

A machine learning study to identify spinodal clumping in high energy nuclear collisions

Jan Steinheimer,^a Long-Gang Pang,^{b,c} Kai Zhou,^a Volker Koch,^c Jørgen Randrup^c and Horst Stoecker^{a,d,e}

^aFrankfurt Institute for Advanced Studies,
Ruth-Moufang-Str. 1, 60438 Frankfurt am Main, Germany

^bPhysics Department, University of California,
Berkeley, CA 94720, U.S.A.

^cNuclear Science Division, Lawrence Berkeley National Laboratory,
Berkeley, CA 94720, U.S.A.

^dInstitut für Theoretische Physik, Goethe Universität Frankfurt,
D-60438 Frankfurt am Main, Germany

^eGSI Helmholtzzentrum für Schwerionenforschung GmbH,
D-64291 Darmstadt, Germany

E-mail: steinheimer@fias.uni-frankfurt.de, lgpang@qq.com,
zhou@fias.uni-frankfurt.de, vkoch@lbl.gov, JRandrup@LBL.gov,
stoecker@fias.uni-frankfurt.de

ABSTRACT: The coordinate and momentum space configurations of the net baryon number in heavy ion collisions that undergo spinodal decomposition, due to a first-order phase transition, are investigated using state-of-the-art machine-learning methods. Coordinate space clumping, which appears in the spinodal decomposition, leaves strong characteristic imprints on the spatial net density distribution in nearly every event which can be detected by modern machine learning techniques. On the other hand, the corresponding features in the momentum distributions cannot clearly be detected, by the same machine learning methods, in individual events. Only a small subset of events can be systematically differentiated if only the momentum space information is available. This is due to the strong similarity of the two event classes, with and without spinodal decomposition. In such scenarios, conventional event-averaged observables like the baryon number cumulants signal a spinodal non-equilibrium phase transition. Indeed the third-order cumulant, the skewness, does exhibit a peak at the beam energy ($E_{\text{lab}} = 3\text{--}4 A \text{ GeV}$), where the transient hot and dense system created in the heavy ion collision reaches the first-order phase transition.

KEYWORDS: Heavy Ion Phenomenology, QCD Phenomenology

ARXIV EPRINT: [1906.06562](https://arxiv.org/abs/1906.06562)

Contents

1	Introduction	1
2	Method	2
3	The equations of state	5
4	Using deep learning	6
4.1	Coordinate space	6
4.2	Momentum space	8
4.3	Using different features	10
4.4	Unsupervised learning	12
5	Conventional observables	13
6	Discussion	16
A	Short introduction to neural networks and the training process	18
B	Technical details on the neural networks used	19
B.1	Convolutional neural network for coordinate space distributions	19
B.2	Convolutional neural network for momentum space distributions	19
B.3	Point-cloud neural network for particle lists	19
B.4	Fully connected neural network for the momentum difference distribution	20
B.5	The principal component analysis (PCA)	21

1 Introduction

The possible phase transition between a confined chirally broken phase of hadrons and a deconfined phase of quarks and gluons where chiral symmetry is restored has evaded experimental discovery for several decades now. At vanishing net baryon density the transition appears as a smooth crossover, as shown by lattice-QCD simulations [1–3]. At large baryochemical potentials the situation is less clear as direct calculations on the lattice are not possible due to the sign-problem [4]. Effective model predictions range from a strong first-order transition to a smooth crossover for the large density domain (or an even more complex phase structure with several transitions) [5–8]. In order to determine the phase structure of QCD, several heavy-ion experiments are currently performed or in preparation at RHIC, SPS-CERN, GSI, FAIR, NICA and JPARC-HI.

The proposed signals for a phase transition can be roughly split into two categories:

1. Effects of the softening of the equation of state (EoS) where the appearance of a phase transition leads to a local minimum in the speed of sound. This softening then can be related to changes in the collective flow due to the decreased pressure gradients in the early evolution of the system. Several observables have been suggested ranging from the mean transverse mass to several orders of the azimuthal anisotropies generated in heavy-ion collisions [9–20].
2. Effects from the non-equilibrium features and critical phenomena. In addition to effects of the softening of the equation of state, it is well known that systems can show signals that are related to the appearance of multi-particle correlations. For example, at the critical point the correlation length (in infinite systems) will diverge, leading to characteristic changes in the particle-number fluctuations [21–27]. For systems that undergo a phase transition, formation of baryon clusters can occur due to the spinodal decomposition associated with the mechanically unstable region of the phase diagram [28–35]. See also [36] for other works on effects of clumping due to a phase transition.

Recently it was suggested that a new approach based on modern machine-learning methods may offer a new promising venue. As modern neural networks are powerful tools for extracting information from complex datasets, it was suggested to use them to circumvent the biased ‘handcrafting’ of observables. Instead, the neural network should itself select the appropriate features within the data which are most sensitive to the properties of the equation of state. Indeed, foundational work [37] has shown that this is feasible, at least in a state-of-the-art relativistic fluid-dynamical approach. The machine learning methods can be employed for the classification of single events, which is different to previous attempts using conventional event-averaged observables. This makes these methods interesting for the analysis of specific features in these individual events.

In the present paper we will extend this method to try to identify special phase space features of a first order phase transition, namely features that should appear through instabilities in domains away from phase equilibrium, which are expected to occur in nuclear collisions. In particular we will show that the machine learning methods are able to find these features in the coordinate space but not in momentum space distributions of individual events. The further analysis, using unsupervised learning, supports the idea that these features should be identifiable by event-averaged statistical quantities like the cumulants of the net baryon number distribution.

2 Method

First we define a framework that is capable of correctly reproducing the underlying physics of the conjectured spinodal decomposition at a QCD phase transition and identifying the appropriate physical observables. In the present work, we employ relativistic fluid dynamics augmented with a gradient term to ensure the proper dispersion relation as expected for

spinodal decomposition. In addition we implement an equation of state that is mechanically unstable in the phase-coexistence region at large densities. Such a model has been presented in previous works [35, 38–40]. Simulations with this model have shown significant baryon clumping due to the spinodal decomposition during the passage of the unstable region in the phase diagram.

The time evolution of the system is based on the equations of relativistic ideal fluid dynamics, namely local four-momentum conservation,

$$\partial_\mu T^{\mu\nu} = 0, \tag{2.1}$$

and local flavor conservation,

$$\partial_\mu j^\mu = 0. \tag{2.2}$$

In this paper we include only the net baryon number current which is expected to carry most of the spinodal instability strength. These equations can be solved numerically on a 3+1 dimensional Cartesian lattice [41]. In order to take into account the effects of finite-range interactions (which, for example, are responsible for the presence of a surface tension), a gradient term is included. It modifies the equation of state $p(e, n)$ locally,

$$p(\mathbf{r}) = p_0(\varepsilon(\mathbf{r}), \rho(\mathbf{r})) - a^2 \frac{\varepsilon_s}{\rho_s^2} \rho(\mathbf{r}) \nabla^2 \rho(\mathbf{r}), \tag{2.3}$$

where $p_0(\varepsilon, \rho)$ is the equation of state in equilibrium, i.e. the pressure in uniform matter characterized by the energy density ε and by the net baryon density ρ . Furthermore, $\rho_s = 0.153/\text{fm}^3$ is the nuclear saturation density and $\varepsilon_s \approx m_N \rho_s$ is the associated ground state energy density. The strength of the gradient term is conveniently governed by the length parameter $a = 0.033 \text{ fm}$ [35]. As discussed in detail [38], this choice of a results in the formation of spinodal clusters of a characteristic size of 1–2 fm.

This model describes nuclear collisions at various incident beam energies. As strong deviations from thermal equilibrium appear in the initial penetrating phase, the ideal fluid dynamical description is supplemented by a non-equilibrium description for the initial state. For this, the non-equilibrium transport model UrQMD, which has been shown to successfully describe a wealth of data [42–44] but does not include any effects of a phase transition, is used.

In our calculation the spinodal instabilities are seeded by the fluctuations generated by the initial UrQMD evolution. In principle there are additional seeds from thermal noise, which, however, are small corrections to the dominant effects from the UrQMD fluctuations, as discussed in [25, 39]. Since the purpose of the paper is the detectability of these fluctuations rather than their precise determination, we omit thermal noise in our calculation.

Once the local density of particles reaches a certain value (usually below the coexistence density) the fluid dynamical description loses its validity and the system undergoes freeze-out. At this point we employ the Cooper-Frye procedure, based of the following integration over a predefined hypersurface Σ [45],

$$E \frac{dN}{d^3p} = \int_\Sigma f(x, p) p^\mu d\Sigma_\mu. \tag{2.4}$$

At this point, a decision must be made about how to implement the Cooper-Frye freeze-out transformation. In most studies, a random sampling of particles on the hypersurface is performed, where the local particle densities, calculated by eq. (2.4), are interpreted as probabilities for producing particles. Because the particle numbers per hypersurface element are usually significantly smaller than one, it is assumed that the probability for finding N particles in a cell follows a Poisson distribution, thus ignoring local correlations within a cell. In such a scenario one can enforce global conservation of flavors, which leads to a multinomial probability distribution for finding N particles in a finite number of cells [46].

Alternatively, the Cooper-Frye distribution can be implemented for each hypersurface element exactly. In such an approach, a non-integer particle number will then be emitted from any hypersurface element. Even though this is impractical, it is the only way to conserve all quantum numbers locally exactly (within each single hypersurface cell). In a recent paper it was shown that this local exact conservation reduces significantly the observed fluctuations [40].

Other methods have also been discussed recently [47], where the relevant flavors are conserved over clusters of fluid-cells of varying size.

It is not yet clear which of these methods should be favored, and how flavors are conserved locally in heavy-ion collisions. As shown in [40] the strength of the fluctuations depend on whether at the freeze-out flavors are conserved locally or globally. In case of one test particle, flavors are conserved only globally, while with increasing number of test particles flavor conservation becomes more local, as discussed in [40]. To address both limits of flavor conservation we consider the freeze-out with one and twenty testparticles. In either case we do *not* evolve the system after freeze-out with hadronic transport. Thus, we will compare results obtained with a varying number of test particles in the Cooper-Frye procedure. In this way, we will be able to take into account different scenarios of hadron production:

1. One test particle per real particle, TP=1: only global flavor conservation.
2. Twenty test particles per real particle, TP=20: this scenario comes close to local conservation of flavors, although non-integer particle numbers occur.

In the present work all (test) particles are sampled from the Cooper-Frye equation (2.4) and stored in an array. The list of all possible hadrons that can be sampled includes all stable ground-state hadrons, as well as an extensive list of unstable hadronic resonances. After all particles are sampled, the decays of these unstable resonances are calculated within the UrQMD transport model, using the complete Particle Data Group tables as input for the properties of the hadrons and their microscopic decays. Each test-particle resonance then decays exactly as a regular resonance, where the decay products carry only a fractional flavor.

With this model, heavy-ion collisions can be modeled at any beam energy which provides sufficient compression and heating to allow for a coexistence of confined and deconfined matter. To find out where the conditions for observable signals is best, the beam energy can be varied.




Figure 1. Pressure as a function of the net baryon density at a fixed temperature $T = 100$ MeV, for both the mechanically unstable spinodal equation of state (solid) and the corresponding stable equation of state (dashed) where the instabilities are removed by means of the standard Maxwell construction.

In the following section, we discuss the equation of state, the most important physics ingredient in the present simulations.

3 The equations of state

As we seek to identify signals for the predicted spinodal decomposition, or baryon clumping, we will focus on two equations of state that differ only with respect to the instabilities associated with the phase transition. Hence, they are identical outside of the spinodal region of the phase diagram, but within that phase coexistence region they differ significantly.

The spinodal EoS has a mechanically unstable region with a negative square of the isothermal speed of sound $c_s^2 < 0$. The stable partner EoS is obtained by means of a Maxwell construction (which has no effect in the already stable phase regions). The two equations of state are illustrated in figure 1 and more details on the construction of these equations of state may be found in ref. [48]. Well within the confined and deconfined phase regions, these equations of state describe a gas of interacting nucleons and pions and a gas of free two-flavor quarks and gluons in a bag, respectively.

Previous work has shown that the spinodal EoS and its Maxwell partner EoS lead to similar collective radial expansions [38]. By virtue of its construction, the Maxwell EoS produces the same amount of work during the expansion, which is proportional to $\int p dV$, as does the spinodal EoS, hence the amount of energy transformed to collective motion is exactly the same in both EoS cases.

In principle, there are many model predictions for the high baryon density equilibrium equation of state and here we adopt that of ref. [35]. As the present paper focuses on the detectability of the spinodal clumping associated with a first order phase transition, any equation of state that exhibits a co-existence between hadronic and quark matter and its associated spinodal region would serve this purpose.

We expect that any dynamical differences between the two scenarios are to be associated to the coordinate space correlations, in particular to the degree of baryon clumping during the phase separation. Indeed previous studies have shown that the two equations of state yield statistically significant differences in the baryon number distributions [40]. On the other hand, it has not been established whether the spinodal decomposition mechanism leads to clumping in essentially every event or in only a few. The main unanswered question is whether the clumping in coordinate space will actually yield a measurable significant signal in momentum space.

4 Using deep learning

In this paper, several popular machine- and deep-learning methods are applied in order to determine whether it is possible to discriminate between those events that are generated by the fluid dynamical model through the non-equilibrium spinodal EoS and those events that are generated by the Maxwell EoS corresponding to an equilibrium phase transition. To accomplish this goal, we first compare two different neural network architectures, which represent different supervised learning approaches. The first one is a convolution neural network (CNN), where event-by-event images of the density distribution in coordinate as well as momentum space serve as input. CNNs are used successfully in pattern recognition tasks in image applications. The second model is a point cloud network (PCN) [49], whose inputs are lists of discrete particle properties, e.g. particle four-momenta for every individual particle in a single event. The PCN is well suited for dealing with particles from collision experiments because it can use the momentum information for discrete particles as direct input. For a short introduction on neural networks and terminology we refer to appendix A.

The last section presents a unsupervised learning approach, i.e. a principal component analysis (PCA), which is used to extract the principal components of a given analysis feature, namely the two-particle momentum difference distributions (see appendix B.5 for details on the PCA). This feature is fed to a fully connected neural network (NN) to identify the EoS. The PCA yields a slightly improved accuracy.

4.1 Coordinate space

In a first step, we test the neural network for the coordinate clumping as expected from the spinodal equation of state: about 20 000 Pb+Pb collision events are generated at a (typical FAIR/GSI) beam energy of $E_{\text{lab}} = 3.5 A$ GeV, for each EoS. We know from previous studies of the moments of the density distribution [35] that the density fluctuations in coordinate space are strongest at $t = 3 \text{ fm}/c$,¹ at this beam energy and subside after another $3 \text{ fm}/c$. Thus we stop the time evolution of the system at the point in time where the density fluctuations are expected to be strongest, at $t = 3 \text{ fm}/c$. From each event an ‘image’ is then generated, containing information on the net baryon density distribution in the transverse spatial $X - Y$ plane for $Z = 0$. A naive way to classify the two scenarios is to just compare the maximum density for all the images and assume/postulate that the

¹The precise value of the optimal energy, $E_{\text{lab}} = 3.5 \text{ GeV}$, depends somewhat on the specific equation of state employed, so others would yield different values.




Figure 2. Example images of the normalized transverse density distributions. 20 000 of these individual Pb+Pb collision events are generated with the spinodal EoS (left) and Maxwell EoS (right). These images are used as input to train a convolutional neural network.

largest densities can be reached only in the spinodal case. However, to avoid such a trivial comparison and to make the task more challenging for the neural network, we renormalized the event-by-event density distributions by their maximum value for each event separately. Examples of the resulting, single event, distributions are shown in figures 2. Each image has a dimension of 100×100 pixels which corresponds to the number of fluid-cells shown in these figures.

For the training stage, a CNN with three convolutional layers and two pooling layers in between is used (See appendix B.1 for more details on the network structure). The output of this network is a binary classification on whether an input picture is from a spinodal or from a Maxwell event. Figure 3 shows the resulting accuracy during the training stage for the training dataset and for an independent validation set. It is obvious that the training as well as the testing sets accuracy increase quickly as the network learns the most important features for both cases. At later times the network overfits the data, indicated by the ever increasing training accuracy. In any case, already for the simple network, a 95% accuracy is obtained for the classification of events. This important finding has a consequence:




Figure 3. Training and validation accuracy for the Convolutional Neural Network (CNN) with coordinate-space input. After about 5 epochs the network starts over-fitting. Nevertheless, a good accuracy of about 95% can be achieved without optimization of the network structure.

it suggests that the spinodal EoS creates characteristic features in almost every event in coordinate space, which can be discriminated from features in the Maxwell case.

We speculate that these features correspond to the actual density clumping, however the network does not reveal how it reaches its results. Given the underlying physics, our speculation seems to be a reasonable explanation.

Note that this finding confirms previous findings. Indeed spinodal instabilities lead to characteristic structures in coordinate space [40, 50]. It is most important to point out that we have now verified that these clumping structures appear in nearly every sampled event.

4.2 Momentum space

Next let us focus on the event-by-event baryon distributions in momentum space. Here the connection of the final momentum space distributions to the baryon clumping during the early compression phase evolution is less obvious. In order to obtain the ‘final’ information on the momenta of all produced particles we need to run the fluid dynamical simulation until a later time. Although one may naively argue that the momentum space correlations should be largest at the point in time when the coordinate space clumping is large, this is actually not the case. In order to transform the coordinate clusters to momentum space correlations, these clusters need to expand fluid dynamically. Furthermore, the dense clusters are composed of dense quark matter which cannot be detected by experiment. In a realistic scenario, baryons will in fact be produced on an isoenergy density hyper-surface that is below the coexistence density at $e = 4e_0$ with e_0 being the nuclear saturation energy density. On this isoenergy hypersurface, the baryon density is approximately constant. and thus the signal in coordinate space has vanished completely. It is however possible that the coordinate space clumping has been transformed to momentum space correlations by the fluid dynamical evolution. Consequently, at the point of particle production, the clumping




Figure 4. Validation accuracy for the CNN and the PCN, comparing the spinodal EoS and the Maxwell construction for TP=1 and TP=20. The PCN shows a better performance than the CNN, but both network structures display only a very low validation accuracy just above 50%.

in coordinate space has disappeared and we have to rely on correlations in momentum space to discriminate between the two event classes.

The baryons are produced by a random sampling of the Cooper-Frye hypersurface as discussed in section 2, either sampling the actual number of particles (TP=1) or a large number of test particles (TP=20) per real particle. For both cases, all resonance decays are performed and the resulting phase space information for all produced particle (real and test particle) is saved in a list.

This list can be used to create a two-dimensional histogram of all baryons, using equal-sized bins in p_x and p_y . The histograms dimension are 20×20 bins of width 200 MeV in the range of $-2 < p_{x,y} < 2$ GeV. This histogram then gives an image, similar to the coordinate space image, but more coarse grained due to the finite number of particles. For the case of TP=20, 120 000 training events and 20 000 test events were generated.

Again a CNN (see appendix B.2 for the network structure) was trained to classify the equation of state on an event-by-event basis. The results on the training and testing accuracy are shown in figure 4 as black and blue lines. It is clear that the network fails to find significant features that would allow it to discern the spinodal from the Maxwell class on an event-by-event basis. While for TP=20 an accuracy of 52% can be reached, the accuracy for TP=1 is almost as low as 50% which would mean that the network result is no better than random guessing. One may argue that the 20×20 bin size input is too coarse grained to capture the relevant correlations, however, increasing the number of bins does not lead to a better result as it significantly increases the noise on the event-wise input data.

The CNN is a state-of-the-art pattern recognition method when the input is an image-like instance. In this particular case, the input ‘images’ are created from the momentum density distribution of baryons. More precise momentum information about the baryons could be lost during the data processing and one would prefer to actually use all particle information as training input.

To use the full information about all the discrete particles, we also employ a point cloud network architecture taking as input points in a five-dimensional feature space (for more details on the network structure and the hyper-parameters used, see appendix B.3). The five features $(E, p_x, p_y, p_z, \text{mass})$ define one point for each particle. Thus the PCN is more convenient and better suited for dealing with lists of particle information as it is obtained from experiments and event generators. An important advantage of this network is that the particle lists of all events have a permutation symmetry, i.e. changing the order of particles in an event does not affect the final result.

With the PCN, the testing accuracy can now reach 52% for the TP=1 case and 54% for TP=20, as shown in figure 4. Even though the PCN is more convenient for dealing with lists of particles and does an overall slightly better job at the event classification, the accuracy is still rather poor when only momentum-space information is taken into account.

4.3 Using different features

An important result of the last section is that the CNN as well as the PCN are not able to efficiently extract distinguishing features from the momentum space information, i.e. the plain p_x - p_y spectrum or the baryon momentum vectors. As a result the classification accuracy is very low, only 52%.

Since the underlying hope behind the consideration of spinodal clumping is that the clusters will produce correlations in momentum space, it may be useful to construct a distribution of momentum differences instead of the pure spectrum. This feature engineering may yield better results than leaving the full engineering to the neural network.

The two-particle momentum-difference distribution $dN_{\text{pairs}}/d\Delta p$ can be calculated event-by-event by binning the momentum difference between all pairs of baryons a and b in the event,

$$\begin{aligned} \Delta p_{ab} &\equiv \frac{1}{2} |\mathbf{p}^a - \mathbf{p}^b| \\ &= \frac{1}{2} \left[(p_x^a - p_x^b)^2 + (p_y^a - p_y^b)^2 + (p_z^a - p_z^b)^2 \right]^{\frac{1}{2}}. \end{aligned} \tag{4.1}$$

The quantity Δp_{ab} is equal to the momentum of each of the two particles as measured in the rest frame of the pair. An earlier study [51] considered a somewhat similar observable, namely the average kinetic energy of each particle in an N -body cluster and found that the signal-to-background grows stronger as the cluster size N is increased but, at the same time, the counting rate decreases progressively. In the present exploratory study, we stick to just two-baryon clusters.

The resulting event-averaged distributions for the spinodal and Maxwell cases are shown in the upper panel of figure 5, while the relative difference of these distributions is shown in the lower panel of figure 5. It should be noted that this distribution gives identical results for different numbers of test particles, but the amount of noise is larger when the number of test particles is small. The relative difference of these baryon pair distributions is small, but appears systematic. In the spinodal case, intermediate-momentum differences ($\Delta p < 0.5 \text{ GeV}$) are preferred, while large-momentum differences are suppressed. In order




Figure 5. Panel *a*: the distribution of the momentum difference between baryon pairs after the Cooper-Frye hadronization procedure; both, the spinodal and Maxwell distributions, are very similar. Panel *b*: the small relative difference between the distributions.

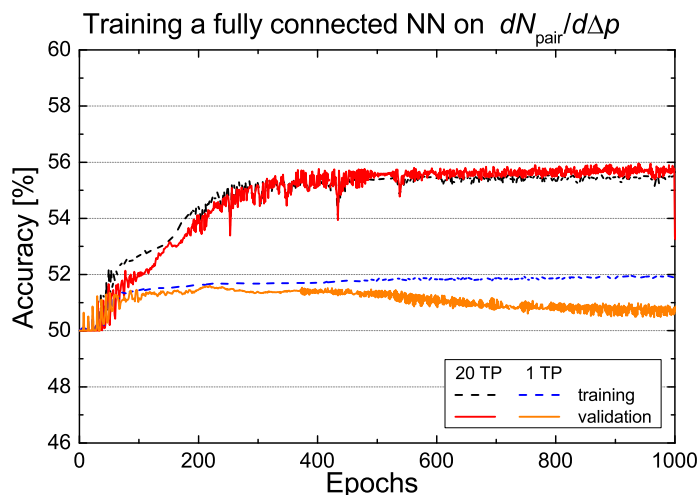


Figure 6. Training and validation accuracy as function of the number of training epochs for the fully connected neural network, training with event-by-event distributions of the momentum difference of baryon pairs. We compare results for TP=1 and TP=20.

to find out whether these small differences can be used to distinguish the event classes, we will employ a fully connected neural network for classification. This network structure is chosen due to the simpler input data, namely the $dN/d\Delta p$ distribution, which is a 200-bin dataset. After training both the TP=1 and TP=20 datasets, we find that indeed the neural network performs better than on the pure $\{p_x, p_y\}$ spectra. However only an accuracy of 55% for TP=20 and 52% for TP=1 is reached, as shown in figure 6. Even though this is better than random, it is not sufficient to claim the discovery of a two class distribution.




Figure 7. Scatter plot of the first two components of a PCA of the distribution of the momentum difference for baryon pairs. The red crosses indicate the events with the spinodal instabilities, which dominate the southern hemisphere, while the blue pluses indicate events with a Maxwell construction. The large symbols with error bars near the center indicate the mean values with their standard deviation. The spinodal EoS creates a clear crescent of crosses in the southern hemisphere, $x_2 < 0$. Also the mean value is shifted downwards slightly.

4.4 Unsupervised learning

In order to better understand why only a relatively poor individual event discrimination can be obtained, we will analyze the $dN/d\Delta p$ distribution using a simple unsupervised learning tool, because the distribution is only one dimensional. Here we will use a principal component analysis (PCA) which extracts the most relevant features of a function. To understand the results, one should keep in mind that for a machine-learning analysis the $dN/d\Delta p$ distribution is a 200-dimensional vector, i.e. each Δp bin corresponds to a dimension in the training feature vector. The PCA analysis will try to reduce the number of dimensions of the feature vector by transforming the basis of the vector such that the variance in the first components of the vector (for example a two-dimensional vector) is maximized. It is trying to preserve the maximum amount of information. Such principal components might be related to the peak position or width of the distribution or some combination of these.

Figure 7 shows a plot of the first two principal components x_1 and x_2 for all training events with TP=20. Each point corresponds to a single event and the color indicates whether the event corresponds to the spinodal EoS (red crosses) or the Maxwell construction (blue pluses). Note that in the PCA this information was not given. The large square symbols denote the average values of the components of a specific event class. In addition, the standard deviation of the extracted components are shown as error bars.

The figure indicates that the mean values are similar with only a small systematic shift. It was checked that this shift is statistically significant (the error of the mean is much smaller than the difference of the means). A closer examination reveals that the red crosses (spinodal) are more frequent for negative values of x_2 . This indicates that the

spinodal EoS more likely produces events that ‘look’ different than the average event. On the other hand, these events are rare which explains why an individual event classification analysis did not produce the desired results. In other words, only few events display characteristic features, while most events appear indistinguishable.

Fortunately there is a way to verify this suspicion. The output of the neural network, before the final *Softmax* function, are actual values that are proportional to the probability that a given input event belongs to either class Maxwell or spinodal. Thus the network assigns each input a probability that it belongs to either class, that is then transformed into a binary decision. By studying how these probabilities are assigned, with the help of the unsupervised learning approach, we can learn about the networks decision process. First, we selected those 2000 events to which the fully connected neural network had assigned the largest probabilities for belonging to either class, spinodal or Maxwell. For these events we plotted in figure 8 only those that are correct classifications as spinodal (red crosses) and Maxwell (green pluses). The blue points refer to the total dataset. It is obvious that the neural network separates the events mostly according to the x_2 feature also found in the PCA. It turns out that the accuracy for those events shown (the 2000 events with the largest probabilities) is around 70%.² This is considerably better than for the total dataset, which was around 54%. This finding confirms the intuitive suspicion that there is a strong overlap in the features of the spinodal and Maxwell events. The network thus focuses on events that are outliers, having the strongest features, namely x_2 in our case. Consequently, only a relatively poor accuracy for all events can be obtained.

We note that the fact that the network has the highest accuracy for the events in the tails of the x_2 distribution of the PCA (green and red points of figure 8 does not mean that the physics for these events is different than that for those in the same event class with a smaller value of x_2 . As shown in section 4.1 almost all events with the spinodal EOS show characteristic features in coordinate space. However, the correlations, which are clearly visible in coordinate space, are rather weak in momentum space, so that the Δp distribution is dominated by noise. This is illustrated in figure 9, where we show the x_2 distribution for both event classes. While the peaks are clearly separated their widths are very large so that only the events in the respective tails can be uniquely associated with a given event class.

5 Conventional observables

In this final part we use ‘conventional’ statistical methods to distinguish spinodal events from Maxwell events. In particular, we construct the cumulants of the baryon number distribution function, as they are also measured in several heavy ion experiments. The advantage of these cumulants is that they are sensitive also to the outlying events. The first three cumulants are given by

$$K_1 = M = \langle N \rangle , \tag{5.1}$$

$$K_2 = \sigma^2 = \langle (\delta N)^2 \rangle , \tag{5.2}$$

$$K_3 = S\sigma^3 = \langle (\delta N)^3 \rangle , \tag{5.3}$$

²Again, accuracy refers to the likelihood that the network makes the correct decision, while the probability refers to how likely the network thinks an event belongs to either class.




Figure 8. Scatter plot of the first two components of a PCA of the distributions of the momentum difference for baryon pairs. The blue points correspond to all 70 000 events, Maxwell and spinodal. The red crosses correspond to those events that were identified correctly as being in the spinodal class from among those 1000 events that had the highest probability for being spinodal events, according to the neural network. Similarly, the green pluses show the correctly identified Maxwell events among the 1000 event having largest probability of belonging to that class. According to this distribution it is reasonable to assume that the neural network would assign points (distributions) in the norther hemisphere a larger probability to be from the Maxwell EoS and points in the southern hemisphere to be spinodal events. Therefore it is clear that the x_2 variable serves as discriminator for the two event classes and that the network result is dominated by the small shift in an x_2 -like feature.




Figure 9. Distribution functions of the second principal component x_2 for both classes, *Maxwell* (blue dashed) and *spinodal* (red solid). One can clearly see the small shift in the mean x_2 and the effect on the tails of the distributions.




Figure 10. The scaled variance of the net baryon number as a function of the rapidity window size around mid-rapidity. Different scenarios are shown: the results for TP=1 and TP=20 with spinodal and Maxwell EoS, as well as a binomial baseline.

where $\delta N = N - \langle N \rangle$, N is the number of particles in a given experimental acceptance window and the brackets denote an event average. These cumulants, and ratios thereof, have been measured by the STAR, ALICE, NA61 and HADES experiments [52–54]. Usually, convenient ratios of the cumulants are presented:

$$\omega = K_2/K_1, \quad (5.4)$$

$$S\sigma = K_3/K_2. \quad (5.5)$$

In the following we show results from our model simulations for the same events that were used in the machine-learning analysis. We have calculated the net baryon number cumulants in intervals in rapidity, around mid-rapidity, integrating over the entire transverse momentum distribution. The errors are estimated using the delta-theorem and a random distribution, which should give an upper estimate of the error (for more details on the error see e.g. [55, 56]).

Figure 10 shows the scaled variance ω of the baryon number as a function of the rapidity window size for the two scenarios with TP=1 and TP=20. A clear difference between the two cases is observed, but the scaled variance for the two cases, spinodal and Maxwell are almost identical.

Figure 11 shows K_3/K_2 , for the same rapidity window size. Again the results show a clear dependence on the number of test particles used, whether twenty (TP=20) or one (TP=1). Importantly, for this observable the two EoS scenarios considered lead to distinguishable results, the difference being quite significant for TP=20 and still visible for TP=1.

Finally, the beam energy dependence of the skewness is presented in figure 12. A clear enhancement is seen in this observable at the beam energy with the strongest clustering ($E_{\text{lab}} = 3.5$ A GeV). The peak is strongest for TP=20, but the enhancement is still visible for TP=1. We propose that this peak presents an observable indication of spinodal decomposition in nuclear collisions due to the occurrence of a first-order phase transition.




Figure 11. The normalized skewness of the net baryon number as function of the rapidity window size around mid-rapidity. Different scenarios are shown: the results for TP=1 and TP=20 with spinodal and Maxwell EoS, as well as a binomial baseline.




Figure 12. The normalized skewness of the net baryon number distribution in the rapidity window of $-0.5 < y < 0.5$ for several incident beam energies (in the lab frame). Results for TP=1 and TP=20 are compared. A peak is found for the beam energy that produces the largest effect of the spinodal decomposition.

6 Discussion

We have presented a detailed investigation of how phase-space clumping of the baryon number, due to a first order phase transition, can or cannot be observed in heavy ion collisions. Employing state-of-the-art machine learning methods we were able to show that the QCD phase transition will lead to systematic features, in the coordinate space baryon density distribution, in essentially every event, given that the beam energy is tuned so that the created system will lie in the unstable region of the phase diagram. It was also demonstrated that the translation from coordinate space clumping to momentum

space clumping is far from trivial. Using Cooper-Frye sampling of particles as well as test particles, while globally conserving the baryon number, the systematic distinguishable features of the spinodal clumping, in individual events, almost entirely disappeared in momentum space (which is the only kind of data available to heavy ion experiments).

Both the CNN and the PCN were able to reach only slightly better than random accuracy, discriminating between single spinodal and Maxwell events, on the basis of momentum space correlations. Similarly, no measurable discrimination was obtained on the basis of individual event net baryon rapidity distributions. Furthermore, when the two-baryon correlation function was explicitly calculated and used as training input for a fully connected neural network, small effects on an individual events were found. The accuracy of this method was similar to that of the PCN. The deep learning methods presented in this work were not able to reliably identify the effects of spinodal decomposition when given single event momentum distributions as input. Clearly, it seems very difficult to identify these effects in the momentum distributions of individual events.

It is important to note that we did indeed find a systematic shift of the event properties from Maxwell to spinodal events when a PCA analysis was performed. We have found that the PCA is most sensitive to the momentum difference distribution around 500–600 MeV. At the same momentum difference also the average distributions show the largest absolute difference, as seen in figure 5(b). Therefore it is reasonable to assume that the difference in physics lies in a small shift of the mean of the two proton momentum difference distributions, and those events that have the largest difference in the Δp_T distribution around 500 MeV are most likely to be correctly identified.

Consequently, observables that are sensitive to the tails of the event distributions, are more promising than individual event observables. Such observables are for example higher-order cumulants of the net-baryon distribution.

We predict that the transition through the QCD phase transition is visible as a maximum of the skewness of the net-baryon number multiplicity distribution in a central rapidity window of $-0.5 < y < 0.5$, which is accessible with current and future experiments.

We also expect that higher-order cumulants, such as the kurtosis, should also show effects of the spinodal decomposition, but within the current calculational statistics we could not make reliable predictions. Thus, creating a fluid dynamical model that can incorporate the effects of spinodal decomposition, but has a significantly reduced computational time would be an important future task.

Acknowledgments

JS and KZ thank the Samson AG and the BMBF through the ErUM-Data project for funding. JS and HS thank the Walter Greiner Gesellschaft zur Förderung der physikalischen Grundlagenforschung e.V. for her support. This work was supported by the DAAD through a PPP exchange grant. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under contract number DE-AC02-05CH11231. L.G. Pang is supported by NSF under Grant No. ACI-1550228 within the JETSCAPE Collaboration. Computational resources were provided by the

NVIDIA Corporation with the donation of two NVIDIA TITAN Xp GPUs and the Frankfurt Center for Scientific Computing (Goethe-HLR).

A Short introduction to neural networks and the training process

In simple terms one can understand a neural network as some kind of mapping function, which maps an n -dimensional input vector to an m -dimensional output vector. In our specific case the output vector is two dimensional, the value of each dimension determining whether the input vector belongs to class 1 (Spinodal) or class 2 (Maxwell). Besides these so-called input and output layers a neural networks consists of a varying number of hidden layers. These hidden layers consist of neurons who themselves perform a non-linear transformation on their input. This non-linear transformation usually consists of a linear transformation $y = ax + b$, where x is the input of the neuron, y is the output of the neuron and a and b are parameters of that specific neuron (in the case of many neurons and many layers a and b take the form of multi dimensional matrices). The output y then serves as argument of a so called *activation function* which can take on different forms (often sigmoid or Relu functions are used). Depending on the structure of the network the neurons of one single layer can be connected to any number of neurons in the next layer. For example in a fully connected neural network, the output of the j^{th} neuron in the $i + 1^{\text{th}}$ layer is the sum of all outputs of the i^{th} layer $y_{i+1,j} = f(\sum_k a_{k,i+1}x_k + b)$ where k is the number of neurons in the i^{th} layer and $f()$ is an activation function. As one can see such a network can easily have a large number of parameters to be determined. The determination of these parameters is done during the training phase of the network. During this training, a number of pre-labeled data-points is fed to the network. Using an error-function which measures how well the network output agrees with the true training label (in our case the *cross entropy*), the parameter values are changed using a gradient descent method in order to minimize the error with respect to the training data. Since the gradient descent method only changes the parameters incrementally, many passes through the whole training data sets are necessary in order for the gradient descent to converge. Such a single pass through the training data-set is called *epoch*. Once the network is trained it can be validated using an independent set of data.

Due to the large number of parameters a neural network is prone to *overfit* the data, i.e. the error on the training data decreases as the error on the validation data increases. A popular regularization method to avoid overfitting is the *dropout*.

For the *dropout* an additional layer of i neurons is added between the hidden layers. This layer performs an identity transformation, i.e. it only passes through all the information $y_i = x_i$, where x_i is the output of the previous layers neurons. However, the neurons of the *dropout* layer have a finite probability to for their output to be set to zero, i.e. there is a finite probability that $y_i := 0$ for any neuron i in that layer. Since the neurons that are *dropped out* are randomly chosen during each feed-forward training pass, this method not only reduces the overfitting but also introduced stochasticity in the training and often makes the trained model more robust and generalizable.

For a much more in-depth explanation on neural networks we refer the interested reader to [57].

B Technical details on the neural networks used

In the following, the details on the neural networks and PCA method used are summarized.

B.1 Convolutional neural network for coordinate space distributions

The network employed for the binary classification of the density distributions in the x - y plane, as discussed in section 4.1, is structured as shown in figure 13. It corresponds to a two dimensional convolutional network, using 3 convolutional layers of different size. Between these layers Pooling layers are introduced, which decrease the dimensionality of the convolutional layers.

The network is constructed and trained using the *Tensorflow* library and *PYTHON*, using the *cross entropy* to calculate the loss function and the *Adam optimizer* for the gradient descent.

B.2 Convolutional neural network for momentum space distributions

The network employed for the binary classification on the momentum density distributions in the $p_x - p_y$ plane, as discussed in section 4.2, is structured as shown in figure 13.

The network is constructed and trained using the *Keras* library and *PYTHON*. Batch normalization and L2 regularization (0.0001) are used. The *Adam optimizer* is used with learning rate 0.00005. These hyperparameters as well as the dropout rates are varied to minimize the overfitting while still allowing the network to capture the important features of the data set.

B.3 Point-cloud neural network for particle lists

This network is employed to do the binary classification using only a list of particle information, i.e. the momentum 4-vector and mass as the coordinate space information cannot be measured directly by the experiments. The input, the network structure and other hyper parameters are listed below and a schematic view is shown in figure 14.

- For each event, the input is a list of particles. Each list contains 380 to 420 particles, corresponding to the variation of the number of participant baryons in each event, i.e. the length of the input list has to be 420 to allow every event to fit. Each particle has 5 features $(E, p_x, p_y, p_z, \text{mass})$, these features are scaled to approximately lie within the range $[0, 1]$ using the minimum and the maximum values of the first event.
- In the point-cloud network, for events with less than 420 baryons, we pad the remaining input array with five-zero vectors.
- The *cross entropy* is used to calculate the loss function.
- An *SGD optimizer* is used for the gradient descent. For this optimizer, the learning rate is set to 0.0005, the momentum is set to 0.9 and the nesterov is turned on.

The network is constructed and trained using the *Keras* framework and *PYTHON*.




Figure 13. Structure of the CNN used in the coordinate space analysis (left) and momentum space analysis (right).

B.4 Fully connected neural network for the momentum difference distribution

The network employed for the binary classification on the momentum-difference distributions $dN_{pairs}/\Delta p$, as discussed in section 4.3, is structured in the following way:

- Input are the 200 bin momentum-difference distributions as function of Δp .
- First fully connected hidden layer of 200 neurons and a *leakyReLU* activation function.
- Second fully connected hidden layer of 100 neurons and a *leakyReLU* activation function.




Figure 14. Structure of the Point Cloud Network used in this analysis.

- An output layer with 2 neurons and a softmax activation function.
- The learning rate is set to 0.0001 and the *cross entropy* is used as loss function for the training.

The network is constructed and trained using the *Tensorflow* library and *PYTHON*.

B.5 The principal component analysis (PCA)

The PCA used in our analysis is a linear dimensionality reduction using Singular Value Decomposition of the data to project it to a lower dimensional space [58], as provided by the *sklearn* library of *PYTHON*. The resulting x_n vectors of the first n components are the principal axes in feature space, representing the directions of the maximum variance in the data.

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](https://arxiv.org/abs/hep-lat/0611014)] [[INSPIRE](https://inspirehep.net/literature/571983)].
- [2] S. Borsányi et al., *The QCD equation of state with dynamical quarks*, *JHEP* **11** (2010) 077 [[arXiv:1007.2580](https://arxiv.org/abs/1007.2580)] [[INSPIRE](https://inspirehep.net/literature/876103)].
- [3] HOTQCD collaboration, *Deconfinement and chiral transition with the highly improved staggered quark (HISQ) action*, *J. Phys. Conf. Ser.* **230** (2010) 012014 [[arXiv:1005.1131](https://arxiv.org/abs/1005.1131)] [[INSPIRE](https://inspirehep.net/literature/867111)].

- [4] P. de Forcrand, *Simulating QCD at finite density*, [PoS\(LAT2009\)010](#) (2009) [[arXiv:1005.0539](#)] [[INSPIRE](#)].
- [5] M.A. Stephanov, K. Rajagopal and E.V. Shuryak, *Signatures of the tricritical point in QCD*, [Phys. Rev. Lett.](#) **81** (1998) 4816 [[hep-ph/9806219](#)] [[INSPIRE](#)].
- [6] L. McLerran and R.D. Pisarski, *Phases of cold, dense quarks at large N_c* , [Nucl. Phys. A](#) **796** (2007) 83 [[arXiv:0706.2191](#)] [[INSPIRE](#)].
- [7] M.G. Alford, A. Schmitt, K. Rajagopal and T. Schäfer, *Color superconductivity in dense quark matter*, [Rev. Mod. Phys.](#) **80** (2008) 1455 [[arXiv:0709.4635](#)] [[INSPIRE](#)].
- [8] A. Motornenko, J. Steinheimer, V. Vovchenko, S. Schramm and H. Stoecker, *Equation of state for hot QCD and compact stars from a mean field approach*, [arXiv:1905.00866](#) [[INSPIRE](#)].
- [9] J. Hofmann, H. Stoecker, U.W. Heinz, W. Scheid and W. Greiner, *Possibility of Detecting Density Isomers in High Density Nuclear MACH Shock Waves*, [Phys. Rev. Lett.](#) **36** (1976) 88 [[INSPIRE](#)].
- [10] H. Stoecker and W. Greiner, *High-Energy Heavy Ion Collisions: Probing the Equation of State of Highly Excited Hadronic Matter*, [Phys. Rept.](#) **137** (1986) 277 [[INSPIRE](#)].
- [11] D.H. Rischke, Y. Pursun, J.A. Maruhn, H. Stoecker and W. Greiner, *The Phase transition to the quark-gluon plasma and its effects on hydrodynamic flow*, [Acta Phys. Hung. A](#) **1** (1995) 309 [[nucl-th/9505014](#)] [[INSPIRE](#)].
- [12] J. Brachmann et al., *Antiflow of nucleons at the softest point of the EoS*, [Phys. Rev. C](#) **61** (2000) 024909 [[nucl-th/9908010](#)] [[INSPIRE](#)].
- [13] J. Brachmann, A. Dumitru, H. Stoecker and W. Greiner, *The Directed flow maximum near $c(s) = 0$* , [Eur. Phys. J. A](#) **8** (2000) 549 [[nucl-th/9912014](#)] [[INSPIRE](#)].
- [14] L.P. Csernai and D. Rohrlich, *Third flow component as QGP signal*, [Phys. Lett. B](#) **458** (1999) 454 [[nucl-th/9908034](#)] [[INSPIRE](#)].
- [15] Y.B. Ivanov, E.G. Nikonov, W. Noerenberg, A.A. Shanenko and V.D. Toneev, *Directed flow of baryons in heavy ion collisions*, [Acta Phys. Hung. A](#) **15** (2002) 117 [[nucl-th/0011004](#)] [[INSPIRE](#)].
- [16] H. Stoecker, *Collective flow signals the quark gluon plasma*, [Nucl. Phys. A](#) **750** (2005) 121 [[nucl-th/0406018](#)] [[INSPIRE](#)].
- [17] L.P. Csernai et al., *The 3rd flow component as a QGP signal*, [Acta Phys. Hung. A](#) **22** (2005) 181 [[hep-ph/0405277](#)] [[INSPIRE](#)].
- [18] Y. Nara, H. Niemi, J. Steinheimer and H. Stöcker, *Equation of state dependence of directed flow in a microscopic transport model*, [Phys. Lett. B](#) **769** (2017) 543 [[arXiv:1611.08023](#)] [[INSPIRE](#)].
- [19] Y. Nara, H. Niemi, A. Ohnishi, J. Steinheimer, X. Luo and H. Stöcker, *Enhancement of elliptic flow can signal a first order phase transition in high energy heavy ion collisions*, [Eur. Phys. J. A](#) **54** (2018) 18 [[arXiv:1708.05617](#)] [[INSPIRE](#)].
- [20] Y. Nara, J. Steinheimer and H. Stoecker, *The enhancement of v_4 in nuclear collisions at the highest densities signals a first-order phase transition*, [Eur. Phys. J. A](#) **54** (2018) 188 [[arXiv:1809.04237](#)] [[INSPIRE](#)].
- [21] K. Paech, H. Stoecker and A. Dumitru, *Hydrodynamics near a chiral critical point*, [Phys. Rev. C](#) **68** (2003) 044907 [[nucl-th/0302013](#)] [[INSPIRE](#)].

- [22] M.A. Stephanov, *QCD phase diagram and the critical point*, *Prog. Theor. Phys. Suppl.* **153** (2004) 139 [[hep-ph/0402115](#)] [[INSPIRE](#)].
- [23] M.A. Stephanov, *Non-Gaussian fluctuations near the QCD critical point*, *Phys. Rev. Lett.* **102** (2009) 032301 [[arXiv:0809.3450](#)] [[INSPIRE](#)].
- [24] M. Nahrgang, S. Leupold, C. Herold and M. Bleicher, *Nonequilibrium chiral fluid dynamics including dissipation and noise*, *Phys. Rev. C* **84** (2011) 024912 [[arXiv:1105.0622](#)] [[INSPIRE](#)].
- [25] C. Herold, M. Nahrgang, I. Mishustin and M. Bleicher, *Chiral fluid dynamics with explicit propagation of the Polyakov loop*, *Phys. Rev. C* **87** (2013) 014907 [[arXiv:1301.1214](#)] [[INSPIRE](#)].
- [26] C. Plumberg and J.I. Kapusta, *Hydrodynamic fluctuations near a critical endpoint and Hanbury-Brown-Twiss interferometry*, *Phys. Rev. C* **95** (2017) 044910 [[arXiv:1702.01368](#)] [[INSPIRE](#)].
- [27] F. Li and C.M. Ko, *Spinodal instabilities of baryon-rich quark-gluon plasma in the Polyakov-Nambu-Jona-Lasinio model*, *Phys. Rev. C* **93** (2016) 035205 [[arXiv:1601.00026](#)] [[INSPIRE](#)].
- [28] O. Scavenius, A. Dumitru, E.S. Fraga, J.T. Lenaghan and A.D. Jackson, *First order chiral phase transition in high-energy collisions: Can nucleation prevent spinodal decomposition?*, *Phys. Rev. D* **63** (2001) 116003 [[hep-ph/0009171](#)] [[INSPIRE](#)].
- [29] L.F. Palhares and E.S. Fraga, *Droplets in the cold and dense linear σ -model with quarks*, *Phys. Rev. D* **82** (2010) 125018 [[arXiv:1006.2357](#)] [[INSPIRE](#)].
- [30] C. Herold, M. Nahrgang, I. Mishustin and M. Bleicher, *Formation of droplets with high baryon density at the QCD phase transition in expanding matter*, *Nucl. Phys. A* **925** (2014) 14 [[arXiv:1304.5372](#)] [[INSPIRE](#)].
- [31] F. Li and C.M. Ko, *Spinodal instabilities of baryon-rich quark matter in heavy ion collisions*, *Phys. Rev. C* **95** (2017) 055203 [[arXiv:1606.05012](#)] [[INSPIRE](#)].
- [32] P. Chomaz, M. Colonna and J. Randrup, *Nuclear spinodal fragmentation*, *Phys. Rept.* **389** (2004) 263 [[INSPIRE](#)].
- [33] J. Randrup, *Spinodal decomposition during the hadronization stage at RHIC?*, *Phys. Rev. Lett.* **92** (2004) 122301 [[hep-ph/0308271](#)] [[INSPIRE](#)].
- [34] C. Sasaki, B. Friman and K. Redlich, *Density fluctuations in the presence of spinodal instabilities*, *Phys. Rev. Lett.* **99** (2007) 232301 [[hep-ph/0702254](#)] [[INSPIRE](#)].
- [35] J. Steinheimer and J. Randrup, *Spinodal amplification of density fluctuations in fluid-dynamical simulations of relativistic nuclear collisions*, *Phys. Rev. Lett.* **109** (2012) 212301 [[arXiv:1209.2462](#)] [[INSPIRE](#)].
- [36] K.-J. Sun, L.-W. Chen, C.M. Ko, J. Pu and Z. Xu, *Light nuclei production as a probe of the QCD phase diagram*, *Phys. Lett. B* **781** (2018) 499 [[arXiv:1801.09382](#)] [[INSPIRE](#)].
- [37] 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](#)].
- [38] J. Steinheimer and J. Randrup, *Spinodal density enhancements in simulations of relativistic nuclear collisions*, *Phys. Rev. C* **87** (2013) 054903 [[arXiv:1302.2956](#)] [[INSPIRE](#)].

- [39] J. Steinheimer, J. Randrup and V. Koch, *Non-equilibrium phase transition in relativistic nuclear collisions: Importance of the equation of state*, *Phys. Rev. C* **89** (2014) 034901 [[arXiv:1311.0999](#)] [[INSPIRE](#)].
- [40] J. Steinheimer and V. Koch, *Effect of finite particle number sampling on baryon number fluctuations*, *Phys. Rev. C* **96** (2017) 034907 [[arXiv:1705.08538](#)] [[INSPIRE](#)].
- [41] D.H. Rischke, S. Bernard and J.A. Maruhn, *Relativistic hydrodynamics for heavy ion collisions. 1. General aspects and expansion into vacuum*, *Nucl. Phys. A* **595** (1995) 346 [[nucl-th/9504018](#)] [[INSPIRE](#)].
- [42] S.A. Bass et al., *Microscopic models for ultrarelativistic heavy ion collisions*, *Prog. Part. Nucl. Phys.* **41** (1998) 255 [[nucl-th/9803035](#)] [[INSPIRE](#)].
- [43] 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](#)].
- [44] H. Petersen, J. Steinheimer, G. Burau, M. Bleicher and H. Stocker, *A Fully Integrated Transport Approach to Heavy Ion Reactions with an Intermediate Hydrodynamic Stage*, *Phys. Rev. C* **78** (2008) 044901 [[arXiv:0806.1695](#)] [[INSPIRE](#)].
- [45] F. Cooper and G. Frye, *Comment on the Single Particle Distribution in the Hydrodynamic and Statistical Thermodynamic Models of Multiparticle Production*, *Phys. Rev. D* **10** (1974) 186 [[INSPIRE](#)].
- [46] P. Huovinen and H. Petersen, *Particlization in hybrid models*, *Eur. Phys. J. A* **48** (2012) 171 [[arXiv:1206.3371](#)] [[INSPIRE](#)].
- [47] D. Oliinychenko and V. Koch, *Microcanonical Particlization with Local Conservation Laws*, *Phys. Rev. Lett.* **123** (2019) 182302 [[arXiv:1902.09775](#)] [[INSPIRE](#)].
- [48] J. Randrup, *Phase transition dynamics for baryon-dense matter*, *Phys. Rev. C* **79** (2009) 054911 [[arXiv:0903.4736](#)] [[INSPIRE](#)].
- [49] C.R. Qi, H. Su, K. Mo and L.J. Guibas, *PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation*, [arXiv:1612.00593](#).
- [50] K. Paech and A. Dumitru, *Density inhomogeneities in heavy ion collisions around the critical point*, *Phys. Lett. B* **623** (2005) 200 [[nucl-th/0504003](#)] [[INSPIRE](#)].
- [51] J. Randrup, *Probes of phase decomposition in high-energy nuclear collisions*, *Acta Phys. Hung. A* **22** (2005) 69 [[nucl-th/0406031](#)] [[INSPIRE](#)].
- [52] STAR collaboration, *Higher Moments of Net-proton Multiplicity Distributions at RHIC*, *Phys. Rev. Lett.* **105** (2010) 022302 [[arXiv:1004.4959](#)] [[INSPIRE](#)].
- [53] STAR collaboration, *Beam energy dependence of moments of the net-charge multiplicity distributions in Au+Au collisions at RHIC*, *Phys. Rev. Lett.* **113** (2014) 092301 [[arXiv:1402.1558](#)] [[INSPIRE](#)].
- [54] ALICE collaboration, *Event-by-Event Identified Particle Ratio Fluctuations in Pb-Pb Collisions with ALICE using the Identity Method*, *Nucl. Phys. A* **956** (2016) 870 [[arXiv:1512.03372](#)] [[INSPIRE](#)].
- [55] X. Luo and N. Xu, *Search for the QCD Critical Point with Fluctuations of Conserved Quantities in Relativistic Heavy-Ion Collisions at RHIC: An Overview*, *Nucl. Sci. Tech.* **28** (2017) 112 [[arXiv:1701.02105](#)] [[INSPIRE](#)].
- [56] A. Bzdak and V. Koch, *Mapping the QCD phase diagram with statistics friendly distributions*, *Phys. Rev. C* **100** (2019) 051902 [[arXiv:1811.04456](#)] [[INSPIRE](#)].

- [57] P. Mehta et al., *A high-bias, low-variance introduction to Machine Learning for physicists*, *Phys. Rept.* **810** (2019) 1 [[arXiv:1803.08823](#)] [[INSPIRE](#)].
- [58] T.P. Minka, *Automatic Choice of Dimensionality for PCA*, *NIPS* **13** (2000) 598.