## AI-based Monte Carlo event generator for electron-proton scattering

Y. Alanazi,<sup>1</sup> P. Ambrozewicz,<sup>2</sup> M. P. Kuchera,<sup>3</sup> Y. Li,<sup>1</sup> T. Liu,<sup>2</sup> R. E. McClellan,<sup>2</sup>

W. Melnitchouk,<sup>2</sup> E. Pritchard,<sup>3</sup> M. Robertson,<sup>4</sup> N. Sato,<sup>2</sup> R. Strauss,<sup>4</sup> and L. Velasco<sup>5</sup>

<sup>1</sup>Department of Computer Science, Old Dominion University, Norfolk, Virginia 23529, USA

<sup>2</sup> Jefferson Lab, Newport News, Virginia 23606, USA

<sup>3</sup>Department of Physics, Davidson College, Davidson, North Carolina 28035, USA

<sup>4</sup>Department of Mathematics and Computer Science,

<sup>5</sup>Department of Physics, University of Dallas, Irving, Texas 75062, USA

(Dated: August 3, 2020)

We present a novel new strategy using artificial intelligence (AI) to build the first AI-based Monte Carlo event generator (MCEG) capable of faithfully generating final state particle phase space in high-energy reactions. We show a blueprint for integrating machine learning strategy with calibrated detector simulations to build a vertex-level, AI-based MCEG, free of theoretical assumptions about femtometer scale physics. As the first steps towards this goal, we present a case study for inclusive electron-proton scattering using synthetic data from the PYTHIA MCEG for testing and validation purposes. Our quantitative results validate our proof of concept and demonstrate the predictive power of the trained models. The work suggests new venues for data preservation to enable future QCD studies of hadrons structure, and the developed technology can boost the science output of physics programs at facilities such as Jefferson Lab and the future Electron-Ion Collider.

Since the early 1970s, Monte Carlo event generators (MCEGs) have played a vital role in facilitating studies of high-energy particle scattering. From the experimental perspective, MCEGs are crucial for modeling the responses of the complex arrays of detectors that measure the energies and momenta of final state particles. The development of modern MCEGs, such as PYTHIA [1], HERWIG [2], or SHERPA [3], has been driven by a combination of high-precision experimental data and theoretical inputs. The latter have involved a mix of perturbative QCD methods, describing the dynamics of quarks and gluons at short distances, and phenomenological models that map the transition from quarks and gluons to observable hadrons, as well as nonperturbative inputs such as parton distribution functions for applications involving hadrons in the initial state. In this work we present an alternative approach to building an event generator based on machine learning (ML), free of theoretical assumptions about the femtometer-scale physics that controls the production of particles at high energy.

An MCEG can be viewed as a form of a data compatification utility, encapsulating large amounts of data collected from multiple experiments which can be regenerated from the MCEG. On the other hand, the reliance of existing MCEGs on theoretical assumptions of factorization and hadronization models limits their ability to capture the full range of possible correlations between the produced particles' momenta and spins. To date, for instance, no MCEG is able to reproduce all possible single-spin asymmetries in inclusive or semi-inclusive electron-nucleon deep-inelastic scattering (DIS).

Having an MCEG that faithfully simulates particle reactions is also important for the development of theoretical models. In practice, experimental data analyses are tailored to specific physical observables, which limits the possibility of accessing some other observable that was not conceived prior to the analysis. While it is possible to reconstruct other observables by repeating the analysis, the relatively high cost of transforming detector-level events into physical observables that can be studied in QCD is often prohibitive.

This scenario poses difficulty for developing theory, as the generic trends from particle distributions provide important insights into how to improve the theory using QCD factorization methods. For instance, in semiinclusive DIS different regions of phase space are expected to be controlled by different physical mechanisms. Having access to generic gross features of the data can provide hints about such mechanisms, and can motivate corresponding theoretical development. Such generic features of the data cannot be accessed from theory-based MCEGs as they are, by construction, biased towards specific physical pictures to which its parameters are tuned.

In this Letter, we suggest a new strategy for constructing MCEGs using machine learning (ML) methods involving generative adversarial networks (GANs) [4], which have been utilized in high-energy physics as a tool for fast Monte Carlo simulations [5–11]. A crucial feature of GANs is their ability to generate synthetic data by learning from real samples without knowing the underlying physical laws of the original system. Using particle momenta as event representations, we present a case study of inclusive electron-proton (ep) DIS. The results provide a new venue for experimental data analysis by building an AI-based MCEG that serves as a data compactification tool, which has the potential to boost the scientific discoveries at existing and future facilities such as Jefferson Lab [12] and the Electron-Ion Collider [13].

Davidson College, Davidson, North Carolina 28035, USA



FIG. 1. Schematic blueprint of the AI-based event-level data analysis framework "ETHER", comprising a vertex-level GAN (red boxes), which includes an event generator paired with a discriminator that receives inputs from the detector-level events, and a secondary folding GAN (green boxes) to fold the vertex-level events with detector response to mimic real detector-level events. The folding GAN must be trained separately using calibrated detector simulators with synthetic event samples to faithfully mimic the detector effects.

One of the challenges of any analysis of experimental data is that data, in the form of final state particle momenta, are affected by distortions introduced by particle detectors. At present, such corrections are taken into account using unfolding procedures that attempt to correct for the detector effects at the histogram level, and need to be developed for each type of observable.

On the other hand, the GAN technology allows us to define an analysis framework in which the unfolding procedure can be realized at the event level. The novelty of this approach is that once the AI-based MCEG is trained along with detector response, arbitrary physical observables can be reconstructed, free of detector effects as well as free of theoretical biases. A blueprint for such an AI-based event-level data analysis framework is illustrated in Fig. 1, which we refer to as the "empirically trained hadronic event regenerator" (ETHER). The blueprint utilizes two GANs: (i) a GAN as a vertex-level event generator, and (ii) a GAN for the event-level folding procedure to include detector effects for the vertexlevel events. Since the latter carries out the opposite of the unfolding procedure, it still has the same goal for accessing vertex-level distributions. While the full realization of the blueprint in Fig. 1 is our ultimate goal, in the current work we focus on an idealized situation in which detector effects are absent.

For our idealized AI-based MCEG, the GAN model is composed of a generator and a discriminator. The generator transforms random noise through a deep neural network to produce candidate samples, while the discriminator learns through another deep neural network to differentiate the true samples from ones produced by the



FIG. 2. Design of the ETHER in the absence of detector effects. The GAN (red boxes) is composed of a generator that creates synthetic events and a discriminator that trains the generator to mimic true particle production. The FAT procedure enhances the discriminator sensitivity to achieve maximal fidelity for the production of synthetic data by the generator. The Wasserstein and maximum mean discrepancy (MMD) [11] loss functions provide a smooth metric to match the distributions of fake and true events.

generator. The GAN training evolves as the generator and discriminator compete adversarially, each updating their parameters during the training process. Eventually, the GAN is able to produce synthetic samples such that the discriminator can no longer distinguish them from the real samples.

Although GANs have demonstrated impressive results in applications such as generating near-realistic images [14], music [15], and videos [16], training a successful GAN model is known to be notoriously difficult. Many GAN models suffer from major problems, including a mode collapse, non-convergence, a model parameter oscillation, destabilization, vanishing gradient, and overfitting due to unbalanced training of the generator and discriminator. Approaches and techniques to address these general problems have been proposed and discussed recently in the literature [17–19].

Unlike common GAN applications, such as the generation of realistic and high resolution images, the success of our GAN application as an MCEG relies on its ability to faithfully reproduce the aggregated representation of events given by histograms or cross sections. The latter typically span over several orders of magnitude, with complex correlations among the particle momenta, and its reproduction via synthetic events generated by the GAN requires custom-designed architectures and training strategies beyond pre-existing and commonly used architectures.

In Fig. 2 we present our ML prototype for the AI-based MCEG built upon the GAN framework and extended with a feature-augmented and transformed (FAT) strategy, which we refer to as the FAT-GAN. Our approach is based on the following scheme:

- We design the generator to produce a minimal set of event features, and reconstruct the rest of the features using momentum conservation and particle on-shell conditions. By construction we therefore enforce the generator to obey basic physics laws, and prioritize the learning on particle momentum distributions and particle correlations.
- We use a suitable transformation of variables for the feature space to flatten the particle distributions and reinforce energy and momentum conservation in cases where the training is on inclusive particle samples.
- To increase the discriminator sensitivity to particle correlations, the feature set is augmented with additional features such as combinatorial products among particle momenta.

We emphasize that the success of training the FAT-GAN relies on a careful choice of the feature sets, parameter initialization, and selection of hyper-parameters.

In our case study we apply the FAT-GAN to inclusive electron-proton scattering,  $ep \rightarrow e'X$ , where the scattered electron e' is detected with momentum k' and energy E', we use event samples from the PYTHIA [1] MCEG as a surrogate for real events. While our ultimate goal is to train on real experimental events using the ML design in Fig. 1, the case study provides unique insights into the GAN capabilities for generating event "images", as well as identifying challenges in formulating a suitable feature space to be learned by the FAT-GAN.

Since the generator only produces inclusive electron samples, momentum conservation cannot be imposed exactly. We therefore employ a series of transformations on the momentum vectors to generate a set of transformed features, which help the generator to produce the phase space for the electrons within physical limits. The feature set to be used as input to the discriminator is further augmented by including the scattered electron's transverse momentum,  $k'_{\rm T}$ , the energy E' (derived from the on-shell condition) and the momentum correlations  $k'_x k'_y$ ,  $k'_x k'_z$ and  $k'_y k'_z$  as additional features.

The FAT-GAN architecture in Fig. 2 is designed with the following characteristics:

— The generator: The input to the generator is a 100-dimensional white noise array centered at 0 with unit standard deviation. The generator network consists of 5 hidden dense layers, each with 512 neurons, activated by a leaky Rectified Linear Unit (ReLU) function. The last hidden layer is fully connected to a 3-neuron output, activated by a linear function representing the generated features.

— Feature augmentation and transformation: A customized Lambda layer is incorporated to calculate the augmented features from the generated features.

— **Discriminator**: The output of the feature augmentation and transformation is passed as an input for the discriminator which consists on 5 hidden dense layers, each of which has 512 neurons and activated by a leaky ReLU function. To avoid overfitting, a 10% dropout rate is applied to each hidden layer. The last hidden layer is fully connected to a single-neuron output, activated by a sigmoid function, where "1" indicates a true event and "0" is a fake event.

— Loss functions: The discriminator D is trained to give  $D(\mathbf{F}) = 1$  for each training sample  $\mathbf{F}$  generated by PYTHIA, and  $D(\widetilde{\mathbf{F}}) = 0$  for each sample  $\widetilde{\mathbf{F}}$  produced by the generator. The discriminator is optimized using the Wasserstein loss with gradient penalty [20] to improve training stability and reduce the likelihood for mode collapse. The loss function  $L_D$  of the discriminator is defined as

$$L_D = \left( \mathbb{E}[D(\widehat{F})] - \mathbb{E}[D(F)] \right) + \lambda \mathbb{E}_{\widehat{F} \sim P_{\widehat{F}}} \left[ (\|\nabla_{\widehat{F}} D(\widehat{F})\|_2 - 1)^2 \right],$$
(1)

where  $\mathbb{E}$  denotes the expectation value. The first term in (1) measures the Wasserstein distance [21], while the second term is the gradient penalty, where  $\hat{F}$  is a random sample from  $P_{\hat{F}}$ , defined by a uniform distribution along the straight lines between pairs of samples from the training samples and the generator's output. The coefficient  $\lambda$  is a harmonic parameter to balance the Wasserstein distance and the gradient penalty.

To ensure that the distributions of the event features created by the generator match well with the true distributions, we incorporate a two-sample test based on kernel MMD in ETHER. To compare two distributions, the MMD employs a kernel-based statistical test method to determine if the two samples were drawn from different distributions. As a result, the loss function  $L_G$  of the generator includes a Wasserstein distance term from the discriminator and an MMD term [22],

$$L_G = -\mathbb{E}[D(\widetilde{F})] + \eta \,\mathrm{MMD}^2(F, \widetilde{F}), \qquad (2)$$

where  $\eta$  is the balancing hyperparameter. The MMD term is defined as

$$MMD^{2}(\boldsymbol{F}, \widetilde{\boldsymbol{F}}) = \mathbb{E}_{\boldsymbol{F}_{a}, \boldsymbol{F}_{a'} \sim P_{\boldsymbol{F}}}[k(\boldsymbol{F}_{a}, \boldsymbol{F}_{a'})] \\ + \mathbb{E}_{\boldsymbol{F}_{b}, \boldsymbol{F}_{b'} \sim P_{\widetilde{\boldsymbol{F}}}}[k(\boldsymbol{F}_{b}, \boldsymbol{F}_{b'})] \\ - 2 \mathbb{E}_{\boldsymbol{F}_{a} \sim P_{\boldsymbol{F}}, \boldsymbol{F}_{b} \sim P_{\widetilde{\boldsymbol{F}}}}[k(\boldsymbol{F}_{a}, \boldsymbol{F}_{b})],$$
(3)

where  $k(\mathbf{F}_a, \mathbf{F}_b)$  is a positive definite kernel function. In this analysis we select a Gaussian kernel such that  $k(\mathbf{F}_a, \mathbf{F}_b) = \exp[-(\mathbf{F}_a - \mathbf{F}_b)^2/2\sigma^2]$ , where  $\sigma$  is the hyperparameter determining the MMD resolution, tuned to the same order of magnitude as the event feature width.

The combined network is trained adversarially for 100k epochs, where an epoch is defined as one pass through the training dataset. A batch of 10k events is employed



FIG. 3. Phase space distributions (with yield normalized to unity) for inclusive ep scattering from PYTHIA (black) and ETHER (red with yellow bands) trained on the electron momenta  $k'_i$  (i=x, y, z, T) and energy E', with predictions for the reconstructed variables  $\theta$ ,  $\phi$ ,  $Q^2$  and  $x_{bj}$  (vertical scales for  $\theta$  and  $x_{bj}$  are given on the right hand sides of the panels). The ratios of the ETHER distributions to PYTHIA are shown at the bottom of each panel, with the uncertainty bands for ETHER generated via bootstrap.

to ensure that there are sufficient samples to calculate a stable MMD value in each batch, which contains random examples from the training samples. For the optimizer we use Adam [23] with a  $10^{-4}$  learning rate,  $\beta_1 = 0.5$ , and  $\beta_2 = 0.9$ . To balance the generator and discriminator training, the training ratio is set to 5.

In Fig. 3 we compare the inclusive ep phase space distributions for the scattered electron, with yields normalized to unity, from the PYTHIA training samples and those generated from ETHER with the same statistics, at a center of mass energy of 100 GeV. The uncertainty bands were generated by training 10 independent FAT-GANs, where for each training the samples were prepared using the bootstrapping procedure (taking random samples with replacement). Each of the features on which the generator is trained, namely the components  $k'_i$  (i = x, y, z) of the electron momentum, the electron transverse momentum  $k'_{\rm T} = \sqrt{k'^2_x + k'^2_y}$ , and the energy  $E' = \sqrt{k_T'^2 + k_z'^2 + m_e^2}$  display sharp peak structures with tails that fall rapidly within the phase space. Despite variations over several orders of magnitude, ETHER is able to learn the underlying hidden distributions directly from the training events. Its predictability worsens in regions where the event rates are low, as reflected in the



FIG. 4. Comparison of the neutral current inclusive ep (reduced) cross section  $\sigma_r^{e^+p}$  (scaled by a factor  $2^i$ ) from the HERA collider [24] (black points) with data generated from PYTHIA [1] (yellow solid lines) and the trained ETHER (red dashed lines).

ratio plots, consistent with statistical expectations.

While the momenta  $k'_i$  and energy E' were part of the feature sets in the training, the polar and azimuthal scattering angles  $\theta$  and  $\phi$  are derived quantities. Even when these features were not included in the training set. Fig. 3 shows that ETHER is able to predict correctly the true distribution within uncertainties. Also included are predictions for the standard DIS variables, namely, the fourmomentum transfer squared  $Q^2 = -(k - k')^2$  and the Bjorken scaling variable  $x_{\rm bj} = Q^2/2P \cdot (k-k')$ , which are found to be in excellent agreement with PYTHIA. Since the structure functions and parton distribution functions extracted from DIS measurements depend explicitly on  $x_{\rm bi}$  and  $Q^2$ , the results in Fig. 3 indicate that physical observables can be accessed a posteriori from the eventlevel distributions even if they are not included in the original data analysis.

As an example of ETHER's ability to replicate real data, in Fig. 4 we compare the  $e^+p$  neutral current inclusive DIS (reduced) cross section  $\sigma_r^{e^+p}$  from the HERA collider [24] with data generated from PYTHIA [1] and with ETHER trained on PYTHIA samples at HERA

kinematics. Note that the high- $x_{\rm bj}$  and high- $Q^2$  region is expected to have larger statistical errors, reflecting the smaller event rates relative to those in the low- $x_{\rm bj}$ and low- $Q^2$  region. The cross section reconstructed with 1M PYTHIA samples indeed exhibits rather noisy behavior at high  $x_{\rm bj}$  and  $Q^2$ . While ETHER was trained on the same 1M PYTHIA samples, the reconstructed ETHER cross sections in Fig. 4 were generated from 100M trained samples. Consequently, the phase space with 100M events populates the high- $x_{\rm bj}$  and high- $Q^2$  region more densely, providing statistical convergence for the reconstructed cross sections.

More strikingly, the  $Q^2$  dependence of the synthetic cross sections displays the scaling behavior consistent with the HERA data and theoretical expectations from perturbative QCD [25]. The fact that ETHER was trained with low statistics relative to the number of simulated synthetic samples highlights the super-resolution capabilities of our generator, as encountered in earlier ML applications for images. This demonstrates that training an AI-based MCEG allows the reconstruction of observables in regions of phase space that have low rates. The ETHER technology therefore offers the possibility of boosting the scientific discoveries beyond the maximum machine luminosity attainable for a given experiment.

While the long term goal remains to build a generator capable of producing vertex-level events for QCD studies that can be trained directly on experimental data, the present analysis is an important proof of concept that demonstrates the viability of an AI-based MCEG, free of theoretical bias about the underlying particle dynamics. The promising results found with the ETHER simulations of synthetic as well as real data suggests potentially important applications of the AI-based technology to physical processes beyond inclusive *ep* scattering. As obvious improvements of this technology, we envisage implementation of the event-level folding tool based on GANs and trained on calibrated detector simulators, as well as integration of the folding with the FAT-GAN framework, where the generator produces vertex-level events while the discriminator is trained on detector-level events.

Future extensions will include dedicated closure tests of the super-resolution capabilities of GANs in phenomenological applications, such as the extraction of parton densities in global QCD analysis [26]. The use of the trained AI-based MCEG as a data compactification utility will also allow future exploration of QCD beyond the physical observables that were conceived of during the data taking, thus boosting the scientific discovery potential of experimental physics programs at current and planned facilities such as Jefferson Lab [12] and the Electron-Ion Collider [13]. Acknowledgments We thank M. Battaglieri and J. Qiu for helpful discussions. This work was supported by the Jefferson Lab LDRD project No. LDRD19-13 and No. LDRD20-18.

- T. Sjostrand, S. Mrenna and P. Skands, Comput. Phys. Commun. 178, 852 (2008).
- [2] M. Bahr et al., Eur. Phys. J. C 58, 639 (2008).
- [3] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert and J. Winter, JHEP 02, 007 (2009).
- [4] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville and Y. Bengio, Neural Information Processing Systems **32** (NIPS 2014), arXiv:1406.2661 [stat.ML].
- [5] L. de Oliveira, M. Paganini and B. Nachman, Comput. Softw. Big Sci. 1, 4 (2017).
- [6] M. Paganini, L. de Oliveira and B. Nachman, Phys. Rev. Lett. **120**, 042003 (2018).
- [7] M. Paganini, L. de Oliveira and B. Nachman, Phys. Rev. D 97, 014021 (2018).
- [8] P. Musella and F. Pandolfi, Comput. Softw. Big Sci. 2, 8 (2018).
- [9] B. Hashemi, N. Amin, K. Datta, D. Olivito and M. Pierini, arXiv:1901.05282 [hep-ex].
- [10] A. Butter, T. Plehn and R. Winterhalder, SciPost Phys. 7, 075 (2019).
- [11] A. Gretton, K. Borgwardt, M. J. Rasch, B. Scholkopf and A. J. Smola, Neural Information Processing Systems 19 (NIPS 2006).
- [12] V. D. Burkert, Ann. Rev. Nucl. Part. Sci. 68, 405 (2018).
- [13] A. Accardi et al., Eur. Phys. J. A 52, 268 (2016).
- [14] T. Karras, S. Laine and T. Aila, Conference on Computer Vision and Pattern Recognition (CVPR 2019, arXiv:1812.04948 [cs.NE].
- [15] O. Mogren, Constructive Machine Learning Workshop (NIPS 2016), arXiv:1611.09904 [cs.AI].
- [16] A. Clark et al., arXiv:1907.06571 [cs.CV].
- [17] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford and X. Chen, arXiv:1606.03498 [cs.LG].
- [18] S. Arora and Y. Zhang, arXiv:1706.08224 [cs.LG].
- [19] M. Arjovsky and L. Bottou, arXiv:1701.04862 [stat.ML].
- [20] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, A. C. Courville, Neural Information Processing Systems **30** (NIPS 2017), arXiv:1704.00028 [cs.LG].
- [21] M. Arjovsky, S. Chintala and L. Bottou, arXiv:1701.07875 [stat.ML].
- [22] C. Li, W. Chang, Y. Cheng, Y. Yang and B. Poczos, Neural Information Processing Systems **31** (NIPS 2017), arXiv:1705.08584 [cs.LG].
- [23] D. P. Kingma and J. Ba, International Conference for Learning Representations, San Diego (2015), arXiv:1412.6980 [cs.LG].
- [24] H. Abramowicz et al., Eur. Phys. J. C 75, 580 (2015).
- [25] J. C. Collins, D. E. Soper and G. F. Sterman, Adv. Ser. Direct. High Energy Phys. 5, 1 (1989).
- [26] N. Sato, C. Andres, J. J. Ethierand W. Melnitchouk, Phys. Rev. D 101, 074020 (2020).