





Interpretable machine learning methods applied to jet background subtraction in heavy-ion collisions

Tanner Mengel , Patrick Steffanic , Charles Hughes , Antonio Carlos Oliveira da Silva , and Christine Nattrass 
University of Tennessee, Knoxville, Tennessee 37996, USA

 (Received 14 March 2023; accepted 9 August 2023; published 22 August 2023)

Jet measurements in heavy ion collisions can provide constraints on the properties of the quark gluon plasma, but the kinematic reach is limited by a large, fluctuating background. We present a novel application of symbolic regression to extract a functional representation of a deep neural network trained to subtract background from jets in heavy ion collisions. We show that the deep neural network is approximately the same as a method using the particle multiplicity in a jet. This demonstrates that interpretable machine learning methods can provide insight into underlying physical processes.

DOI: [10.1103/PhysRevC.108.L021901](https://doi.org/10.1103/PhysRevC.108.L021901)

Introduction. The quark gluon plasma (QGP) is a hot, dense, strongly interacting liquid of quarks and gluons that is created briefly in high energy heavy ion collisions [1–4]. Measurements of jets produced by hard scatterings between partons in heavy ion collisions can be used to investigate the properties of the QGP [5]. Quantitative comparisons between jet measurements and physics models can provide further constraints on these properties [6,7]. However, heavy ion events are dominated by a fluctuating background of soft particles not due to hard scatterings. The details of these fluctuations are sensitive to correlations from hydrodynamical flow and the shape of the single-particle spectra [8], and as such are unlikely to be exactly the same in data and models. Mixed events are able to successfully describe the background in measurements of hadron-jet correlations by the STAR Collaboration [9] at the Relativistic Heavy Ion Collider (RHIC). Studies of the background at the Large Hadron Collider (LHC) by the ALICE Collaboration found that the distribution of background energy density in random cones is well described by a random background with correlations due to hydrodynamical flow and Poissonian fluctuations [10]. A better understanding of this background will facilitate more precise jet measurements for comparisons between data and models.

Measurement precision and kinematic range is limited by the ability to correct for this background and its fluctuations. Background correction in jet measurements requires subtraction of contributions from soft particles within the jet, and suppression of fluctuations which have been reconstructed as combinatorial jets. At low momenta, combinatorial jets limit the kinematic reach of the measurement. Improved background subtraction methods would increase measurements'

sensitivity to partonic energy loss. Measurements of jet spectra which extend to low momenta primarily use the area method [11] for background subtraction. This method was initially proposed to correct for the underlying event in $p+p$ collisions in high pile-up conditions [11] and has also been applied to heavy ion collisions [12–15].

The complexity of jet background subtraction makes it an interesting environment to apply machine learning techniques. However, application of machine learning methods to background subtraction should be handled with care since models are not able to fully reproduce background fluctuations in heavy ion collisions [8]. Nuclear physics has prioritized the continued advancement in machine learning analysis techniques with a focus on interpretable methods that are robust, provide clear uncertainty quantification, and are explainable [16]. Applications of noninterpretable machine learning methods are insufficient when models available for training may be inaccurate, when it may be necessary to understand the method to interpret the results, or when a result is needed outside of the training space.

Application of a deep neural network, i.e., a neural network with multiple hidden layers, to jet background subtraction in heavy ion collisions has demonstrated significant improvements compared to the area method, particularly at low jet momenta [17,18]. Deep neural networks are susceptible to model bias because their predictions risk being unreliable outside the domain of their training space. These methods may break down when they are extrapolated beyond this space, and due to their opaque nature, offer little indication where and why this break down occurs. In addition, one cannot validate the technique against data because we do not know the true jet momenta in data.

Increased performance of machine learning methods over traditional methods is an indication that there is information accessible to the machine learning that accounts for this improvement. We present an interpretable machine learning technique that allows us to understand why a deep neural network improves the jet momentum resolution in heavy ion

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

collisions. We empirically derive an alternate method based on the background described in [10,19], we call the multiplicity method. We compare the widths of the fluctuations of the jet momenta for this method to the area and neural network methods and estimate the impact of the methods on the kinematic range. We apply symbolic regression to determine a functional form describing the mapping learned by the neural network, which was trained using TENNGEN [8] for the background and PYTHIA [20] for the signal. We compare this functional description of the neural network to the form of the multiplicity method.

Simulation. TENNGEN [21,22] generates heavy ion collisions with π^\pm , K^\pm , p , and \bar{p} hadrons with yields [23], momentum distributions [24,25], and azimuthal anisotropies [26,27] matched to published data. TENNGEN was updated to simulate collision energies per nucleon of $\sqrt{s_{NN}} = 200$ GeV collisions as well as $\sqrt{s_{NN}} = 2.76$ TeV, including multiplicity fluctuations, and improved computational efficiency. Proton-proton collisions at $\sqrt{s} = 200$ GeV were simulated with the PYTHIA 8.307 [20] Monash 2013 tune [28] in 25 p_T^{hard} bins starting at 5 GeV, with 10^6 $p+p$ events in each bin. Only final state charged particles from PYTHIA are mixed with a TENNGEN background event. Charged particles from both PYTHIA and TENNGEN are required to have a minimum p_T of 150 MeV and be within pseudorapidity $|\eta| < 0.9$.

Jets are clustered using the anti- k_T algorithm with FASTJET 3.4.0 [29] with jet resolution parameters $R = 0.2, 0.4,$ and 0.6 . To determine the true momentum, jets are reconstructed separately in both PYTHIA and the combined event. Jets in the combined PYTHIA and TENNGEN event are geometrically matched to a PYTHIA jet if $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.1$, where $\Delta\eta$ and $\Delta\phi$ are the differences in η and ϕ between the jets and there is a bijective match. Reconstructed jets are required to have $p_T > 5$ GeV and be within pseudorapidity $|\eta_{jet}| < 0.9 - R$. The momentum of the PYTHIA jet is taken as the truth momentum, $p_{T,\text{Jet}}^{\text{Truth}} \equiv p_{T,\text{Jet}}^{\text{PYTHIA}}$.

Area and multiplicity methods. For area-based background subtraction, the jet area [30] is estimated through the use of “ghost” particles, jets are reconstructed using the k_T jet finder [31]. The corrected jet momentum is then estimated as

$$p_{T,\text{Jet}}^{\text{Corr},A} = p_{T,\text{Jet}}^{\text{tot}} - \rho A, \quad (1)$$

where A is the jet area, ρ is the background momentum density per unit area, and $p_{T,\text{Jet}}^{\text{tot}}$ is the total momentum in the jet. The ρ in an event is approximated as the median $p_{T,\text{Jet}}^{\text{tot}}/A$ for k_T jets because k_T jets are dominated by background.

To a good approximation, the standard deviation of the momentum residual $\delta p_T = p_T^{\text{Corr}} - p_T^{\text{Truth}}$ with the area method is given by

$$\sigma_{\delta p_T} = \sqrt{N\sigma_{p_T}^2 + \left(N + 2N^2 \sum_{n=1}^{\infty} v_n^2\right) \langle p_T \rangle^2}, \quad (2)$$

where N is the number of background particles in the jet, σ_{p_T} is the standard deviation of the single track momentum distribution, v_n are the coefficients of the azimuthal anisotropies of the single particle distributions, and $\langle p_T \rangle$ is the average momentum of background particles [10]. This is derived by

assuming each of the N particles is drawn from a single track momentum distribution, which is approximately a gamma distribution, giving rise to the first term [19]. The second term is from Poissonian fluctuations in the number of background particles and the third term is from fluctuations in the number of particles due to hydrodynamical flow. Deviations of the single track momentum distribution from a gamma distribution and momentum dependence of the v_n lead to slightly larger widths [8].

The area method is usually used instead of iterative background subtraction methods [32–34] for measurements of jets at lower momenta. Iterative methods may suppress the fluctuations described in Eq. (2) by estimating the local background and suppress combinatorial jets by requiring high momentum or energy constituents. At low momenta, these requirements may impose a bias on the surviving jets. Fluctuations and the contribution from combinatorial jets are generally higher with the area method, but with less bias.

We propose a multiplicity-based method as an alternative to the area method:

$$p_{T,\text{Jet}}^{\text{Corr},N} = p_{T,\text{Jet}}^{\text{tot}} - \rho_{\text{Mult}}(N_{\text{tot}} - N_{\text{signal}}), \quad (3)$$

where N_{tot} is the total number of particles in the jet, N_{signal} is the number of particles in the signal, and $N = N_{\text{tot}} - N_{\text{signal}}$. This leverages the fact that the natural variable in the background fluctuations is the number of particles, largely eliminating the second and third terms in Eq. (2). ρ_{Mult} in an event is the mean transverse momentum per background particle, which is approximated as the median $p_{T,\text{Jet}}^{\text{tot}}/N_{\text{tot}}$ for k_T jets. N_{signal} is roughly described by models [35] and therefore can be estimated. Measurements of γ - h correlations [36] and reconstructed jets [37,38] indicate that there are around 0.5 additional particles for $p_T^{\text{jet}} \approx 10$ GeV/ c and 1.0 additional particles for $p_T^{\text{jet}} \approx 100$ GeV/ c in heavy ion collisions. If this were applied as an additional uncertainty, it would be proportional to $\sigma_{N_{\text{signal}}} \times \rho_{\text{Mult}}$, or around 0.25–0.5 GeV/ c .

Machine learning methods. A sufficiently complex neural network can interpolate any function, at the cost of transparency to the user. This poses an obstacle to application of deep neural networks in physics where understanding predictions and identifying their potential biases is crucial. Our approach to addressing this challenge is through symbolic regression, one example of interpretable machine learning, to extract mathematical expressions from trained deep neural networks. The resulting equations provide an effective description of the neural network’s mapping between the input and output. By constraining the types of operations available, we can impose complexity and smoothness requirements.

We train a deep neural network to predict the corrected jet momentum from the following input features: the uncorrected jet momentum, jet area, jet angularity, number of jet constituents, and seven leading constituent momenta. The architecture and input features of the network are motivated by previous application of neural networks to proton-proton jets with a thermal background [17]. The deep neural network is implemented with TENSORFLOW 2.10.0 [39]. The deep neural network has three hidden layers consisting of 100, 100, and 50 nodes, each activated by a rectified linear unit (ReLU) [40]

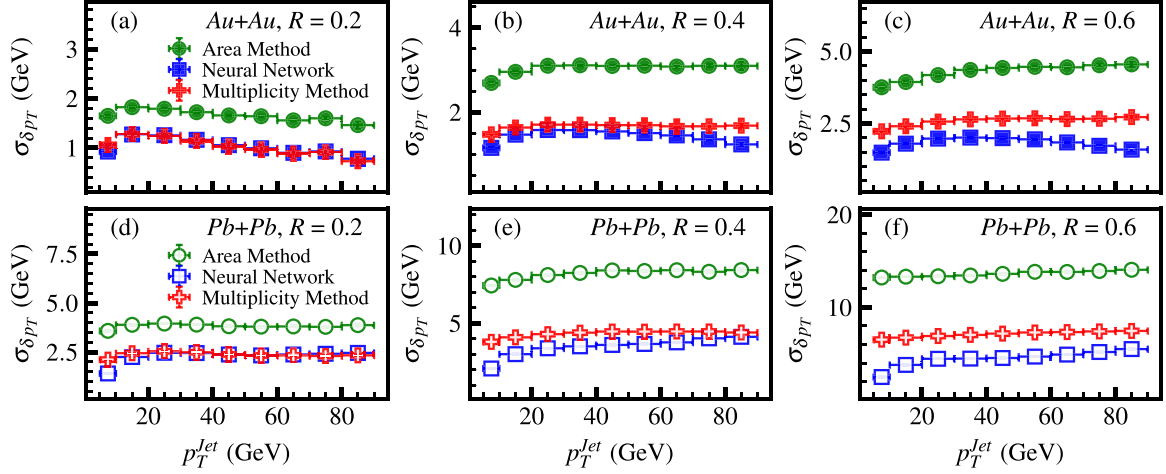


FIG. 1. Comparisons of jet p_T residual width for each background subtraction method as a function of reconstructed jet momentum for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for jet resolution parameters $R = 0.2, 0.4,$ and 0.6 .

function. The model is optimized with ADAM [40] and the loss function is a modified mean squared error,

$$\mathcal{L} = \langle \left\| p_{T,\text{Jet}}^{\text{Truth}} - p_{T,\text{Jet}}^{\text{DNN}} \right\|^2 \rangle + \lambda \sum_{l=1}^L \|\mathbf{W}_l\|^2, \quad (4)$$

where $p_{T,\text{Jet}}^{\text{DNN}}$ is the predicted jet momentum, $p_{T,\text{Jet}}^{\text{Truth}}$ is the truth momentum, and the last term is an $L^2(\lambda)$ regularization where $\lambda = 0.001$, \mathbf{W}_l is the weight matrix of layer l , and the sum is over the L layers. The regularization term penalizes redundancy and encourages sparsity in the final trained network. The network is trained using 50% of the simulated jets while the remaining 50% are reserved for testing.

Once the neural network is trained, it represents an approximate mapping between the input jet features and the truth jet momentum. We apply a genetic algorithm to symbolically regress a functional form which describes this mapping using the PYSR 0.11.11 [41] package. The PYSR model samples the phase space of analytic expressions defined by operators, input features, and constants for minimization through genetic programming. The input features are comparable to those of the neural network, and the pool of operations are arithmetic, exponential, trigonometric, and exponentiation. The model mutates over 50 generations of 20 different population samples, with each population containing 33 individuals. The loss function for the PYSR model,

$$\mathcal{L} = \langle \left\| p_{T,\text{Jet}}^{\text{DNN}} - p_{T,\text{Jet}}^{\text{PYSR}} \right\|^2 \rangle, \quad (5)$$

is the mean squared error between the prediction from PySR $p_{T,\text{Jet}}^{\text{PYSR}}$ and the corrected jet momentum predicted by the neural network. PYSR evaluates expressions based on a score S that rewards minimizing the loss function \mathcal{L} and penalizes equation complexity C :

$$S = -\frac{\delta \ln \mathcal{L}}{\delta C}, \quad (6)$$

where the equation complexity C is defined as the total number of operations, variables, and constants used in an equation [42]. The simulated jets, designated for testing,

are used to sufficiently sample the neural network outputs throughout the possible input feature space. The highest scoring PYSR expression is a functional representation of the mapping from input jet features to corrected jet momentum learned by the deep neural network.

Unfolding. The lower threshold for unfolding is typically set to be between 2–5 times the width of the jet momentum resolution to suppress effects of combinatorial jets on the unfolded results [12,18]. We unfold the reconstructed jet momentum spectra using five iterations of the Bayesian unfolding method [43] in ROOUNFOLD 2.0.0 [44]. We construct a response matrix using PYTHIA jets (truth jets) matched to PYTHIA + TENNGEN jets (reconstructed jets). The momentum of the PYTHIA jet is taken as the truth momentum, $p_{T,\text{jet}}^{\text{Truth}} \equiv p_{T,\text{jet}}^{\text{PYTHIA}}$. We then unfold our reconstructed jet spectra. The reconstructed spectra has no matching criteria between the PYTHIA + TENNGEN jets and PYTHIA jets and no kinematic cuts to suppress combinatorial jet contributions. We use reconstructed jet spectra including combinatorial background to investigate the sensitivity of the lower momentum threshold to combinatorial background.

Results. Figure 1 shows the width of the jet momentum residual distributions as a function of jet momentum for each background subtraction method in both Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The $\sigma_{\delta p_T}$ increases with increasing jet resolution parameter, as expected because there is more background when the jet is larger. The $\sigma_{\delta p_T}$ also increases with $\sqrt{s_{NN}}$ because the particle multiplicity increases. As seen in [17], the deep neural network reconstructs the momentum considerably more accurately than the area method. The performance of the multiplicity method is comparable to that of the deep neural network in Au+Au collisions and small jet resolution parameter. The ability of each method to sufficiently suppress contributions from combinatorial jets at low p_T is demonstrated with the ratios of the reconstructed jet spectra to the true jet spectra, shown in Fig. 2. The contributions from combinatorial jets decreases with increasing jet momentum for all methods, with all jet resolution parameters, and for

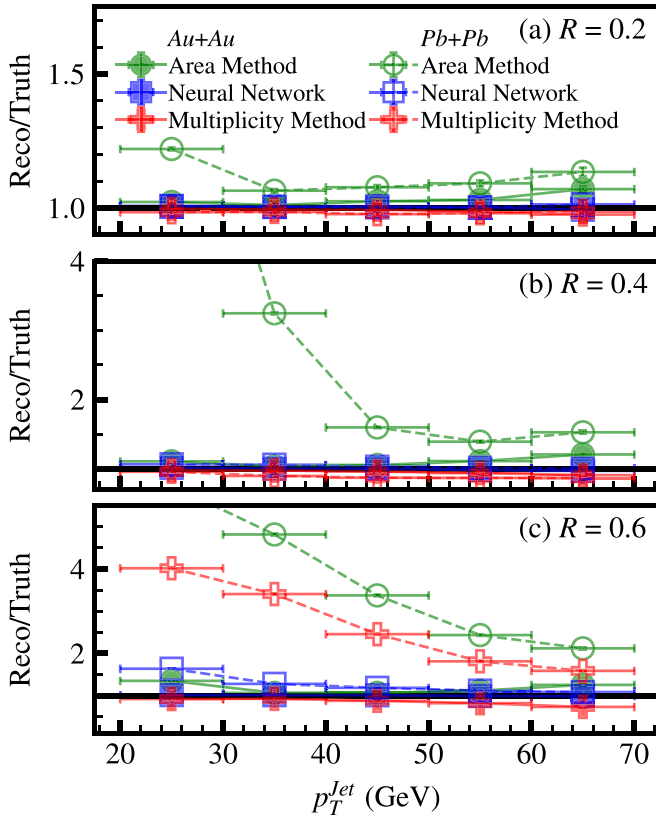


FIG. 2. Ratio of the reconstructed jet spectrum over the truth spectrum for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for jet resolution parameters $R = 0.2$, 0.4 , and 0.6 . Low momentum points for the area method for LHC energies at $R = 0.4$ and $R = 0.6$ are off scale.

both collision energies. The ratios for the deep neural network and multiplicity methods are both lower than those of the area method.

The ratios of the unfolded spectra to the true jet spectra are shown in Fig. 3. Fluctuations from one at lower jet momenta are where the method becomes unstable due to overwhelming contributions from combinatorial background. Reconstructed jet spectra have no kinematic cuts to suppress combinatorial jet contributions, therefore any extension in the lower kinematic range is due to the momentum resolution of the background subtraction method. The multiplicity and deep neural network methods are stable to at least 10 GeV/c lower in momentum than the area method for all jet resolution parameters and collisions systems.

For all jet resolution parameters and both collision energies, the symbolic regression found that the best description of the deep neural network has the functional form

$$p_{T,\text{jet}}^{\text{Corr,PYSR}} = p_{T,\text{jet}}^{\text{tot}} - C_1(N_{\text{tot}} - C_2), \quad (7)$$

where the two parameters, C_1 and C_2 , are optimization constants defined by PYSR. These parameters are plotted in Fig. 4 and compared to the average value of the parameters used in the multiplicity method. We find that the symbolic regression parameters C_1 and C_2 are comparable to the averages of those for the multiplicity method, $\langle \rho_{\text{Mult}} \rangle$ and $\langle N_{\text{signal}} \rangle$, respectively,




FIG. 3. Ratio of unfolded jet spectrum over truth spectrum for (a) Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and (b) Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for jet resolution parameter $R = 0.4$.

with greater deviations at LHC energies and larger R . This indicates that the deep neural network is using a relationship similar to the multiplicity method to predict jet momenta. This approach to machine learning enables use of domain knowledge. The optimization parameters from PYSR would otherwise not have a clear physical interpretation. Since these parameters are understood in the multiplicity method, it is possible to assign a physically motivated uncertainty to them. Assumptions inherent in the method can then be understood.

Conclusions. We have shown that interpretable machine learning methods can learn an underlying physical correlation, such as the multiplicity dependence for jet background, that was previously overlooked. The dependence of the neural network on jet multiplicity, rather than other input features, is easily understood since the fluctuations in the background are well described by Eq. (2) [8,10,19], where the multiplicity is the dominant variable in the standard deviation. Using the multiplicity instead of the the area removes the second two terms in the standard deviation because these arise purely from fluctuations in the number of particles. The systematic uncertainty on N_{signal} in the multiplicity method can be constrained by existing measurements and is therefore data driven. The multiplicity method achieves similar performance without the model dependence of the neural network.

We previously showed that, when we used a random forest to classify jets as either combinatorial or signal, the optimal selection was on the leading hadron momentum [45], already

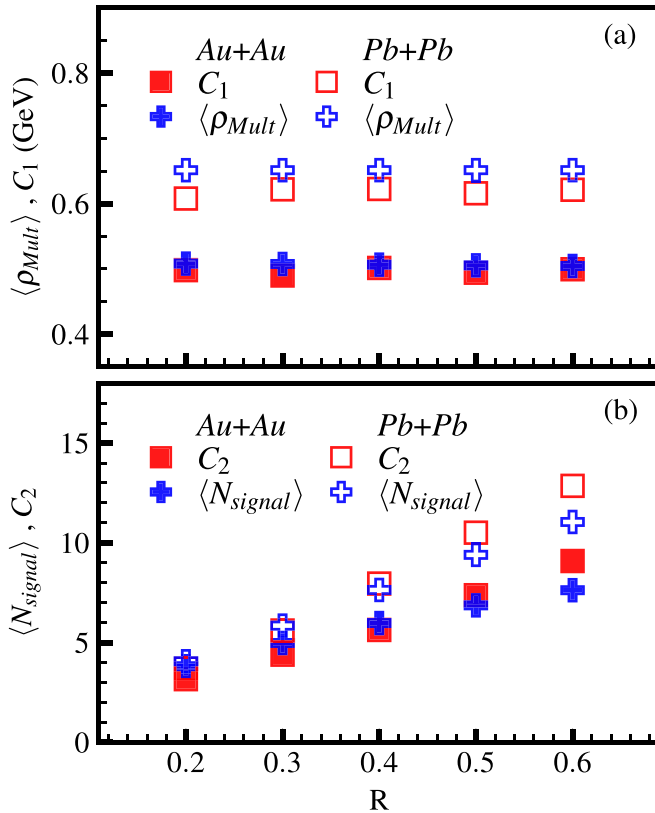


FIG. 4. PySR optimization constants compared to average value of multiplicity method parameters versus jet resolution parameter for Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for jet resolution parameters $R = 0.2, 0.4, \text{ and } 0.6$.

used as a standard technique [12–15]. We argue that applying machine learning to scientific problems requires methods that are interpretable. The definition of interpretability is often ambiguous or underspecified, but [46] presents several definitions of interpretability to guide our selection of machine learning methods. We argue that for a machine learning method to be interpretable (1) it should be applicable equivalently to data and simulation, (2) the method’s output can be understood outside the range of the training set, and (3) a measurement uncertainty can be calculated. We argue that an uncertainty on the method is not a proxy for a measurement uncertainty. These stricter criteria are consistent with those outlined in [16]. Symbolic regression satisfies these requirements because the output is a formula. The convergence of the empirically based multiplicity method and the formula produced through symbolic regression is a clear indication of the usefulness of an interpretable method. Machine learning should be used to gain knowledge about the underlying physical processes that drive the relationships in our data. We must interpret the details of any method in terms of these underlying physical processes.

We are grateful to F. Bock, H. Bossi, A. D. Maestro, J. Nagle, K. Read, and A. Schmier for useful discussions and feedback on the manuscript. This work was supported in part by funding from the Division of Nuclear Physics of the U.S. Department of Energy under Grant No. DE-FG02-96ER40982. This work was performed on the computational resources at the Infrastructure for Scientific Applications and Advanced Computing (ISAAC) supported by the University of Tennessee.

- [1] K. Adcox *et al.* (PHENIX Collaboration), Formation of dense partonic matter in relativistic nucleus nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration, *Nucl. Phys. A* **757**, 184 (2005).
- [2] J. Adams, Jr. *et al.* (STAR Collaboration), Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR collaboration’s critical assessment of the evidence from RHIC collisions, *Nucl. Phys. A* **757**, 102 (2005).
- [3] B. B. Back *et al.*, The PHOBOS perspective on discoveries at RHIC, *Nucl. Phys. A* **757**, 28 (2005).
- [4] I. Arsene *et al.* (BRAHMS Collaboration), Quark gluon plasma an color glass condensate at RHIC? The perspective from the BRAHMS experiment, *Nucl. Phys. A* **757**, 1 (2005).
- [5] M. Connors, C. Nattrass, R. Reed, and S. Salur, Jet measurements in heavy ion physics, *Rev. Mod. Phys.* **90**, 025005 (2018).
- [6] K. M. Burke *et al.* (JET Collaboration), Extracting the jet transport coefficient from jet quenching in high-energy heavy-ion collisions, *Phys. Rev. C* **90**, 014909 (2014).
- [7] S. Cao *et al.* (JETSCAPE Collaboration), Determining the jet transport coefficient q from inclusive hadron suppression measurements using Bayesian parameter estimation, *Phys. Rev. C* **104**, 024905 (2021).
- [8] C. Hughes, A. C. Oliveira Da Silva, and C. Nattrass, Model studies of fluctuations in the background for jets in heavy ion collisions, *Phys. Rev. C* **106**, 044915 (2022).
- [9] L. Adamczyk *et al.* (STAR Collaboration), Jet-Hadron Correlations in $\sqrt{s_{NN}} = 200$ GeV $p + p$ and Central Au + Au Collisions, *Phys. Rev. Lett.* **112**, 122301 (2014).
- [10] B. Abelev *et al.* (ALICE Collaboration), Measurement of Event Background Fluctuations for Charged Particle Jet Reconstruction in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *J. High Energy Phys.* **03** (2012) 053.
- [11] M. Cacciari and G. P. Salam, Pileup subtraction using jet areas, *Phys. Lett. B* **659**, 119 (2008).
- [12] B. Abelev *et al.* (ALICE Collaboration), Measurement of charged jet suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *J. High Energy Phys.* **03** (2014) 013.
- [13] J. Adam *et al.* (ALICE Collaboration), Measurement of jet suppression in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *Phys. Lett. B* **746**, 1 (2015).
- [14] S. Acharya *et al.* (ALICE Collaboration), Measurements of inclusive jet spectra in pp and central Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, *Phys. Rev. C* **101**, 034911 (2020).
- [15] J. Adam *et al.* (STAR Collaboration), Measurement of inclusive charged-particle jet production in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV, *Phys. Rev. C* **102**, 054913 (2020).
- [16] P. Achenbach *et al.*, The present and future of QCD, arXiv:2303.02579.

- [17] R. Haake and C. Loizides, Machine Learning based jet momentum reconstruction in heavy-ion collisions, *Phys. Rev. C* **99**, 064904 (2019).
- [18] ALICE Collaboration, Measurement of the radius dependence of charged-particle jet suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, [arXiv:2303.00592](https://arxiv.org/abs/2303.00592).
- [19] M. Tannenbaum, The distribution function of the event-by-event average pt for statistically independent emission, *Phys. Lett. B* **498**, 29 (2001).
- [20] T. Sjöstrand, S. Mrenna, and P. Z. Skands, A Brief Introduction to PYTHIA 8.1, *Comput. Phys. Commun.* **178**, 852 (2008).
- [21] T. Mengel and C. Hughes, *TennGen*, Github (2023).
- [22] T. Mengel and C. Hughes, *tenngen200*, Github (2022), <https://github.com/tmengel/tenngen200>.
- [23] K. Aamodt *et al.* (ALICE Collaboration), 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**, 032301 (2010).
- [24] A. Adare *et al.* (PHENIX Collaboration), Spectra and ratios of identified particles in Au+Au and $d + Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV, *Phys. Rev. C* **88**, 024906 (2013).
- [25] B. Abelev *et al.* (ALICE Collaboration), Centrality dependence of π , K , p production in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *Phys. Rev. C* **88**, 044910 (2013).
- [26] A. Adare *et al.* (PHENIX Collaboration), Measurement of the higher-order anisotropic flow coefficients for identified hadrons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV, *Phys. Rev. C* **93**, 051902 (2016).
- [27] J. Adam *et al.* (ALICE Collaboration), Higher harmonic flow coefficients of identified hadrons in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *J. High Energy Phys.* **09** (2016) 164.
- [28] P. Skands, S. Carrazza, and J. Rojo, Tuning PYTHIA 8.1: the Monash 2013 Tune, *Eur. Phys. J. C* **74**, 3024 (2014).
- [29] M. Cacciari, G. P. Salam, and G. Soyez, FastJet User Manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [30] M. Cacciari, G. P. Salam, and G. Soyez, The catchment area of jets, *J. High Energy Phys.* **04** (2008) 005.
- [31] S. D. Ellis and D. E. Soper, Successive combination jet algorithm for hadron collisions, *Phys. Rev. D* **48**, 3160 (1993).
- [32] J. A. Hanks *et al.*, Method for separating jets and the underlying event in heavy ion collisions at the BNL Relativistic Heavy Ion Collider, *Phys. Rev. C* **86**, 024908 (2012).
- [33] V. Khachatryan *et al.* (CMS Collaboration), Measurement of inclusive jet cross sections in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *Phys. Rev. C* **96**, 015202 (2017).
- [34] G. Aad *et al.* (ATLAS Collaboration), Measurement of the jet radius and transverse momentum dependence of inclusive jet suppression in lead-lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector, *Phys. Lett. B* **719**, 220 (2013).
- [35] B. B. Abelev *et al.* (ALICE Collaboration), Charged jet cross sections and properties in proton-proton collisions at $\sqrt{s} = 7$ TeV, *Phys. Rev. D* **91**, 112012 (2015).
- [36] U. Acharya *et al.* (PHENIX Collaboration), Measurement of jet-medium interactions via direct photon-hadron correlations in Au+Au and $d + Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV, *Phys. Rev. C* **102**, 054910 (2020).
- [37] S. Chatrchyan *et al.* (CMS Collaboration), Measurement of jet fragmentation in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV, *Phys. Rev. C* **90**, 024908 (2014).
- [38] M. Aaboud *et al.* (ATLAS Collaboration), Measurement of jet fragmentation in Pb+Pb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector, *Phys. Rev. C* **98**, 024908 (2018).
- [39] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from <https://tensorflow.org>.
- [40] I. J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016), <http://www.deeplearningbook.org>.
- [41] M. Cranmer, PySR: Fast & parallelized symbolic regression in Python/Julia, *Zenodo* (2020).
- [42] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, and S. Ho, Discovering symbolic models from deep learning with inductive biases, [arXiv:2006.11287](https://arxiv.org/abs/2006.11287).
- [43] G. D'Agostini, A multidimensional unfolding method based on Bayes' theorem, *Nucl. Instrum. Methods Phys. Res., Sect. A* **362**, 487 (1995).
- [44] L. Brenner, R. Balasubramanian, C. Burgard, W. Verkerke, G. Cowan, P. Verschuur, and V. Croft, Comparison of unfolding methods using RooFitUnfold, *Int. J. Mod. Phys. A* **35**, 2050145 (2020).
- [45] P. Steffanic, C. Hughes, and C. Nattrass, Separating signal from combinatorial jets in a high background environment, [arXiv:2301.09148](https://arxiv.org/abs/2301.09148).
- [46] Z. C. Lipton, The myths of model interpretability: In machine learning, the concept of interpretability is both important and slippery, *Queue* **16**, 31 (2018).