

Toward a foundation model for heavy-ion collision experiments based on point-cloud diffusion

Manjunath Omana Kuttan^{1,2,*}, Kai Zhou^{1,3,†}, Jan Steinheimer^{1,4}, and Horst Stoecker^{1,4,5}

¹*Frankfurt Institute for Advanced Studies, Ruth-Moufang-Straße 1, D-60438 Frankfurt am Main, Germany*

²*Xidian-FIAS International Joint Research Center, Giersch Science Center, D-60438 Frankfurt am Main, Germany*

³*School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, People's Republic of China*

⁴*GSI Helmholtzzentrum für Schwerionenforschung GmbH, Planckstraße 1, D-64291 Darmstadt, Germany*

⁵*Institut für Theoretische Physik, Goethe Universität Frankfurt, Max-von-Laue-Straße 1, D-60438 Frankfurt am Main, Germany*



(Received 26 February 2025; revised 17 October 2025; accepted 22 October 2025; published 24 November 2025)

A novel point-cloud diffusion model for relativistic heavy-ion collisions, capable of ultrafast generation of complete, event-by-event collision output, is introduced. When trained on ultrarelativistic quantum molecular dynamics cascade simulations, the model generates realistic collision event output containing 26 distinct hadron species, as a list of particle momentum vectors along with their particle ID. From solving inverse problems to accelerating model calculations or detector simulations, the model can be a promising general-purpose tool for heavy-ion collisions beneficial to both theoretical studies and experimental applications.

DOI: [10.1103/6ndd-d1nl](https://doi.org/10.1103/6ndd-d1nl)

Relativistic nucleus-nucleus collisions create short-lived, strongly interacting systems characterized by high temperatures and/or densities, providing a unique terrestrial experimental probe of the conjectured phase diagram of quantum chromodynamics (QCD) [1–4]. Although several observables [1,5–9] have been shown to be sensitive to the properties and dynamics of the matter created in these collisions, essential details of the phase diagram, including the exact locations of phase boundaries, the nature of the transitions between them, and the presence of any possible critical points, remain poorly understood. Currently, first-principles lattice QCD calculations [10–15] are possible only at vanishing and small net-baryon densities, necessitating QCD-inspired effective models alongside direct experimental searches to accurately map the regions of high net-baryon density of the QCD phase diagram.

Numerous experimental programs are currently dedicated to investigate QCD matter at high baryon density [16–23]. These experiments will generate copious amounts of high-precision data at beam energies of $\sqrt{s_{NN}} \leq 15$ GeV. This wealth of data offers a unique opportunity to investigate baryon-rich QCD matter in previously uncharted ways, employing rare and novel probes.

Significant progress has also been made in modeling moderate- to intermediate-energy collisions with various

approaches emerging. Transport models [24–29] provide a microscopic, nonequilibrium description of the collisions, while hybrid models [30] incorporate a hydrodynamic description for the hot, dense, intermediate phase that depends on the unknown high-density and -temperature equation of state (EOS). Extensive efforts have also been made to implement realistic density- and momentum-dependent EOS within transport descriptions [31–35].

Despite the substantial advancements in realistic descriptions of heavy-ion collisions, these models suffer from their high computational expense. Even the fastest simulation models run orders of magnitude slower than the rate of experimental data collection. This computational bottleneck limits our ability to perform modern model-data analysis and to fully capitalize on the data to be produced by these next-generation experiments.

Extracting fundamental properties of the system such as the EOS or the location of a possible critical point in the phase diagram from the experimental data would require comparisons to model calculations for multiple observables such as collective flow, fluctuations, and correlations. For multiparameter fits or Bayesian analysis, fast emulators are developed using machine learning techniques to quickly map model parameters to relevant observables [35–41]. The main shortcoming of these approaches is that a new emulator needs to be trained for every new observable included in the inference. This would quickly become unfeasible when a large number of observables such as multidifferential flow and n -body correlation spectra are to be included in the analysis.

In addition to model simulations, experiments also rely heavily on expensive detector simulations for physics analyses, detector calibration, efficiency correction, uncertainty estimation, etc. Frameworks like GEANT4 [42], which are often used to model various experimental effects, can be substantially slower than experimental data rates, limiting

*Contact author: manjunath@fias.uni-frankfurt.de

†Contact author: zhoukai@cuhk.edu.cn

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

their applicability in real-time analyses, online event selection, and event characterization. As a result, the computational demands of detector simulations present an additional bottleneck, which poses a significant challenge to fully exploiting experimental data.

Deep learning (DL) methods have emerged as a novel approach for fast and accurate analysis of experimental data in high-energy heavy-ion collisions (HICs) [43–55]. However, these models are often designed only for a specific task (e.g., predicting the flow coefficients) within a given experiment, limiting their flexibility as a versatile artificial intelligence (AI) tool. A more general approach can be considered using deep generative models. As a class of machine learning algorithms capable of learning the underlying distribution of the training data, generative models can produce new data samples, offering a flexible alternative to DL models that perform only a specific task.

A recent work [56] applied generative models to generate p_T - ϕ spectra (as a 64×64 pixel image) of charged particles produced in collisions. Similarly, in Ref. [57], it was shown that a simulation chain with an event generator and detector simulation of the sPHENIX [58] experiment can be replaced using a generative model that generates the final calorimeter response. In Ref. [59], a deep learning model was trained to predict the final energy density and flow velocity profiles in relativistic hydrodynamic simulations. Although these methods are promising, they are currently limited to producing partial information in the form of histograms of the event, lacking the capability to generate complete event-level data, i.e., a complete list of all hadrons emitted from a given event, including information on their 4-momenta and hadron type. Generative models that generate complete event-level data will be able to emulate an entire simulation chain for an experiment with remarkable speed, offering pathways to build a robust and adaptable AI analysis framework for both heavy-ion collision experiments and theory.

In this work, we introduce, for the first time, an event-by-event generative model for heavy-ion collisions, which generates collision events as lists of particle vectors, named HEIDI (Heavy-ion Events through Intelligent Diffusion). With its strong ability to generate complete event information in a flexible representation, this model lays the groundwork for a future foundation model for HICs. Foundation models are general-purpose DL models that can be easily adapted for different tasks without requiring extensive retraining. Furthermore, the model's adaptability extends beyond HICs, offering potential applications in accelerating cascade simulations of cosmic ray showers or detector simulations in particle and astroparticle physics.

In HEIDI, we employ a conditional diffusion probabilistic model based on Ref. [60], to generate a point cloud of final-state particles in a collision event. Diffusion probabilistic models [61,62] are generative models that are inspired by nonequilibrium thermodynamics. It comprises a forward diffusion process and a reverse antidiffusion process. During the forward diffusion process, the point cloud of final-state particles evolves stochastically over several time steps until it becomes indistinguishable from pure noise, akin to a nonequilibrium system in contact with a heat bath. The generation

process is then achieved by learning to reverse this diffusion process. Starting from random samples from a simple well-defined distribution, the model iteratively predicts the noise added in the previous step and then denoises it, ultimately transforming it into a clean, structured point cloud.

Diffusion-based generative models have been applied in high-energy physics for various tasks [63–76]. In particle physics, such point-cloud generative models have primarily been used for simulating jets or calorimeter responses and typically involve small point clouds with limited numbers of particle types. These models are mainly intended for tasks such as performance benchmarking, detector calibration, or algorithm development. Our work addresses a more complex problem of simulating complete heavy-ion events containing over a thousand particles and dozens of hadron species with the additional but important goal of sophisticated downstream tasks such as addressing inverse problems and online event analyses. Given that heavy-ion physics strongly rely on model-to-data comparisons, an easily adaptable and fast, full event generation framework will significantly enhance the physics potential of current and future experiments.

The present work realizes a point-cloud diffusion model for heavy-ion collisions by implementing a deep generative emulator of the microscopic ultrarelativistic quantum molecular dynamics (UrQMD) cascade model [24,25]. UrQMD is widely used in the description of nucleus-nucleus, hadron-nucleus, and hadron-hadron collisions over a wide range of energies. It is based on the covariant propagation of hadrons and provides an effective solution of the relativistic Boltzmann equation, with n -body correlations.

In this work, as proof of concept, the model was trained to generate Au-Au collision events at 10A GeV using 18 000 UrQMD events, each with impact parameter $b = 1$ fm. Each event output generated by UrQMD is represented as a point cloud $\mathbf{X}^{(0)} = \{\mathbf{x}_i^{(0)}\}_{i=1}^{1084}$, where each point is a particle described by its momentum vector and particle information (ID), i.e., $\mathbf{x}_i^{(0)} = \{\mathbf{p}_i^{(0)}, \text{ID}_i^{(0)}\}$, where $\mathbf{p}_i^{(0)} = (p_{x_i}^{(0)}, p_{y_i}^{(0)}, p_{z_i}^{(0)})$. The number of points in an event is fixed to be 1084, a number larger than the highest multiplicity event in the training data, and events with fewer particles are padded with zeros. The superscript indicates the time step in the diffusion process. The particle IDs are one-hot encoded and the diffusion model is trained to generate 26 distinct types of hadrons that make up more than 99% of particles produced at this energy.

To capture the various correlations present in the training data, the diffusion model takes a latent vector z as an additional input. A PointNet-based encoder and a normalizing flow-based decoder [77], trained end-to-end with the diffusion model, learn to map each UrQMD event into a latent vector z that encodes various global features and correlations of the event. During generation, the normalizing flow decoder produces z , which becomes the input condition for the diffusion model. This conditioning facilitates the generation of physically realistic point clouds that preserve essential multiparticle correlations and global event properties that would otherwise be absent, since the diffusion model generates particles independently to maintain the permutation invariance of the point-cloud representation.

The forward diffusion process is realized through a Markov chain that progressively adds noise to each point in the point clouds over several time steps until the data are transformed into complete noise. The conditional probability distribution $q(\mathbf{x}_i^{(1:T)}|\mathbf{x}_i^{(0)})$ for a sequence of states $\mathbf{x}_i^{(1:T)} = \{\mathbf{x}_i^{(1)}, \mathbf{x}_i^{(2)}, \dots, \mathbf{x}_i^{(T)}\}$ generated through the diffusion steps starting from the initial point $\mathbf{x}_i^{(0)}$ in the point cloud is modeled as the product of probabilities of all consecutive diffusion steps:

$$q(\mathbf{x}_i^{(1:T)}|\mathbf{x}_i^{(0)}) = \prod_{t=1}^T q(\mathbf{x}_i^{(t)}|\mathbf{x}_i^{(t-1)}), \quad (1)$$

where $q(\mathbf{x}_i^{(t)}|\mathbf{x}_i^{(t-1)}) = \mathcal{N}(\mathbf{x}_i^{(t)}|\sqrt{1-\beta_t}\mathbf{x}_i^{(t-1)}, \beta_t\mathbf{I})$ is the diffusion kernel which introduces noise to the point in the previous time step, and is modeled as a Gaussian with a mean of $\sqrt{1-\beta_t}\mathbf{x}_i^{(t-1)}$ and variance of $\beta_t\mathbf{I}$. Here β_t is a hyperparameter controlling the amount of noise added at each time step and $\mathbf{I}$ is the identity matrix with dimensions matching the particle vector $\mathbf{x}_i$.

The process of generating point clouds involves a reverse Markov chain, which takes as input the latent vector z generated by the normalizing flow model, along with random samples from a simple noise distribution $\tilde{q}(\mathbf{x}_i^{(T)})$ which approximates $q(\mathbf{x}_i^{(T)})$. The Markov chain then sequentially denoises the random samples to generate states $\mathbf{x}_i^{(T-1)}, \mathbf{x}_i^{(T-2)}, \dots, \mathbf{x}_i^{(0)}$, ultimately recovering the original point cloud $\mathbf{x}_i^{(0)}$. The probability $\tilde{q}_\theta(\mathbf{x}_i^{(0:T)}|z)$ for a sequence of states $\mathbf{x}_i^{(0:T)} = \{\mathbf{x}_i^{(0)}, \mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(T)}\}$, given a latent vector z , is then given by the product of the probabilities of all consecutive reverse diffusion steps and the initial noise distribution $\tilde{q}(\mathbf{x}_i^{(T)})$:

$$\tilde{q}_\theta(\mathbf{x}_i^{(0:T)}|z) = \tilde{q}(\mathbf{x}_i^{(T)}) \prod_{t=1}^T \tilde{q}_\theta(\mathbf{x}_i^{(t-1)}|\mathbf{x}_i^{(t)}, z). \quad (2)$$

The reverse diffusion kernel $\tilde{q}_\theta(\mathbf{x}_i^{(t-1)}|\mathbf{x}_i^{(t)}, z)$ is parametrized by a neural network with parameters θ . This network is trained to learn the noise added in the previous time step and to remove it at each step in the reverse diffusion process. The reverse diffusion kernel is thus also modeled as a Gaussian distribution given by

$$\tilde{q}_\theta(\mathbf{x}_i^{(t-1)}|\mathbf{x}_i^{(t)}, z) = \mathcal{N}(\mathbf{x}_i^{(t-1)}|\mu_\theta(\mathbf{x}_i^{(t)}, t, z), \beta_t\mathbf{I}), \quad (3)$$

where μ_θ is the denoised mean as predicted by the network based on the current state $\mathbf{x}_i^{(t)}$, the time step t , and the latent vector z . This event generation process in HEIDI is illustrated in Fig. 1. For more details on HEIDI's network structure, see Ref. [78].

After training, 50 000 events were generated using both UrQMD and HEIDI to evaluate the performance of the generative model. These tests aim to quantify the model's ability in capturing the underlying physics of the HIC events and produce realistic point clouds that are consistent with the features of UrQMD-generated events.

HEIDI generates a list of particle vectors with 26 different hadrons species, representing the complete output of a collision event. The mean multiplicities of various hadrons




FIG. 1. Visualization of the generation process in HEIDI. Starting from random initial values for the momentum and ID at time step $t = 300$, the reverse diffusion process progressively denoises the samples, ultimately producing realistic particles at $t = 0$. The absolute deviation for the p_x - p_z probability densities of HEIDI from those of a random Gaussian probability density for various time steps is shown. Note that HEIDI generates particles, not distributions. The distributions presented are constructed from the particles in 2000 HEIDI events.

from events generated by HEIDI are compared to those from UrQMD in Fig. 2. Evidently, the model accurately learns the differences between the multiplicities of various hadron types in the training data.




FIG. 2. Mean event multiplicity of various hadrons for Au-Au collisions at 10A GeV with $b = 1$ fm. The results from the diffusion model are shown as orange bars, while the blue bars correspond to multiplicities from the UrQMD cascade model. The HEIDI model can also generate $\bar{p}$, $\bar{n}$, $\bar{\Lambda}$, $\bar{\Sigma}^-$, $\bar{\Sigma}^0$, $\bar{\Sigma}^+$, $\bar{\Xi}^0$, $\bar{\Xi}^+$, $\bar{\Omega}^+$, and $\bar{\Omega}^-$ in an event. However, these hadrons have been excluded from the plot due to their very small multiplicities ($< 1/\text{event}$). The particles marked 0 are the fake particles used to maintain a constant multiplicity to the event while n' and p' are spectator neutrons and protons, respectively. Inset: Multiplicity distributions of selected hadrons at mid-rapidity ($|y| < 0.1$). The results from the diffusion model are shown as solid lines, while the dashed lines correspond to multiplicities from the UrQMD cascade model.




FIG. 3. Rapidity distributions of selected hadrons for Au-Au collisions at 10A GeV with $b = 1$ fm. The solid lines show the diffusion model results while the dotted lines denote UrQMD results.

The multiplicity distributions of the selected hadrons at mid-rapidity are shown in Fig. 2 (inset). Although HEIDI generates events with 26 distinct hadron types, for clarity and conciseness, in Fig. 2 (inset) and in subsequent plots, we show the results only for protons (p), charged pions (π^- , π^+), kaons (K^+), and lambdas (Λ^0), which represent the most abundant hadrons in the system. It shows that the model, besides learning to reproduce mean multiplicities, can also effectively capture the probability distributions of multiplicities for various hadron species produced in an event. From less abundant particles, such as Ξ or Σ , with multiplicities of ≈ 1 –3 per event, to the most abundant particles, like protons and pions, HEIDI accurately reproduces the UrQMD multiplicity distributions. Figure 3 compares the rapidity distributions of the selected hadrons in the generated events with the corresponding distributions from UrQMD. The rapidity distributions of particles generated from HEIDI are in good agreement with those simulated from the UrQMD for all hadron types. For high-multiplicity particles like nucleons and pions, a small deviation appears around rapidities close to zero, where the generative model slightly overestimates the number of particles. This deviation is primarily due to a slight excess of the very low-momentum particles in the data generated by the diffusion model.

This effect is also noticeable in the transverse momentum (p_T) distributions at mid-rapidity ($|y| < 0.1$), as shown in Fig. 4. The particles generated by the generative model closely reproduces the p_T distribution across most momentum ranges, but at lower p_T values, $p_T < 400$ MeV, HEIDI generates an excess of low-momentum particles over the UrQMD model. However, this large deviation at small p_T values results only in a small difference to the total yield. This discrepancy suggests that the model does not fully capture the correlations in the very low-momentum region of phase space. While this issue could be mitigated using a more diverse training dataset, such as minimum-bias data, it is of less concern for experimental applications. Detectors typically have limited acceptance and




FIG. 4. Transverse momentum distributions of selected hadrons at mid-rapidity ($|y| < 0.1$) for Au-Au collisions at 10A GeV with $b = 1$ fm. The results of the generative model are shown as solid lines while the dotted lines denote UrQMD results. The generative model successfully reproduces the UrQMD p_T spectra across most momentum ranges, with the exception of the deviation at very low p_T values.

efficiency for low-momentum particles [79,80]. Hence, a low-momentum cutoff is usually applied to the simulation data for analyses [81,82].

Overall, the present model does an impressive job in learning the probability distributions of various hadrons spanning four orders of magnitude, from abundant lower-momentum particles to the exponentially less likely high-momentum particles. To achieve this level of accuracy with a limited training dataset in point-cloud format is not trivial.

In addition to single-particle-specific quantities like momenta and multiplicities of various hadrons, it is also important to check how well global, event-by-event quantities like total energy (E), total baryon number (B), and total charge (Q), as well as dynamic correlations like collective flow, are learned by the generative model. These results are presented in Fig. 5, where the distributions of the E , B , and Q values of all generated hadrons at mid-rapidity per event are compared




FIG. 5. Global features of generated events. The distributions of the total energy, baryon number, and charge of all mid-rapidity particles ($|y| < 0.1$) are shown in the left-hand, middle, and right-hand panels, respectively. Results from the generative model are shown in red, while the corresponding UrQMD distributions are shown in blue.

to those obtained from UrQMD calculations. Notably, such global features are also well captured by the diffusion model. The AI-generated events and UrQMD-generated events have very similar distributions for E , B , and Q . Furthermore, the generative model also accurately reproduces the vanishing elliptic flow for protons and pions, consistent with the UrQMD results at mid-rapidity for central, $b = 1$ fm collisions at 10A GeV. The model provides sufficiently accurate outputs for most event-averaged observables. However, as evident in Fig. 5, observables sensitive to strict conservation laws such as event-by-event fluctuation measurements may require further refinements to the model or additional conditioning strategies. For a detailed analysis of HEIDI's performance, see Ref. [78].

While accurately capturing the physics and correlations observed in UrQMD is essential, a primary motivation for implementing deep learning models is its significant speedup, particularly when the model is deployed on GPUs. On an Nvidia A100 GPU with 40 GB of memory, the model generates an event in about 30 ms, whereas the UrQMD model takes about 3 s to generate a single event. This represents a speedup of two orders of magnitude. It is important to note that here, UrQMD in cascade mode, the fastest version, was used for demonstration. The inclusion of potential interactions in UrQMD will increase the computational time to 1 minute, and the hybrid version of UrQMD, which includes an intermediate hydrodynamic stage, can take up to 1 hour per event at SIS-100 energies. When trained on such data, the deep learning model will deliver a speedup of at least five orders of magnitude.

The capability of HEIDI to not only distinguish a wide range of hadron species but also to learn well their multiplicity and momentum distributions, while correctly capturing various global features of the event, marks a significant advancement in generative modeling of heavy-ion collisions. Nevertheless, practical use in downstream tasks such as Bayesian inference, detector-level simulations, and data-driven physics analyses requires the model to generalize across a broader set of conditions, including varying centralities, beam energies, collision systems, EOS, and various detector effects. The network architecture [78], the point-cloud data format, and the limited training dataset were chosen to support this adaptability.

Extending HEIDI to incorporate and generalize across various conditions or physical parameters is straightforward. Any additional conditioning variable can be introduced to the model by appending it to the latent vector z , which is generated by the normalizing flow-based decoder. A proof-of-principle demonstration of this adaptability, without requiring any changes to the network architecture, is presented in the Supplemental Material [83]. This highlights that extending the model to other physical parameters is straightforward, requiring only the inclusion of those parameters in the conditioning vector and appropriate training data.

Such an ultrafast point-cloud diffusion model for heavy-ion collisions could eliminate the need for training specific surrogate models to predict individual observables, as the model is capable of generating events at high speed and directly

calculating any observable. This will make it an indispensable tool for comprehensive Bayesian analysis. Furthermore, a deep generative model can serve as a differentiable version of the physics model which can be directly used for parameter estimation tasks or to perform comprehensive global fits directly on raw, particle-level experimental data. This opens the door to novel gradient-based inference methods in heavy-ion physics for the first time. This approach eliminates the need to train deep learning models for specific tasks or inputs and improves interpretability, as the result can then be directly verified on the original model.

The point-cloud structure of electronically collected data in sensors also offers exciting potential for point-cloud diffusion models in biomedical imaging and diagnostics, natural hazard early warning systems, autonomous navigation, robotics, etc. In high-energy physics, the presented model is flexible to be adapted for any event-by-event model calculations, experimental simulations, or a combination of both. It is more precise and avoids the loss of information typically associated with histograms or image-based representations. Owing to the compact representation of point clouds, these models can also be more resource efficient as compared to generating images or voxel grids. Both the theoretical and experimental communities can rely on point-cloud diffusion models to quickly generate large numbers of events for initial-level analysis, the results of which can later be verified using actual physics models.

Even though the presented results are highly promising, extensive tests on the reliability and consistency of the model are still necessary before HEIDI can be applied to physics analyses. In particular, the sources of the deviations observed at small transverse momentum and rapidities need to be better understood, and further developments are required to mitigate these discrepancies. Moreover, it would be valuable to compare and quantitatively benchmark the performance of HEIDI against other point-cloud-based generative models developed for particle physics tasks. Such comparisons could help identify the limitations of PointNet-based architectures in capturing complex higher-order correlations in the data. However, these efforts lie beyond the scope of the present work and represent important directions for future research. While not yet a complete foundation model, the presented HEIDI framework represents a key step toward building one for heavy-ion physics through a fast, general-purpose, and extensible full event-by-event generation method adaptable for a wide range of downstream tasks.

Acknowledgments. This work is supported by the the BMBF under the KISS project (M.O.K., K.Z.), the SAMSON AG (J.S., K.Z.), the CUHK-Shenzhen university development fund under Grants No. UDF01003041 and No. UDF03003041 (K.Z.), the Shenzhen Peacock fund under Grant No. 2023TC0179 (K.Z.), and the Walter Greiner Gesellschaft zur Förderung der physikalischen Grundlagenforschung e.V. through the Judah M. Eisenberg Laureatus Chair at the Goethe Universität Frankfurt am Main (H.S.).

Data availability. The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

- [1] H. Stoecker and W. Greiner, *Phys. Rep.* **137**, 277 (1986).
- [2] K. Fukushima and T. Hatsuda, *Rep. Prog. Phys.* **74**, 014001 (2011).
- [3] K. Rajagopal and F. Wilczek, in *At the Frontier of Particle Physics: Handbook of QCD*, edited by M. Shifman and B. Ioffe (World Scientific, Singapore, 2000), Vol. 1–3, pp. 2061–2151.
- [4] M. A. Halasz, A. D. Jackson, R. E. Shrock, M. A. Stephanov, and J. J. M. Verbaarschot, *Phys. Rev. D* **58**, 096007 (1998).
- [5] J. Hofmann, H. Stoecker, U. W. Heinz, W. Scheid, and W. Greiner, *Phys. Rev. Lett.* **36**, 88 (1976).
- [6] H. Stoecker, *Nucl. Phys. A* **750**, 121 (2005).
- [7] M. A. Stephanov, K. Rajagopal, and E. V. Shuryak, *Phys. Rev. Lett.* **81**, 4816 (1998).
- [8] Y. Hatta and M. A. Stephanov, *Phys. Rev. Lett.* **91**, 102003 (2003); **91**, 129901(E) (2003).
- [9] K. Zhou, N. Xu, Z. Xu, and P. Zhuang, *Phys. Rev. C* **89**, 054911 (2014).
- [10] Y. Aoki, G. Endrodi, Z. Fodor, S. D. Katz, and K. K. Szabo, *Nature (London)* **443**, 675 (2006).
- [11] S. Borsanyi, G. Endrodi, Z. Fodor, A. Jakovac, S. D. Katz, S. Krieg, C. Ratti, and K. K. Szabo, *J. High Energy Phys.* **11** (2010) 077.
- [12] A. Bazavov *et al.* (HotQCD Collaboration), *Phys. Rev. D* **90**, 094503 (2014).
- [13] M. Cheng *et al.*, *Phys. Rev. D* **77**, 014511 (2008).
- [14] A. Bazavov *et al.*, *Phys. Rev. D* **85**, 054503 (2012).
- [15] S. Aoki *et al.* (PACS-CS Collaboration), *Phys. Rev. D* **79**, 034503 (2009).
- [16] G. Agakishiev *et al.* (HADES Collaboration), *Eur. Phys. J. A* **41**, 243 (2009).
- [17] C. Yang (STAR Collaboration), *Nucl. Phys. A* **967**, 800 (2017).
- [18] M. M. Aggarwal *et al.* (STAR Collaboration), *arXiv:1007.2613*.
- [19] L. Adamczyk *et al.* (STAR Collaboration), *Phys. Rev. C* **96**, 044904 (2017).
- [20] K. C. Meehan (STAR Collaboration), *Nucl. Phys. A* **956**, 878 (2016).
- [21] T. Ablyazimov *et al.* (CBM Collaboration), *Eur. Phys. J. A* **53**, 60 (2017).
- [22] J. C. Yang *et al.*, *Nucl. Instrum. Methods Phys. Res. B* **317**, 263 (2013).
- [23] V. Golovatyuk, V. Kekelidze, V. Kolesnikov, O. Rogachevsky, and A. Sorin, *Eur. Phys. J. A* **52**, 212 (2016).
- [24] S. A. Bass *et al.*, *Prog. Part. Nucl. Phys.* **41**, 255 (1998).
- [25] M. Bleicher *et al.*, *J. Phys. G: Nucl. Part. Phys.* **25**, 1859 (1999).
- [26] W. Cassing and E. L. Bratkovskaya, *Nucl. Phys. A* **831**, 215 (2009).
- [27] Y. Nara, N. Otuka, A. Ohnishi, K. Niita, and S. Chiba, *Phys. Rev. C* **61**, 024901 (1999).
- [28] O. Buss, T. Gaitanos, K. Gallmeister, H. van Hees, M. Kaskulov, O. Lalakulich, A. B. Larionov, T. Leitner, J. Weil, and U. Mosel, *Phys. Rep.* **512**, 1 (2012).
- [29] Z.-W. Lin, C. M. Ko, B.-A. Li, B. Zhang, and S. Pal, *Phys. Rev. C* **72**, 064901 (2005).
- [30] H. Petersen, J. Steinheimer, G. Burau, M. Bleicher, and H. Stöcker, *Phys. Rev. C* **78**, 044901 (2008).
- [31] M. Omana Kuttan, A. Motornenko, J. Steinheimer, H. Stoecker, Y. Nara, and M. Bleicher, *Eur. Phys. J. C* **82**, 427 (2022).
- [32] J. Steinheimer, M. Omana Kuttan, A. Motornenko, A. Sorensen, Y. Nara, V. Koch, M. Bleicher, and H. Stoecker, *EPJ Web Conf.* **276**, 01021 (2023).
- [33] J. Steinheimer, T. Reichert, Y. Nara, and M. Bleicher, *J. Phys. G: Nucl. Part. Phys.* **52**, 035103 (2025).
- [34] Y. Nara, T. Maruyama, and H. Stoecker, *Phys. Rev. C* **102**, 024913 (2020).
- [35] D. Oliinychenko, A. Sorensen, V. Koch, and L. McLerran, *Phys. Rev. C* **108**, 034908 (2023).
- [36] J. E. Bernhard, J. S. Moreland, S. A. Bass, J. Liu, and U. Heinz, *Phys. Rev. C* **94**, 024907 (2016).
- [37] D. Everett *et al.* (JETSCAPE Collaboration), *Phys. Rev. C* **103**, 054904 (2021).
- [38] S. Pratt, E. Sangaline, P. Sorensen, and H. Wang, *Phys. Rev. Lett.* **114**, 202301 (2015).
- [39] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, *Phys. Rev. C* **103**, 054909 (2021).
- [40] M. Omana Kuttan, J. Steinheimer, K. Zhou, and H. Stoecker, *Phys. Rev. Lett.* **131**, 202303 (2023).
- [41] S. Huth *et al.*, *Nature (London)* **606**, 276 (2022).
- [42] S. Agostinelli *et al.* (GEANT4 Collaboration), *Nucl. Instrum. Methods Phys. Res. A* **506**, 250 (2003).
- [43] K. Zhou, L. Wang, L.-G. Pang, and S. Shi, *Prog. Part. Nucl. Phys.* **135**, 104084 (2024).
- [44] M. Omana Kuttan, J. Steinheimer, K. Zhou, A. Redelbach, and H. Stoecker, *Phys. Lett. B* **811**, 135872 (2020).
- [45] M. Omana Kuttan, K. Zhou, J. Steinheimer, A. Redelbach, and H. Stoecker, *J. High Energy Phys.* **10** (2021) 184.
- [46] M. Omana Kuttan, J. Steinheimer, K. Zhou, A. Redelbach, and H. Stoecker, *Proc. Sci. FAIRness2022*, 040 (2023).
- [47] L.-G. Pang, K. Zhou, N. Su, H. Petersen, H. Stöcker, and X.-N. Wang, *Nat. Commun.* **9**, 210 (2018).
- [48] Y.-L. Du, K. Zhou, J. Steinheimer, L.-G. Pang, A. Motornenko, H.-S. Zong, X.-N. Wang, and H. Stöcker, *Eur. Phys. J. C* **80**, 516 (2020).
- [49] Y.-L. Du, D. Pablos, and K. Tywoniuk, *J. High Energy Phys.* **03** (2021) 206.
- [50] W.-B. He, Y.-G. Ma, L.-G. Pang, H.-C. Song, and K. Zhou, *Nucl. Sci. Tech.* **34**, 88 (2023).
- [51] F. Li, Y. Wang, H. Lü, P. Li, Q. Li, and F. Liu, *J. Phys. G* **47**, 115104 (2020).
- [52] N. Mallick, S. Prasad, A. N. Mishra, R. Sahoo, and G. G. Barnaföldi, *Phys. Rev. D* **105**, 114022 (2022).
- [53] Y.-D. Song, R. Wang, Y.-G. Ma, X.-G. Deng, and H.-L. Liu, *Phys. Lett. B* **814**, 136084 (2021).
- [54] H. Hirvonen, K. J. Eskola, and H. Niemi, *Phys. Rev. C* **108**, 034905 (2023).
- [55] F. Sergeev, E. Bratkovskaya, I. Kisel, and I. Vassiliev, *Int. J. Mod. Phys. A* **35**, 2043002 (2020).
- [56] J.-A. Sun, L. Yan, C. Gale, and S. Jeon, *arXiv:2410.13069*.
- [57] D. Torbunov, Y. Huang, M. Lin, Y. Ren, Y. Go, T. Rinn, H. Yu, B. Viren, and J. Huang, *Phys. Rev. C* **110**, 034912 (2024).
- [58] C. Aidala *et al.*, *arXiv:1207.6378*.
- [59] H. Huang, B. Xiao, Z. Liu, Z. Wu, Y. Mu, and H. Song, *Phys. Rev. Res.* **3**, 023256 (2021).
- [60] S. Luo and W. Hu, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, New York, 2021), pp. 2837–2845.
- [61] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 10850 (2023).
- [62] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, W. Zhang, B. Cui, and M.-H. Yang, *ACM Comput. Surv.* **55**, 1 (2022).
- [63] L. Wang, G. Aarts, and K. Zhou, *J. High Energy Phys.* **05** (2024) 060.

- [64] V. Mikuni and B. Nachman, *Phys. Rev. D* **106**, 092009 (2022).
- [65] O. Amram and K. Pedro, *Phys. Rev. D* **108**, 072014 (2023).
- [66] F. T. Acosta, V. Mikuni, B. Nachman, M. Arratia, B. Karki, R. Milton, P. Karande, and A. Angerami, *J. Instrum.* **19**, P05003 (2024).
- [67] V. Mikuni and B. Nachman, *J. Instrum.* **19**, P02001 (2024).
- [68] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, W. Korcari, K. Krüger, and P. McKeown, *J. Instrum.* **18**, P11025 (2023).
- [69] A. Butter, N. Huetsch, S. Palacios Schweitzer, T. Plehn, P. Sorrenson, and J. Spinner, *SciPost Phys. Core* **8**, 026 (2025).
- [70] M. Leigh, D. Sengupta, J. A. Raine, G. Quétant, and T. Golling, *Phys. Rev. D* **109**, 012010 (2024).
- [71] E. Buhmann, C. Ewen, D. A. Faroughy, T. Golling, G. Kasieczka, M. Leigh, G. Quétant, J. A. Raine, D. Sengupta, and D. Shih, *arXiv:2310.00049*.
- [72] E. Buhmann, G. Kasieczka, and J. Thaler, *SciPost Phys.* **15**, 130 (2023).
- [73] V. Mikuni, B. Nachman, and M. Pettee, *Phys. Rev. D* **108**, 036025 (2023).
- [74] M. Leigh, D. Sengupta, G. Quétant, J. A. Raine, K. Zoch, and T. Golling, *SciPost Phys.* **16**, 018 (2024).
- [75] J. Y. Araz, V. Mikuni, F. Ringer, N. Sato, F. T. Acosta, and R. Whitehill, *Phys. Lett. B* **868**, 139694 (2025).
- [76] P. Devlin, J.-W. Qiu, F. Ringer, and N. Sato, *Phys. Rev. D* **110**, 016030 (2024).
- [77] I. Kobyzev, S. J. Prince, and M. A. Brubaker, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964 (2021).
- [78] M. Omana Kuttan, K. Zhou, J. Steinheimer, and H. Stöcker, companion paper, *Phys. Rev. C* **112**, 054907 (2025).
- [79] M. S. Abdallah *et al.* (STAR Collaboration), *Phys. Lett. B* **827**, 137003 (2022).
- [80] J. Adamczewski-Musch *et al.* (HADES Collaboration), *Phys. Rev. C* **102**, 024914 (2020).
- [81] T. Reichert, M. Omana Kuttan, A. Kittiratpattana, N. Buyukcizmeci, A. Botvina, J. Steinheimer, and M. Bleicher, *arXiv:2504.17389*.
- [82] P. Hillmann, J. Steinheimer, T. Reichert, V. Gaebel, M. Bleicher, S. Sombun, C. Herold, and A. Limphirat, *J. Phys. G: Nucl. Part. Phys.* **47**, 055101 (2020).
- [83] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/6nndd-d1nl> for a demonstration of how the model can be flexibly adapted to simulate collisions at different centralities.