RECEIVED: May 1, 2020

REVISED: June 8, 2020

ACCEPTED: June 15, 2020

PUBLISHED: July 21, 2020

Classification of equation of state in relativistic heavy-ion collisions using deep learning

Yu. Kvasiuk,^{a,b} E. Zabrodin,^{a,c,1} L. Bravina,^a I. Didur^d and M. Frolov^d

^a*Department of Physics, University of Oslo,
PB 1048 Blindern, Oslo N-0316, Norway*

^b*Faculty of Physics, Taras Shevchenko National University of Kyiv,
Akademika Hlushkova 4, Kyiv UA-01033, Ukraine*

^c*Skobeltsyn Institute of Nuclear Physics, Moscow State University,
Vorob'evy Gory, Moscow RU-119991, Russia*

^d*DataRoot Labs, Dmytrivska 80, Kyiv, Ukraine*

E-mail: yurii.kvasiuk@fys.uio.no, evgeny.zabrodin@fys.uio.no,
larissa.bravina@fys.uio.no, didur@datarootlabs.com,
frolov@datarootlabs.com

ABSTRACT: Convolutional Neural Nets, which is a powerful method of Deep Learning, is applied to classify equation of state of heavy-ion collision event generated within the UrQMD model. Event-by-event transverse momentum and azimuthal angle distributions of protons are used to train a classifier. An overall accuracy of classification of 98% is reached for Au+Au events at $\sqrt{s_{NN}} = 11$ GeV. Performance of classifiers, trained on events at different colliding energies, is investigated. Obtained results indicate extensive possibilities of application of Deep Learning methods to other problems in physics of heavy-ion collisions.

KEYWORDS: Heavy Ion Phenomenology, Phenomenological Models

ARXIV EPRINT: [2004.14409](https://arxiv.org/abs/2004.14409)

¹Corresponding author.

Contents

1	Introduction	1
2	Work setup	2
2.1	Features of microscopic model	2
2.2	Hard and soft EoS	3
2.3	Example: directed flow	3
2.4	Model setup	4
3	Deep Convolutional Neural Network	5
4	Results	6
4.1	Classification accuracy for central Au+Au collisions at three different energies	6
4.2	Application of more efficient algorithm — Adam	8
5	Conclusions and discussion	9

1 Introduction

One of the main goals of physics of heavy-ion collisions is to explore properties of exotic state of matter, namely, hot, dense and hard-interacting baryonic matter. It can be recreated in laboratory by colliding heavy nuclei at relativistic energies. Lattice Quantum Chromodynamics (QCD) calculations indicate that a transition from quark-gluon plasma (QGP) to hadron gas is a smooth crossover at high energies and low baryon densities [1]. It is widely believed that a first-order phase transition that ends up with tricritical point takes place within the energy range between $\sqrt{s} = 3$ and 10 GeV, see, e.g., [2] and references therein. Various past and ongoing experiments, such as beam energy scan (BES) and BES II at Relativistic Heavy Ion Collider (RHIC) [3, 4], experiments at Super Proton Synchrotron (SPS) at CERN, are exploring collisions with gold and lead ion beams to find any peculiarities within the aforementioned energy range. However, neither the first-order phase transition, nor the tricritical point has been observed so far. Future experiments, such as Nuclotron-based Ion Collider fAcility (NICA) and Facility for Antiproton and Ion Research (FAIR) are aiming to perform collisions at given energies with higher luminosity, which give us a hope to see something new there. Difficulties in observing the phase transition originate from numerous factors. Some of them are an extremely short, approximately 10^{-24} fm/c, time of existence of the QGP phase, small number of particles in the system, high anisotropy of matter both in coordinate and in momentum space, etc. All valuable information registered by detectors is roughly several thousands of particles with corresponding energies and momenta. Therefore, it is immensely difficult to make any reasonable assumptions about the media they came from.

Recent success of methods of Artificial Intelligence (AI), such as Machine Learning (ML) and Deep Learning (DL), in approximating highly obscure dependencies gives us a justifiable hope that they can benefit in heavy-ion physics objectives. Now these methods are widely employed in various problems that involve pattern recognition, for instance, image classification, natural language processing, and so on. Major advantage of such algorithms is that they do not need to be explicitly programmed. For example, AlphaGo [5], a program developed by Google Deep Mind to play a board game Go, managed to beat human world champion with the score 4:1. Played on 19×19 board, Go has roughly 2.08×10^{170} legal positions which makes it impossible for explicitly-written algorithm due to computational complexity. Moreover, a more advanced versions of the program, AlphaGo Zero and AlphaZero, were trained even without using the human data or guidance, but only by playing between versions of themselves and reached even better results [6]. Recently, algorithms of DL have started to be applied to complex tasks in various fields of physics, namely, high-energy particle physics [7–9], condensed matter physics [10, 11], astroparticle physics [12–14], and quantum information [15]. The work [16] demonstrates how optimal representation variables, describing different states of physical systems, could be learned by neural nets. Generative Adversarial Nets, a generative Deep Learning models have started to be successfully applied for the event simulations [17, 18]. Review of possibilities of application of ML and DL methods to high energy physics can be found elsewhere [19–21]. Prominent results of application of DL to problems of high energy nuclear physics and heavy-ion collisions are described in [22–24].

Present work continues efforts of application of deep learning methods to identification of properties of equation of state (EoS). The paper is organized as follows. Section 2 introduces basic features of the microscopic ultrarelativistic quantum molecular dynamics (UrQMD) model and parameters of hard and soft mean-field potentials used in the calculations. It is outlined that the Deep Learning methods are able to differentiate the calculations made with two different equations of state on event-by-event basis compared to standard methods, which need usually significant statistics of generated events. Principles of the Convolutional Neural Network (CNN) algorithm are sketched in section 3. Section 4 presents main results of our study. Here we show also the ability of model, trained on particular collision energy, to make correct distinctions for the collisions at neighbor energies. Effectiveness of different training algorithms is discussed as well. Conclusions are drawn in section 5.

2 Work setup

2.1 Features of microscopic model

The Ultra-relativistic Quantum Molecular Dynamics model [25, 26] is a Monte-Carlo event generator designed for the description of hh , hA , and $A + A$ collisions in a broad energy range from hundred MeV up to several TeV. The model contains 55 baryon and 32 meson states with masses up to $2.25 \text{ GeV}/c^2$ as independent degrees of freedom, together with their antiparticles and explicit isospin-projected states. At energies below few GeV, the interaction dynamics of hh or $A + A$ collisions can be described via interactions between

Parameter	Hard potential (EoS1)	Soft potential (EoS2)
α	0.25	0.25
t_1	−163	−353
t_γ	125.93	304
γ	1.676	1.167
V_0^{Yuk}	−0.498	1.0038
γ^{Yuk}	1.4	1.4

Table 1. Parameters of the hard and soft mean-field potentials used in the current version of the UrQMD.

the hadrons and their excited states, resonances. At higher energies new processes of multiparticle production come into play. The UrQMD treats the production of new particles via formation and fragmentation of specific colored objects, strings. Strings are uniformly stretched between the quarks, diquarks and their antistates with constant string tension $\kappa \approx 1 \text{ GeV/fm}$. The excited strings are fragmenting into pieces via the Schwinger mechanism of $q\bar{q}$ -pair production, and the distribution of newly produced hadrons is uniform in the rapidity space. The model utilizes Hamiltonian dynamics of particle motions and incorporates particle interaction via geometric cross sections, taken from available experimental data or from quark models.

2.2 Hard and soft EoS

There is a possibility to switch-on potential interaction in the form of Yukawa, Coulomb, and Skyrme potentials

$$V = \sum_{i \neq j} \left(V_0^{\text{Yuk}} \frac{e^{-|\vec{r}_i - \vec{r}_j|/\gamma_Y}}{|\vec{r}_i - \vec{r}_j|} + \frac{Z_i Z_j e^2}{|\vec{r}_i - \vec{r}_j|} \right) + t_1 \rho_j^{\text{int}} + t_\gamma (\gamma + 1)^{-3/2} (\rho_j^{\text{int}})^\gamma, \quad (2.1)$$

where

$$\rho_j^{\text{int}} = \left(\frac{\alpha}{\pi} \right)^{3/2} e^{-\alpha r_j^2}. \quad (2.2)$$

Here $\vec{r}_i$ is a coordinate vector of i -th particle, Z_i is value of electric charge, ρ_j^{int} is interaction density, V_0^{Yuk} , γ_Y , t_1 , t_γ , γ are constants that define stiffness of EoS, see [25–27] for details. In present work we use potential with two different sets of parameters which can be associated with two separate equations of state, hard and soft, of baryonic matter. Parameters of the hard and the soft potentials employed in the calculations are listed in table 1.

2.3 Example: directed flow

It was shown earlier that EoS in UrQMD influences kinematic observables, such as directed flow, v_1 , quite significantly [28]. Recall, that directed flow [29] is an experimental observable

Figure 1. Slope of directed flow of protons coming from UrQMD minimum bias Au+Au collisions generated with hard (stars) and soft (crosses) EoS, respectively. Data are taken from [28].

that characterizes momentum space anisotropy of particle spectrum. It is defined as first coefficient of transverse flow Fourier decomposition

$$E \frac{d^3 N}{d^3 p} = \frac{d^2 N}{\pi d p_T^2 dy} \left\{ 1 + 2 \sum_{n=1}^{\infty} v_n \cos [n(\phi - \Psi_n)] \right\}, \quad (2.3)$$

$$v_n = \langle \cos [n(\phi - \Psi_n)] \rangle. \quad (2.4)$$

Here p_T is particle transverse momentum, y is rapidity, ϕ is the azimuth between the $\vec{p}_T$ and the participant event plane, and Ψ_n is the azimuthal angle of the event plane of n -th flow component. Averaging in eq. (2.4) is done over all particle of a given species from the single event and over all set of events.

Figure 1 shows the midrapidity slopes of directed flow of protons in minimum bias Au+Au collisions at energies $\sqrt{s} = 4, 7.7, 11.5, 19.6$ GeV generated with two different EoS. One can see that the value of the slope is sensitive to the choice of EoS. However, to make the differences in v_1 slopes visible, about one million events for each data point were needed.

2.4 Model setup

It is worth mentioning that application of the Deep Learning methods permits one to differentiate between EoS on *event-by-event* basis. To check the efficiency of the method we generated 5000 Au+Au collisions of each kind at $\sqrt{s} = 11$ GeV with centrality 0–5%. Distribution of proton number $n_p = n_p(p_T, \phi)$ was calculated for each separate event. Transverse

No.	Layer	Parameters
0.	Input	$1 \times 10 \times 10$ tensor
1.	Conv2D	8 kernels 3×3 , ReLU
2.	MaxPool2D	2×2 kernel, stride 2
3.	Conv2D	2 kernels 3×3 , ReLU
4.	MaxPool2D	3×3 kernel, stride 1
5.	FullyConnected	144 neurons, ReLU
6.	FullyConnected	36 neurons, ReLU
7.	Sigmoid	

Table 2. Architecture of Convolutional Neural Network.

Figure 2. Network architecture diagram.

momentum and azimuthal angle were selected in ranges $(0, 1)$ GeV/ c and $(-\pi, \pi)$, respectively. Altogether, the training dataset is consisted of 10000 histograms with dimensions of 10×10 bins and labeled “0” for hard potential and “1” for soft potential. All bin values were scaled to fit the range $(0, 1)$, so they represent probability density distribution.

3 Deep Convolutional Neural Network

Convolutional Neural Network (CNN) [30] has proven to be a powerful deep learning algorithm for tasks dealing with graphical data, such as image recognition, classification, etc. It was inspired by the process of image processing in mammal’s visual cortex. Similar to other deep learning algorithms, CNN consists of layers of artificial neurons. However, there are three main types of layers that are used in deep CNN. Convolutional layers are used to extract useful features. Pooling layers are responsible for reduction of parameters and fully-connected ones are used for making the final prediction. They map extracted features to the desired output, which is either “1” or “0” in our case. Configuration of layers is determined by the complexity of the task. Network architecture employed here is described in table 2 and visualized in figure 2, build with the help of online tool for CNN architecture visualization.¹ Binary cross-entropy, used as a loss function, reads

$$S = -w_n (y_n * \log(x_n) + (1 - y_n) * \log(1 - x_n)) . \quad (3.1)$$

¹An online tool for CNN architecture visualization: <http://alexlenail.me/NN-SVG/LeNet.html>.

Figure 3. Proton densities for (a) soft and (b) hard potentials from a single UrQMD generated Au+Au collision at $\sqrt{s} = 11$ GeV.

Here w_n is a sample weight (set to 1), x_n and y_n are network outputs and true labels, respectively. Index n runs over the batch. The network was trained with a help of stochastic gradient descent (SGD) method with learning rate 0.01 during 800 epochs. Dropout [31] layers ($p = 0.3$) were used before every fully-connected layer to reduce overfitting. Validation loss was recorded after every 10 epochs and model parameters were saved every time it decreased. Model parameters that yielded lowest loss at validation set were used for model evaluation afterwards. Eighty percent of the dataset was used for training and twenty percent for validation. Data acquisition and early preprocessing steps were done within the ROOT framework [32]. Further process was realized in Jupyter-lab [33], which is an interactive Python notebook, with an extensive use of the Python libraries NumPy [34, 35], Matplotlib [36], Pandas [37], and Scikit-learn [38]. The network was realized in PyTorch [39] framework.

4 Results

4.1 Classification accuracy for central Au+Au collisions at three different energies

Typical particle density profiles coming from events generated with both hard and soft potentials are shown in figure 3. We have to emphasize that density profiles from single event are strongly influenced by fluctuations and substantially differ on event-by-event basis. They are hardly distinguishable by naked eye, yet if we consider larger statistics, e.g., 100 events, the differences become more noticeable, as shown in figure 4.

Overall classification accuracy of 98% percent was achieved for both hard and soft potentials. It means that out of 100 items of every kind 98 were classified correctly. Learning curves are depicted in figure 5. One can see that after 100-th epoch both training and

Figure 4. The same as figure 3, but for 100 Au+Au collisions.

Figure 5. Training and validation losses and accuracies for UrQMD generated Au+Au collisions at $\sqrt{s} = 11$ GeV.

validation losses start to abruptly fall, whereas both accuracies rapidly grow. This circumstance indicates that a direction of minimum of loss function in the model parameter space is found. All curves enter the saturation region after 150-th epoch. The fact that training metrics coincides with validation metrics indicates that our model does not overfit and generalizes well.

Then, analogous setup was used for energies of $\sqrt{s} = 7$ and 14 GeV. For the last case overall accuracy of 98% was reached, however, it drops to 94% of accuracy for $\sqrt{s} = 7$ GeV collision energy within the proposed training scheme. It could be attributed to

Training energies $\sqrt{s}$ (GeV)	Efficiency (%)		
	7	11	14
7	94	55	50
11	77	98	90
14	74	96	98

Table 3. Performance of models trained at different energies.

low multiplicity of secondary particles, produced in a single event. Next, it is interesting to study the ability of a model, trained on one particular energy, to make predictions for another energy. These results are listed in table 3. There are three trained models together with their predictions for Au+Au collisions at three energies, $\sqrt{s} = 7, 11$, and 14 GeV. Table 3 reveals an interesting result that a classifier, trained on events at collision energy of $\sqrt{s} = 14$ GeV, shows good performance while classifying the events at lower energies. The same is true for the classifier trained at $\sqrt{s} = 11$ GeV. It could indicate on similarities in hadron spatial distributions in both reactions. In contrast, performance of the classifier trained on $\sqrt{s} = 7$ GeV is significantly worse. Possible explanation could be low particle yields in a single event at this energy which are not enough to make a distinction properly. As it turned out, the smallest amount of data needed for accuracy better than 96% is only 100 images, i.e., only 1% of the whole dataset, for energies $\sqrt{s} = 11$ and 14 GeV. For the energy of 7 GeV the best performance with the least amount of data used is 92% with 400 training instances. However, training time is roughly 5 times longer compared to 4000 epochs for $\sqrt{s} = 11$ GeV.

4.2 Application of more efficient algorithm — Adam

It is possible to achieve very good results, namely, 92% of accuracy on validation set for $\sqrt{s} = 7$ GeV and 96% of accuracy on validation set for $\sqrt{s} = 11$ and 14 GeV, by using the Adam routine [40], which is a more advanced algorithm of loss minimization. Figure 6, figure 7, and figure 8 show overall validation accuracy during 5000 epochs of training for three energy data sets with different amount of training data and different optimizers.

Results obtained for central Au+Au collisions at $\sqrt{s} = 7$ GeV are displayed in figure 6. To reach the desired accuracy above 90% by standard SGD method using just 4% of generated statistics as a train factor, one needs about 2500 epochs. This time can be shortened to approximately 600 epochs if 20% of events are used for the training. In contrast, application of the Adam algorithm with 4% train fraction needs only 30(!) epochs to reach the accuracy about 90%. Peculiarity in application of this method is the slow declining of the validation accuracy after 100 epochs, which means that the model starts to overfit. Recall that training with low amounts of data is unstable with both training algorithms. More time is required by the model to find a path to the minimum of the loss. The problem can be cured by increase of the training statistics. It is important because the multiplicity of produced hadrons in heavy-ion collisions drops with decreasing collision

Figure 6. Validation loss during 5000 epochs of training for different training data fractions and optimizers in Au+Au central collisions generated within UrQMD at $\sqrt{s} = 7$ GeV. Solid and dashed curves indicate the results obtained by SGD method with train fraction 4% and 20%, respectively. Dotted curve denote the results obtained by Adam method with 4% train fraction.

energy. For Au+Au collisions at $\sqrt{s} = 11$ and 14 GeV the multiplicities of secondaries are larger than that at $\sqrt{s} = 7$ GeV.

Here advantages of the Adam algorithm become more obvious, as shown in figure 7 for the collisions at $\sqrt{s} = 11$ GeV. The standard SGD optimization method demonstrates sharp increase of the validation accuracy after 3200 epochs for the train fraction 1%, and after 400 epochs for the train fraction 20%. Training with Adam algorithm converges much faster. We see a rapid increase of accuracy in EoS recognition with the train fraction 1% after 30 epochs already. The same efficiency is obtained for central Au+Au collisions at $\sqrt{s} = 14$ GeV, as depicted in figure 8. Therefore, a classifier trained with the Adam optimization algorithm needs very limiting statistics for training within quite short time to discriminate between the hard and the soft equation of state with accuracy about 95%.

5 Conclusions and discussion

Our main results can be summarized as follows. Convolutional Neural Networks demonstrate excellent performance for equation of state classification in UrQMD generated heavy-ion collisions. Binary classification accuracy of 98% is achieved for events at $\sqrt{s} = 11$ and

Figure 7. The same as figure 6 but for central Au+Au collisions at $\sqrt{s} = 11$ GeV.

14 GeV and 94% for $\sqrt{s} = 7$ GeV, using the data set of primitive experimental observables, such as particle transverse momentum and azimuthal angle. Moreover, it is possible to differentiate between equation of state using the model trained on different neighboring energies.

It turned out that even a small fraction of data is enough to achieve satisfactory results. Training time depends strongly on data fraction used for the training, and on the employed algorithm. Our study shows that to reach the top classification accuracy for central Au+Au collisions at energies about 11 GeV and higher by the standard Stochastic Gradient Descent method, one needs more than 3000 epochs for training data of 1% of the generated statistics, and more than 400 epochs if the training fraction was increased to 20%. Similar accuracy can be achieved by application of the stochastic optimization algorithm Adam trained at statistics of 1% just after 30 epochs. Therefore, training Convolutional Neural Network classifier with the Adam algorithm permits one to reach high accuracy during very short training time with little amount of training data.

It is worth mentioning that the performance of trained CNN routine is very robust against initial state fluctuations and final state interactions, such as decays of resonances. However, similar to other applications of DL technique to high energy physics [22–24], the program has to decide between two discrete choices: calculations either with soft or stiff equation of state. In other words, this is a one-dimensional discrete space. In reality, the space of possible EoS is continuous. The next step, therefore, should be calculations

Figure 8. The same as figure 6 but for central Au+Au collisions at $\sqrt{s} = 14$ GeV.

with a combined EoS containing certain fraction x of the soft EoS and complementary fraction $1 - x$ of the stiff EoS. The model trained on such data would definitely give us fractions of both EoS. It will take more time because of the amount of generated data and longer training procedure. But, as we show in our study, the application of sophisticated algorithms, like Adam, makes solving this problem possible.

After that one can make a transition to multi-dimensional continuous space. This will allow us to extract different parameters of the EoS provided we can fix its functional form. If the limits of values of each parameter, entering the EoS, are determined, we have to generate the events for each set of parameters. Then, we train the model so that it learns the correspondence between (p_T, ϕ) -distributions of hadrons and set of parameters. The whole procedure looks cumbersome but feasible. This interesting problem needs further investigations.

Acknowledgments

Yu.K. acknowledges financial support of the Norwegian Centre for International Cooperation in Education (SIU) under Grant “CPEA-LT-2016/10094 — From Strong Interacting Matter to Dark Matter”. The work of L.B. and E.Z. was supported by Russian Foundation for Basic Research (RFBR) under Grants No. 18-02-40084 and No. 18-02-40085, and by the Norwegian Research Council (NFR) under Grant No. 255253/F50 — “CERN Heavy

Ion Theory”. Computer calculations were made at Abel (UiO, Oslo) and Govorun (JINR, Dubna) computer cluster facilities.

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

References

- [1] Y. Aoki, G. Endrodi, Z. Fodor, S.D. Katz and K.K. Szabo, *The Order of the quantum chromodynamics transition predicted by the standard model of particle physics*, *Nature* **443** (2006) 675 [[hep-lat/0611014](#)] [[INSPIRE](#)].
- [2] K. Fukushima and T. Hatsuda, *The phase diagram of dense QCD*, *Rept. Prog. Phys.* **74** (2011) 014001 [[arXiv:1005.4814](#)] [[INSPIRE](#)].
- [3] G. Odyniec, *The RHIC Beam Energy Scan program in STAR and what’s next...*, *J. Phys. Conf. Ser.* **455** (2013) 012037 [[INSPIRE](#)].
- [4] D. Thusty, *The RHIC Beam Energy Scan Phase II: Physics and Upgrades*, in proceedings of the 13th Conference on the Intersections of Particle and Nuclear Physics (CIPANP 2018), Palm Springs, CA, United States, 29 May–3 June 2018, [arXiv:1810.04767](#) [[INSPIRE](#)].
- [5] D. Silver et al., *Mastering the game of Go with deep neural networks and tree search*, *Nature* **529** (2016) 484.
- [6] D. Silver et al., *Mastering the game of Go without human knowledge*, *Nature* **550** (2017) 354.
- [7] J. Barnard, E.N. Dawe, M.J. Dolan and N. Rajcic, *Parton Shower Uncertainties in Jet Substructure Analyses with Deep Neural Networks*, *Phys. Rev. D* **95** (2017) 014018 [[arXiv:1609.00607](#)] [[INSPIRE](#)].
- [8] P. Baldi, P. Sadowski and D. Whiteson, *Enhanced Higgs Boson to $\tau^+\tau^-$ Search with Deep Learning*, *Phys. Rev. Lett.* **114** (2015) 111801 [[arXiv:1410.3469](#)] [[INSPIRE](#)].
- [9] CMS collaboration, *A deep neural network to search for new long-lived particles decaying to jets*, [arXiv:1912.12238](#) [[INSPIRE](#)].
- [10] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, *Science* **355** (2017) 602.
- [11] P. Mehta et al., *A high-bias, low-variance introduction to Machine Learning for physicists*, *Phys. Rept.* **810** (2019) 1 [[arXiv:1803.08823](#)] [[INSPIRE](#)].
- [12] M. Erdmann, J. Glombitza and D. Walz, *A Deep Learning-based Reconstruction of Cosmic Ray-induced Air Showers*, *Astropart. Phys.* **97** (2018) 46 [[arXiv:1708.00647](#)] [[INSPIRE](#)].
- [13] E.A. Huerta et al., *Enabling real-time multi-messenger astrophysics discoveries with deep learning*, *Nature Rev. Phys.* **1** (2019) 600 [[arXiv:1911.11779](#)] [[INSPIRE](#)].
- [14] C.J. Shallue and A. Vanderburg, *Identifying Exoplanets with Deep Learning: A Five-Planet Resonant Chain Around Kepler-80 and an Eighth Planet Around Kepler-90*, *Astron. J.* **155** (2018) 94 [[arXiv:1712.05044](#)].
- [15] M. Dalgaard, F. Motzoi, J.J. Sørensen and J. Sherson, *Global optimization of quantum dynamics with AlphaZero deep exploration*, *npj Quantum Inf.* **6** (2020) 6.

- [16] R. Iten, T. Metger, H. Wilming, L. del Rio and R. Renner, *Discovering Physical Concepts with Neural Networks*, *Phys. Rev. Lett.* **124** (2020) 010508 [[arXiv:1807.10300](#)].
- [17] R. Di Sipio, M. Fauci Giannelli, S.K. Haghighat and S. Palazzo, *DijetGAN: A Generative-Adversarial Network Approach for the Simulation of QCD Dijet Events at the LHC*, *JHEP* **08** (2019) 110 [[arXiv:1903.02433](#)] [[INSPIRE](#)].
- [18] SHiP collaboration, *Fast simulation of muons produced at the SHiP experiment using Generative Adversarial Networks*, *2019 JINST* **14** P11028 [[arXiv:1909.04451](#)] [[INSPIRE](#)].
- [19] A. Radovic et al., *Machine learning at the energy and intensity frontiers of particle physics*, *Nature* **560** (2018) 41 [[INSPIRE](#)].
- [20] K. Albertsson et al., *Machine Learning in High Energy Physics Community White Paper*, *J. Phys. Conf. Ser.* **1085** (2018) 022008 [[arXiv:1807.02876](#)] [[INSPIRE](#)].
- [21] M. Paganini, *Machine Learning Solutions for High Energy Physics: Applications to Electromagnetic Shower Generation, Flavor Tagging and the Search for di-Higgs Production*, Ph.D. Thesis, Yale University, New Haven U.S.A. (2019) [[arXiv:1903.05082](#)] [[INSPIRE](#)].
- [22] L.-G. Pang, K. Zhou, N. Su, H. Petersen, H. Stöcker and X.-N. Wang, *An equation-of-state-meter of quantum chromodynamics transition from deep learning*, *Nature Commun.* **9** (2018) 210 [[arXiv:1612.04262](#)] [[INSPIRE](#)].
- [23] Y.-L. Du et al., *Identifying the nature of the QCD transition in relativistic collision of heavy nuclei with deep learning*, *Eur. Phys. J. C* **80** (2020) 516 [[arXiv:1910.11530](#)] [[INSPIRE](#)].
- [24] J. Steinheimer, L.-G. Pang, K. Zhou, V. Koch, J. Randrup and H. Stöcker, *A machine learning study to identify spinodal clumping in high energy nuclear collisions*, *JHEP* **12** (2019) 122 [[arXiv:1906.06562](#)] [[INSPIRE](#)].
- [25] S.A. Bass et al., *Microscopic models for ultrarelativistic heavy ion collisions*, *Prog. Part. Nucl. Phys.* **41** (1998) 255 [[nucl-th/9803035](#)] [[INSPIRE](#)].
- [26] M. Bleicher et al., *Relativistic hadron hadron collisions in the ultrarelativistic quantum molecular dynamics model*, *J. Phys. G* **25** (1999) 1859 [[hep-ph/9909407](#)] [[INSPIRE](#)].
- [27] C. Hartnack et al., *Modeling the many body dynamics of heavy ion collisions: Present status and future perspective*, *Eur. Phys. J. A* **1** (1998) 151 [[nucl-th/9811015](#)] [[INSPIRE](#)].
- [28] L. Bravina, Y. Kvasiuk, S. Sivoklov, O. Vitiuk and E. Zabrodin, *Directed Flow in Microscopic Models in Relativistic A + A Collisions*, *Universe* **5** (2019) 69 [[INSPIRE](#)].
- [29] A.M. Poskanzer and S.A. Voloshin, *Methods for analyzing anisotropic flow in relativistic nuclear collisions*, *Phys. Rev. C* **58** (1998) 1671 [[nucl-ex/9805001](#)] [[INSPIRE](#)].
- [30] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, *Nature* **521** (2015) 436.
- [31] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, *Dropout: A Simple Way to Prevent Neural Networks from Overfitting*, *J. Mach. Learn. Res.* **15** (2014) 1929.
- [32] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81 [[INSPIRE](#)] and online at <http://root.cern.ch/>.
- [33] T. Kluyver et al., *Jupyter Notebooks — a publishing format for reproducible computational workflows*, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, IOS Press Ebooks (2016), pp. 87–90.

- [34] S. van der Walt, S.C. Colbert and G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, *Comput. Sci. Eng.* **13** (2011) 22 [[arXiv:1102.1523](#)] [[INSPIRE](#)].
- [35] T.E. Oliphant, *Guide to NumPy*, Continuum Press, Austin Texas U.S.A. (2015).
- [36] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [37] W. McKinney, *Data Structures for Statistical Computing in Python*, in proceedings of the *9th Python in Science Conference (SciPy 2010)*, Austin, Texas, U.S.A., 28 June–3 July 2010.
- [38] F. Pedregosa et al., *Scikit-learn: Machine Learning in Python*, *J. Mach. Learn. Res.* **12** (2011) 2825 [[arXiv:1201.0490](#)] [[INSPIRE](#)].
- [39] A. Paszke et al., *Automatic differentiation in PyTorch*, in proceedings of the *NIPS 2017 Workshop on Machine Learning for the Developing World*, Long Beach, California, U.S.A., 8 December 2017 and online at <https://openreview.net/group?id=NIPS.cc/2017/Workshop/Autodiff>.
- [40] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].