU}$	20 (CHS)	80 (CHS)	140 (CHS)	80 (No CHS)
p_T (SoftKiller)		92.3%	92.3%	92.5%	64.9%
PUPPI weight		94.1%	93.9%	94.4%	65.1%
Fully-connected		95.0%	94.8%	94.8%	68.5%
GRU		94.8%	94.8%	94.7%	68.8%
GGNN		96.1%	96.1%	96.0%	70.1%

Table 1. Area under the curve for the different discriminating variables and models. The highest values, highlighted in bold, are obtained when the GGNN architecture is used.

as an indicator of their expected performances. For this purpose, we allow the PUPPI weight to take on negative values if it is to the left of the median. We observe that the GGNN consistently obtains the best classification performance. Based on these results, we only consider the performance of our GGNN and the state of the art algorithms from here on. We tune $R_1 = 0.3$ and $N_0 = 5$ to maximize the auc. We fix the threshold parameter $p_{cut} = 0.4$ at $\bar{n}_{PU} = 20$ and $p_{cut} = 0.35$ at $\bar{n}_{PU} = 80$ or $\bar{n}_{PU} = 140$, so as to minimize the offset between the reconstructed and the LV observables. We find that minimizing the offset for one observable also approximately minimizes the offset for all other observables.

Figure 4 shows the effect of running our proposed approach on an event at $\bar{n}_{PU} = 20$. The event reconstructed by PUPPIML is shown on the bottom right, with particles represented as filled circles, sized according to their p_T . Dots are colored as orange (dark) if they come from the LV and blue (light) if they originate from pileup interactions. The event is also shown as reconstructed by PUPPI (bottom left) and by *SoftKiller* (bottom center). Moreover, we show the ground truth on the top left and the unprocessed event on the top right. Using the same plotting conventions, Fig. 5 shows the effect of the three PV mitigation algorithms for three jets in events with at $\bar{n}_{PU} = 80$. We note qualitatively that PUPPIML improves on the state-of-the-art approaches, removing some low- p_T pileup particles close to the jet that PUPPI does not (dotted ellipses), and removing some high- p_T particles far away from the jet that *SoftKiller* does not (dashed ellipses).

5.1 Jet variables

Following the methodology of Ref. [1], we run the anti- k_t clustering algorithm [38] using FastJet 3.3.0 [39, 40] with $R = 0.7$ on both the LV-only and full events, after applying PUPPI, *SoftKiller*, and PUPPIML. We consider a jet to match a LV-only one if the distance between the two in the (η, ϕ) plane is less than $\Delta R_{max} = 0.3$. Figure 6 shows the dependence of the jet p_T resolution on $\bar{n}_{PU}$ and on the jet p_T . Figure 7 shows this resolution for jets in the $50 < p_T < 100$ GeV and the $100 < p_T < 200$ GeV region for $\bar{n}_{PU} = 140$. We choose to quantify the resolutions of all indicated variables using the half-difference between the 14th and 86th percentile, $\delta = (P_{86} - P_{14})/2$, so as to avoid statistical fluctuations from potential outliers. Our method shows an improvement in resolution. Furthermore, such

Fig. 4. Depiction of the effect of running the different pileup mitigation algorithms on a random test event at $\bar{n}_{\text{PU}} = 20$. Particles from the LV are shown in orange (dark) and pileup particles are shown in blue (light). We show the ground truth (top left), the event contaminated by the pileup interactions (top right) and the reconstructed event after running PUPPI (bottom left), *SoftKiller* (bottom center) and PUPPIML (bottom right). PUPPIML improves on PUPPI by eliminating some of the low p_T particles close to jets and on *SoftKiller* by eliminating the high p_T particles that are far away from jets. All algorithms are run with CHS.

an improvement is larger for larger values of $\bar{n}_{\text{PU}}$, which makes the proposed solution particularly relevant for future LHC runs, when the average number of pileup interactions is expected to increase substantially. Mainly, at low pileup levels and very high p_T the improvement with respect to the already adopted algorithms is marginal. We also test the resolution worsening when CHS is not applied. In this case, we train the network without the CHS flag in the dataset. Figure 8 shows that, while the resolution gain is smaller, there is still a small advantage in using PUPPIML.

Figure 9 shows how the mitigation algorithms affect the resolution of the jets position in the (η, ϕ) space, quantified through the mean ΔR . PUPPIML shows an improvement of $\sim 15\%$ in resolution at $\bar{n}_{\text{PU}} = 20$, $\sim 25\%$ at $\bar{n}_{\text{PU}} = 80$, and $\sim 30\%$ at $\bar{n}_{\text{PU}} = 140$.

Finally, in order to study the reconstruction of jet shapes, we consider the jet-mass resolution, which is commonly taken as a proxy to evaluate a pileup-removal performance. Figure 10 shows the resolution as a function of $\bar{n}_{\text{PU}}$ and of the jet p_T . We find that a better mass resolution can be obtained when PUPPIML is applied.

Since our network takes the full event as input, we expect it to be able to correct not only jet-related quantities but also global ones. Figure 11 shows that this is indeed the case, when the missing transverse energy (MET) in the event is considered. Since this variable involves summing over all present particles, its sensitivity to pileup contamination is typically stronger.

5.2 Robustness

It is important that the network learns a pileup mitigation strategy that is effective across samples and not overly dependent on the specific features of the training dataset. We probe this dependence by evaluating the network auc performance on (i) different mean pileup levels, (ii) different PYTHIA8 tunes and (iii) different decay processes. Table 2 shows that the performance degrades when evaluating the network at highly different pileup distributions. Nevertheless, a model trained at multiple pileup levels is effective over a wide range. Tables 3 and 4 show that PUPPIML reaches auc values larger than the corresponding *SoftKiller* and PUPPI ones, even when using different tunes and on unseen decay channels respectively. This feature is particularly important in view of the fact that pileup mitigation algorithms are designed and tuned on simulation to then be used on data.

5.3 Computational Performance Considerations

We run all our experiments on a Titan Xp GPU and observe speeds of $\mathcal{O}(100 \text{ events/s})$ at $\bar{n}_{\text{PU}} = 20$, $\mathcal{O}(20 \text{ events/s})$ at $\bar{n}_{\text{PU}} = 80$ and $\mathcal{O}(10 \text{ events/s})$ at $\bar{n}_{\text{PU}} = 140$ for the network forward propagation. This is more than an order of

Fig. 5. Depiction of the effect of running the different pileup mitigation algorithms on three jets at $\bar{n}_{\text{PU}} = 80$. Each circle represent one particle and the size of the circle is a function of the particle p_T . Particles from the LV are shown in orange (dark) and pileup particles are shown in blue (light). The top row shows the the ground truth, before (left) and after (right) adding the pileup contribution. The bottom row shows the output returned by PUPPI (left), *SoftKiller* (center), and PUPPIML (right), all running after applying CHS. PUPPIML improves on PUPPI by eliminating some of the low p_T particles close to jets (dotted ellipses) and on *SoftKiller* by eliminating some of the high p_T pileup particles that are far away from jets (dashed ellipses).

		Trained on $\bar{n}_{\text{PU}}$			
		20	80	140	20+80+180
Evaluated on $\bar{n}_{\text{PU}}$	20	96.1%	95.6%	95.0%	96.1%
	80	95.7%	96.1%	95.9%	96.1%
	140	95.1%	96.0%	96.0%	96.1%

Table 2. Area under the curve for PUPPIML, trained and evaluated on different pileup configurations. We train models at $\bar{n}_{\text{PU}}$ equal to 20, 80 and 140, and on a dataset containing equal splits of the three. While the performance decreases when testing on previously unseen pileup configurations, it is still superior to the *SoftKiller* and PUPPI proxies (see Table 1). A model trained on the combined dataset is capable of learning a strategy that generalizes to a wide range of pileup levels. Some results are highlighted for visual clarity.

	Tune 4C	Tune 1	CUEP8M1
p_T (SoftKiller)	92.5%	91.7%	92.4%
PUPPI weight	94.4%	93.7%	94.2%
PUPPIML	96.0%	95.4%	95.9%

Table 3. Area under the curve (auc) for different pileupmitigation algorithms. The pileup configuration is fixed to $\bar{n}_{\text{PU}} = 140$. The models are trained on $Z \rightarrow \nu\bar{\nu}$ events, generated with PYTHIA 4C tune and then tested on the same process, generated with different PYTHIA8 tunes. For PUPPIML, the most significant drop of performance is observed using Tune 1, yet with an auc value larger than the corresponding *SoftKiller* and PUPPI values, also shown for comparison.

magnitude slower than *SoftKiller* and PUPPI. We note, however, that our implementation is not tuned for computational performance. We verified that reducing the number of layers, the number of edges and size of the hidden representation can result in improvements of a factor of 2 with a reduction in auc of less than 0.2%. Moreover, the network requires a

Fig. 6. Jet p_T resolution as a function of $\bar{n}_{\text{PU}}$ for jets in the range $100 < p_T < 150$ GeV (top) and as a function of the jet transverse momentum at $\bar{n}_{\text{PU}} = 140$ (bottom) when CHS is applied.

Fig. 7. Jet p_T resolution at $\bar{n}_{\text{PU}} = 140$ for jets in the range $50 < p_T < 100$ GeV (top) and $100 < p_T < 200$ GeV (bottom) when CHS is applied.

	$Z \rightarrow \nu\bar{\nu}$	$H \rightarrow b\bar{b}$	$H \rightarrow g\bar{g}$
p_T (SoftKiller)	92.5%	92.0%	92.0%
PUPPI weight	94.4%	94.0%	94.0%
PUPPIML	96.0%	95.7%	95.7%

Table 4. Area under the curve (auc) for different pileupmitigation algorithms. The pileup configuration is fixed to $\bar{n}_{\text{PU}} = 140$. The models are trained on $Z \rightarrow \nu\bar{\nu}$ events, generated with PYTHIA 4C tune and then tested on different processes, generated with the same PYTHIA tune. The PUPPIML model is able to generalize to decays involving multiple jets and including bottom quarks and gluons with only a small decrease in auc performance. The corresponding values for the *SoftKiller* and PUPPI proxies are shown for comparison.

pre-processing step in which the neighbours to each particle are found; while this is slow if done through pairwise brute force (as we did in our study, trading simplicity for computation efficiency), the computational cost can be reduced substantially if implemented through a heap or other more advanced data structures. Similarly to PUPPI, removing a large portion of the pileup particles (e.g., by applying a threshold on the particles' p_T) results in much faster jet clustering downstream, partially compensating for the time spent to perform the pileup mitigation.

We further note that our approach maintains locality, since each particle is connected only to those particles with distance $\Delta R < R_1$ from it. By choosing $R_1 = 0.3$ and with a total of four time-steps across the three GGNN layers, only particles local neighborhood of $\Delta R < 1.2$ can influence the classification of a given particle. Taking advantage of this

Fig. 8. Jet p_T resolution at $\bar{n}_{PU} = 80$ as a function of the jet p_T (top) and for jets in the range $100 < p_T < 200$ GeV (bottom) when CHS is not applied.

Fig. 9. Distribution of $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ between the jet in absence and in presence of pileup effects, for $\bar{n}_{PU} = 80$ (top) and $\bar{n}_{PU} = 140$ (bottom). CHS is applied to all algorithms.

local nature, one could further speed up the PUPPIML inference by parallelizing it. Asynchronous Gated Graph Neural Networks and "towers" in Message Passing Neural Networks are a step in this direction and can handle overlapping parts of the graph.

6 Conclusions and outlook

This paper introduces a new pileup mitigation algorithm, PUPPIML, that extends the commonly-used PUPPI algorithm through Deep Learning via Graph Neural Networks. The proposed algorithm shows resolution improvements in all the considered jet- and event-related quantities. For jet-related quantities, PUPPIML provides a resolution improvement up to $\sim 30\%$ with respect to PUPPI and even more with respect to SoftKiller.

The algorithm is designed to return the probability that a given particle belongs to the leading vertex of the event (as opposed to originating from pileup). Its output can be used to weight each particle by its probability to originate from the leading vertex. On the other hand, it can also be used for pileup removal by tuning a threshold on the probability value returned by the classifier. When tuned on different pileup configurations, this threshold is found to be mildly dependent on the details of the training sample.

By representing an event as a graph of reconstructed particles, PUPPIML allows to avoid any assumptions on the geometry or granularity of a given detector. In this respect, it represents an interesting alternative to the approach presented in Ref. [24].

Fig. 10. Jet mass resolution as a function of $\bar{n}_{\text{PU}}$ for jets in the range $100 < p_T < 150$ GeV (top) and as a function of the jet transverse momentum at $\bar{n}_{\text{PU}} = 80$ (bottom) when CHS is applied.

Fig. 11. Missing transverse energy resolution at $\bar{n}_{\text{PU}} = 20$ (top) and $\bar{n}_{\text{PU}} = 80$ (bottom). One should notice that, due to the sample definition, the events are characterized by real MET, corresponding to the Z p_T .

In terms of future work, we believe that a further exploration of similar locality-preserving architectures, in particular variations of Message Passing Neural Networks, could further improve performances. Adversarial training, requiring the network to correct the full event directly as opposed to simply classifying individual particles, could aid in improving the resolution gain on jet substructure quantities.

Acknowledgments

We are grateful to Caltech and the Kavli Foundation for their support of undergraduate student research in cross-cutting areas of machine learning and domain sciences. This work was conducted at “*iBanks*”, the AI GPU cluster at Caltech. We acknowledge NVIDIA, SuperMicro and the Kavli Foundation for their support of “*iBanks*”. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement n° 772369). This project is partially supported by the United States Department of Energy, Office of High Energy Physics Research under Caltech Contract No. de-sc0011925.

References

1. D. Bertolini, P. Harris, M. Low, N. Tran, JHEP **10**, 059 (2014), 1407.6013

2. G. Aad et al. (ATLAS), JINST **3**, S08003 (2008)
3. S. Chatrchyan et al. (CMS), JINST **3**, S08004 (2008)
4. S. Chatrchyan et al. (CMS), JINST **9**, P10009 (2014), 1405.6569
5. ATLAS Collaboration, ATLAS-CONF-2010-069 (2010)
6. M. Cacciari, J. Rojo, G.P. Salam, G. Soyez, JHEP **12**, 032 (2008), 0810.1304
7. D. Krohn, M.D. Schwartz, M. Low, L.T. Wang, Phys. Rev. **D90**, 065020 (2014), 1309.4777
8. M. Cacciari, G.P. Salam, G. Soyez, Phys. Rev. **D92**, 014003 (2015), 1404.7353
9. P. Berta, M. Spousta, D.W. Miller, R. Leitner, JHEP **06**, 092 (2014), 1403.3108
10. G. Soyez, G.P. Salam, J. Kim, S. Dutta, M. Cacciari, Phys. Rev. Lett. **110**, 162001 (2013), 1211.2811
11. M. Cacciari, G.P. Salam, G. Soyez, Eur. Phys. J. **C75**, 59 (2015), 1407.0408
12. Tech. Rep. CMS-PAS-JME-14-001, CERN, Geneva (2014), <https://cds.cern.ch/record/1751454>
13. Tech. Rep. ATLAS-CONF-2017-065, CERN, Geneva (2017), <https://cds.cern.ch/record/2281055>
14. Y. Li, D. Tarlow, M. Brockschmidt, R.S. Zemel, CoRR **abs/1511.05493** (2015), 1511.05493
15. J. Orbach, Archives of General Psychiatry **7**, 218 (1962), <https://jamanetwork.com/journals/jamapsychiatry/articlep>
16. K. Cho, B. van Merriënboer, D. Bahdanau, Y. Bengio, CoRR **abs/1409.1259** (2014), 1409.1259
17. A. Radovic et al., Nature **560**, 41 (2018)
18. CMS Collaboration, Tech. Rep. CMS-DP-2017-027, CERN, Geneva (2017), <http://cds.cern.ch/record/2275226>
19. ATLAS Collaboration, Tech. Rep. ATL-PHYS-PUB-2017-003, CERN, Geneva (2017), <http://cds.cern.ch/record/2255226>
20. A. Aurisano et al., JINST **11**, P09001 (2016), 1604.01444
21. I. Henrion et al., *Neural message passing for jet physics* (2017), https://cims.nyu.edu/~bruna/Media/nmp_jet.pdf
22. M. Abdughani, J. Ren, L. Wu, J.M. Yang (2018), 1807.09088
23. N. Choma et al., CoRR **abs/1809.06166** (2018)
24. P.T. Komiske, E.M. Metodiev, B. Nachman, M.D. Schwartz, JHEP **12**, 051 (2017), 1707.08600
25. S. Chatrchyan et al. (CMS), JINST **9**, P10009 (2014), 1405.6569
26. M. Aaboud et al. (ATLAS), Eur. Phys. J. **C77**, 466 (2017), 1703.10485
27. A.M. Sirunyan et al. (CMS), JINST **12**, P10003 (2017), 1706.04965
28. T. Sjöstrand et al., Comput. Phys. Commun. **191**, 159 (2015), 1410.3012
29. R. Corke, T. Sjöstrand, JHEP **03**, 032 (2011), 1011.1759
30. J. de Favereau et al. (DELPHES 3), JHEP **02**, 057 (2014), 1307.6346
31. R. Brun, F. Rademakers, Nucl. Instrum. Meth. **A389**, 81 (1997)
32. D.P. Kingma, J. Ba, CoRR **abs/1412.6980** (2014), 1412.6980
33. J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals, G.E. Dahl, CoRR **abs/1704.01212** (2017), 1704.01212
34. K. He, X. Zhang, S. Ren, J. Sun, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
35. F. Chollet et al., *Keras*, <https://keras.io> (2015)
36. M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems* (2015), software available from tensorflow.org, <https://www.tensorflow.org/>
37. Microsoft, *Gated graph neural network samples*, <https://github.com/Microsoft/gated-graph-neural-network-samples> (2017)
38. M. Cacciari, G.P. Salam, G. Soyez, JHEP **04**, 063 (2008), 0802.1189
39. M. Cacciari, G.P. Salam, G. Soyez, Eur. Phys. J. **C72**, 1896 (2012), 1111.6097
40. M. Cacciari, G.P. Salam, Phys. Lett. **B641**, 57 (2006), hep-ph/0512210