

This is the accepted manuscript made available via CHORUS. The article has been published as:

Machine-learning-based jet momentum reconstruction in heavy-ion collisions

Rüdiger Haake and Constantin Loizides

Phys. Rev. C **99**, 064904 — Published 17 June 2019

DOI: [10.1103/PhysRevC.99.064904](https://doi.org/10.1103/PhysRevC.99.064904)

Machine Learning based jet momentum reconstruction in heavy-ion collisions

Rüdiger Haake¹ and Constantin Loizides²

¹*Yale University, Wright Laboratory, New Haven, CT, USA*

²*ORNL, Physics Division, Oak Ridge, TN, USA*

(Dated: June 4, 2019)

The precise reconstruction of jet transverse momenta in heavy-ion collisions is a challenging task. A major obstacle is the large number of (mainly) low- p_T particles overlaying the jets. Strong region-to-region fluctuations of this background complicate the jet measurement and lead to significant uncertainties. In this paper, a novel approach to correct jet momenta (or energies) for the underlying background in heavy-ion collisions is introduced. The proposed method makes use of common Machine Learning techniques to estimate the jet transverse momentum based on several parameters, including properties of the jet constituents. Using a toy model and HIJING simulations, the performance of the new method is shown to be superior to the established standard area-based background estimator. The application of the new method to data promises the measurement of jets down to extremely low transverse momenta, unprecedented thus far in data on heavy-ion collisions.

INTRODUCTION

In ultrarelativistic heavy-ion collisions a new state of nuclear matter is created: the Quark-Gluon Plasma (QGP) [1]. In the QGP, deconfined quarks and gluons interact strongly and form a hot and dense medium that can be approximately described by hydro- and thermodynamics. The regime of strong coupling at large distances, especially in systems of high temperature or large energy densities, is still not well understood in Quantum Chromodynamics (QCD). An ideal self-generated probe to explore the properties of the medium and its interactions are particle jets [2]. Reconstructed in the detector as collimated sprays of color-neutral particles, jets are created in a large-momentum-transfer scattering of partons in the early stage of a high-energy collision. Their production is well understood within the framework of QCD and their rates can be perturbatively calculated in vacuum. In a heavy-ion collision, jets traverse the strongly-interacting medium and interact mainly non-perturbatively and, thus, can serve as valuable probes of the QGP.

The reconstruction of particle jets in heavy-ion collisions is a complex task. The main obstacle is the overwhelmingly large background of particles that do not originate from hard interactions. In ALICE [3], the mean momentum density in 0–10% most central collisions at $\sqrt{s_{NN}} = 2.76$ TeV leads to a contribution to the jet momentum that is already of the order of the typical jet momentum itself. The average charged particle transverse momentum density for particles with momenta above 0.15 GeV/c is $\langle \rho \rangle \approx 138.2$ GeV/c per unit area, while its standard deviation is $\sigma(\rho) = 18.5$ GeV/c [4]. Since jets are rare objects, these numbers provide already a good estimate of the mean background in the selected events. In addition, this background shows large uncorrelated and also correlated region-to-region fluctuations. Uncorrelated fluctuations are due to random Poissonian fluctuations of the number of particles and their momenta. Sources of correlated fluctuations are e.g. physical correlations of the particles from the particle flow or also the non-uniform detector acceptances. These fluctuations have a large impact on the reconstructed jet momentum and on the jet axes by directly affecting the jet

finding algorithm and eventually result in large uncertainties on the final measurements. An approach to at least lower the impact of the background at the expense of a potential fragmentation bias is a higher p_T -cut for constituents used in the jet finding algorithm. This massively reduces the background, which mostly consists of low- p_T particles, but it also discards the low- p_T parts of the jet. The treatment of the background and its fluctuations depends on the observable under study. In this paper, the focus is on the correction of observables based on jet momentum, i.e. the correction of the jet energy scale, without applying a particular constituent cut. The impact of the background on the jet (sub)structure, e.g. by distorting the jet axis, is not discussed here.

In the standard method for jet spectra measurements in ALICE, the background momentum density per unit area is calculated on an *event-by-event* basis. Each jet is then corrected by taking into account the event-averaged density multiplied by the jet area. The area-based method corrects the jet momentum for the average background but leads to large residual fluctuations. These residual fluctuations are then typically corrected for on a statistical basis in an unfolding procedure, see for instance Ref. [5].

The new approach, introduced in this paper, calculates the corrected jet momentum on a *jet-by-jet* basis to reduce the residual fluctuations and to allow a more precise estimate for the jet momentum. As we demonstrate below, this enables the measurement of jets in heavy-ion collisions with transverse momenta much lower than what is currently possible. We apply Machine Learning (ML) techniques, which are widely used in the HEP community [6], to obtain the mapping between jet parameters, e.g. constituent momenta, and the *true* transverse momentum of the jet. The background estimator is trained on a toy model that embeds jets with known (true) transverse momenta in a simulated thermal heavy-ion background. Using this toy model, the performance of the new method is compared to that of the established correction method. To demonstrate the applicability of the new approach on jet populations in heavy-ion collisions, the correction method performance is then also presented when applied onto HIJING [7] simulations.

JET DEFINITION AND TOY DATASETS

The basis of this analysis are reconstructed charged jets.¹ They are reconstructed by the anti- k_T algorithm [8] implemented in Fastjet [9]. Since the background and its fluctuations increase quadratically in R , current jet analyses in heavy-ion physics typically avoid larger resolution parameters. The resolution parameter is chosen deliberately at the relatively high value $R = 0.4$ to test the method. Charged particles above $p_T = 0.15$ GeV/ c are taken into account in the jet reconstruction.

In contrast to an algorithmic correction approach where the correction, i.e. the mapping of raw and corrected jet momentum, is directly described, ML techniques *learn* the mapping from a training dataset. The quality of the training dataset is crucial for the final performance of the correction method on real data. It should only encode well-known physics features and should be as simple as possible. A good dataset allows the method to learn general features connected to the problem and enables the method to generalize to even slightly different data. To judge its performance, the correction method is evaluated on the toy dataset and directly compared to the established correction method (described below). The events in the toy dataset used for the evaluation are not used in the training phase.

To create events with particle jets in a heavy-ion background, PYTHIA-generated events are embedded in a thermal background. The events from PYTHIA [10] are generated at $\sqrt{s} = 2.76$ TeV with particles in a pseudorapidity range of $|\eta| < 0.9$ and in full azimuth.

The PYTHIA part of the events is created with PYTHIA6, using the Perugia 2011 tune [11]. The thermal background is created by randomly distributing charged particles according to realistic particle multiplicity and momentum distributions. The multiplicity distribution is modeled with a Gaussian function with a mean of 1800 and a width of 200. This roughly reproduces the multiplicity in central Pb–Pb events at $\sqrt{s_{NN}} = 2.76$ TeV [12, 13]. The impact of different thermal background multiplicities on the model performance will be evaluated below. The momentum distribution is a modified power law function and defined such that it coincides with the momentum distribution at low p_T [14] but falls much faster for higher particle momenta roughly above 4 GeV/ c . Actual details on the higher- p_T region of the thermal spectrum do not influence the model much, as this region only corresponds to a small part of the dataset.

For the actual application of the correction method to real detector-level heavy-ion events, the method should be trained on a dataset that describes the real, reconstructed data as precise as possible. For instance, the momentum distribution

could be adjusted to better fit the real data at low p_T . Or the entire thermal model could be replaced by mixed events created from real data [15].

For the training, the supervised learning techniques that are applied need a truth value assigned to each sample, i.e. to each jet. In the present model, the truth that will be approximated by the new correction method is the true jet momentum. Here, it is defined as the reconstructed jet momentum multiplied by the momentum fraction that is carried by PYTHIA particles in the jet,

$$p_{T, \text{chjet}}^{\text{true}} = p_{T, \text{chjet}}^{\text{raw}} \cdot \sum_i p_{T, \text{const } i}^{\text{PYTHIA}} / \sum_i p_{T, \text{const } i}. \quad (1)$$

In this equation, $p_{T, \text{chjet}}^{\text{raw}}$ is the reconstructed jet transverse momentum before any background correction. With this definition of the true jet momentum, also the true background is implicitly defined: It consists of all the particles from the thermal model. As an alternative definition, the true jet momentum can also be defined as the momentum of jets reconstructed by PYTHIA particles only and geometrically matched to the reconstructed jets in the full toy event. Since the background influences the jet finding algorithm, these matched jets are conceptually closer to the perfectly corrected jets. On the other hand, this definition is technically more complex since it needs further parameters like the matching radius and is thus less robust. There is a chance for mismatching jets, in particular at low transverse momentum. However, in the end both truth definitions allow to train a model with very similar performance.

The toy model dataset is split into a training and a testing dataset, which is a standard procedure for ML-based approaches. In total, 5M toy events were generated from which roughly 6.5M samples were extracted for further analysis. More samples would be available in the generated events, but only fractions are used to save memory. In subsequent analyses, these fractions are fully taken into account in the normalization. For the training, 10% of the data is used. The remaining 90% form the testing dataset.

For the training of the estimators, only jets with $p_{T, \text{chjet}}^{\text{true}} > 5$ GeV/ c are processed. The inclusion of jets below this threshold in the training distorted the corrected spectrum at higher p_T . The jet production in bins of $p_{T, \text{hard}}$ have not been reweighted before being used for the training. Therefore, the training sample consists of more jets at higher transverse momentum compared to the natural abundances. This has the advantage that the training sample is not dominated by low- p_T jets.

Besides the evaluation of the background estimator concerning its resolution and potential biases, the performance of the correction method for actual heavy-ion-like collision data is also tested. For this purpose, the toy model cannot be used as it represents just one hard jet–jet interaction in a soft background. Instead, a test dataset of 2M events is created by HIJING at the energy of $\sqrt{s_{NN}} = 2.76$ TeV, with jet quenching

¹ An extension of this method to fully-reconstructed jets including neutral particles is straightforward.

switched off in the simulation. HIJING is tuned to roughly reproduce the expected multiplicities. An impact parameter of $b < 3.6$ fm and $p_{T, \text{hard}} > 5.8$ leads to a nearly Gaussian multiplicity distribution with a mean value of $N \approx 2200$ and a width of 200, which is further discussed below.

BACKGROUND ESTIMATORS

As introduced above, the established method to correct the transverse momentum of jets is based on a per-event background estimate [16]. In each event, clusters are reconstructed using the k_T algorithm because it is particularly sensitive to background. After removing the two hardest clusters from the cluster collection, the estimated background density is calculated as the median of the k_T -cluster momentum densities. It is defined by

$$\rho = \text{median} \left(\frac{p_{T,i}}{A_i} \right), \quad (2)$$

where i represents the index of all k_T -clusters of the collection and A_i their areas. The median provides stability against outliers in the underlying distribution, e.g. against outliers from the jet signal clustered by the k_T algorithm. Jets are corrected for this background density on a jet-by-jet basis by subtracting the background density multiplied by the jet area, $p_{T, \text{ch jet}}^{\text{rec}} = p_{T, \text{ch jet}}^{\text{raw}} - \rho \cdot A$.

Residual fluctuations are usually treated in an unfolding procedure. They can be quantified by embedding jets of known transverse momenta into heavy-ion events. These jets are affected by the heavy-ion environment and, therefore, their reconstructed momenta are directly modified by residual fluctuations. δp_T quantifies the fluctuations and is defined by

$$\delta p_T = p_{T, \text{ch jet}}^{\text{rec}} - p_{T, \text{ch jet}}^{\text{true}}, \quad (3)$$

for reconstructed jets matched with a known reference jet. As described in detail in [4], these fluctuations can be approximated by a Γ -function. In an alternative procedure to quantify residual fluctuations, one calculates the background-corrected momentum of particles in jet-sized cones randomly placed in heavy-ion events.

The novel background estimator introduced in this paper follows a different approach. Instead of calculating the background once for an event, the background correction is applied completely on a jet-by-jet basis. The main idea behind this is that the properties of background particles and those belonging to the jet process are very different. While the background is dominated by low- p_T particles, the jet signal should contain sizable higher momentum contributions. The information on the fraction of momentum from background is therefore partially encoded within the jet. However, the relation of the input parameters, including jet constituents, and the true jet momentum is not trivial. Machine Learning techniques are ideal to approximate this mapping by learning from data instead of modeling the relation by hand with expert knowledge.

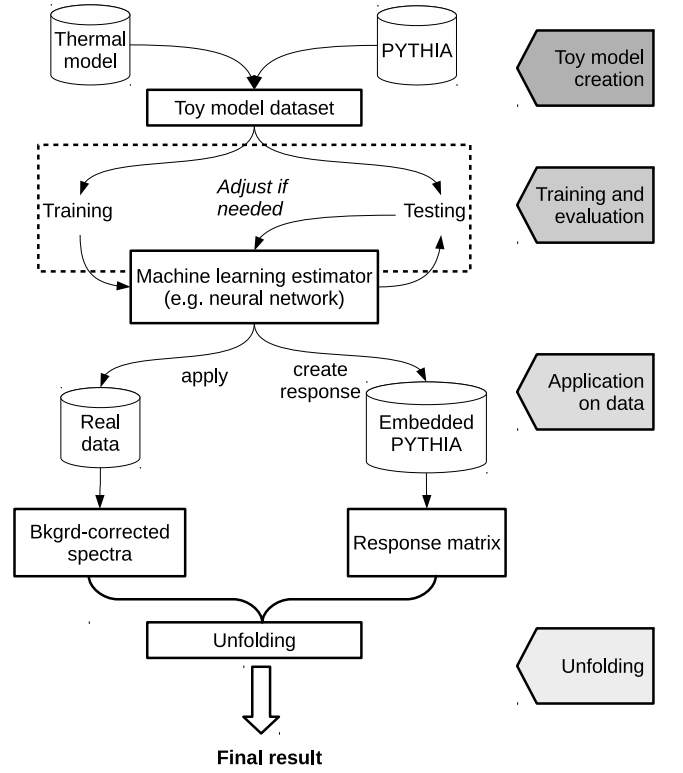


FIG. 1. Illustration of the full proposed analysis strategy. Explanations are given in the text.

Since the task is to predict a numerical value for each sample, a regression algorithm needs to be used.

Several algorithms have been evaluated and will be compared in the next section: shallow neural networks [17], random forests [18], and linear regression. The neural network is implemented as a shallow multi-layer perceptron, consisting of three layers with [100, 100, 50] nodes. The ADAM optimizer [19], an improved stochastic gradient-based optimizer, is used and the nodes/neurons are activated by the ReLU activation function [20]. For the random forest algorithm, an ensemble of 30 decision trees is used. The linear regression model is the simplest model probed. All models are used as implementations within the *scikit-learn* [21] Python module. If a parameter is not explicitly given above, its default setting in *scikit-learn* v0.19.2 is used.

In order to find a suitable combination of input parameters, the analysis was repeated for a large variety of parameter sets. The number of parameters used is kept small to avoid a dependence on data subtleties. Eventually, the following input parameters prove to be useful, discriminative features:

- (1) The uncorrected jet momentum as reconstructed by the jet finding algorithm,
- (2) the jet transverse momentum, corrected by the established area-based method,
- (3) several jet shape observables, namely jet mass, radial momentum, momentum dispersion, and LeSub,

- (4) the number of constituents within the jet,
- (5) mean and median of all constituent transverse momenta,
- (6) the transverse momenta of the first ten leading, i.e. hardest, particles within the jet.

The new ML-based approach defines a background estimator that works on a jet-by-jet basis and that uses the jet constituents as major input. Residual fluctuations – which turn out to be smaller than for the established area-based background correction method – can still be corrected in an unfolding procedure. The corresponding response should be estimated with the embedding method, given in Eq. 3.

The proposed analysis strategy to train, validate, and apply the model to real data is depicted in Fig.1. The first step is the creation of the toy model data, which is described in the previous section. Proton–proton data promise to provide more realistic vacuum jets compared to PYTHIA simulations and can be used instead if available for comparable detector conditions and same energy. For the same reason, as a more accurate alternative to a thermal model, the vacuum jets can also be embedded into heavy-ion events or mixed heavy-ion events. Background estimators trained on different toy model configurations could serve for a systematic uncertainty analysis on real data. The second step is the training of the ML-based estimators and their evaluation on the toy model as will be described in the next section. Training and evaluation dataset are independent subsamples of the full toy model dataset. During this step, the model hyperparameters can be adjusted to obtain a good performance. These first two steps are presented in this paper. In the third step, the trained and commissioned estimator is applied to heavy-ion data to obtain background-corrected spectra, e.g. jet momentum or jet mass spectra. In addition, a response matrix is created to unfold residual fluctuations (and possible detector effects). We propose to create the response matrix by embedding vacuum jets into a realistic background, e.g. given by data. The last step is the unfolding procedure to measure the final spectra.

COMPARISON OF THE BACKGROUND ESTIMATORS

To assess the performance of the new ML-based background estimator, it is applied to data created by the PYTHIA + thermal toy model introduced above and compared to the established area-based method. A direct measure for the resolution and the overall bias of an estimator is the distribution of residuals $p_{T, \text{ch jet}}^{\text{rec}} - p_{T, \text{ch jet}}^{\text{true}}$. The residual distributions are normalized to unity to present a probability density and shown with Gaussian fits. The mean of these fits can be interpreted as the overall bias, where the width σ indicates the resolution of the method.

In the following, some results will only be presented for the neural network estimator. However, in these cases, the other estimators exhibit qualitatively the same results.




FIG. 2. Residual distributions for several background estimators and $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c.

In Fig. 2, the residual distributions of the estimators considered are shown for jets with $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c. It is immediately seen that all ML-based estimators considered exhibit a superior performance compared to the area-based correction method. While, to first order, the ML-based estimators all perform equally well on the toy model, the area-based correction has an inferior resolution, as summarized in Tab. I. The width of the residuals is a factor two smaller for the new background method. These smaller fluctuations also lead to much less purely combinatorial jets at low jet transverse momentum and, therefore, allow to measure jets down approximately factor two lower p_T .

It can also be shown that the new background estimator works similarly for different $p_{T, \text{ch jet}}^{\text{true}}$. The $p_{T, \text{ch jet}}^{\text{true}}$ -dependence of the estimator is depicted in Fig. 3. The estimator has roughly the same resolution and no sizable bias.

An important property of the estimator is how robust it is to modifications in the underlying background or with respect to jet fragmentation or jet type. The estimators performance should not strongly depend on these ingredients of the toy model. In reality, the background is definitely more complex than in the present model and the jets will be different in real data. First, the jets in the toy model are generated by PYTHIA and the PYTHIA-simulated jet fragmentation will be slightly different than in real data. Second, the fraction of quark and gluon jets is different in heavy-ion collisions than in PYTHIA, so a dependence on the jet type could bias the results. And finally, for the same reason, the background estimator should not yield qualitatively different results for quenched and unquenched jets, which may differ in shape and fragmentation.

To estimate the sensitivity to differences in the underlying thermal background, the correction is applied in events with modified thermal background multiplicities. The multiplicities are adjusted by roughly 10%, from the baseline of $N = 1800$ to $N = 1600$ and $N = 2000$. Hydrodynamic particle flow is not part of the thermal toy model but will occur on data. Its influence on the model performance is checked by applying the model to a particle distribution with $N = 1800$ which is modulated with $v_2 = 0.1$ in azimuth ϕ using

$$dN_{\text{flow}}/d\phi = dN/d\phi \cdot (1 + 2 \cdot v_2 \cos(2\phi)). \quad (4)$$




FIG. 3. Residual distributions for the neural network estimator and several (true) jet transverse momenta.




FIG. 4. Comparison of residual distributions for several thermal background settings for $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c.




FIG. 5. Comparison of the widths of the residual distributions as a function of the jet resolution parameter for $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c.

This corresponds to a relative particle modulation of up to 20% in the azimuthal yield distribution. It turns out that these changes have only a negligible effect on the model performance, as seen in Fig. 4. Here, the impact of different thermal backgrounds is presented for $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c, but also for lower jet transverse momenta it is very small.

Since jets are extended objects consisting of correlated particles up to large distances far beyond $R = 0.4$, using a large jet resolution parameter is in general desirable. While the presented study was performed for jets with $R = 0.4$, the performance of larger jet resolution parameter was also tested. Figure 5 presents the dependence of the background estima-

Estimator	Inclusive jets		Quark jets		Gluon jets	
Fit param. (GeV/c)	Mean	σ	Mean	σ	Mean	σ
Area-based method	0.7	7.8	0.8	7.3	0.9	7.7
Neural network	0.2	3.5	1.0	3.1	-0.5	3.6
Random Forest	-0.1	4.0	0.7	3.6	-0.8	3.9
Linear regression	-0.5	4.2	0.8	3.8	-1.4	4.1

TABLE I. Properties of the residual distributions for the considered estimators for $40 \leq p_{T, \text{ch jet}}^{\text{true}} < 60$ GeV/c. Mean and σ represent mean and standard deviation of the Gaussian fits, respectively.

Feature	Score	Feature	Score
Jet p_T (no corr.)	0.1355	$p_{T, \text{const}}^1$	0.0012
Jet mass	0.0007	$p_{T, \text{const}}^2$	0.0039
Jet area	0.0005	$p_{T, \text{const}}^3$	0.0015
Jet p_T (area-based corr.)	0.7876	$p_{T, \text{const}}^4$	0.0011
LeSub	0.0004	$p_{T, \text{const}}^5$	0.0009
Radial moment	0.0005	$p_{T, \text{const}}^6$	0.0009
Momentum dispersion	0.0007	$p_{T, \text{const}}^7$	0.0008
Number of constituents	0.0008	$p_{T, \text{const}}^8$	0.0007
Mean of const. p_T	0.0585	$p_{T, \text{const}}^9$	0.0006
Median of const. p_T	0.0023	$p_{T, \text{const}}^{10}$	0.0007

TABLE II. Random forest feature importances. A higher score corresponds to a higher importance of the feature. $p_{T, \text{const}}^i$ is the transverse momentum of the i^{th} -hardest particles. The four most important features are marked in bold face.

tor on the jet resolution parameter as the R -dependence of the width of the residual distribution, obtained by a Gaussian fit. As expected, the width gets larger, i.e. the resolution worsens, for larger R . However for the ML-based approach, both, the increase with increasing R , as well as the absolute value of the width, are significantly smaller than for the area-based correction.

To analyze the impact of different jet fragmentation or jet types on the background estimator, the residual distributions are calculated for different jet types. As a quite extreme test, quark and gluon jets serve here as proxies for jets with different fragmentation. The residual distributions are presented in Figs. 6 and 7. A small bias can be observed, which is of the order of the inclusive jet bias, see also Tab. I.

Another check on the fragmentation dependence of the background estimator was performed with jets generated by JEWEL [22] with the *no recoils*-setting. JEWEL is a Monte Carlo generator that simulates the jet evolution within a medium created in heavy-ion collisions and was used to produce samples of jets affected or not affected by medium effects. The residual distribution for the neural network estimator can be found in Fig. 8. Like for the quark/gluon jet comparison, a small bias can also be observed for JEWEL jets in the medium. The bias is of similar magnitude as the quark/gluon bias, though slightly smaller.

It should be emphasized that even a small bias can have a sizeable effect on the jet production yields due to the steeply falling jet spectrum. It might lead to deviations of up to approximately 20% at low transverse momentum around 20 GeV/c for the worst-case of pure quark or gluon samples.

For more relativistic jet collections, the bias will be significantly lower. In an analysis on real data, a systematic uncertainty should be evaluated taking this effect into account. Note though, that this bias will at least partially be corrected in an unfolding procedure.

The random forest estimator in *scikit-learn* directly calculates the mean decrease in impurity in the tree ensemble, also called Gini importance [24], for each input parameter/feature. In the Gini definition, a feature is more important when it is used in more tree splittings of more samples. Feature importances are simple measures (potentially) indicating how the random forest estimator learns from the data. Since all estimators perform approximately equally well, the importances should also give a rough estimate for the features' importances in the other estimators. The values can be found in Tab. II. The most important features are all related to variables that characterize the “hardness” of the jet. The most important feature turns out to be the area-based corrected jet momentum, indicating that the area-based correction already gives a reasonable estimate for the true jet momentum. Interestingly, it is not the hardest particle within the jet which has the highest importance but the second-hardest.

The performance evaluation on the toy model data is helpful in gaining insights on how precisely the true momentum of a jet in a heavy-ion-like background is reconstructed by the correction method. Unfortunately, the corrected jet spectra in the toy model are not realistic heavy-ion-like jet spectra, since they are made from one PYTHIA spectrum, embedded in thermal background. This changes drastically the amount of expected purely combinatorial jets as well as the jet spectra itself. Therefore, the toy model jet spectra are not presented here.

Instead, an independent study is performed on HIJING simulated data, shown in Fig. 10. Central HIJING events are simulated and the new correction method (without retraining) is applied as if the HIJING data were real data. The corrected jet spectra are then compared to N_{coll} -scaled PYTHIA spectra to judge the performance of the background estimators. The spectra comparison reveals that the new background estimators are also superior for realistic heavy-ion-like events down to very low transverse momenta of about 20 GeV/c. The spectra are closer to the expected truth and, therefore, an unfolding procedure of residual fluctuations will be possible with much higher precision and smaller uncertainties than for the standard method. As a further cross check, several HIJING spectra were analyzed for different minimum $p_{T, \text{hard}}$ -cuts and different impact parameter ranges, leading to the same conclusion.

In order to test how the measurement of the jet mass is affected by the different jet energy scales introduced by the background estimators, the jet mass has been measured for the toy model data for area-based and neural network background estimators. Figure 9 shows the jet mass for jets with $40 \leq p_{T, \text{ch jet}} < 60$ GeV/c, where the definition of $p_{T, \text{ch jet}}^{\text{rec}}$ – and therefore the jet energy scale – differs for the shown distributions. The jet mass itself is also background-corrected [23].




FIG. 6. Comparison of residual distributions for gluon, quark, and inclusive jets for the neural network estimator.




FIG. 7. Comparison of residual distributions for gluon, quark, and inclusive jets for the area-based method.




FIG. 8. Comparison of residual distributions for jets created by JEWEL for the neural network estimator.

It turns out that the reconstructed jet masses coincide well for the true and ML-based jet energy scale. Therefore, the higher precision of the jet energy scale might allow a jet mass measurement at lower jet momenta than currently possible. The jet mass for the area-based method energy scale seems to be significantly overestimated. Further tests show that these general observations are also valid for higher transverse jet momenta, although the bias of the area-based method decreases.




FIG. 9. Jet mass distribution compared for different jet energy scales. A background correction for the jet mass itself was also applied [23].




FIG. 10. Reconstructed charged jet spectra in HIJING events and the ratio to (N_{coll} -scaled) PYTHIA jet spectra.

SUMMARY AND OUTLOOK

A novel method to correct for the large and strongly-fluctuating background in central heavy-ion collisions using Machine Learning techniques is introduced. In contrast to the established correction method, the background is estimated on a jet-by-jet basis rather than event-by-event. A toy-model analysis indicates that the new background estimator is superior to the established method in terms of precision. Studies also show that the residual fluctuations do not depend on the jet transverse momentum above $p_{T, \text{ch jet}}^{\text{true}} \geq 20$ GeV/c. The method was evaluated on several different underlying background models to investigate its susceptibility to differences compared to real data, revealing that the performance does not strongly depend on subtleties of the background. In addition, quark and gluon jets as well as JEWEL-generated jets are used as proxies for a different jet fragmentation, with a small bias observed. An analysis on real data should take into account a possible systematic uncertainty that could arise from this bias. It has also been shown that the jet mass reconstructed with the ML-based jet energy scale coincides well with the

mass in the true jet energy scale. Hence, the application of the new method on real experimental data may allow considerably more precise measurements in heavy-ion collisions, in particular at low transverse momentum.

ACKNOWLEDGMENTS

We would like to thank the ALICE Collaboration for providing computational resources for this study. We are grateful to L. Cunqueiro and J. Harris for useful discussions and comments to the manuscript. R. Haake is supported in part by a Feodor Lynen Research Fellowship from the Alexander von Humboldt Foundation and by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Grant number DE-SC004168. C. Loizides is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under contract number DE-AC05-00OR22725.

- [1] **HotQCD** Collaboration, T. Bhattacharya, M. I. Buchoff, N. H. Christ, *et al.*, “QCD Phase Transition with Chiral Quarks and Physical Quark Masses,” *Phys. Rev. Lett.* **113** (Aug, 2014) 082001, arXiv:1402.5175 [hep-lat]. <http://link.aps.org/doi/10.1103/PhysRevLett.113.082001>.
- [2] M. Gyulassy and M. Plumer, “Jet Quenching in Dense Matter,” *Phys.Lett.* **B243** (1990) 432–438.
- [3] **ALICE** Collaboration, K. Aamodt *et al.*, “The ALICE experiment at the CERN LHC,” *JINST* **3** (2008) S08002.
- [4] **ALICE** Collaboration, B. Abelev *et al.*, “Measurement of event background fluctuations for charged particle jet reconstruction in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV,” *JHEP* **1203** (2012) 053, arXiv:1201.2423 [hep-ex].
- [5] **ALICE** Collaboration, B. Abelev *et al.*, “Measurement of charged jet suppression in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV,” *JHEP* **1403** (2014) 013, arXiv:1311.0633 [nucl-ex].
- [6] K. Albertsson *et al.*, “Machine Learning in High Energy Physics Community White Paper,” arXiv:1807.02876 [physics.comp-ph].
- [7] X.-N. Wang, and M. Gyulassy, “HIJING: A Monte Carlo model for multiple jet production in pp, pA and AA collisions,” *Phys. Rev. D* **44** (1991) 3501.
- [8] M. Cacciari, G. P. Salam, and G. Soyez, “The Anti-k(t) jet clustering algorithm,” *JHEP* **0804** (2008) 063, arXiv:0802.1189 [hep-ph].
- [9] M. Cacciari, G. P. Salam, and G. Soyez, “FastJet User Manual,” *Eur.Phys.J.* **C72** (2012) 1896, arXiv:1111.6097 [hep-ph].
- [10] T. Sjostrand, S. Mrenna, and P. Z. Skands, “PYTHIA 6.4 Physics and Manual,” *JHEP* **0605** (2006) 026, arXiv:hep-ph/0603175 [hep-ph].

- [11] P. Z. Skands, “Tuning Monte Carlo generators: The Perugia tunes,” *Phys.Rev.* **D82** (2010) 074018, arXiv:1005.3457 [hep-ph].
- [12] ALICE Collaboration, K. Aamodt *et al.*, “Charged-particle multiplicity density at mid-rapidity in central Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” *Phys.Rev.Lett.* **105** (2010) 252301, arXiv:1011.3916 [nucl-ex].
- [13] ALICE Collaboration, K. Aamodt *et al.*, “Centrality dependence of the charged-particle multiplicity density at mid-rapidity in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” *Phys.Rev.Lett.* **106** (2011) 032301, arXiv:1012.1657 [nucl-ex].
- [14] ALICE Collaboration, K. Aamodt *et al.*, “Suppression of charged particle production at large transverse momentum in central Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” *Phys.Lett.* **B696** (2011) 30–39, arXiv:1012.1004 [nucl-ex].
- [15] STAR Collaboration, L. Adamczyk *et al.*, “Measurements of jet quenching with semi-inclusive hadron+jet distributions in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV,” *Phys. Rev.* **C96** no. 2, (2017) 024905, arXiv:1702.01108 [nucl-ex].
- [16] M. Cacciari and G. P. Salam, “Pileup subtraction using jet areas,” *Phys. Lett.* **B659** (2008) 119–126, arXiv:0707.1378 [hep-ph].
- [17] S. Haykin, *Neural networks: A comprehensive foundation*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2nd ed., 1998.
- [18] L. Breiman, “Random forests,” *Machine Learning* **45** no. 1, (Oct, 2001) 5–32.
- [19] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR* (2014), arXiv:1412.6980.
- [20] V. Nair and G. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML’10, pp. 807–814. Omnipress, USA, 2010.
- [21] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research* **12** (2011) 2825–2830.
- [22] K. C. Zapp, “JEWEL 2.0.0: directions for use,” *Eur.Phys.J.* **C74** no. 2, (2014) 2762, arXiv:1311.0048 [hep-ph].
- [23] G. Soyez, G. P. Salam, J. Kim, S. Dutta, and M. Cacciari, “Pileup subtraction for jet shapes,” *Phys. Rev. Lett.* **110** no. 16, (2013) 162001, arXiv:1211.2811 [hep-ph].
- [24] G. Louppe, L. Wehenkel, A. Suter, and P. Geurts, “Understanding variable importances in forests of randomized trees,” in *Advances in Neural Information Processing Systems* 26, pp. 431–439. Curran Associates, Inc., 2013.