Neural-network-enhanced Bayesian global analysis of relativistic heavy-ion collisions

Jussi Auvinen [1,2,3,*] Kari J. Eskola [2,3] Henry Hirvonen [4,5] and Harri Niemi [2,3]

[1] *Incubator of Scientific Excellence, Centre for Simulations of Superdense Fluids, University of Wrocław, plac Maksa Borna 9, PL-50204 Wrocław, Poland*

[2] *University of Jyväskylä, Department of Physics, P. O. Box 35, FI-40014 University of Jyväskylä, Finland*

[3] *Helsinki Institute of Physics, P. O. Box 64, FI-00014 University of Helsinki, Finland*

[4] *Department of Mathematics, Vanderbilt University, Nashville, Tennessee 37240, USA*

[5] *Department of Physics and Astronomy, Vanderbilt University, Nashville, Tennessee 37240, USA*

(Received 31 March 2026; accepted 27 May 2026; published 10 July 2026)

We introduce a novel deep convolutional neural-network-(NN) enhanced Bayesian global analysis of bulk observables in highest-energy heavy-ion collisions using relativistic 2 + 1D second-order viscous hydrodynamics with a dynamical freeze-out and with perturbative quantum chromodynamics (QCD) and saturation-based initial conditions from the event-by-event EKRT model. Our analysis has 13 + 2 free parameters for the QCD-matter properties + initial state, which are constrained by the experimental data from $\sqrt{s_{NN}} = 200$ GeV Au + Au collisions at Relativistic Heavy Ion Collider and 2.76 TeV Pb + Pb, 5.02 TeV Pb + Pb, and 5.44 TeV Xe + Xe collisions at the Large Hadron Collider. We replace the computationally demanding hydrodynamical simulations by NNs, which predict bulk observables directly from the initial energy density profiles, event by event, and account for the QCD-matter properties. With the NN output, we train the Gaussian process emulators for obtaining centrality-class averaged observables and their uncertainties. The NNs reduce the computing time significantly, enabling us to include also statistics-hungry flow observables like v_4 and the normalized symmetric cumulant NSC(4, 2) in the analysis. In this paper, we demonstrate the feasibility of the NN-based Bayesian global analysis. We find the data favoring a specific shear viscosity η/s with a minimum-value plateau at temperatures $150 \lesssim T \lesssim 230$ MeV, with $0.12 \lesssim (\eta/s)_{\min} \lesssim 0.18$. The bulk viscous coefficient ζ/s is nonzero at $200 \lesssim T \lesssim 300$ MeV. The Knudsen number at the freeze-out is 0.8–2.3, while the ratio of the mean free path to the system size at freeze-out is in the range 0.3–1.2, implying that the freeze-out indeed happens at the expected limit of the applicability of hydrodynamics.

DOI: 10.1103/13dy-4g64

I. INTRODUCTION

According to quantum chromodynamics (QCD), the strongly interacting matter at high temperatures ($T \gtrsim 150$ MeV) and vanishing baryochemical potential takes the form of the quark-gluon plasma (QGP), a fluid of deconfined quarks and gluons [1–4]. The properties of hot QCD matter can be currently studied in ultrarelativistic heavy-ion collisions at the CERN Large Hadron Collider (LHC) and at the BNL Relativistic Heavy Ion Collider (RHIC), where such extreme temperatures are reachable.

A cornerstone in modeling QCD matter evolution is dissipative relativistic fluid dynamics, "hydrodynamics," where the matter properties are embedded in the transport coefficients and QCD equation of state, which includes the crossover transition from the QGP to a hadron gas. Besides the matter properties, an essential part of the hydrodynamical models are the initial conditions that can be either parametrized, as done, e.g., in TRENTo [5], or computed from QCD-based initial-state models such as color glass condensate–based IP-Sat (Impact parameter dependent saturation) [6–8], the EKRT (Eskola-Kajantie-Ruuskanen-Tuominen) model [9–15], and EPOS (Energy conservation + Parallel scattering + factOrization + Saturation) [16–21]. The hydrodynamical evolution can be further supplemented by additional modeling for prehydrodynamic evolution such as KøMPøST [22] and ALPACA [23] and for late hadronic evolution that can be modeled using hadron cascade simulations, such as UrQMD [24,25] or SMASH [26]. Alternatively, one may describe the whole evolution using hydrodynamics, the approach taken in this work, but supplement it with dynamical freeze-out conditions and chemical freeze-out [27]. Hydrodynamics-based approaches have provided a consistent description of the measured bulk (low transverse momentum) observables and consequently a very strong support for

*Contact author: jussi.a.m.auvinen@jyu.fi

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International 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.

the formation of a hot collectively behaving QGP in these collisions [28].

The determination of the QCD matter properties, such as the temperature dependencies of the specific shear and bulk viscosities, $\eta/s(T)$ and $\zeta/s(T)$, from the LHC and RHIC data is a very active field of research currently. The effective model frameworks mentioned above, which are necessary for describing the full collision event, typically come altogether with a relatively large number of free parameters, $\mathcal{O}(10–20)$, which must be fitted from the experimental data; the matter properties, together with the unknowns in the initial-state dynamics and in the decoupling dynamics, are then expressed in terms of these parameters. Furthermore, many of the fit parameters are correlated in such a way that various different parameter combinations can produce similar effects on the final-state observables. For this reason—to understand these correlations and to have controlled error estimates—a Bayesian global analysis has become a standard tool in the determination of QCD matter properties and their uncertainties [29–47]. For a recent review, see Ref. [48].

One of the major challenges in these analyses is that typically the final experimental results are averages of millions of collision events, and a meaningful data comparison with the models can require a similar amount of simulated collision events. This is combined with the fact that the parameters for the QCD matter properties, together with those for the decoupling conditions and the initial state, form a multidimensional parameter space. The Bayesian analysis then requires that we can compute experimentally measured observables fast and easily at an arbitrary point in the parameter space. By a direct model computation this is clearly an impossible task as it would require a new set of at least tens of thousands of model runs, or even significantly more for rare observables, for each new parameter combination. Even with tools like Gaussian process (GP) emulators [49–51] the task is still computationally very demanding, as it may still require millions of model runs to produce sufficiently comprehensive training data for these emulators.

In this proof-of-principle study, we introduce a novel framework that can speed up this process potentially by several orders of magnitude. The main new feature in our approach is that instead of computing millions of hydrodynamical simulations, we use the hydrodynamical simulations to train deep convolutional neural networks (NN) that can predict the final experimental observables event by event (EbyE) directly from the initial energy density profile. We can then replace the time-consuming hydrodynamical simulations by NNs and use them to generate the training data for the Gaussian process emulators. While the training of the NNs still requires a considerable amount of real hydrodynamical simulations, it is orders of magnitude less than what would be needed for generating data directly to the training of GP emulators using hydrodynamics. The more simulated events are required to obtain meaningful results for a given observable, the more significant is the difference in the computational time. Besides our previous works [52,53], the use of generative neural networks as a fast way to emulate EbyE hydrodynamics simulations has recently been studied in Refs. [54–56]. However, to our knowledge, this is the first time that any of these methods have been applied in the context of Bayesian global analysis.

Our main physics goal here is an improved determination of the QCD matter viscosities $\eta/s(T)$ and $\zeta/s(T)$ and an understanding how these properties are correlated with the EKRT initial-state calculation and its uncertainties. Our NNs, introduced in Ref. [52], are now further developed to accept also the QCD matter parameters as input [53]. These new NNs enable a very fast EbyE computation of flow coefficients, average transverse momenta, and multiplicities. With the centrality-class averaged results, we then train the GP emulators, which are employed in the Bayesian analysis.

For the EbyE training of the NNs, we employ 2 + 1D relativistic shear and bulk viscous hydrodynamics from Ref. [27], with initial conditions computed from the perturbative QCD + saturation-based EbyE-EKRT model [13] and with the decoupling conditions determined dynamically as introduced in Ref. [27]. Instead of coupling our hydrodynamics to a hadron cascade afterburner, we run hydrodynamics in partial chemical equilibrium until the final kinetic freeze-out. In doing so, we avoid further assumptions for the various unknown hadronic cross sections and unphysical discontinuities, e.g., in the shear viscosity. Strong and electromagnetic hadron decays, however, we still account for after the freeze-out.

Our Bayesian inference setup has 15 input parameters in total: 10 for the $\eta/s(T)$ and the $\zeta/s(T)$, 1 for the chemical decoupling temperature, 2 for the dynamical freeze-out conditions, and 2 for the initial-state calculation. Exploring such a high-dimensional parameter space with full EbyE 2 + 1D hydrodynamic simulations is computationally expensive. Even when utilizing statistical model emulators such as Gaussian processes, full EbyE simulations still need to be run for $\mathcal{O}(1000)$ different parameter combinations, and for each of these the EbyE statistics required by the observables needs to be generated. With the NNs trained in these $\mathcal{O}(1000)$ parameter points, we can, however, easily generate such a large number of events that inclusion of more statistics-hungry flow observables in the analysis, such as here the Fourier coefficient v_4 and the normalized symmetric cumulant NSC(4, 2), becomes possible.

This paper is structured as follows: We describe the EbyE-EKRT + viscous hydrodynamics model in Sec. II and the neural-network framework for predicting model output in Sec. III. We go through the statistical analysis procedure in Sec. IV and present the analysis results in Sec. V. A summary of our findings can be found in Sec. VI.

II. MODEL DESCRIPTION

The EbyE-EKRT + viscous hydrodynamics model with dynamical freeze-out utilized in this study has been described in detail in Ref. [27]. Here we summarize the main features and parameters relevant to the current study.

A. EbyE-EKRT initial state

The initial energy density profiles of QCD matter are obtained from the EbyE-EKRT minijet saturation model [13], which has been developed from Refs. [9–12]. The model is

based on a collinearly factorized next-to-leading-order (NLO) perturbative QCD calculation of transverse energy (E_T) of minijets that are produced with transverse momenta (p_T) above a cutoff scale p_0 and rapidities in a midrapidity interval $\Delta y = 1$ [11,13,57–59]. Saturation of the minijet processes, emerging from QCD nonlinearities that become dominant at high parton densities [9,60], is conjectured to suppress the primary parton production towards lower values of p_T. For the transverse density of E_T in an AA collision of an impact parameter $\mathbf{b}$, the saturation limit is obtained from a local saturation condition [11],

$$\frac{dE_T}{d^2\mathbf{r}}(T_A(\mathbf{r}_1,\sigma_n), T_A(\mathbf{r}_2,\sigma_n), p_0, \sqrt{s_{NN}}, \Delta y, \beta) = \frac{K_{\rm sat}}{\pi} p_0^3 \Delta y, \tag{1}$$

where $\mathbf{r}$ is the transverse coordinate, $\mathbf{r}_{1,2} = \mathbf{r} \pm \mathbf{b}/2$, T_A's are the nuclear thickness functions, σ_n is the width parameter of the nucleon density distribution, $\sqrt{s_{NN}}$ is the nucleon-nucleon center-of-mass system collision energy, β is a parameter encoding the freedom in the infrared-collinear safe NLO definition of the E_T (see Refs. [11,13] for the details), and $K_{\rm sat}$ is a parameter controlling the onset of saturation.

The nuclear collision geometry enters the calculation via the nuclear overlap density,

$$\rho_{AA}(\mathbf{b}, \mathbf{r}, \sigma_n) = T_A(\mathbf{r}_1, \sigma_n) T_A(\mathbf{r}_2, \sigma_n), \tag{2}$$

and also via the T_A-dependent nuclear effects in the employed EPS09s nuclear parton distribution functions [61]. The T_As here fluctuate due to the EbyE fluctuating configurations of nucleons in the colliding nuclei. The nucleon positions $\mathbf{r}_{i1,2}$ around the centers of the nuclei are sampled from the Woods-Saxon densities, with thicknesses $d = 0.55$, 0.55, 0.54 fm and radii $R = 6.38$, 6.7, 5.49 fm for the Au, Pb, Xe nuclei, correspondingly, and for Xe with fixed deformation parameters $\beta_2 = 0.162$ and $\beta_4 = -0.003$ [62]. The T_As are obtained as sums of individual nucleon thickness functions, which are assumed to be Gaussians of a width σ_n,

$$T_A(\mathbf{r}_1, \sigma_n) = \sum_{i=1}^{A} \frac{1}{2\pi\sigma_n^2} \exp\left(-\frac{|\mathbf{r}_1 - \mathbf{r}_{i1}|^2}{2\sigma_n^2}\right), \tag{3}$$

and similarly for $T_A(\mathbf{r}_2, \sigma_n)$. We leave σ_n here as a free fit parameter to be determined from the experimental data.

The saturation scale that regulates the parton production at low p_T is obtained as a solution of Eq. (1),

$$p_{\rm sat} = p_{\rm sat}(\rho_{AA}(\mathbf{b}, \mathbf{r}; \sigma_n), \sqrt{s_{NN}}, \Delta y, \beta, K_{\rm sat}). \tag{4}$$

As noticed in Ref. [13], to a very good precision $p_{\rm sat}$ does not depend individually on $T_A(\mathbf{r}_1)$ and $T_A(\mathbf{r}_2)$ but effectively only on the value of the overlap density ρ_{AA} (as the above list of arguments to $p_{\rm sat}$ is to indicate). Thus, for fixed A and $\sqrt{s_{NN}}$, it has been possible to parametrize the ρ_{AA} dependence of $p_{\rm sat}$ with $K_{\rm sat}$- and β-dependent coefficients [13,63,64]. To minimize the number of free parameters in our Bayesian inference study, we fix $\beta = 0.8$ as in our earlier works [12,13,27,63,64] but leave $K_{\rm sat}$ as another free initial-state parameter to be constrained by the data.

Assuming the formation time of the minijet system to be $\tau_{\rm sat} = 1/p_{\rm sat}$, the EbyE-EKRT model gives the local energy density at the midrapidity unit, for each nucleon configuration, impact parameter and center-of-mass energy, and fixed β, as [13]

$$e(\mathbf{r}, \tau_{\rm sat}(\mathbf{r})) = \frac{K_{\rm sat}}{\pi} [p_{\rm sat}(\rho_{AA}(\mathbf{b}, \mathbf{r}; \sigma_n), K_{\rm sat})]^4, \tag{5}$$

where now the dependence on the fit parameters σ_n and $K_{\rm sat}$ is indicated. Concretely, for a fast computation of the EbyE fluctuating energy densities for the present Bayesian inference study, we employ the parametrizations of $p_{\rm sat}$ introduced in Refs. [13,63,64].

Following our earlier EbyE-EKRT studies, in the transverse coordinate region where $p_{\rm sat} \geqslant p_{\rm sat}^{\rm min} = 1$ GeV, we evolve the energy densities from $\tau_{\rm sat}(\mathbf{r})$ to $\tau_0 = 1/p_{\rm sat}^{\rm min} \approx 0.2$ fm using $0+1$D Bjorken hydrodynamics. In the region where $p_{\rm sat}(\mathbf{r}) < p_{\rm sat}^{\rm min}$, we connect the energy density profile at the time τ_0 smoothly to a binary collision profile ($\propto \rho_{AA}$); see again Ref. [13] for the details. We then use the obtained full energy density profile $e(\mathbf{r}, \tau_0)$, assuming no initial transverse flow, as the initial condition for our $2+1$D viscous hydrodynamics described in the next section.

For triggering the nuclear collisions here, we require that between the colliding two randomly sampled nucleon configurations the transverse distance of at least one pair of nucleons is less than $\sqrt{\sigma_{NN}^{\rm in}/\pi}$, where $\sigma_{NN}^{\rm in}$ is the inelastic nucleon-nucleon cross section. As in Ref. [27], we use here the values $\sigma_{NN}^{\rm in} = 42$, 64, 70, 72 mb for $\sqrt{s_{NN}} = 200$ GeV Au + Au, 2.76 TeV Pb + Pb, 5.02 TeV Pb + Pb, and 5.44 TeV Xe + Xe collisions, correspondingly.

B. 2 + 1D viscous hydrodynamics

The fluid dynamical equations describe local conservation of energy and momentum, $\partial_\mu T^{\mu\nu} = 0$. The energy-momentum tensor can be decomposed in the Landau frame as

$$T^{\mu\nu} = e u^\mu u^\nu - P\Delta^{\mu\nu} + \pi^{\mu\nu}, \tag{6}$$

where u^μ is the fluid 4-velocity, e is the local energy density, P is the isotropic pressure, and $\pi^{\mu\nu}$ is the shear-stress tensor. Here $P = -\frac{1}{3}\Delta_{\mu\nu}T^{\mu\nu}$, with $\Delta^{\mu\nu} = g^{\mu\nu} - u^\mu u^\nu$. The bulk viscous pressure is defined as $\Pi = P - P_0$, where P_0 is the equilibrium pressure provided by the equation of state (EoS), $P_0 = P_0(e)$, with conserved charges neglected. For the EoS we use the *s95p* parametrization [65] with a variable chemical freeze-out temperature $T_{\rm chem}$ [66].

The evolution of the dissipative quantities is given by second-order transient fluid dynamics [67–69]

$$\tau_\Pi \frac{d}{d\tau}\Pi + \Pi = -\zeta\theta - \delta_{\Pi\Pi}\Pi\theta + \lambda_{\Pi\pi}\pi^{\mu\nu}\sigma_{\mu\nu}, \tag{7}$$

$$\begin{aligned}\tau_\pi \frac{d}{d\tau}\pi^{\langle\mu\nu\rangle} + \pi^{\mu\nu} &= 2\eta\sigma^{\mu\nu} + 2\tau_\pi \pi_\alpha^{\langle\mu}\omega^{\nu\rangle\alpha} \\ &\quad - \delta_{\pi\pi}\pi^{\mu\nu}\theta - \tau_{\pi\pi}\pi_\alpha^{\langle\mu}\sigma^{\nu\rangle\alpha} \\ &\quad + \varphi_7 \pi_\alpha^{\langle\mu}\pi^{\nu\rangle\alpha} + \lambda_{\pi\Pi}\Pi\sigma^{\mu\nu},\end{aligned} \tag{8}$$

where η is the shear viscosity, ζ is the bulk viscosity, $\tau_\Pi = [15(\frac{1}{3} - c_s^2)^2(e + P_0)]^{-1}$ is the bulk relaxation time with c_s denoting the speed of sound, $\tau_\pi = \frac{5\eta}{e+P_0}$ is the shear relaxation

time, $\theta = \nabla_\mu u^\mu$ is the expansion rate, $\sigma^{\mu\nu} = \nabla^{\langle\mu} u^{\nu\rangle}$ is the strain-rate tensor, and $\omega^{\mu\nu} = \frac{1}{2}(\nabla^\mu u^\nu - \nabla^\nu u^\mu)$ is the vorticity tensor. The angular brackets denote the symmetric and traceless projection that is orthogonal to the 4-velocity, and the spatial gradient is defined as $\nabla^\mu = \Delta^{\mu\nu}\partial_\nu$. The second-order transport coefficients are obtained from the 14-moment approximation for massless gas [68–70] and bulk related coefficients from Ref. [71], $\delta_{\Pi\Pi} = \frac{2}{3}\tau_\Pi$, $\lambda_{\Pi\pi} = \frac{8}{5}(\frac{1}{3} - c_s^2)\tau_\Pi$, $\delta_{\pi\pi} = \frac{4}{3}\tau_\pi$, $\tau_{\pi\pi} = \frac{10}{7}\tau_\pi$, $\varphi_7 = \frac{9}{70P_0}$, and $\lambda_{\pi\Pi} = \frac{6}{5}\tau_\pi$.

The temperature dependence of the specific shear viscosity is parametrized as [27]

$$(\eta/s)(T) = \begin{cases} (\eta/s)_{\min} + S_H T\left(\left(\frac{T}{T_H}\right)^{-P_H} - 1\right), & T < T_H \\ (\eta/s)_{\min}, & T_H \leqslant T \leqslant T_Q \\ (\eta/s)_{\min} + S_Q(T - T_Q), & T > T_Q, \end{cases} \tag{9}$$

where s is the equilibrium entropy density and $T_Q = T_H + W_{\min}$. The parametrization allows both wide and narrow constant-η/s region near the QCD transition temperature and also a rapid increase in η/s towards the lower temperatures in the hadronic matter. The latter is particularly important in connection with the dynamical freeze-out, discussed below, that relates decoupling and the temperature dependence of the hadronic shear viscosity.

The temperature dependence of the specific bulk viscosity is parametrized as [27]

$$(\zeta/s)(T) = \frac{(\zeta/s)_{\max}}{1 + \left(\frac{T - T_{\max}^{\zeta/s}}{w(T)}\right)^2}, \tag{10}$$

where

$$w(T) = \frac{2(\zeta/s)_{\text{width}}}{1 + \exp\left(\frac{a_{\zeta/s}(T - T_{\max}^{\zeta/s})}{(\zeta/s)_{\text{width}}}\right)}. \tag{11}$$

The parametrization gives a maximum of ζ/s at temperature $T_{\max}^{\zeta/s}$, and the width of the peak is controlled by $w(T)$. The peak is allowed to be asymmetric (regulated by $a_{\zeta/s}$) so that the most important part of the bulk viscous effects at temperatures below the QCD transition is given by the chemical freeze-out [72–74].

Thus, the $\eta/s(T)$- and $\zeta/s(T)$-related 10 free parameters to be determined from the experimental data are $(\eta/s)_{\min}$, T_H, $W_{\min}$, P_H, S_H, S_Q, $(\zeta/s)_{\max}$, $T_{\max}^{\zeta/s}$, $(\zeta/s)_{\text{width}}$, and $a_{\zeta/s}$.

C. Dynamical freeze-out conditions

In order to model the system-size dependence of freeze-out, while still treating hadronic QCD-matter evolution using fluid dynamics, we employ a dynamical decoupling condition, where the decoupling surface is determined by imposing simultaneously two different conditions.[1] We take the microscopic scale to be given by the shear relaxation time τ_π and assume that the fluid decouples when the local Knudsen number exceeds a limit C_{Kn}. Thus, at decoupling

$$\text{Kn} = \tau_\pi \max(\sqrt{|\nabla_\mu e \nabla^\mu e|}/e, \theta) = C_{\text{Kn}}. \tag{12}$$

In addition, we require that the system decouples when the mean free path, assumed to be proportional to τ_π, becomes larger than the system size. This gives a nonlocal system-size decoupling condition

$$\frac{\gamma\tau_\pi}{R} = C_R, \tag{13}$$

where the Lorentz-γ factor accounts for τ_π being in the fluid rest frame while the system size is determined in the center-of-momentum frame of the nuclear collision, and $R = \sqrt{\frac{A_{\text{Kn}}}{\pi}}$, with A_{Kn} denoting the area in the transverse-coordinate plane where $\text{Kn} < C_{\text{Kn}}$. In the cases where there are multiple separate areas of the fluid, the system size is calculated separately for each region. Both C_{Kn} and C_R are free parameters to be determined from the experimental data.

A clear advantage of treating also the hadronic QCD-matter evolution with fluid dynamics is that it is then possible to avoid artificial discontinuities in the temperature dependence of transport coefficients and at the same time obtain constraints from the experimental data also for the transport coefficients below the QCD transition temperature.

III. PREDICTING HYDRODYNAMIC SIMULATION OUTPUT WITH NEURAL NETWORKS

Some observables, such as higher flow harmonics $v_3, v_4, \ldots$, and their correlations, require a lot of events to constrain statistical errors, which makes them computationally expensive. We use neural networks to obtain quick estimates for single-event model output for any given energy density profile and parameter combination. The performance gain is significant; we can run approximately two full simulations in one CPU-hour, while with neural networks we can produce ∼100 events *per second* using a GPU (NVIDIA V100).

Following Refs. [53,75], the neural-network architecture consists of two main parts. The first part extracts the most relevant information from the initial energy density profile at midrapidity, which is discretized into a 256×256 matrix in the transverse xy plane with a resolution of 0.07 fm. This segment of the network is built from a series of convolutional layers connected via dense connections. This design, known as the Densely Connected Convolutional Network (DenseNet), was originally developed for computer vision tasks [76]. Dense connections preserve information from all preceding convolution operations within a block by concatenating the output channels of each layer with those of all previous layers in that block. Formally, for n successive convolution layers, the output of one densely connected block is as follows:

$$\begin{aligned} \mathbf{y} = [\mathbf{x}, C_1(\mathbf{x}), C_2([\mathbf{x}, C_1(\mathbf{x})]), \ldots, \\ C_n([\mathbf{x}, C_1(\mathbf{x}), C_2([\mathbf{x}, C_1(\mathbf{x})]), \ldots]], \end{aligned} \tag{14}$$

where C_i denotes a convolution operation and $\mathbf{x}$ is the input to the block. Using a convolutional design allows the network to

[1] As an additional constraint, the minimum and maximum temperatures for the decoupling are $T_{\min} = 75$ MeV and $T_{\max} = 150$ MeV.

TABLE I. The structure of a neural network with initial energy density and model parameter inputs. The numbers in parentheses next to the input variable describe the dimensions of the input. The output of the initial energy density part of the network is connected to the model parameter input branch via concatenation.

Initial energy density (256x256)

Block	Output size	Layers
Convolution	128x128x64	7x7 conv, stride 2
Pooling	64x64x64	3x3 max pool, stride 2
Dense Block	64x64x256	[1x1 conv, 3x3 conv] x 6
Transition Layer	67x67x128	1x1 conv
	32x32x128	2x2 average pooling, stride 2
Dense Block	32x32x512	[1x1 conv, 3x3 conv] x 12
Transition Layer	33x33x256	1x1 conv
	16x16x256	2x2 average pooling, stride 2
Dense Block	16x16x896	[1x1 conv, 3x3 conv] x 20
Transition Layer	16x16x448	1x1 conv
	8x8x448	2x2 average pooling, stride 2
Dense Block	8x8x1216	[1x1 conv, 3x3 conv] x 24
Pooling	1216	8x8 global average pooling

Additional inputs

Layer	Outputs size		
Input parameters	η/s (6)	ζ/s (4)	Decoupling (3)
Fully Connected	128	128	128
Concatenation	384		
Fully Connected	1024		

Concatenation, output size: 2240

Layer	Outputs size
Fully Connected	512
Fully Connected	256
Fully Connected	N_{out}

learn local features hierarchically, with smaller-scale patterns building up to recognize larger structures. Furthermore, dense connectivity enhances training stability and mitigates the vanishing gradient problem [76].

As in Refs. [52,53], we use a DenseNet-BC variant with a growth rate $k = 32$ and compression parameter $\theta = 0.5$, which reduces the number of feature maps by a factor of 2 after each densely connected block. The DenseNet part contains a total of four densely connected blocks. Before each block, the spatial size is reduced by a factor of two by performing average pooling, except for the first block, where 3×3 average pooling is used. Each convolution is followed by batch normalization and a rectified linear unit (ReLU) activation function, $F(x) = \max(0, x)$. The DenseNet section concludes with a global average pooling layer, after which the condensed information is passed forward.

The second part of the network processes the model parameters (excluding those related to the initial state, as their information is already encoded in the energy density profile). This segment consists of two fully connected layers. The output of the second segment is then concatenated with the output from the DenseNet, followed by two additional dense layers and a final output layer. A ReLU activation is applied after each fully connected layer. The detailed structure of the full neural network is shown in Table I.

Separate neural networks are trained to produce different types of p_T-integrated event-by-event observables. For this purpose, the observables are divided into the following classes for which the networks are trained:

(1) Particle multiplicities: protons, kaons, pions, charged particles
(2) Mean transverse momenta $\langle p_T \rangle$: protons, kaons, pions, charged particles
(3) Charged-particle flow harmonics v_2, v_3, and v_4, each in a separate class.

Each class contains observables with multiple p_T integration ranges matched to the kinematic cuts of various experiments. Furthermore, each observable is normalized such that the typical values are $\mathcal{O}(1)$.

Network inputs are normalized based on the training data. The initial energy density profiles are scaled to have zero mean and unit standard deviation. Other inputs are scaled to the range [0, 1].

For training, we run full simulations for 40 events per training point for each investigated collision system. Thirty-six of these events are used for the training, and the remaining four are used in validation.

The training points have been sampled using the following iterative procedure: First, we sample 950 points uniformly from the full input parameter space using Latin hypercube sampling and perform initial training of the neural networks. The training was done using the Adam optimizer [77] for a period of 120 epochs. After each epoch, the network performance was tested against the validation set, and the network was saved if it outperformed the previously best-performing network. A logarithmic root-mean-squared (RMS) error loss is used for multiplicities, while the standard RMS error is used for other observables. During the training, the training data are augmented by applying random rotations, flips, and shifts.

Next, we perform the global analysis as described in the following section. After this initial analysis, we sample an

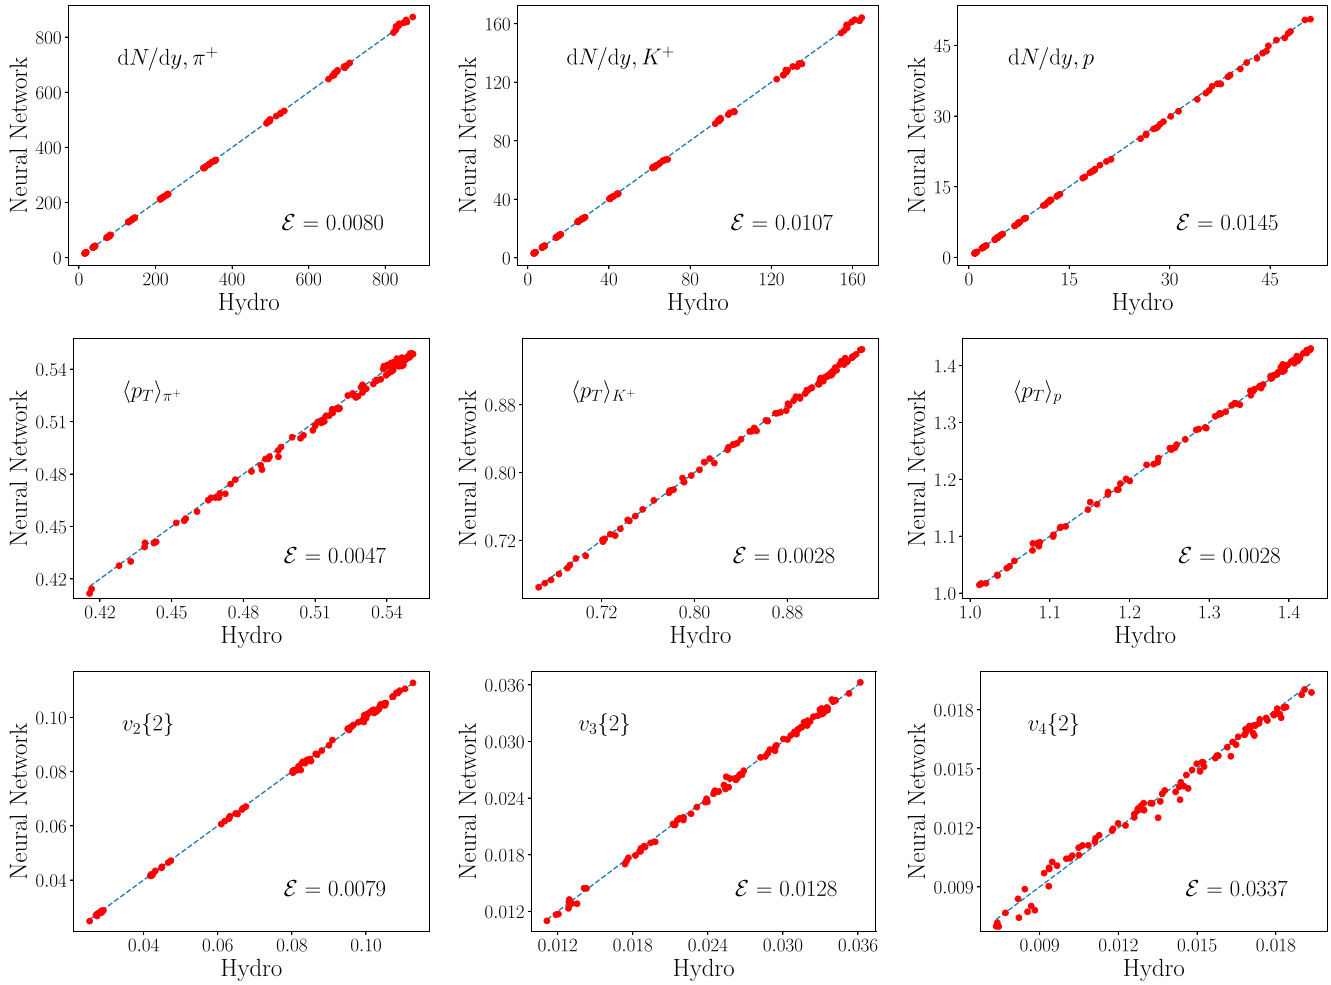

FIG. 1. Validation of neural-network predictions against full-simulation results for various observables in $\sqrt{s_{NN}} = 5.02$ TeV Pb + Pb collisions. The predictions are compared with hydrodynamical simulation results for 10 parameter points across the centrality classes listed in Table III. The straight line is added as a reference to indicate where the two results match exactly. $\mathcal{E}$ is the root-mean-square (RMS) error of the observable, defined in Eq. (A1).

additional 100 training points from the obtained posterior distribution and add these to the training set. We then repeat the procedure by retraining the neural network with this enhanced training set by an additional 20 epochs to obtain more accurate estimates in the more relevant region of the parameter space.

The neural networks are used solely to increase event statistics for individual parameter points, while interpolation across the parameter space is handled by a Gaussian process emulator, which will be described in the next section. For each of the (950 + 100) parameter sets, we generate 10^5 events per collision system utilizing the trained NNs.

To validate the neural network's performance for the event-averaged observables, we selected 10 parameter points from the second training set and performed 5000 hydrodynamic simulations for 5.02 TeV Pb + Pb collisions at each point. A comparison with neural-network predictions is shown in Fig. 1. The accuracy is very good, with an average deviation of ∼1% of the observable value. This error is subleading relative to the error from parameter-space interpolation when the training data for the Gaussian process are sparsely distributed in the parameter space (see Appendix A).

Unfortunately, 5000 collision events are not enough to do the same test for the normalized symmetric cumulants. Therefore, to validate the neural-network performance in this case, we chose the training parameter point with the highest value of the likelihood function and ran 50 000 hydrodynamic simulations. The results for this test for NSC(4, 2), NSC(3, 2), and NSC(4, 3) are shown in Fig. 2. The accuracy of the neural network is of the order 1–10% for the NSC(4, 2) in the 10–70% centrality range. In the most central collision systems and for other normalized symmetric cumulants, the relative errors can be slightly larger, but these data were not included in the analysis.

IV. STATISTICAL ANALYSIS

The goal of a Bayesian global analysis is to obtain the posterior probability distribution of the model parameters. According to Bayes's theorem, the posterior probability P for x_0 to be the "true" value of parameter x, on the condition of the observed data y_{exp}, is proportional to the prior probability

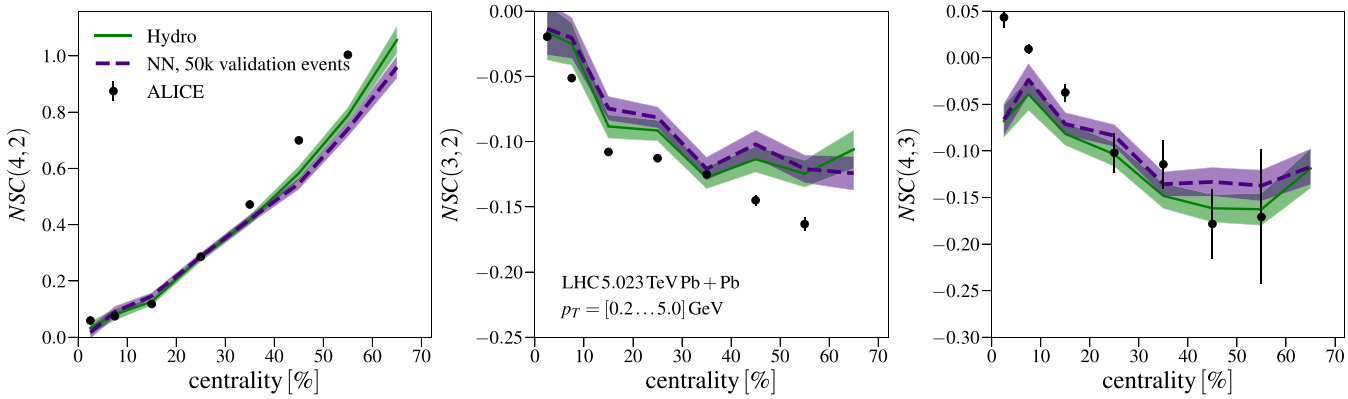

FIG. 2. Validation of the neural-network prediction for normalized symmetric cumulants against hydrodynamic simulations with 50 000 simulated events. Note that only NSC(4, 2) is included in our Bayesian analysis.

p of x_0 and the likelihood of observation y_{exp} if $x = x_0$,

$$P(x = x_0|y_{\text{exp}}) \propto p(x = x_0) \cdot \mathcal{L}(y_{\text{exp}}|x_0). \tag{15}$$

In the case of a uniform prior probability distribution, determining the posterior amounts to determining the likelihood function.

A. Input parameters

Our EbyE-EKRT + (2 + 1)D viscous hydrodynamics with dynamical freeze-out model has altogether 15 tunable parameters, which were introduced in Sec. II. To focus the sampling in the physically most meaningful areas of the input parameter space, we do the following replacements when constructing our uniform prior:

(1) $S_H \to \Delta(\eta/s)_H \equiv (\eta/s)(T = 80\ \text{MeV}) - (\eta/s)_{\text{min}}$,
(2) $C_{\text{Kn}} \to C_{\text{Kn}}/(\eta/s)_{\text{min}}$,
(3) $C_R \to C_R/(\eta/s)_{\text{min}}$.

The final 15 fit parameters in our Bayesian inference study are summarized in Table II, together with their assumed prior ranges. Furthermore, for the model emulation and Markov-chain Monte Carlo (MCMC), the prior ranges have been scaled into a unit hypercube.

B. Observables

Simulations have been performed for four collision systems: Au + Au at $\sqrt{s_{NN}} = 200$ GeV, Pb + Pb at $\sqrt{s_{NN}} = 2.76$ TeV and 5.02 TeV, and Xe + Xe at $\sqrt{s_{NN}} = 5.44$ TeV. For each collision system, the following centrality classes have been included: 0–5%, 5–10%, 10–20%, 20–30%, 30–40%, 40–50%, 50–60%, 60–70%, and 70–80%. For each centrality class, the following 15 observables have been determined:

(1) Charged-particle multiplicity $dN_{\text{ch}}/d\eta|_{|\eta|<0.5}$
(2) Pion multiplicity $dN(\pi^+)/dy|_{|y|<0.5}$
(3) Kaon multiplicity $dN(K^+)/dy|_{|y|<0.5}$
(4) Proton multiplicity $dN(p)/dy|_{|y|<0.5}$
(5) Charged-particle mean transverse momentum $\langle p_T\rangle(\text{ch})$
(6) Pion mean transverse momentum $\langle p_T\rangle(\pi^+)$
(7) Kaon mean transverse momentum $\langle p_T\rangle(K^+)$
(8) Proton mean transverse momentum $\langle p_T\rangle(p)$
(9) Charged particle $v_2\{2\} \cdots v_4\{2\}$, p_T integrated from 0.2 to 3.0 GeV
(10) Charged particle $v_2\{2\} \cdots v_4\{2\}$, p_T integrated from 0.2 to 5.0 GeV
(11) Normalized symmetric cumulant $\text{NSC}(4, 2) = \frac{\langle v_4^2 v_2^2\rangle}{\langle v_4^2\rangle\langle v_2^2\rangle} - 1$

We provide a detailed list of the included experimental data in Table III.

C. Principal component analysis

The dimension of the observable space can be reduced considerably via principal component analysis, where the goal is to find the linear combination of observables with the largest variance (so the largest sensitivity to the change in input). For the variance comparison to make sense, the observables are

TABLE II. Free parameters of the model and their prior ranges.

Parameter	Range
Initial state	
K_{sat}	0.2–1.0
σ_n	0.35–0.60 fm
QCD matter	
T_H	120–250 MeV
$(\eta/s)_{\text{min}}$	0.08–0.22
P_H	2.0–10.0
$\Delta(\eta/s)_H$	0.1–1.5
S_Q	0.0–3.0 GeV^{-1}
W_{min}	0–300 MeV
$(\zeta/s)_{\text{max}}$	0.0–0.15
$T^{\zeta/s}_{\text{max}}$	180–350 MeV
$a_{\zeta/s}$	−2.0–0.0
$(\zeta/s)_{\text{width}}$	10–100 MeV
Decoupling	
T_{chem}	135–165 MeV
$C_{\text{Kn}}/(\eta/s)_{\text{min}}$	2.0–15.0
$C_R/(\eta/s)_{\text{min}}$	0.5–7.5

TABLE III. Summary of experimental data. For Au+Au, the proton target multiplicity is an average of protons and antiprotons and identified particle multiplicities at 10–20% are averages of 10–15% and 15–20% centrality classes; for $\langle p_T \rangle$, the 10–20% centrality class is not included. For Pb+Pb, in cases where $v_n\{2\}$ data have narrower centrality bins than the model output, an average over bins has been taken.

	Au+Au 200 GeV		Pb+Pb 2760 GeV		Pb+Pb 5020 GeV		Xe+Xe 5440 GeV	
Observable	Centrality (%)	Ref.	Centrality (%)	Ref.	Centrality (%)	Ref.	Centrality (%)	Ref.
$dN_{\rm ch}/d\eta\vert_{\vert\eta\vert<0.5}$	0–5, . . . , 50–60	[78]	0–5, . . . , 60–70	[79]	0–5, . . . , 60–70	[80]	0–5, . . . , 70–80	[81]
$dN(\pi^+)/dy\vert_{\vert y\vert<0.5}$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$dN(K^+)/dy\vert_{\vert y\vert<0.5}$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$dN(p)/dy\vert_{\vert y\vert<0.5}$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$\langle p_T \rangle$(ch)	–	–	–	–	–	–	0–5, . . . , 70–80	[81]
$\langle p_T \rangle(\pi^+)$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$\langle p_T \rangle(K^+)$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$\langle p_T \rangle(p)$	0–5, . . . , 70–80	[82]	0–5, . . . , 70–80	[83]	0–5, . . . , 70–80	[84]	–	–
$v_2\{2\}$	0–5, . . . , 70–80	[85]	0–5, . . . , 70–80	[86]	0–5, . . . , 70–80	[86]	0–5, . . . , 60–70	[87]
$v_3\{2\}$	0–5, . . . , 70–80	[85]	0–5, . . . , 60–70	[86]	0–5, . . . , 70–80	[86]	0–5, . . . , 50–60	[87]
$v_4\{2\}$	0–5, . . . , 70–80	[85]	0–5, . . . , 60–70	[86]	0–5, . . . , 70–80	[86]	0–5, . . . , 50–60	[87]
NSC(4, 2)	10−20, . . . , 40–50	[88]	10 − 20, . . . , 40–50	[89]	10 − 20, . . . , 40–50	[90]	–	–

centered so that their mean is 0 and scaled with their sample standard deviation to make them dimensionless and of similar scale with regards to the variance.

Given a matrix $\hat{Y}$ of (centered and scaled) model data, where each row corresponds to one point in input parameter space, with the number of rows corresponding to the number of parameter sets that we will use for the training of the GPs (900 in the first analysis and 980 in the second, see Sec. IV D below) and each column corresponds to one of m data points, we find an eigenvalue decomposition of the covariance matrix $\hat{Y}^T\hat{Y}$,

$$\hat{Y}^T\hat{Y} = V\Lambda V^T, \tag{16}$$

where Λ is the diagonal matrix containing the eigenvalues $\lambda_1, \ldots, \lambda_m$ and V is an orthogonal matrix containing the eigenvectors of the covariance matrix. Since $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_m$, the fraction of the total variance explained by the principal components above some index k will be negligible. This allows us to define a lower-rank approximation of the original transformed data matrix $Z = \hat{Y}V$ as $Z_k = \hat{Y}V_k$, where V_k contains the first k columns of V. The transformation of a vector $\vec{y}$ from the space of observables to a vector $\vec{z}$ in the reduced-dimension principal component space with k components is then defined as

$$\vec{z} = \vec{y}\,V_k. \tag{17}$$

For each collision system we include enough principal components to cover 99% of the data variance:

(1) In Au + Au collisions, 89 data points are reduced to 15 principal components
(2) In Pb + Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, 91 data points are reduced to 12 principal components
(3) In Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, 93 data points are reduced to 12 principal components
(4) In Xe + Xe collisions, 40 data points are reduced to 8 principal components

D. Gaussian process emulators

Given a list of model data (X, Z_k) for a principal component $z(\vec{x}) \in Z_k$ evaluated at parameter points $\vec{x} \in X$, we obtain the Gaussian process predictive mean and uncertainty for z at a point $\vec{x}_0$ as follows:

$$z_{\rm GP}(\vec{x}_0) = \mathcal{C}_{0,Z}\mathcal{C}_{Z,Z}^{-1}Z, \tag{18}$$

$$\sigma^2_{\rm GP}(\vec{x}_0) = \mathcal{C}_{0,0} - \mathcal{C}_{0,Z}\mathcal{C}_{Z,Z}^{-1}\mathcal{C}_{Z,0}, \tag{19}$$

where the covariance matrix has the form

$$\mathcal{C} = \begin{pmatrix} \mathcal{C}_{0,0} & \mathcal{C}_{0,Z} \\ \mathcal{C}_{Z,0} & \mathcal{C}_{Z,Z} \end{pmatrix} \tag{20}$$

further defined by the covariance function also known as GP kernel k: $\mathcal{C}_{0,0} = k(\vec{x}_0, \vec{x}_0), \mathcal{C}_{Z,Z} = k(X, X)$. In this work we use the radial basis function with an additional white noise term for the kernel,

$$k(\vec{x}_i, \vec{x}_j) = A \exp\left(-\frac{1}{2}\sum_{k=1}^{15}\frac{(x_{i,k} - x_{j,k})^2}{\ell_k^2}\right) + \delta_{\rm noise}. \tag{21}$$

The white noise term $\delta_{\rm noise}$ applies only on the diagonal elements of the covariance matrix, $k(\vec{x}_i, \vec{x}_i)$, and accounts for the statistical and neural network uncertainties in the event averages. The kernel hyperparameters $A, \vec{\ell}, \delta_{\rm noise}$ are initially tuned using 900 points from the neural-network training dataset, while the remaining 50 were used for validation. On the second round, 80 of the additional 100 training points are used for hyperparameter tuning and remaining 20 are added to the validation set.

To compare the emulator predictions against physical observables, the prediction $\vec{z}_{\rm GP}$ in the principal component space is transformed back to the space of observables:

$$\vec{y}_{\rm GP} = \vec{z}_{\rm GP}\,V_k^T. \tag{22}$$

One must also remember to invert the scaling and centering of the model data which was performed in preparation for the principal component analysis.

We utilize the SCIKIT-LEARN [91] Python library for both the principal component analysis and the Gaussian process regression. A summary of the emulator quality checks can be found in Appendix A.

E. Likelihood function

For each collision energy, the likelihood function is defined as

$$\mathcal{L}(\vec{z}_{\rm exp}|\vec{x}) = \frac{1}{\sqrt{|2\pi\,\Sigma|}} e^{-d_M^2/2}, \tag{23}$$

where d_M is the Mahalanobis distance

$$d_M^2 \equiv (\vec{z}_{\rm GP}(\vec{x}) - \vec{z}_{\rm exp})^T \Sigma^{-1} (\vec{z}_{\rm GP}(\vec{x}) - \vec{z}_{\rm exp}). \tag{24}$$

The covariance matrix Σ contains three different sources of uncertainty: the GP prediction errors (19), the experimental uncertainties, and an additional theoretical uncertainty $\sigma_{\rm th}$ reflecting the remaining modeling imperfections. Such a theoretical uncertainty was included also in Refs. [30,46]; we, however, do not treat $\sigma_{\rm th}$ as a free parameter but assume it to be $\sim\mathcal{O}(10\%)$ of the measured value for each observable.[2] We investigate three possible error values, 10%, 20%, and 30%, to get a clearer view on its effect on the posterior distribution.[3] Such a study is relevant since a theoretical model, like EbyE-EKRT here, can be too rigid in some of its predictions. If the model is not matching perfectly with data for all observables, then the rigidity easily leads to an artificial bias in an automated global analysis. There is also some degree of uncertainty in the numerical simulations due to the discretization, which is now effectively accounted for by $\sigma_{\rm th}$.

The total likelihood at input parameter point $\vec{x}$ is then the product of the likelihoods from all collision systems:

$$\mathcal{L}_{\rm total} = \mathcal{L}_{200} \times \mathcal{L}_{2760} \times \mathcal{L}_{5020} \times \mathcal{L}_{5440}. \tag{25}$$

We use the Markov-chain Monte Carlo code EMCEE [92] to draw samples from the likelihood function. We run an ensemble of 300 walkers for 40 000 steps and discard the first 30 000 steps to ensure that all samples used in results are drawn after MCMC has converged to the target distribution.

V. RESULTS

A. Posterior distribution

We present the summary statistics of the posterior probability distributions in Table IV, i.e., the median values and 90% credible intervals for the marginal distributions of each parameter.

For the initial-state parameters, we find the saturation parameter has a median value $K_{\rm sat} \approx 0.5$ and the full 90% credible range considering all three theoretical uncertainties is $0.38 < K_{\rm sat} < 0.67$, while the nucleon width favors the lowest allowed values, 0.35 fm $< \sigma_n <$ 0.47 fm. These values for $K_{\rm sat}$ and σ_n are consistent with the earlier studies [13,27]. Interestingly, the σ_n obtained here is quite close to the value 0.43 ± 0.01 fm, extracted from the HERA data for exclusive electroproduction of J/Ψ [93].

TABLE IV. Summary of the median values and 90% credible intervals from the marginal posterior distributions of the model parameters, for three different values of theoretical uncertainty.

	Theoretical uncertainty $\sigma_{\rm th}$		
Parameter	10%	20%	30%
$K_{\rm sat}$	$0.57^{+0.10}_{-0.09}$	$0.53^{+0.12}_{-0.11}$	$0.51^{+0.16}_{-0.13}$
σ_n [fm]	$0.38^{+0.03}_{-0.02}$	$0.38^{+0.05}_{-0.03}$	$0.40^{+0.07}_{-0.04}$
T_H [MeV]	181^{+61}_{-59}	178^{+61}_{-55}	178^{+63}_{-54}
$(\eta/s)_{\rm min}$	$0.14^{+0.01}_{-0.02}$	$0.15^{+0.02}_{-0.02}$	$0.15^{+0.03}_{-0.03}$
P_H	$8.4^{+1.5}_{-3.9}$	$8.4^{+1.5}_{-3.7}$	$8.3^{+1.5}_{-3.8}$
$\Delta(\eta/s)_H$	$0.3^{+0.3}_{-0.2}$	$0.3^{+0.6}_{-0.2}$	$0.4^{+0.8}_{-0.3}$
S_Q [GeV^{-1}]	$1.2^{+1.6}_{-1.1}$	$1.3^{+1.5}_{-1.2}$	$1.4^{+1.5}_{-1.2}$
$W_{\rm min}$ [MeV]	225^{+68}_{-130}	183^{+104}_{-151}	165^{+120}_{-144}
$(\zeta/s)_{\rm max}$	$0.09^{+0.05}_{-0.05}$	$0.09^{+0.05}_{-0.06}$	$0.08^{+0.06}_{-0.06}$
$T^{\zeta/s}_{\rm max}$ [MeV]	215^{+66}_{-31}	223^{+59}_{-37}	230^{+76}_{-43}
$a_{\zeta/s}$	$-0.5^{+0.4}_{-1.1}$	$-0.7^{+0.7}_{-1.1}$	$-0.8^{+0.7}_{-1.0}$
$(\zeta/s)_{\rm width}$ [MeV]	57^{+38}_{-43}	60^{+36}_{-43}	61^{+35}_{-43}
$T_{\rm chem}$ [MeV]	147^{+3}_{-4}	148^{+5}_{-5}	149^{+7}_{-6}
$C_{\rm Kn}$	$1.2^{+0.6}_{-0.4}$	$1.4^{+0.7}_{-0.5}$	$1.5^{+0.8}_{-0.6}$
C_R	$0.7^{+0.3}_{-0.4}$	$0.7^{+0.3}_{-0.4}$	$0.8^{+0.4}_{-0.4}$

The chemical freeze-out temperature $T_{\rm chem}$ is constrained within the range 143–156 MeV, with the increase in theoretical uncertainty effecting mainly the upper bound. The current posterior distribution is extending to a slightly lower range than the previously obtained 150–158 MeV for *s95p* EoS from the global analysis presented in Ref. [31].

The Knudsen number for freeze-out is 0.8–2.3, while the ratio of mean free path to system size at freeze-out is in the range 0.3–1.2 with the median around $C_R \approx 0.7$–0.8. Thus, according to the experimental data, the freeze-out takes place, as expected, at the very limit of the applicability of hydrodynamics.

For the transport coefficients, rather than investigating each parameter individually, we illustrate their combined effect by plotting the actual temperature dependence of η/s and ζ/s based on samples from the posterior distribution. Even if the individual parameters themselves are not very tightly constrained, they are not independent from each other, and the net effect is that in particular shear viscosity itself is much more tightly constrained than the constraints on the individual parameters would suggest. We also note that the final constraints for the transport properties of QCD matter are very robust against the variations in the assumed theoretical error, which is an encouraging result. The most significant difference is seen in bulk viscosity, with the smallest 10% error there is

[2]Note that the $\mathcal{O}(10\%)$ theoretical errors are of the same order as the errors in the computed charged multiplicities that arise from the uncertainties of nuclear parton distributions as estimated in Ref. [11].

[3]The initial analysis (with 950 parameter combinations sampled from the prior) is done only once, with $\sigma_{\rm th} = 0$, so the GP emulators use the same training set and hyperparameters for all three scenarios.

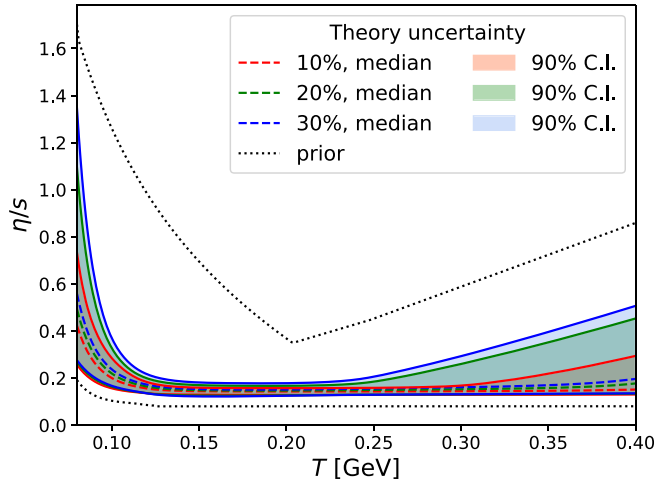

FIG. 3. The median and 90% credible interval for the specific shear viscosity η/s with respect to temperature T.

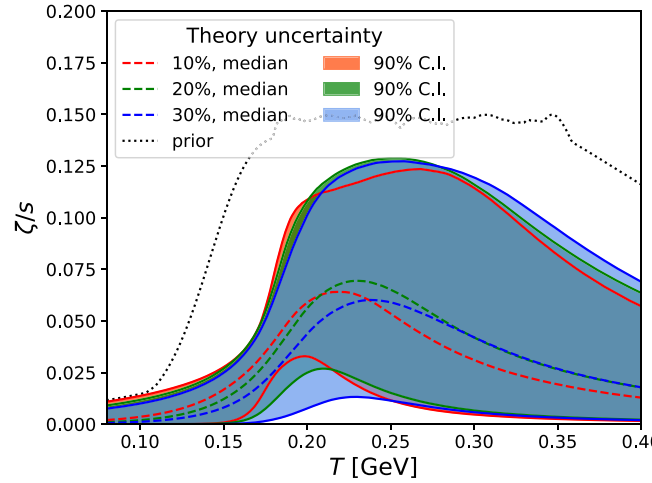

FIG. 4. The median and 90% credible interval for the bulk viscous coefficient ζ/s with respect to temperature T. Note that the lower limit of the prior is zero.

a clear lower limit, but this diminishes with larger theoretical errors. The upper limit for bulk is much less sensitive. On the other hand, the limits for the shear viscosity are practically independent of the assumed error.

We see from Fig. 3 that η/s has a strong temperature dependence below $T \approx 150$ MeV (parameter P_H leaning strongly towards the maximum of the prior range), while the minimum-value plateau is between 150 and 220 MeV with $0.12 < (\eta/s)_{\min} < 0.18$. For temperatures higher than 200 MeV, a rising trend is allowed but not favored; for all three cases the median remains close to constant all the way to 350 MeV. This is a clear difference from the previous global analysis with EKRT initial state [31], where the model with *s95p* equation of state had η/s increasing with temperature beyond $T > 250$ MeV. However, the two analyses are not directly comparable, as the previous analysis was done using averaged initial conditions and without bulk viscosity.

Comparing our result for η/s against other recent Bayesian analyses, we obtain quite a good agreement with the JETSCAPE collaboration (Ref. [33], Fig. 6) for the temperature range investigated in that study, $T = 150$–350 MeV. While the JETSCAPE parametrization for $\eta/s(T)$ did not explicitly include a constant-value plateau, the slope at high temperatures was allowed to be either positive or negative and no strong preference for the sign of the slope was found; a later Bayesian study with the same parametrization but wider prior and IP-Glasma initial state also found temperature-independent η/s to be consistent with data [42]. Our result also fits within the credible intervals from the various Trajectum analyses presented in Ref. [37]; indeed, our 90% credible interval for the 10% theoretical uncertainty is fully consistent with the Trajectum results for η/s. Meanwhile, the rising trend observed by Virta *et al.* in Ref. [46] cannot be fully accommodated within the uncertainties of the present analysis; this discrepancy likely stems from the more restricted parametrization of $(\eta/s)(T)$ in Ref. [46] which does not allow a constant-value plateau or a decrease of η/s with increasing temperature.

For the bulk viscosity in this study, we obtain much weaker constraints compared to η/s; see Fig. 4. We can say with 90% credibility that ζ/s is nonzero in the temperature range $T = 200$–300 MeV, but the variation in the exact value remains considerable. This temperature range for nonzero ζ/s is slightly higher than the range 150–200 MeV obtained in the analyses of JETSCAPE [33] and Virta *et al.* [46]. On the other hand, both the IP-Glasma study [42] and the Trajectum analysis with the total hadronic cross section σ_{AA} included in the fit data and no additional weighting of observables [37] find a nonzero minimum bound for ζ/s in a temperature range similar to our result.

We note that bulk viscosity remains small at and below the chemical freeze-out temperature even if the prior range would allow larger values. This is on par with the discussion in Ref. [27], where the parametrization was constructed so that the main part of the bulk viscous effect in hadron gas is described by chemical freeze-out.

B. Comparison to data

In the following, for each of the three investigated scenarios, we sample 5000 parameter combinations from the posterior distribution and compare the distribution of emulator predictions for these points against the data.

The centrality dependence of charged-particle multiplicities for the investigated collision systems are presented in Fig. 5, along with the average transverse momentum in 5.44 TeV Xe + Xe collisions. The agreement with the multiplicity data is very satisfactory across the whole collision energy range, from 200 GeV Au + Au up to 5.44 TeV Xe + Xe collisions, with larger theoretical uncertainty merely increasing the variation in the predictions while the median distance from the data remains very similar. However, in the case of Xe + Xe in particular, we see that the centrality dependence comes off a bit too steep; the multiplicity falls slightly too strongly towards peripheral collisions. This overly strong centrality dependence applies also for the mean transverse mo-

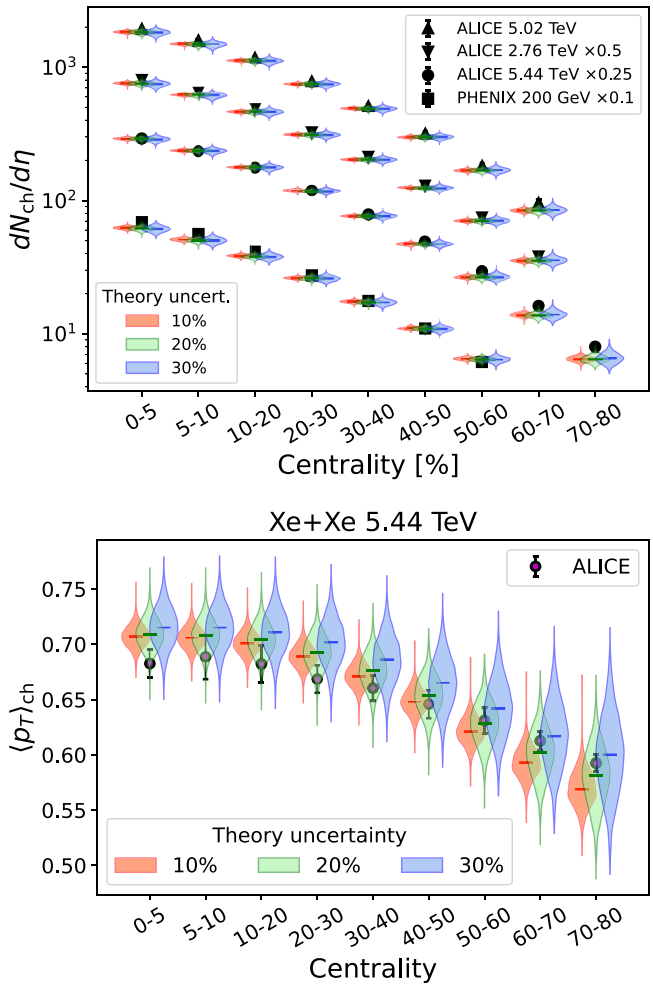

FIG. 5. Posterior distribution–based emulator predictions for the centrality dependence of charged-particle multiplicity for all four collision systems (top panel) and average transverse momentum in Xe + Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV (bottom panel). The violin plots reflect the distribution of emulator predictions, with more points concentrated on the broader parts of the violin. The horizontal line within the violin indicates the mean value of predictions. PHENIX data is from Ref. [78] and ALICE data from Refs. [79] (2.76 TeV), [80] (5.02 TeV), and [81] (5.44 TeV).

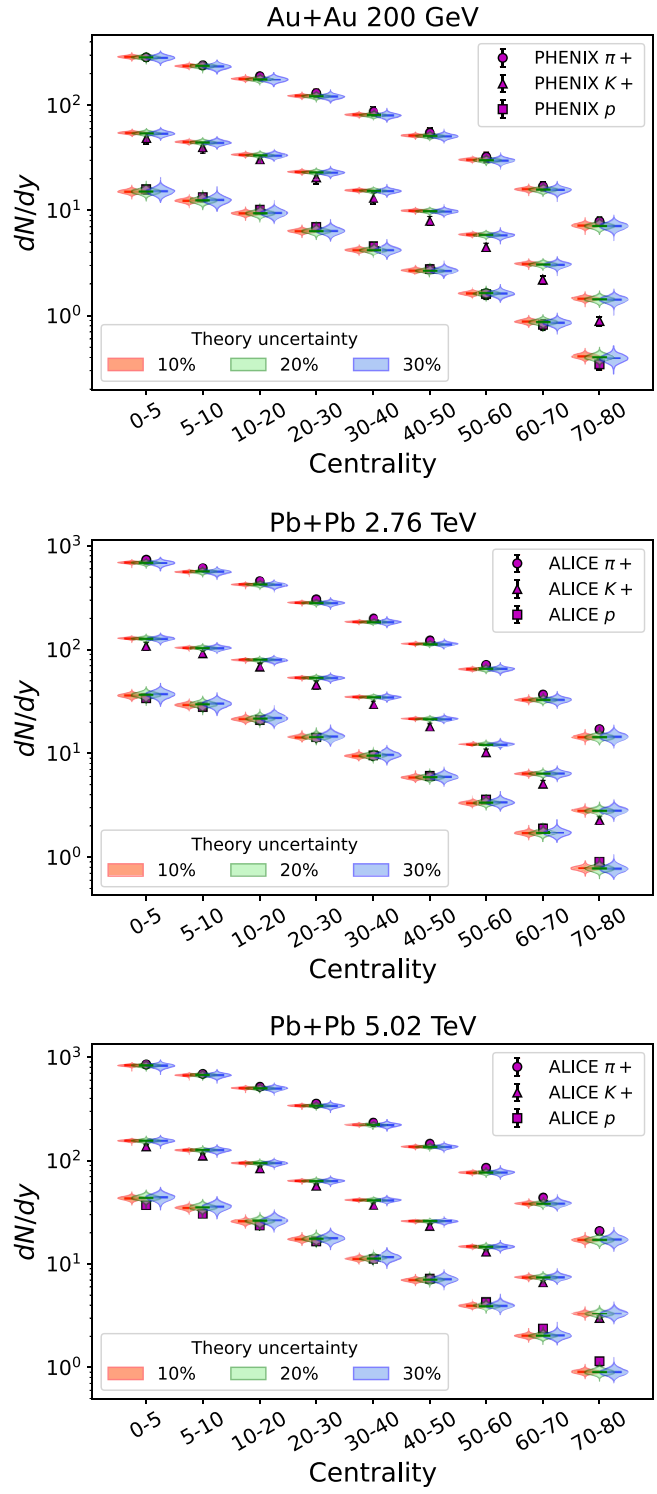

FIG. 6. Posterior distribution emulator predictions for identified particle multiplicities in 200 GeV Au + Au (top panel), 2.76 TeV Pb + Pb (middle panel), and 5.02 TeV Pb + Pb (bottom panel) collisions. PHENIX data from [82]. ALICE data from [83] (2.76 TeV) and [84] (5.02 TeV).

mentum of charged particles; in addition, here the uncertainties tend to increase more towards higher values of $\langle p_T \rangle$, and consequently the increase in theoretical uncertainty also increases the median value of $\langle p_T \rangle_{ch}$ predictions.

Here we also want to emphasize that none of the model parameters depend on the collision energy $\sqrt{s_{NN}}$ or nuclear mass number A and that the $\sqrt{s_{NN}}$, A, and centrality dependencies of multiplicity are mainly predictions of the EKRT model.

We show the data comparison of emulator predictions for pion, kaon, and proton multiplicities in Au + Au and Pb + Pb collisions in Fig. 6. We obtain good descriptions of pion and proton multiplicities in all cases; however, the kaon production is systematically overestimated. As was already seen in the case of charged-particle multiplicity, the centrality dependence tends to be too steep at the LHC energies, suggesting perhaps a collision energy-dependent nucleon width σ_n and/or possibly also that nuclear substructure ("hotspots") might be required for a better fit.

The data comparison for the transverse momenta of pions, kaons, and protons is presented in Fig. 7. Here the overall

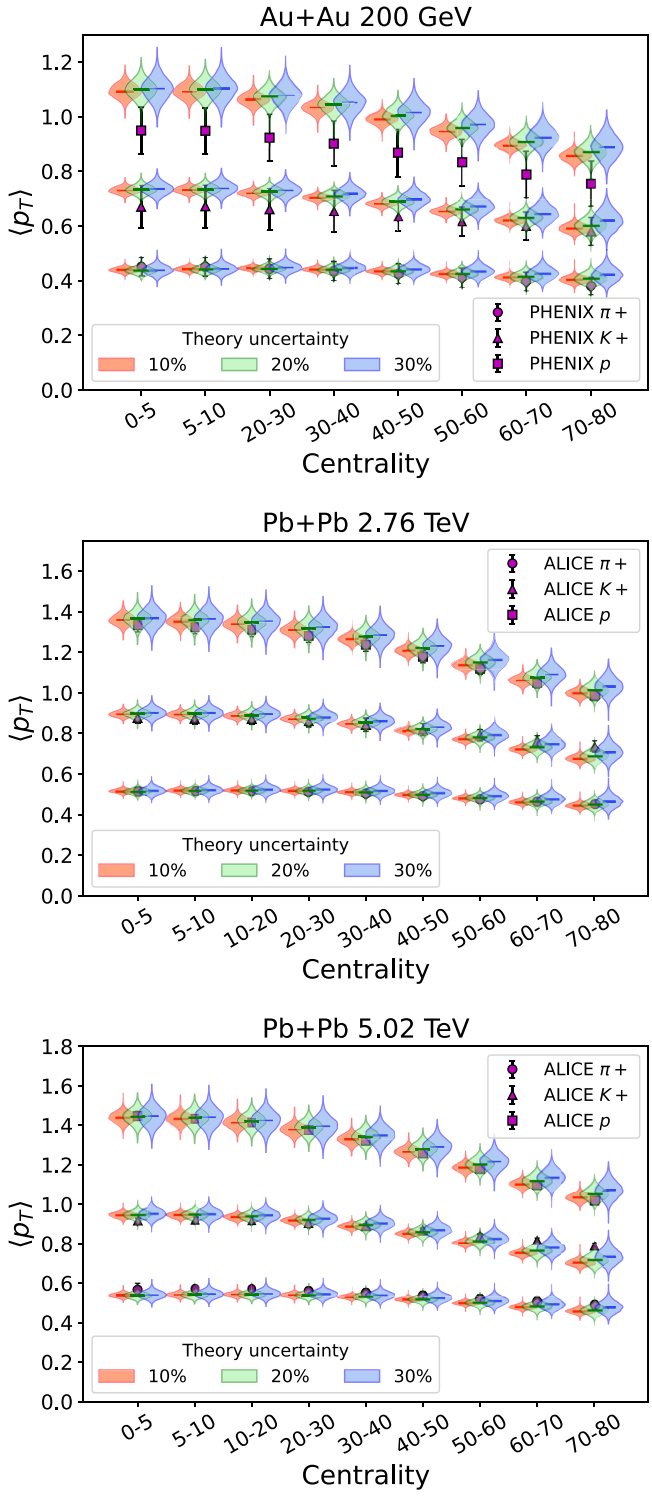

FIG. 7. Posterior distribution emulator predictions for identified particle transverse momenta in 200 GeV Au + Au (top panel), 2.76 TeV Pb + Pb (middle panel), and 5.02 TeV Pb + Pb (bottom panel) collisions. PHENIX data from Ref. [82]. ALICE data from Ref. [83] (2.76 TeV) and Ref. [84] (5.02 TeV).

agreement with the data is very good for all three particle species, the main problem being the proton $\langle p_T \rangle$ at 200 GeV Au + Au collisions which is systematically overestimated. Similarly to the charged particle $\langle p_T \rangle$ in Xe + Xe collisions, it seems easier to increase $\langle p_T \rangle$ rather than decrease it within the posterior distribution, so the increase in theoretical uncertainty allows more variation towards higher $\langle p_T \rangle$ values.

TABLE V. Notable Pearson correlations ($|r| > 0.3$ for at least two of the three scenarios of different theoretical uncertainty) between parameters in the posterior distribution.

Parameter		Th. uncertainty		
A	B	10%	20%	30%
$K_{\rm sat}$	σ_n	−0.41	−0.44	−0.47
$K_{\rm sat}$	T_H	0.42	0.30	0.25
$K_{\rm sat}$	$W_{\rm min}$	0.31	0.40	0.35
σ_n	$T_{\rm chem}$	0.46	0.35	0.33
T_H	P_H	0.60	0.45	0.33
$(\eta/s)_{\rm min}$	$C_{\rm Kn}$	0.57	0.65	0.71
$(\zeta/s)_{\rm max}$	$T^{\zeta/s}_{\rm max}$	0.35	0.38	0.32
$T^{\zeta/s}_{\rm max}$	$(\zeta/s)_{\rm width}$	0.65	0.58	0.47

The flow harmonics v_n for $n = 2, 3, 4$ are shown in Fig. 8. The overall agreement with the data over multiple collision systems and three different harmonics is quite good. However, the centrality dependence is again slightly off; there is too little flow in the most central collisions and too much v_2 produced in more peripheral collisions. The underestimation of flow in the most central collisions may necessitate implementing nuclear deformation effects for Au + Au and Pb + Pb collisions and further deformations in the case of Xe + Xe. Including nucleon substructure would likely improve the v_3/v_2 ratio, as we have demonstrated in a recent study [15]. We note that in particular v_3 and v_4 in Au + Au collisions are in a significantly better agreement with the measurements than in our previous non-Bayesian analysis using the same model setup [27].

To complete the data comparison, we show the emulator predictions for the normalized symmetric cumulant NSC(4, 2) in Fig. 9. We get a good description of the data in the 10–20% centrality class, but the rise towards more peripheral collisions is too weak compared to measurements.

C. Correlations

1. *Correlations between parameters and observables*

Correlations between the parameter-parameter pairs and parameter-observable pairs can serve as a useful guide when investigating the posterior distribution. Strong correlations between parameters can indicate either redundancies in the parametrization (multiple combinations are producing the same effect) or particularly strong constraints from the data, limiting the posterior parameter space and "forcing" a correlation. Correlations between parameters and observables provide further guidance where to look for these data constraints.

In Table V we list the notable Pearson correlations between the parameters within the posterior distribution,

$$r_{ij} = \frac{\sum_k (x_{i,k} - \langle x_i \rangle)(x_{j,k} - \langle x_j \rangle)}{\sqrt{\sum_k (x_{i,k} - \langle x_i \rangle)^2 \sum_k (x_{j,k} - \langle x_j \rangle)^2}}, \tag{26}$$

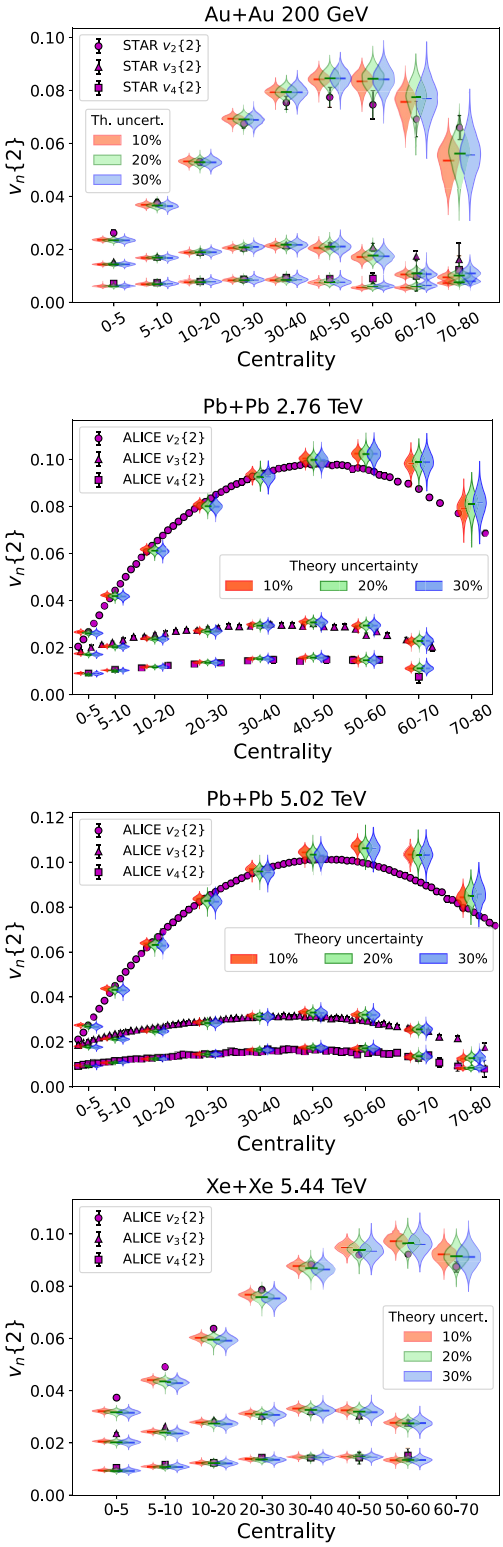

FIG. 8. Posterior distribution emulator predictions for charged-particle flow at RHIC Au + Au (top panel), LHC Pb + Pb (middle panels), and Xe + Xe (bottom panel) collisions. STAR data from Ref. [85]. ALICE data from Ref. [86] (Pb + Pb) and Ref. [87] (Xe + Xe).

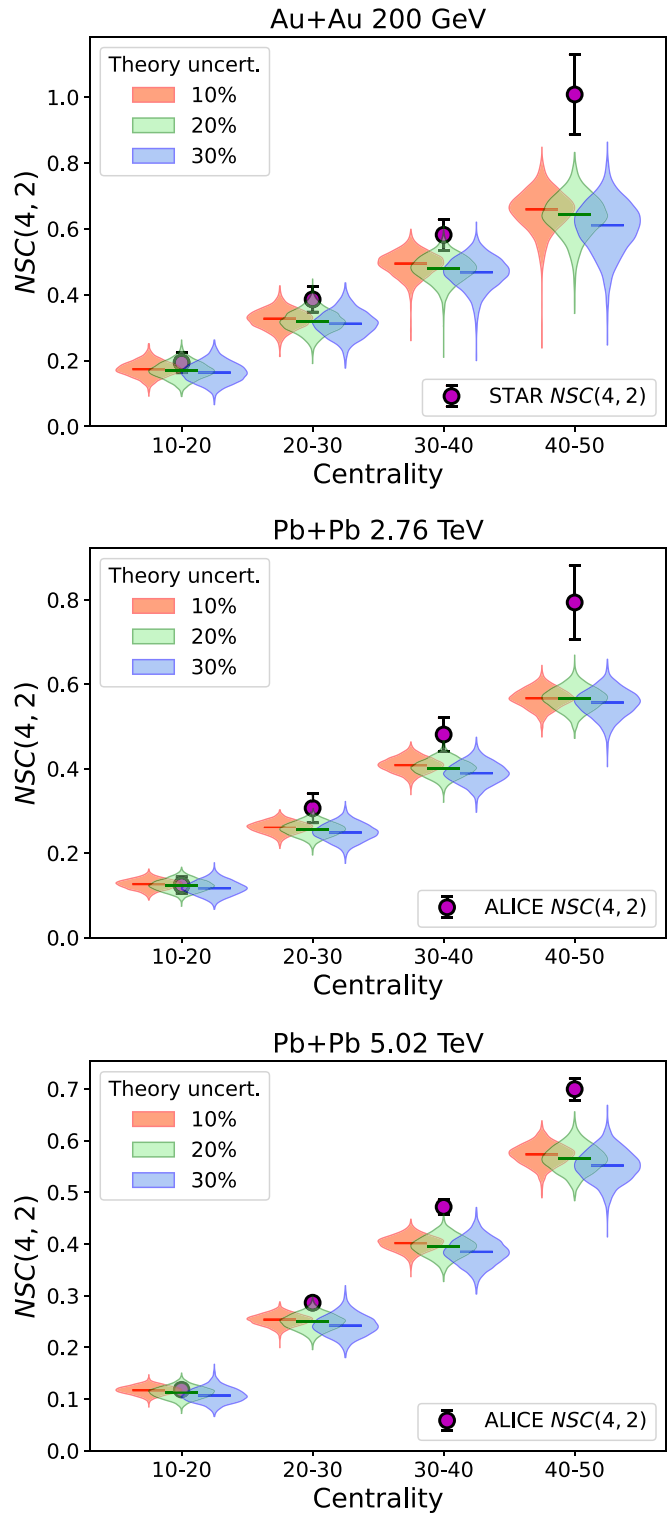

FIG. 9. Posterior distribution emulator predictions for normalized symmetric cumulant NSC(4, 2) in 200 GeV Au + Au (top panel), 2.76 TeV Pb + Pb (middle panel), and 5.02 TeV Pb + Pb (bottom panel) collisions. STAR data from [88]. ALICE data from Ref. [89] (2.76 TeV) and Ref. [90] (5.02 TeV).

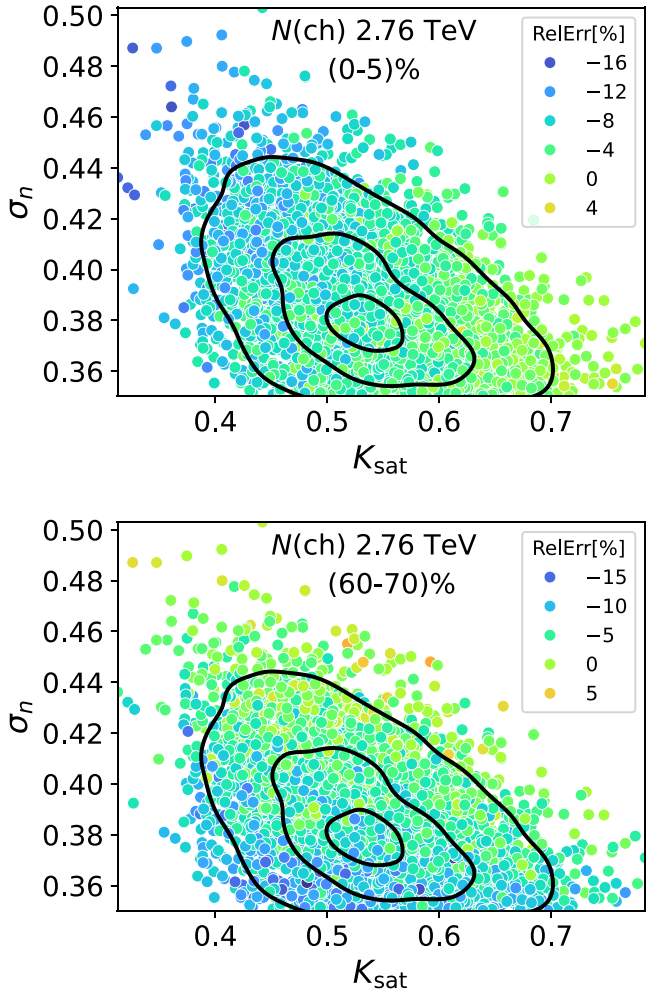

FIG. 10. Relative deviation from the measured value of charged-particle multiplicity $N(\mathrm{ch})$ in Pb + Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for different combinations of initial-state parameters $K_{\rm sat}$ and σ_n, sampled from the posterior distribution with 20% theoretical uncertainty. Black contours indicate the 10%, 50%, and 90% levels of point density, such that the highest 10% of point density is within the innermost curve. Top panel: 0–5% centrality. Bottom panel: 60–70% centrality.

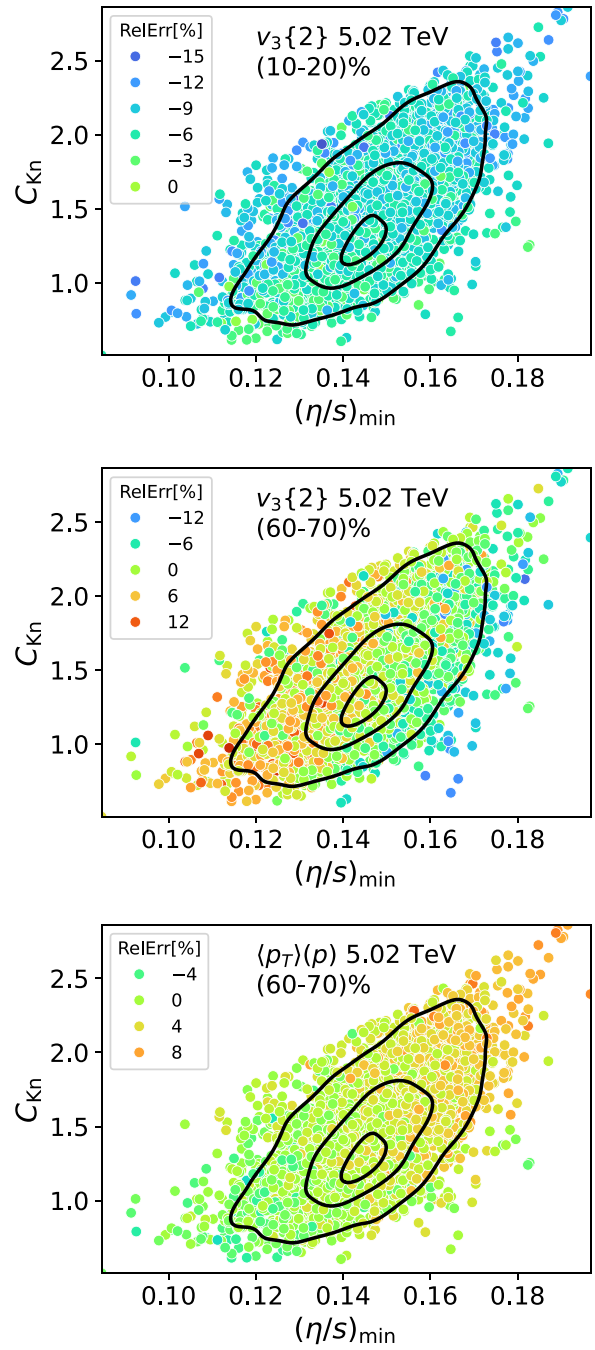

FIG. 11. Relative deviation from the measured value of charged particle $v_3\{2\}$ (top and middle panel) and proton $\langle p_T \rangle$ (bottom panel) in Pb + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV for different combinations of $(\eta/s)_{\rm min}$ and $C_{\rm Kn}$.

where "notable" means $|r| > 0.3$ for at least two of the three investigated scenarios of varying theoretical uncertainty. In addition, summary tables for all notable correlations between model parameters and final-state observables for 20% theoretical uncertainty can be found in Appendix C.

There are three pairs of parameters with particularly strong correlations which appear regardless of the magnitude of the theoretical uncertainty. The initial-state parameters $K_{\rm sat}$ and σ_n are anticorrelated. This emerges from the tension between multiplicity data in central and peripheral collisions, as illustrated in Fig. 10 for the Pb + Pb collisions at 2.76 TeV. Central collisions favor the combination of large $K_{\rm sat}$ and small σ_n, while for the multiplicities in peripheral collisions, smaller values of $K_{\rm sat}$ combined with larger values of σ_n are preferred.

Minimum value of the specific shear viscosity $(\eta/s)_{\rm min}$ correlates with the freeze-out Knudsen number $C_{\rm Kn}$. Here the allowed combinations are limited by mean transverse momentum of protons, which imposes constrains to the upper limit of $(\eta/s)_{\rm min}$ in particular, and the centrality dependence of v_3; we demonstrate this in Fig. 11 for 5.02 TeV Pb + Pb collisions.

The temperature where the bulk viscosity reaches its peak $T^{\zeta/s}_{\rm max}$ is correlated with the width of the peak $(\zeta/s)_{\rm width}$. This correlation does not stem from data constraints, as these parameters have hardly any notable correlations with any of the

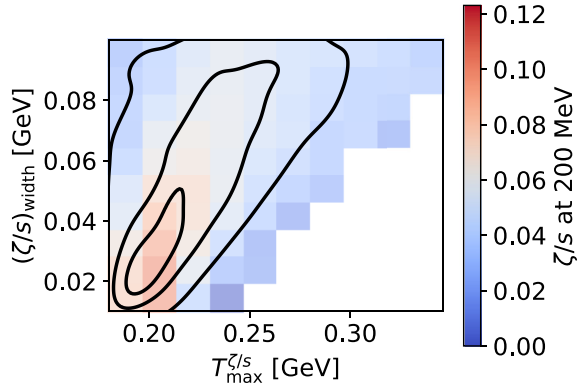

FIG. 12. Value of bulk viscosity coefficient ζ/s at $T = 200$ MeV for different combinations of the peak temperature $T^{\zeta/s}_{\text{max}}$ and peak width $(\zeta/s)_{\text{width}}$ from the posterior distribution.

observables. Rather, the analysis is driven to maximize ζ/s at $T \approx 200$ MeV which makes the combination of narrow peak located at high temperature unpreferable; see Fig. 12.

Further investigating the correlations between model parameters and observables, there is a particularly strong correlation between chemical freeze-out temperature T_{chem} and proton multiplicity $N(p)$. We find notable tension in the $N(p)$ data regarding the optimal value of T_{chem}, both in terms of collision energy and centrality, as shown in Fig. 13. This tension sets strong constraints on T_{chem} and also suggests that the chemical freeze-out temperature might depend on the collision energy.

We already noted above that $(\eta/s)_{\text{min}}$ is constrained by $\langle p_T \rangle(p)$ and v_3, and it is not surprising that $\langle p_T \rangle(K)$ and v_4 are also correlated with $(\eta/s)_{\text{min}}$. However, it is interesting to note that $\langle p_T \rangle(\pi)$ and v_2 do *not* have notable correlations with $(\eta/s)_{\text{min}}$. Instead, both show sensitivity to the initial-state parameters, on the one hand, and the conditions at freeze-out, on the other; as a correlation to $\Delta(\eta/s)_H$ in the case of $\langle p_T \rangle(\pi)$ and as a correlation to C_{Kn} in the case of v_2.

As v_2 is correlated with σ_n and C_{Kn}, while v_4 is correlated with $(\eta/s)_{\text{min}}$, the question arises if these correlations are reflected in NSC(4, 2). This turns out to be true only partially, and the dependence on centrality and collision energy is quite different, as shown in Fig. 14. For v_2, the σ_n correlation is strongest in more central collisions, while for NSC(4, 2) it is notable in more peripheral collisions. For v_4, the correlation with $(\eta/s)_{\text{min}}$ is strongest at higher collision energies for NSC(4, 2) 200 GeV has the strongest correlation. The consistently strong correlation between v_2 and C_{Kn} is greatly diminished for NSC(4, 2) and survives only for the 2.76 TeV Pb + Pb collisions at 10–20% centrality.

2. Correlations between observables and viscous coefficients at different temperatures

While there are some notable correlations between observables and parameters related to the extrema of the shear viscosity coefficient, we notice that many of the parameters controlling the temperature dependence of η/s and ζ/s have

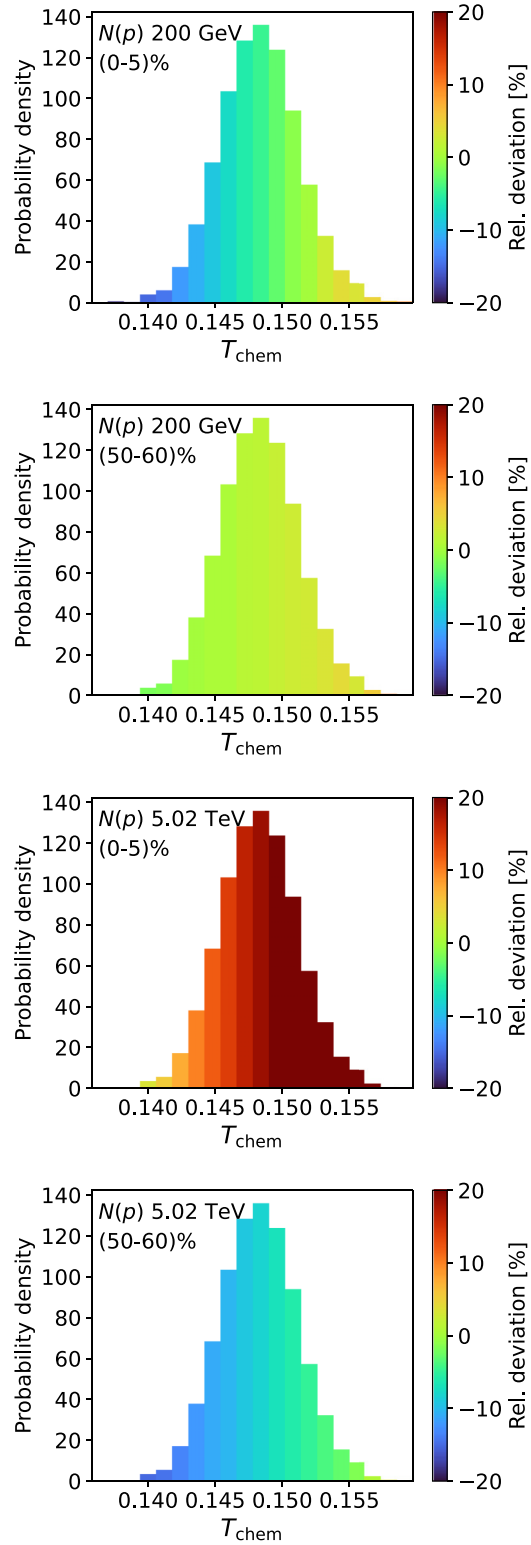

FIG. 13. The posterior distribution of T_{chem} (same in all panels). The color indicates the average relative deviation of proton multiplicity $N(p)$ from the measured value in 200 GeV Au + Au (top panels) and 5.02 TeV Pb + Pb collisions (bottom panels) in 0–5% and 50–60% centrality classes for different values of T_{chem}. Average is taken over the posterior samples within each bin.

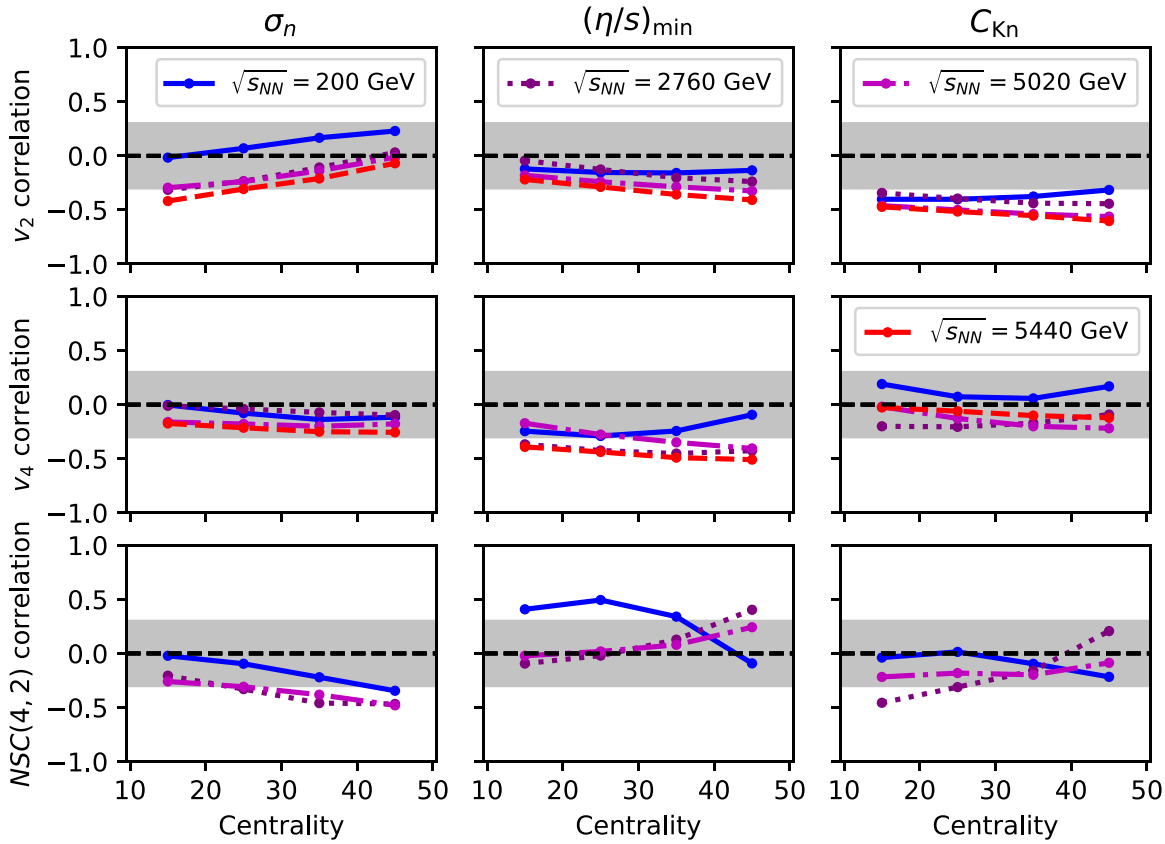

FIG. 14. Centrality dependence of Pearson correlations for $v_2\{2\}$ (top row), $v_4\{2\}$ (middle row), and NSC(4, 2) (bottom row) with parameters σ_n, $(\eta/s)_{\min}$, and C_{Kn} (left, middle, and right column, respectively). Gray band indicates negligible correlation strength.

no systematic correlations with observables. In the following, we investigate whether the observables are correlated with η/s and ζ/s in any particular temperature regions. In addition, tabulated correlation data for a selection of temperatures can be found in Appendix D.

For the shear viscosity, the observables with notable correlations are the mean transverse momenta $\langle p_T \rangle$ of pions, kaons, and protons and the flow observables $v_n\{2\}$ and NSC(4, 2). We show the correlations between $\langle p_T \rangle$ and η/s in the temperature range 80–400 MeV in Fig. 15. Here the findings are consistent with what we already saw in the correlations between model parameters and observables; pion $\langle p_T \rangle$ correlates with η/s at low temperatures, as was suggested by the observed correlation between $\langle p_T \rangle(\pi)$ and $\Delta(\eta/s)_H$, while $\langle p_T \rangle(K)$ and $\langle p_T \rangle(p)$ correlate with η/s in the plateau range $\approx$130–230 MeV, as suggested by their correlations with $(\eta/s)_{\min}$. These correlations have a very weak dependence on collision energy, and also the centrality dependence is notable only at the lowest temperatures, effecting mostly the correlation between $\langle p_T \rangle(\pi)$ and η/s, which is strongest at the most central collisions and becomes insignificant around 40–50% centrality.

The flow observables, however, show more variable behavior; see Fig. 16. While at RHIC the correlations tend to vanish beyond temperatures $T > 200$ MeV, at LHC energies they persist up to 300 MeV. At these high temperatures, the correlation strengths for v_3 and v_4 seem to first increase with collision energy but then abruptly decrease again when moving from 5.02 TeV Pb + Pb collisions to 5.44 TeV Xe + Xe collisions, which may hint to a system-size dependence. We also observe similar weakening of the correlations at high temperatures when moving to more peripheral collisions, supporting the idea that the flow coefficients are not probing these temperatures in smaller systems. It also turns out that the correlator NSC(4, 2) is not sensitive to η/s at high temperatures either, having the strongest correlations with η/s around $T \approx 100$ MeV.

Finally, we move on to the bulk viscosity ζ/s whose temperature dependence remains poorly constrained. Indeed, it only shows notable correlations with mean transverse momenta, which we present in Fig. 17. Here all three investigated particle species tell a remarkably similar story: The correlation is the strongest in the most peripheral collisions and peaks between 150 and 200 MeV, which is the temperature range where the lower limit of our 90% credible interval for ζ/s becomes nonzero. The correlation peak does move towards higher temperatures in more central collisions, but the overall strength of the correlation also weakens, suggesting that no observable in the dataset used in this analysis is sensitive to ζ/s at temperatures beyond $T \gtrsim 250$ MeV.

VI. SUMMARY AND CONCLUSIONS

We have introduced a novel neural-network-enhanced framework for performing a Bayesian global analysis to de-

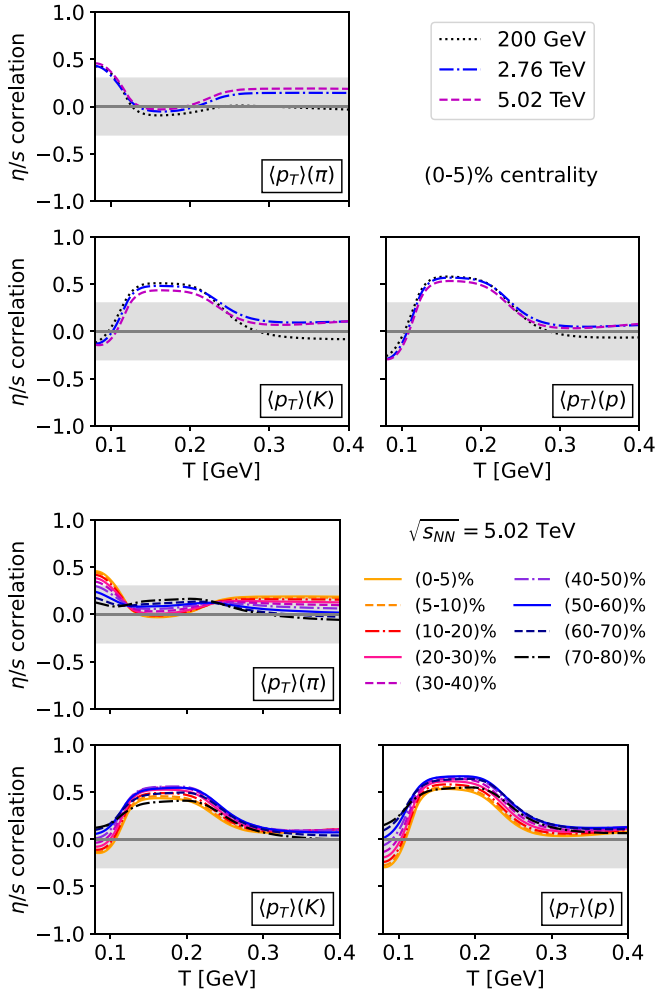

FIG. 15. Temperature dependence of Pearson correlations between mean transverse momenta and η/s. Top three panels: Collision energy dependence of correlations at 0–5% centrality. Bottom three panels: Centrality dependence of correlations at $\sqrt{s_{NN}} = 5.02$ TeV collision energy.

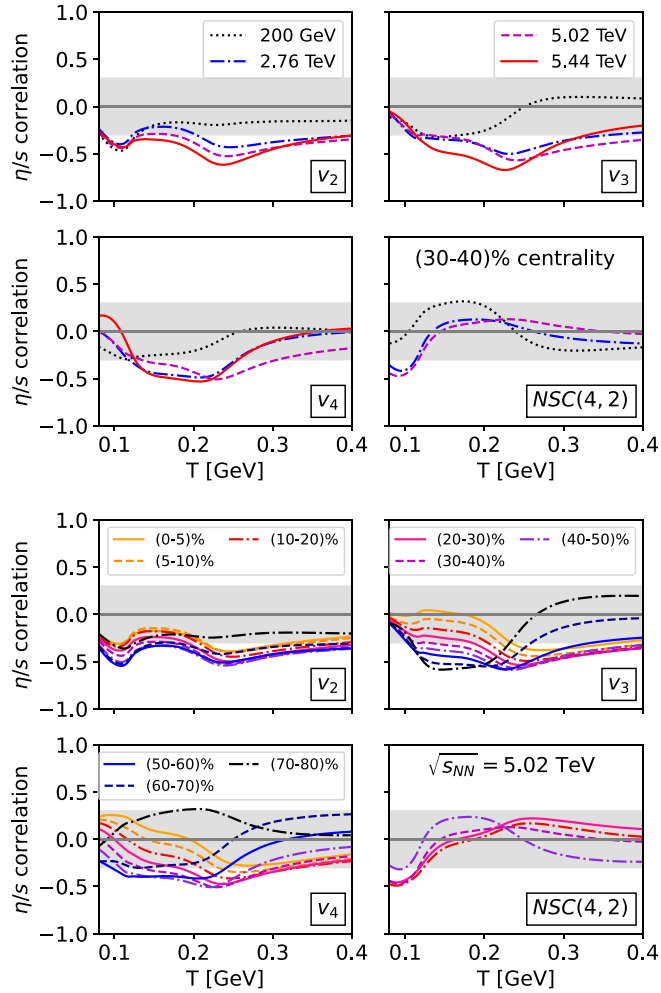

FIG. 16. Temperature dependence of Pearson correlations between flow observables and η/s. Top four panels: Collision energy dependence of correlations at 30–40% centrality. Bottom four panels: Centrality dependence of correlations at $\sqrt{s_{NN}} = 5.02$ TeV collision energy.

termine the QCD matter properties from the high-energy heavy-ion collision data. The main feature of the framework is the replacement of the computationally demanding hydrodynamical simulations with neural networks that can predict the experimental observables directly from the initial energy density profile event by event. The new feature of the networks compared to our earlier work [52] is that they now take also the matter properties as an input, so that it is no longer necessary to train separate networks for each viscosity parametrization and dynamical freeze-out conditions. This makes them very suitable for global analysis. The NN-generated simulation data are then applied to the training of Gaussian process emulators, which in turn are used to evaluate likelihoods in the Bayesian analysis. Even if the training of the networks still requires a considerable amount of full hydrodynamical simulations, the number of simulated events [here 40 events × four collision systems × 1050 parameter sets = $\mathcal{O}(10^5)$ events in total] is orders of magnitude less than what would be needed in computing simulations directly for the training of GP emulators [$10^5 \times 4 \times 1050 = \mathcal{O}(10^8)$ events in total].

Here we have demonstrated the feasibility of the framework by performing a Bayesian analysis on the parameters of the EbyE-EKRT + (2 + 1)D viscous hydrodynamics model with dynamical freeze-out conditions. We incorporate data from four collision systems and, thanks to the neural networks, are able to include statistically demanding observables such as quadrangular flow v_4 and the normalized symmetric cumulant NSC(4, 2). We obtain tightest constraints for the initial-state saturation strength $K_{\rm sat}$ and the nucleon width σ_n, chemical freeze-out temperature $T_{\rm chem}$, and, most importantly, as the QCD matter property, for the temperature dependence of η/s. In particular, we find the data favoring a relatively narrow nucleon width $0.35 < \sigma_n < 0.47$ fm, chemical freeze-out in the range 143–156 MeV, and η/s with minimum-value plateau extending from 150 MeV to at least 230 MeV, with $0.12 < (\eta/s)_{\rm min} < 0.18$. Bulk viscous coefficient ζ/s is nonzero at $T = 200$–300 MeV but otherwise

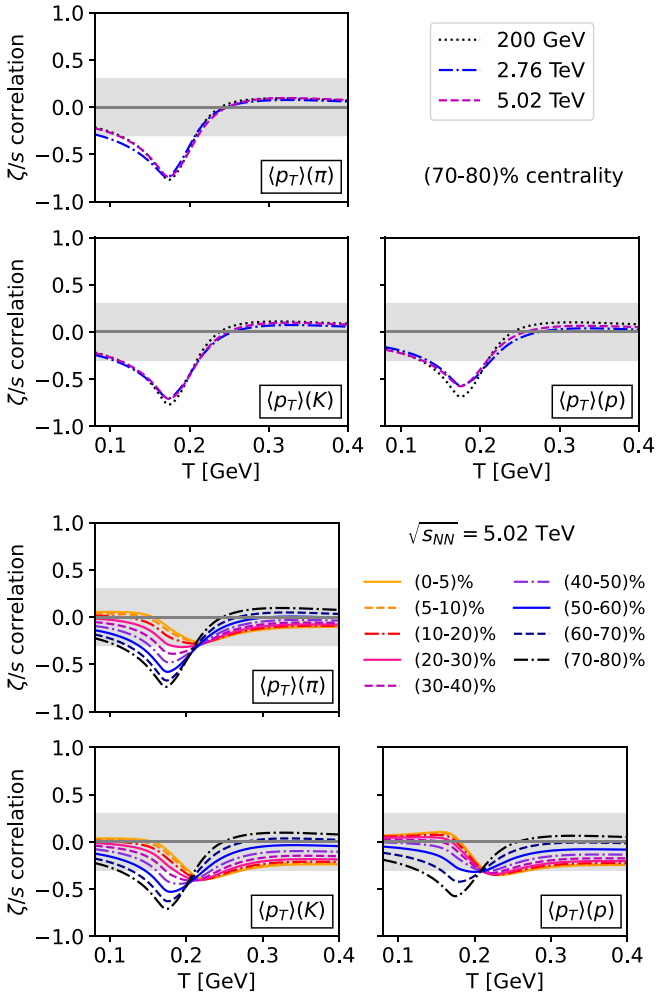

FIG. 17. Temperature dependence of Pearson correlations between mean transverse momenta and ζ/s. Top three panels: Collision energy dependence of correlations at 70–80% centrality. Bottom three panels: Centrality dependence of correlations at $\sqrt{s_{NN}} = 5.02$ TeV collision energy.

remains unconstrained. Overall, our findings for $(\eta/s)(T)$ and $(\zeta/s)(T)$ are consistent with the other recent Bayesian analyses [33,37,42,46]. Our results are also robust against the variations in theoretical uncertainty between 10% and 30% of the experimental data values.

The uncertainties are also reflected in the correlations between parameters and observables, as the most constrained parameters also have most notable correlations with the observables included in the analysis. Investigating the correlations of observables with transport coefficients η/s and ζ/s at different temperatures, we find most notable correlations at temperatures ≈150–250 MeV. For η/s, the strongest constraints in this temperature range appear to come from the mean transverse momenta of kaons and protons, as well as the flow observables $v_n\{2\}$, while the v_2-v_4 correlator NSC(4, 2) and mean transverse momentum of pions are more sensitive to η/s at low temperatures. For ζ/s, the constraining observable is $\langle p_T \rangle$, which shows a very similar correlation peak between 150 and 200 MeV for all three investigated particle species.

There remains some room for improvement both in the model and the analysis. For a better agreement with the data, implementing further nuclear deformation effects is likely required, as well as updating the initial-state model to the MC-EKRT model [14] with hotspots [15]. To obtain further constraints in the future analyses, we aim to incorporate even more statistically demanding observables such as NSC(4, 3) or mixed harmonic cumulants, and consequently we may need to further improve the architecture of the neural networks and the quality of the training data, for example, utilizing active learning approach [34].

ACKNOWLEDGMENTS

J.A. acknowledges support by the program Excellence Initiative–Research University of the University of Wrocław of the Ministry of Education and Science. The project is cofinanced by the Polish National Agency for Academic Exchange, Proposal No. BNI/ULM/2024/1/00179/DEC/I. H.H. is supported by Vanderbilt University and by the U.S. Department of Energy, Office of Science, under Awards No. DE-SC-0024347 and No. DE-SC-0024711, and in part by the National Science Foundation under Grant No. DMS-2406870. We acknowledge the financial support from the Research Council of Finland Project No. 330448 (K.J.E.). This research was funded as a part of the Center of Excellence in Quark Matter of the Research Council of Finland (Projects No. 346325 and No. 364192). This research is part of the European Research Council Project No. ERC-2018-ADG-835105 YoctoLHC. We acknowledge the computation resources from the Finnish IT Center for Science (CSC), Project No. jyy2580, and from the Finnish Computing Competence Infrastructure (FCCI), persistent identifier urn:nbn:fi:research-infras-2016072533.

DATA AVAILABILITY

The code package used for the neural-network training and the event generation is available [94].

APPENDIX: EMULATOR VALIDATION

Following Ref. [95], we define the emulator prediction error as

$$\mathcal{E} \equiv \sqrt{\left\langle \left(\frac{\text{prediction} - \text{truth}}{\text{truth}} \right)^2 \right\rangle}, \tag{A1}$$

and the emulator error estimate accuracy as

$$\mathcal{H} \equiv \ln \left(\sqrt{\left\langle \left(\frac{\text{prediction} - \text{truth}}{\text{GP uncertainty}} \right)^2 \right\rangle} \right), \tag{A2}$$

where the angle brackets indicate the average over test points and the GP uncertainty σ_{GP} is defined in Eq. (19).

We summarize the results of emulator validation for the 20 points sampled from the initial posterior distribution in Table VI. The quality of emulator predictions is good for all the investigated observables. The RMS errors $\mathcal{E}$ presented here

TABLE VI. Summary of centrality-averaged emulator quality parameters for a test set of 20 points sampled from the initial posterior distribution. $\mathcal{E}$ reflects the emulator prediction error relative to the true value, while $\mathcal{H}$ measures the accuracy of the uncertainty estimate reported by the Gaussian process. The negative and positive values of $\mathcal{H}$ indicate overestimation and underestimation of error, respectively.

Observable	$\langle\mathcal{E}\rangle_{200}$	$\langle\mathcal{H}\rangle_{200}$	$\langle\mathcal{E}\rangle_{2760}$	$\langle\mathcal{H}\rangle_{2760}$	$\langle\mathcal{E}\rangle_{5023}$	$\langle\mathcal{H}\rangle_{5023}$	$\langle\mathcal{E}\rangle_{5440}$	$\langle\mathcal{H}\rangle_{5440}$
N(ch)	0.018	−0.261	0.014	−0.306	0.014	−0.293	0.017	−0.567
$N(\pi)$	0.024	−0.155	0.016	−0.275	0.016	−0.255	–	–
$N(K)$	0.027	0.001	0.017	−0.155	0.017	−0.186	–	–
$N(p)$	0.035	−0.185	0.030	−0.114	0.027	−0.200	–	–
$\langle p_T\rangle$(ch)	–	–	–	–	–	–	0.008	−0.327
$\langle p_T\rangle(\pi)$	0.014	−0.222	0.015	−0.054	0.016	−0.057	–	–
$\langle p_T\rangle(K)$	0.010	−0.361	0.009	−0.434	0.009	−0.328	–	–
$\langle p_T\rangle(p)$	0.013	−0.164	0.010	−0.341	0.012	−0.145	–	–
v_2	0.029	−0.206	0.017	−0.096	0.017	0.040	0.017	−0.094
v_3	0.053	−0.106	0.025	−0.066	0.026	−0.173	0.022	−0.141
v_4	0.084	−0.120	0.036	−0.139	0.040	−0.001	0.034	−0.257
NSC(4, 2)	0.088	−0.203	0.052	−0.108	0.053	−0.014	–	–

are on average ≈40% smaller compared to the initial analysis with 900 training points sampled from the prior. Notably the emulator error is systematically overestimated for this set; this happens because in the hyperparameter tuning one has to account also for the edges of the prior, where the quality of predictions degrade due to sparsity of training points.

APPENDIX B: CLOSURE TESTS

Closure tests serve as a consistency check for the full analysis procedure. For the target data, we take the model output at the same 20 points which were used for the GP emulator quality tests and check if the analysis is able to recover the parameter values used to produce the data. Here the relative uncertainties associated with the target data are taken to be the same as in the corresponding experimental data, while the theoretical uncertainty can be considered negligible when comparing to model data. We summarize the results in Table VII.

We obtain the narrowest credible intervals (<25% of the prior width on average) for $K_{\rm sat}$, σ_n, $(\eta/s)_{\rm min}$, $T_{\rm chem}$, and $C_{\rm Kn}/(\eta/s)_{\rm min}$. These intervals are particularly conservative for $K_{\rm sat}$ and σ_n, which have their true values within 20% credible interval in 40% of the trials. On the other extreme is the parameter T_H which has a very poor inference accuracy; the marginal posterior distribution is both wide (the 90% credible interval is ≈70% of the prior width on average) and imprecise (true value is within 90% credible interval in only 65% of the cases).

TABLE VII. Probabilities for finding the closure test truth value in a specified credible interval after 20 trials. The last column shows the average width of the 90% credible interval as a fraction of the prior range of parameter values.

	P(Truth in credible interval)			90% C.I.
Parameter	20% C.I.	68% C.I.	90% C.I.	Avg % of prior
$K_{\rm sat}$	40%	75%	100%	16.6%
σ_n	45%	85%	95%	18.0%
T_H	10%	40%	65%	69.4%
$(\eta/s)_{\rm min}$	25%	85%	100%	12.9%
P_H	10%	40%	70%	27.7%
$\Delta(\eta/s)_H$	5%	60%	85%	36.0%
S_Q	15%	65%	90%	77.6%
$W_{\rm min}$	25%	55%	75%	63.4%
$(\zeta/s)_{\rm max}$	10%	35%	90%	44.6%
$T^{\zeta/s}_{\rm max}$	15%	45%	85%	53.4%
$a_{\zeta/s}$	20%	60%	85%	75.7%
$(\zeta/s)_{\rm width}$	20%	65%	90%	77.1%
$T_{\rm chem}$	15%	75%	100%	17.0%
$C_{\rm Kn}/(\eta/s)_{\rm min}$	5%	60%	85%	23.9%
$C_R/(\eta/s)_{\rm min}$	15%	55%	85%	44.9%

APPENDIX C: NOTABLE PEARSON CORRELATIONS BETWEEN MODEL PARAMETERS AND OBSERVABLES

Tables VIII–XI summarize notable Pearson correlations ($|r| > 0.3$ in at least one centrality class) between the model parameters and observables for the posterior distribution from the analysis with 20% theoretical uncertainty. Upper and lower values in a cell represent the highest and lowest correlations, respectively, among all the investigated centrality classes for each observable.

TABLE VIII. Notable Pearson correlations ($|r| > 0.3$) in at least one centrality class between the model parameters and observables at 200 GeV Au+Au collisions for the posterior distribution from the analysis with 20% theoretical uncertainty. Upper and lower values in a cell represent the highest and lowest correlations, respectively, among all the investigated centrality classes for each observable.

Au+Au 200 GeV	$K_{\rm sat}$	σ_n	T_H	$(\eta/s)_{\rm min}$	P_H	$\Delta(\eta/s)_H$	S_Q	$W_{\rm min}$	$(\zeta/s)_{\rm max}$	$T^{\zeta/s}_{\rm max}$	$a_{\zeta/s}$	$(\zeta/s)_{\rm width}$	$T_{\rm chem}$	$C_{\rm Kn}$	C_R
N(ch)	0.66	0.07	–	–	–	–	–	–	–	–	–	–	−0.33	–	–
	0.41	−0.48	–	–	–	–	–	–	–	–	–	–	−0.49	–	–
$N(\pi)$	0.64	0.39	–	–	–	–	–	–	–	–	–	–	−0.37	–	–
	0.07	−0.49	–	–	–	–	–	–	–	–	–	–	−0.55	–	–
$N(K)$	0.67	0.29	–	–	–	–	–	–	–	–	–	–	−0.19	–	–
	0.06	−0.47	–	–	–	–	–	–	–	–	–	–	−0.46	–	–
$N(p)$	–	–	–	–	–	–	–	–	–	–	–	–	0.70	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	−0.42	–	–
$\langle p_T\rangle(\pi)$	0.52	0.43	–	–	–	0.43	–	–	–	0.30	–	–	0.38	–	–
	0.15	−0.14	–	–	–	0.12	–	–	–	0.12	–	–	−0.34	–	–
$\langle p_T\rangle(K)$	–	0.41	–	0.60	–	–	–	–	–	0.32	–	–	0.40	0.39	–
	–	−0.09	–	0.34	–	–	–	–	–	−0.01	–	–	−0.27	0.14	–
$\langle p_T\rangle(p)$	–	–	–	0.67	–	–	–	–	–	–	–	–	0.32	0.53	0.38
	–	–	–	0.45	–	–	–	–	–	–	–	–	−0.25	0.22	−0.18
v_2	0.37	–	–	–	–	0.03	–	–	–	–	–	–	0.33	−0.03	0.57
	−0.18	–	–	–	–	−0.31	–	–	–	–	–	–	−0.10	−0.40	−0.29
v_3	–	–	–	0.56	–	–	–	–	–	–	–	–	0.48	0.39	–
	–	–	–	−0.33	–	–	–	–	–	–	–	–	−0.21	−0.07	–
v_4	–	–	–	0.65	–	–	–	0.01	–	–	–	–	0.55	0.48	0.31
	–	–	–	−0.29	–	–	–	−0.31	–	–	–	–	−0.20	0.06	−0.31
NSC(4, 2)	0.30	−0.02	–	0.49	–	−0.07	–	–	–	–	–	–	0.01	–	–
	−0.02	−0.35	–	−0.09	–	−0.42	–	–	–	–	–	–	−0.31	–	–

TABLE IX. Same as Table VIII but for 2.76 TeV Pb+Pb collisions.

Pb+Pb 2.76 TeV	$K_{\rm sat}$	σ_n	T_H	$(\eta/s)_{\rm min}$	P_H	$\Delta(\eta/s)_H$	S_Q	$W_{\rm min}$	$(\zeta/s)_{\rm max}$	$T_{\rm max}^{\zeta/s}$	$a_{\zeta/s}$	$(\zeta/s)_{\rm width}$	$T_{\rm chem}$	$C_{\rm Kn}$	C_R
N(ch)	0.63	0.55	–	–	–	–	–	–	–	–	–	–	−0.20	–	–
	−0.03	−0.42	–	–	–	–	–	–	–	–	–	–	−0.32	–	–
$N(\pi)$	0.62	0.62	–	–	–	–	–	–	–	–	–	–	−0.21	–	–
	−0.16	−0.46	–	–	–	–	–	–	–	–	–	–	−0.42	–	–
$N(K)$	0.62	0.59	–	–	–	–	–	–	–	–	–	–	–	–	–
	−0.18	−0.42	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(p)$	–	0.32	–	–	–	–	–	–	0.09	–	–	–	0.82	–	–
	–	0.07	–	–	–	–	–	–	−0.36	–	–	–	−0.03	–	–
$\langle p_T\rangle(\pi)$	0.41	0.43	–	–	–	0.44	–	–	–	–	–	–	0.40	–	–
	0.17	−0.05	–	–	–	0.12	–	–	–	–	–	–	−0.34	–	–
$\langle p_T\rangle(K)$	–	0.38	–	0.56	–	–	–	–	−0.08	–	–	–	0.39	0.33	–
	–	0.01	–	0.39	–	–	–	–	−0.31	–	–	–	−0.26	0.17	–
$\langle p_T\rangle(p)$	–	–	–	0.66	–	0.09	–	–	−0.14	–	–	–	–	0.47	–
	–	–	–	0.53	–	−0.32	–	–	−0.31	–	–	–	–	0.29	–
v_2	0.64	0.30	–	–	–	–	–	–	–	–	–	–	–	−0.06	–
	−0.23	−0.44	–	–	–	–	–	–	–	–	–	–	–	−0.44	–
v_3	0.52	–	–	−0.04	–	–	–	–	–	–	–	–	–	−0.08	–
	−0.18	–	–	−0.47	–	–	–	–	–	–	–	–	–	−0.39	–
v_4	0.33	–	–	0.06	–	–	–	–	–	–	–	–	–	–	–
	−0.12	–	–	−0.45	–	–	–	–	–	–	–	–	–	–	–
NSC(4, 2)	–	−0.21	–	0.40	–	−0.09	–	–	–	–	–	–	−0.22	0.21	–
	–	−0.47	–	−0.09	–	−0.45	–	–	–	–	–	–	−0.32	−0.46	–

TABLE X. Same as Table VIII but for 5.02 TeV Pb+Pb collisions.

Pb+Pb 5.02 TeV	$K_{\rm sat}$	σ_n	T_H	$(\eta/s)_{\rm min}$	P_H	$\Delta(\eta/s)_H$	S_Q	$W_{\rm min}$	$(\zeta/s)_{\rm max}$	$T_{\rm max}^{\zeta/s}$	$a_{\zeta/s}$	$(\zeta/s)_{\rm width}$	$T_{\rm chem}$	$C_{\rm Kn}$	C_R
N(ch)	0.59	0.60	–	–	–	–	–	–	–	–	–	–	–	–	–
	−0.05	−0.41	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(\pi)$	0.58	0.66	–	–	–	–	–	–	–	–	–	–	−0.18	–	–
	−0.17	−0.45	–	–	–	–	–	–	–	–	–	–	−0.40	–	–
$N(K)$	0.58	0.64	–	–	–	–	–	–	–	–	–	–	–	–	–
	−0.20	−0.40	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(p)$	0.13	0.45	–	–	–	–	–	–	0.07	–	–	–	0.82	–	–
	−0.34	0.08	–	–	–	–	–	–	−0.33	–	–	–	0.10	–	–
$\langle p_T\rangle(\pi)$	0.39	0.41	–	–	–	0.46	–	–	–	0.31	–	–	0.37	–	–
	0.22	−0.04	–	–	–	0.12	–	–	–	−0.01	–	–	−0.31	–	–
$\langle p_T\rangle(K)$	–	0.36	–	0.56	–	–	–	–	−0.07	0.30	–	–	0.39	0.33	–
	–	−0.02	–	0.40	–	–	–	–	−0.37	−0.10	–	–	−0.30	0.16	–
$\langle p_T\rangle(p)$	–	–	–	0.66	–	0.12	–	–	−0.09	–	–	–	–	0.56	–
	–	–	–	0.53	–	−0.32	–	–	−0.35	–	–	–	–	0.33	–
v_2	0.63	0.45	–	−0.15	–	−0.20	–	0.30	–	–	–	–	–	−0.12	–
	−0.12	−0.35	–	−0.33	–	−0.34	–	0.19	–	–	–	–	–	−0.56	–
v_3	0.48	0.27	–	0.01	–	–	–	–	–	–	–	–	–	−0.11	–
	−0.27	−0.32	–	−0.56	–	–	–	–	–	–	–	–	–	−0.38	–
v_4	0.31	0.03	–	0.30	–	–	–	–	–	–	–	–	–	0.37	–
	−0.19	−0.34	–	−0.40	–	–	–	–	–	–	–	–	–	−0.22	–
NSC(4, 2)	–	−0.26	–	–	–	−0.29	–	–	–	–	–	–	–	–	–
	–	−0.48	–	–	–	−0.48	–	–	–	–	–	–	–	–	–

TABLE XI. Same as Table VIII but for 5.44 TeV Xe+Xe collisions.

Xe+Xe 5.44 TeV	$K_{\rm sat}$	σ_n	T_H	$(\eta/s)_{\rm min}$	P_H	$\Delta(\eta/s)_H$	S_Q	$W_{\rm min}$	$(\zeta/s)_{\rm max}$	$T^{\zeta/s}_{\rm max}$	$a_{\zeta/s}$	$(\zeta/s)_{\rm width}$	$T_{\rm chem}$	$C_{\rm Kn}$	C_R
N(ch)	0.60	0.74	–	–	–	–	–	–	–	–	–	–	–	–	–
	−0.26	−0.43	–	–	–	–	–	–	–	–	–	–	–	–	–
$\langle p_T \rangle$(ch)	0.33	0.43	–	–	–	–	–	–	−0.09	–	–	–	0.49	–	–
	0.09	0.10	–	–	–	–	–	–	−0.31	–	–	–	0.01	–	–
v_2	0.67	0.24	–	−0.11	0.30	−0.17	–	–	–	–	–	–	–	−0.43	–
	0.12	−0.43	–	−0.47	0.03	−0.36	–	–	–	–	–	–	–	−0.61	–
v_3	0.52	0.06	–	−0.29	–	–	–	–	–	–	–	–	–	−0.27	–
	0.06	−0.31	–	−0.58	–	–	–	–	–	–	–	–	–	−0.41	–
v_4	–	–	–	−0.34	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	−0.51	–	–	–	–	–	–	–	–	–	–	–

APPENDIX D: NOTABLE PEARSON CORRELATIONS BETWEEN OBSERVABLES AND TRANSPORT COEFFICIENTS AT VARIOUS TEMPERATURES

Tables XII and XIII summarize notable Pearson correlations ($|r| > 0.3$ in at least one centrality class) between model parameters and values of specific shear viscosity η/s and bulk viscosity ζ/s, respectively, at various temperatures, for the posterior distribution from the analysis with 20% theoretical uncertainty.

Similarly, Tables XIV and XV summarize notable Pearson correlations between observables and viscous coefficients at different temperatures. Upper and lower values in a cell represent the highest and lowest correlations, respectively, among all the investigated centrality classes for each observable.

TABLE XII. Notable Pearson correlations between η/s at various temperatures and parameters for the posterior distribution from the analysis with 20% theoretical uncertainty.

	$(\eta/s)(T)$ at $T = X$ MeV				
Parameter	100	150	200	250	300
$K_{\rm sat}$	–	–	–	−0.32	−0.46
σ_n	–	–	–	–	–
T_H	–	–	–	–	–
$(\eta/s)_{\rm min}$	0.31	0.97	0.93	–	–
P_H	–	–	–	–	–
$\Delta(\eta/s)_H$	0.81	–	–	0.32	–
S_Q	–	–	–	–	–
$W_{\rm min}$	–	–	–	−0.38	−0.60
$(\zeta/s)_{\rm max}$	–	–	−0.30	–	–
$T^{\zeta/s}_{\rm max}$	–	–	–	–	–
$a_{\zeta/s}$	–	–	–	–	–
$(\zeta/s)_{\rm width}$	–	–	–	–	–
$T_{\rm chem}$	–	–	–	–	–
$C_{\rm Kn}$	0.51	0.66	0.66	0.41	–
C_R	0.31	–	–	–	–

TABLE XIII. Notable Pearson correlations between ζ/s at various temperatures and parameters for the posterior distribution from the analysis with 20% theoretical uncertainty.

	$(\zeta/s)(T)$ at $T = X$ MeV				
Parameter	100	150	200	250	300
$K_{\rm sat}$	–	–	–	–	–
σ_n	–	−0.37	−0.32	–	–
T_H	–	–	–	–	–
$(\eta/s)_{\rm min}$	–	–	−0.30	–	–
P_H	–	–	–	–	–
$\Delta(\eta/s)_H$	–	–	–	–	–
S_Q	–	–	–	–	–
$W_{\rm min}$	–	–	–	–	–
$(\zeta/s)_{\rm max}$	–	–	0.53	0.69	0.56
$T^{\zeta/s}_{\rm max}$	0.46	–	−0.35	0.66	0.84
$a_{\zeta/s}$	0.63	0.44	–	–	–
$(\zeta/s)_{\rm width}$	0.51	0.56	−0.47	0.34	0.58
$T_{\rm chem}$	−0.37	−0.50	–	–	–
$C_{\rm Kn}$	–	–	–	–	–
C_R	–	–	–	–	–

TABLE XIV. Notable Pearson correlations ($|r| > 0.3$ in at least one centrality class) between specific shear viscosity η/s at various temperatures and observables for the posterior distribution from the analysis with 20% theoretical uncertainty. Upper and lower values in a cell represent the highest and lowest correlations, respectively, among all the investigated centrality classes for each observable.

	Au+Au 200 GeV $(\eta/s)(T)$ at $T=X$ MeV					Pb+Pb 2.76 TeV $(\eta/s)(T)$ at $T=X$ MeV					Pb+Pb 5.02 TeV $(\eta/s)(T)$ at $T=X$ MeV					Xe+Xe 5.44 TeV $(\eta/s)(T)$ at $T=X$ MeV				
	100	150	200	250	300	100	150	200	250	300	100	150	200	250	300	100	150	200	250	300
N(ch)	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(\pi)$	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(K)$	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(p)$	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$\langle p_T\rangle$(ch)	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	0.30	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	0.25	–	–
$\langle p_T\rangle(\pi)$	0.35	–	–	–	–	0.33	–	–	–	–	0.37	–	–	–	–	–	–	–	–	–
	0.15	–	–	–	–	0.14	–	–	–	–	0.09	–	–	–	–	–	–	–	–	–
$\langle p_T\rangle(K)$	–	0.60	0.59	–	–	–	0.56	0.56	–	–	–	0.55	0.55	0.31	–	–	–	–	–	–
	–	0.33	0.35	–	–	–	0.38	0.39	–	–	–	0.38	0.41	0.18	–	–	–	–	–	–
$\langle p_T\rangle(p)$	–	0.68	0.65	–	–	–	0.65	0.64	0.31	–	–	0.65	0.66	0.38	–	–	–	–	–	–
	–	0.45	0.46	–	–	–	0.52	0.53	0.22	–	–	0.53	0.49	0.17	–	–	–	–	–	–
v_2	−0.02	–	–	–	–	−0.16	–	−0.06	−0.12	−0.09	−0.31	−0.15	−0.23	−0.23	−0.19	−0.33	−0.11	−0.23	−0.46	−0.27
	−0.47	–	–	–	–	−0.44	–	−0.32	−0.43	−0.38	−0.52	−0.34	−0.42	−0.53	−0.44	−0.50	−0.47	−0.54	−0.60	−0.48
v_3	–	0.53	0.56	0.32	–	–	−0.05	−0.13	−0.19	0.01	–	0.03	−0.09	−0.12	0.13	–	−0.27	−0.40	−0.43	−0.17
	–	−0.33	−0.29	−0.04	–	–	−0.46	−0.48	−0.48	−0.41	–	−0.58	−0.55	−0.57	−0.46	–	−0.56	−0.64	−0.62	−0.43
v_4	0.33	0.65	0.63	0.33	–	–	0.05	0.07	0.07	–	0.26	0.26	0.32	0.22	0.18	–	−0.32	−0.39	−0.08	–
	−0.25	−0.31	−0.24	−0.03	–	–	−0.44	−0.49	−0.39	–	−0.32	−0.40	−0.47	−0.47	−0.35	–	−0.49	−0.53	−0.39	–
NSC(4, 2)	−0.04	0.44	0.44	–	–	−0.04	0.38	0.36	–	–	−0.30	–	–	–	–	–	–	–	–	–
	−0.30	−0.10	−0.14	–	–	−0.52	−0.14	−0.08	–	–	−0.47	–	–	–	–	–	–	–	–	–

TABLE XV. Like Table XIV but for ζ/s instead of η/s.

	Au+Au 200 GeV $(\zeta/s)(T)$ at T=X MeV					Pb+Pb 2.76 TeV $(\zeta/s)(T)$ at T=X MeV					Pb+Pb 5.02 TeV $(\zeta/s)(T)$ at T=X MeV					Xe+Xe 5.44 TeV $(\zeta/s)(T)$ at T=X MeV				
	100	150	200	250	300	100	150	200	250	300	100	150	200	250	300	100	150	200	250	300
N(ch)	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(\pi)$	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(K)$	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
$N(p)$	0.31	0.51	–	–	–	0.32	0.41	–	0.09	–	–	0.30	–	–	–	–	–	–	–	–
	−0.19	−0.24	–	–	–	−0.22	−0.35	–	−0.30	–	–	−0.32	–	–	–	–	–	–	–	–
$\langle p_T\rangle$(ch)	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	−0.12	−0.24	−0.41	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	−0.32	−0.58	−0.51	–	–
$\langle p_T\rangle(\pi)$	–	−0.00	−0.16	–	–	0.07	0.03	−0.23	–	–	–	0.02	−0.22	–	–	–	–	–	–	–
	–	−0.54	−0.44	–	–	−0.34	−0.58	−0.41	–	–	–	−0.53	−0.46	–	–	–	–	–	–	–
$\langle p_T\rangle(K)$	–	−0.06	−0.23	–	–	–	−0.05	−0.29	–	–	–	0.01	−0.28	0.01	–	–	–	–	–	–
	–	−0.53	−0.46	–	–	–	−0.53	−0.47	–	–	–	−0.51	−0.47	−0.35	–	–	–	–	–	–
$\langle p_T\rangle(p)$	–	0.03	−0.14	–	–	–	0.05	−0.20	–	–	–	0.10	−0.18	−0.02	–	–	–	–	–	–
	–	−0.46	−0.44	–	–	–	−0.39	−0.43	–	–	–	−0.41	−0.39	−0.33	–	–	–	–	–	–
v_2	–	–	0.09	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	−0.31	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
v_3	–	−0.01	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	−0.34	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
v_4	–	0.08	0.22	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	−0.42	−0.32	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
NSC(4, 2)	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–

[1] A. Bazavov *et al.* (HotQCD Collaboration), Equation of state in $(2+1)$-flavor QCD, Phys. Rev. D **90**, 094503 (2014).

[2] A. Bazavov, P. Petreczky, and J. H. Weber, Equation of state in $2+1$ flavor QCD at high temperatures, Phys. Rev. D **97**, 014510 (2018).

[3] S. Borsanyi, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, and K. K. Szabo, Full result for the QCD equation of state with $2+1$ flavors, Phys. Lett. B **730**, 99 (2014).

[4] S. Borsanyi, G. Endrodi, Z. Fodor, A. Jakovac, S. D. Katz, S. Krieg, C. Ratti, and K. K. Szabo, The QCD equation of state with dynamical quarks, J. High Energy Phys. 11 (2010) 077.

[5] J. S. Moreland, J. E. Bernhard, and S. A. Bass, Alternative ansatz to wounded nucleon and binary collision scaling in high-energy nuclear collisions, Phys. Rev. C **92**, 011901 (2015).

[6] B. Schenke, P. Tribedy, and R. Venugopalan, Fluctuating glasma initial conditions and flow in heavy ion collisions, Phys. Rev. Lett. **108**, 252301 (2012).

[7] C. Gale, S. Jeon, B. Schenke, P. Tribedy, and R. Venugopalan, Event-by-event anisotropic flow in heavy-ion collisions from combined Yang-Mills and viscous fluid dynamics, Phys. Rev. Lett. **110**, 012302 (2013).

[8] B. Schenke, S. Jeon, and C. Gale, (3+1)D hydrodynamic simulation of relativistic heavy-ion collisions, Phys. Rev. C **82**, 014903 (2010).

[9] K. J. Eskola, K. Kajantie, P. V. Ruuskanen, and K. Tuominen, Scaling of transverse energies and multiplicities with atomic number and energy in ultrarelativistic nuclear collisions, Nucl. Phys. B **570**, 379 (2000).

[10] K. J. Eskola, K. Kajantie, and K. Tuominen, Centrality dependence of multiplicities in ultrarelativistic nuclear collisions, Phys. Lett. B **497**, 39 (2001).

[11] R. Paatelainen, K. J. Eskola, H. Holopainen, and K. Tuominen, Multiplicities and p_T spectra in ultrarelativistic heavy ion collisions from a next-to-leading order improved perturbative QCD + saturation + hydrodynamics model, Phys. Rev. C **87**, 044904 (2013).

[12] R. Paatelainen, K. J. Eskola, H. Niemi, and K. Tuominen, Fluid dynamics with saturated minijet initial conditions in ultrarelativistic heavy-ion collisions, Phys. Lett. B **731**, 126 (2014).

[13] H. Niemi, K. J. Eskola, and R. Paatelainen, Event-by-event fluctuations in a perturbative QCD + saturation + hydrodynamics model: Determining QCD matter shear viscosity in ultrarelativistic heavy-ion collisions, Phys. Rev. C **93**, 024907 (2016).

[14] M. Kuha, J. Auvinen, K. J. Eskola, H. Hirvonen, Y. Kanakubo, and H. Niemi, Monte Carlo event generator MC-EKRT with saturated minijet production for initializing $(3+1)D$ fluid dynamics in high-energy nuclear collisions, Phys. Rev. C **111**, 054914 (2025).

[15] H. Hirvonen, M. Kuha, J. Auvinen, K. J. Eskola, Y. Kanakubo, and H. Niemi, Effects of saturation and fluctuating hotspots for flow observables in ultrarelativistic heavy-ion collisions, Phys. Rev. C **110**, 034911 (2024).

[16] T. Pierog and K. Werner, EPOS model and ultra high energy cosmic rays, Nucl. Phys. B Proc. Suppl. **196**, 102 (2009).

[17] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner, EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider, Phys. Rev. C **92**, 034906 (2015).

[18] K. Werner, B. Guiot, I. Karpenko, and T. Pierog, Analyzing radial flow features in *p*-Pb and *p*-*p* collisions at several TeV by studying identified-particle production with the event generator EPOS3, Phys. Rev. C **89**, 064903 (2014).

[19] K. Werner, Revealing a deep connection between factorization and saturation: New insight into modeling high-energy proton-proton and nucleus-nucleus scattering in the EPOS4 framework, Phys. Rev. C **108**, 064903 (2023).

[20] K. Werner, Core-corona procedure and microcanonical hadronization to understand strangeness enhancement in proton-proton and heavy ion collisions in the EPOS4 framework, Phys. Rev. C **109**, 014910 (2024).

[21] K. Werner and B. Guiot, Perturbative QCD concerning light and heavy flavor in the EPOS4 framework, Phys. Rev. C **108**, 034904 (2023).

[22] A. Kurkela, A. Mazeliauskas, J. F. Paquet, S. Schlichting, and D. Teaney, Effective kinetic description of event-by-event pre-equilibrium dynamics in high-energy heavy-ion collisions, Phys. Rev. C **99**, 034910 (2019).

[23] A. Kurkela, R. Törnkvist, and K. Zapp, AMY Lorentz invariant parton cascade: The thermal equilibrium case, Eur. Phys. J. C **84**, 74 (2024).

[24] S. A. Bass, M. Belkacem, M. Bleicher, M. Brandstetter, L. Bravina, C. Ernst, L. Gerland, M. Hofmann, S. Hofmann, J. Konopka, *et al.*, Microscopic models for ultrarelativistic heavy ion collisions, Prog. Part. Nucl. Phys. **41**, 255 (1998).

[25] M. Bleicher, E. Zabrodin, C. Spieles, S. A. Bass, C. Ernst, S. Soff, L. Bravina, M. Belkacem, H. Weber, H. Stoecker, *et al.*, Relativistic hadron-hadron collisions in the ultra-relativistic quantum molecular dynamics model, J. Phys. G: Nucl. Part. Phys. **25**, 1859 (1999).

[26] J. Weil *et al.* (SMASH Collaboration), Particle production and equilibrium properties within a new hadron transport approach for heavy-ion collisions, Phys. Rev. C **94**, 054905 (2016).

[27] H. Hirvonen, K. J. Eskola, and H. Niemi, Flow correlations from a hydrodynamics model with dynamical freeze-out and initial conditions based on perturbative QCD and saturation, Phys. Rev. C **106**, 044913 (2022).

[28] U. Heinz and B. Schenke, Hydrodynamic description of the quark-gluon plasma, arXiv:2412.19393.

[29] J. Novak, K. Novak, S. Pratt, J. Vredevoogd, C. Coleman-Smith, and R. Wolpert, Determining fundamental properties of matter created in ultrarelativistic heavy-ion collisions, Phys. Rev. C **89**, 034917 (2014).

[30] J. E. Bernhard, J. S. Moreland, and S. A. Bass, Bayesian estimation of the specific shear and bulk viscosity of quark–gluon plasma, Nat. Phys. **15**, 1113 (2019).

[31] J. Auvinen, K. J. Eskola, P. Huovinen, H. Niemi, R. Paatelainen and P. Petreczky, Temperature dependence of η/s of strongly interacting matter: Effects of the equation of state and the parametric form of $(\eta/s)(T)$, Phys. Rev. C **102**, 044911 (2020).

[32] D. Everett *et al.* (JETSCAPE Collaboration), Phenomenological constraints on the transport properties of QCD matter with data-driven model averaging, Phys. Rev. Lett. **126**, 242301 (2021).

[33] D. Everett *et al.* (JETSCAPE Collaboration), Multisystem Bayesian constraints on the transport coefficients of QCD matter, Phys. Rev. C **103**, 054904 (2021).

[34] R. Ehlers *et al.* (JETSCAPE Collaboration), Bayesian inference analysis of jet quenching using inclusive jet and hadron suppression measurements, Phys. Rev. C **111**, 054913 (2025).

[35] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, Transverse momentum differential global analysis of heavy-ion collisions, Phys. Rev. Lett. **126**, 202301 (2021).

[36] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, Bayesian analysis of heavy ion collisions with the heavy ion computational framework Trajectum, Phys. Rev. C **103**, 054909 (2021).

[37] G. Nijs and W. van der Schee, Hadronic nucleus-nucleus cross section and the nucleon size, Phys. Rev. Lett. **129**, 232301 (2022).

[38] G. Nijs and W. van der Schee, A generalized hydrodynamizing initial stage for Heavy Ion Collisions, arXiv:2304.06191.

[39] G. Giacalone, G. Nijs, and W. van der Schee, Determination of the neutron skin of ^{208}Pb from ultrarelativistic nuclear collisions, Phys. Rev. Lett. **131**, 202302 (2023).

[40] H. Mäntysaari, B. Schenke, C. Shen, and W. Zhao, Bayesian inference of the fluctuating proton shape, Phys. Lett. B **833**, 137348 (2022).

[41] M. R. Heffernan, C. Gale, S. Jeon, and J. F. Paquet, Bayesian quantification of strongly interacting matter with color glass condensate initial conditions, Phys. Rev. C **109**, 065207 (2024).

[42] M. R. Heffernan, C. Gale, S. Jeon, and J. F. Paquet, Early-times Yang-Mills dynamics and the characterization of strongly interacting matter with statistical learning, Phys. Rev. Lett. **132**, 252301 (2024).

[43] S. A. Jahan, H. Roch, and C. Shen, Bayesian analysis of $(3+1)D$ relativistic nuclear dynamics with the RHIC beam energy scan data, Phys. Rev. C **110**, 054905 (2024).

[44] J. E. Parkkila, A. Onnerstad, and D. J. Kim, Bayesian estimation of the specific shear and bulk viscosity of the quark-gluon plasma with additional flow harmonic observables, Phys. Rev. C **104**, 054904 (2021).

[45] J. E. Parkkila, A. Onnerstad, S. F. Taghavi, C. Mordasini, A. Bilandzic, M. Virta, and D. J. Kim, New constraints for QCD matter from improved Bayesian parameter estimation in heavy-ion collisions at LHC, Phys. Lett. B **835**, 137485 (2022).

[46] M. Virta, J. Parkkila, and D. J. Kim, Enhancing Bayesian parameter estimation by adapting to multiple energy scales in heavy-ion collisions at RHIC and at the LHC, Phys. Rev. C **111**, 044903 (2025).

[47] N. Götz, I. Karpenko, and H. Elfner, Bayesian analysis of a $(3+1)D$ hybrid approach with initial conditions from hadronic transport, Phys. Rev. C **112**, 014910 (2025).

[48] J. F. Paquet, Applications of emulation and Bayesian methods in heavy-ion physics, J. Phys. G: Nucl. Part. Phys. **51**, 103001 (2024).

[49] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, MA, 2005).

[50] J. E. Bernhard, Bayesian parameter estimation for relativistic heavy-ion collisions, Ph.D. thesis, Duke University, 2018.

[51] D. S. Everett, Quantifying the Quark Gluon Plasma, Ph.D. thesis, The Ohio State University, 2021.

[52] H. Hirvonen, K. J. Eskola, and H. Niemi, Deep learning for flow observables in ultrarelativistic heavy-ion collisions, Phys. Rev. C **108**, 034905 (2023).

[53] H. Hirvonen, K. J. Eskola, and H. Niemi, Deep learning for flow observables in high energy heavyion collisions, EPJ Web Conf. **296**, 02002 (2024).

[54] J. A. Sun, L. Yan, C. Gale, and S. Jeon, End-to-end generative diffusion model for heavy-ion collisions, Phys. Rev. C **112**, L051903 (2025).

[55] M. Omana Kuttan, K. Zhou, J. Steinheimer, and H. Stoecker, Ultrafast, event-by-event heavy-ion simulations for next-generation experiments, Phys. Rev. C **112**, 054907 (2025).

[56] D. Stewart and J. Putschke, Fast prediction of the hydrodynamic QGP evolution in ultrarelativistic heavy-ion collisions using Fourier neural operators, Phys. Rev. C **113**, 014904 (2026).

[57] K. J. Eskola and K. Tuominen, Production of transverse energy from minijets in next-to-leading order perturbative QCD, Phys. Lett. B **489**, 329 (2000).

[58] K. J. Eskola and K. Tuominen, Transverse energy from minijets in ultrarelativistic nuclear collisions: A next-to-leading order analysis, Phys. Rev. D **63**, 114006 (2001).

[59] K. J. Eskola, K. Kajantie, and J. Lindfors, Quark and gluon production in high energy nucleus-nucleus collisions, Nucl. Phys. B **323**, 37 (1989).

[60] L. V. Gribov, E. M. Levin, and M. G. Ryskin, Semihard processes in QCD, Phys. Rept. **100**, 1 (1983).

[61] I. Helenius, K. J. Eskola, H. Honkanen, and C. A. Salgado, Impact-parameter dependent nuclear parton distribution functions: EPS09s and EKS98s and their applications in nuclear hard processes, J. High Energy Phys. 07 (2012) 073.

[62] P. Möller, A. J. Sierk, T. Ichikawa, and H. Sagawa, Nuclear ground-state masses and deformations: FRDM(2012), At. Data Nucl. Data Tables **109-110**, 1 (2016).

[63] H. Niemi, K. J. Eskola, R. Paatelainen, and K. Tuominen, Predictions for 5.023 TeV Pb + Pb collisions at the CERN Large Hadron Collider, Phys. Rev. C **93**, 014912 (2016).

[64] K. J. Eskola, H. Niemi, R. Paatelainen, and K. Tuominen, Predictions for multiplicities and flow harmonics in 5.44 TeV Xe + Xe collisions at the CERN Large Hadron Collider, Phys. Rev. C **97**, 034911 (2018).

[65] P. Huovinen and P. Petreczky, QCD equation of state and hadron resonance gas, Nucl. Phys. A **837**, 26 (2010).

[66] P. Huovinen, Chemical freeze-out temperature in the hydrodynamical description of Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV, Eur. Phys. J. A **37**, 121 (2008).

[67] W. Israel and J. M. Stewart, Transient relativistic thermodynamics and kinetic theory, Ann. Phys. **118**, 341 (1979).

[68] G. S. Denicol, T. Koide, and D. H. Rischke, Dissipative relativistic fluid dynamics: A new way to derive the equations of motion from kinetic theory, Phys. Rev. Lett. **105**, 162501 (2010).

[69] G. S. Denicol, H. Niemi, E. Molnar, and D. H. Rischke, Derivation of transient relativistic fluid dynamics from the Boltzmann equation, Phys. Rev. D **85**, 114047 (2012); **91**, 039902(E) (2015).

[70] E. Molnár, H. Niemi, G. S. Denicol, and D. H. Rischke, Relative importance of second-order terms in relativistic dissipative fluid dynamics, Phys. Rev. D **89**, 074010 (2014).

[71] G. S. Denicol, S. Jeon, and C. Gale, Transport coefficients of bulk viscous pressure in the 14-moment approximation, Phys. Rev. C **90**, 024912 (2014).

[72] K. Paech and S. Pratt, Origins of bulk viscosity in relativistic heavy ion collisions, Phys. Rev. C **74**, 014901 (2006); **93**, 059902(E) (2016).

[73] K. Dusling and T. Schäfer, Bulk viscosity, particle spectra, and flow in heavy-ion collisions, Phys. Rev. C **85**, 044909 (2012).

[74] J. B. Rose, J. M. Torres-Rincon, and H. Elfner, Inclusive and effective bulk viscosities in the hadron gas, J. Phys. G: Nucl. Part. Phys. **48**, 015005 (2021).

[75] H. Hirvonen, Probing properties of quark-gluon plasma using machine learning, Ph.D. thesis, University of Jyväskylä, 2024.

[76] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, in *Proceedings of 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Los Alamitos, CA, 2017), pp. 2261–2269.

[77] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980.

[78] S. S. Adler *et al.* (PHENIX Collaboration), Systematic studies of the centrality and $\sqrt{s_{NN}}$ dependence of the $dE_T/d\eta$ and $dN_{\rm ch}/d\eta$ in heavy ion collisions at midrapidity, Phys. Rev. C **71**, 034908 (2005); **71**, 049901(E) (2005).

[79] K. Aamodt *et al.* (ALICE Collaboration), Centrality dependence of the charged-particle multiplicity density at midrapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, Phys. Rev. Lett. **106**, 032301 (2011).

[80] J. Adam *et al.* (ALICE Collaboration), Centrality dependence of the charged-particle multiplicity density at midrapidity in Pb − Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, Phys. Rev. Lett. **116**, 222302 (2016).

[81] S. Acharya *et al.* (ALICE Collaboration), Transverse momentum spectra and nuclear modification factors of charged particles in Xe–Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV, Phys. Lett. B **788**, 166 (2019).

[82] S. S. Adler *et al.* (PHENIX Collaboration), Identified charged particle spectra and yields in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV, Phys. Rev. C **69**, 034909 (2004).

[83] B. Abelev *et al.* (ALICE Collaboration), Centrality dependence of π, K, and p production in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, Phys. Rev. C **88**, 044910 (2013).

[84] S. Acharya *et al.* (ALICE Collaboration), Production of charged pions, kaons, and (anti-)protons in Pb-Pb and inelastic pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV, Phys. Rev. C **101**, 044907 (2020).

[85] L. Adamczyk *et al.* (STAR Collaboration), Harmonic decomposition of three-particle azimuthal correlations at energies available at the BNL Relativistic Heavy Ion Collider, Phys. Rev. C **98**, 034918 (2018).

[86] S. Acharya *et al.* (ALICE Collaboration), Energy dependence and fluctuations of anisotropic flow in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ and 2.76 TeV, J. High Energy Phys. 07 (2018) 103.

[87] S. Acharya *et al.* (ALICE Collaboration), Anisotropic flow in Xe–Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV, Phys. Lett. B **784**, 82 (2018).

[88] J. Adam *et al.* (STAR Collaboration), Correlation measurements between flow harmonics in Au + Au collisions at RHIC, Phys. Lett. B **783**, 459 (2018).

[89] J. Adam *et al.* (ALICE Collaboration), Correlated event-by-event fluctuations of flow harmonics in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, Phys. Rev. Lett. **117**, 182301 (2016).

[90] S. Acharya *et al.* (ALICE Collaboration), Measurements of mixed harmonic cumulants in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, Phys. Lett. B **818**, 136354 (2021).

[91] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine Learning in Python, J. Mach. Learn. Res. **12**, 2825 (2011).

[92] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Emcee: The MCMC hammer, Publ. Astron. Soc. Pac. **125**, 306 (2013).

[93] S. Chekanov *et al.* (ZEUS Collaboration), Exclusive photoproduction of J/ψ mesons at HERA, Eur. Phys. J. C **24**, 345 (2002).

[94] H. Hirvonen, J. Auvinen, H. Niemi, K. J. Eskola, Zenodo, 2026, doi: 10.5281/zenodo.20499599.

[95] H. Roch, S. A. Jahan, and C. Shen, Model emulation and closure tests for $(3+1)$D relativistic heavy-ion collisions, Phys. Rev. C **110**, 044904 (2024).