Mapping properties of the quark gluon plasma in Pb-Pb and Xe-Xe collisions at energies available at the CERN Large Hadron Collider

L. Vermunt [1,2,*] Y. Seemann,[1,†] A. Dubla [2,‡] S. Floerchinger,[3,§] E. Grossi,[4,∥]
A. Kirchner [5,¶] S. Masciocchi,[1,2,**] and I. Selyuzhenkov [2,††]

[1]*Physikalisches Institut, Universität Heidelberg, 69120 Heidelberg, Germany*
[2]*GSI Helmholtzzentrum für Schwerionenforschung, 64291 Darmstadt, Germany*
[3]*Theoretisch-Physikalisches Institut, Friedrich-Schiller-Universität Jena, 07743 Jena, Germany*
[4]*Dipartimento di Fisica, Università di Firenze and INFN Sezione di Firenze, 50019 Sesto Fiorentino, Italy*
[5]*Institut für Theoretische Physik Heidelberg, 69120 Heidelberg, Germany*

A phenomenological analysis of the experimental measurements of transverse momentum spectra of identified charged hadrons and strange hyperons in Pb-Pb and Xe-Xe collisions at the CERN Large Hadron Collider (LHC) is presented. The analysis is based on the relativistic fluid dynamics description implemented in the numerically efficient FLUID*u*M approach. Building on our previous work, we separate in our treatment the chemical and kinetic freeze-outs, and incorporate the partial chemical equilibrium to describe the late stages of the collision evolution. This analysis makes use of Bayesian inference to determine key parameters of the quark-gluon plasma (QGP) evolution and its properties including the shear and bulk viscosity to entropy ratios, the initialization time, the initial entropy density, and the freeze-out temperatures. The physics parameters and their posterior probabilities are extracted using a global search in multidimensional space with modern machine learning tools, such as ensembles of neural networks. We employ our newly developed fast framework to assess systematic uncertainties in the extracted model parameters by systematically varying key components of our analysis.

DOI: [10.1103/PhysRevC.108.064908](https://doi.org/10.1103/PhysRevC.108.064908)

I. INTRODUCTION

Heavy-ion collisions at ultrarelativistic energies have been at the forefront of modern physics research for several decades. These collisions, studied at facilities such as the BNL Relativistic Heavy Ion Collider (RHIC) and the CERN Large Hadron Collider (LHC), create an extreme state of matter known as the quark-gluon plasma (QGP) [1–4]. This fluid is of great interest because it is described by a renormalizable and fundamental quantum field theory at the microscopic level, i.e., quantum chromodynamics (QCD). While the macroscopic fluid properties remain challenging to calculate from first principles and a limited number of computations exist to

———————
*luuk.vermunt@cern.ch
†yannis.seemann@gmail.com
‡a.dubla@cern.ch
§stefan.floerchinger@uni-jena.de
∥eduardo.grossi@unifi.it
¶kirchner@thphys.uni-heidelberg.de
**s.masciocchi@gsi.de
††ilya.selyuzhenkov@gmail.com

date, an increasing number of experimental results motivate phenomenological and theoretical studies. One of the key features of the QGP is its collective behavior, which is thought to arise from the fluidlike properties of the system. However, recent experimental observations of collective behavior in proton-nucleus and proton-proton collision systems [5–8] have challenged the uniqueness of the fluidlike response of the QGP. The resolution of the origins of collective behavior in heavy-ion collisions will likely rely on a quantitative rather than just qualitative agreement between data and models.

Traditionally, physical properties of QCD matter have been determined by comparing experimental data with model calculations of event-averaged and predefined observables. However, recent developments have shown promise in extracting more information from the final-state particles produced in heavy-ion collisions. Two such methods are Bayesian analysis [9–20] and deep learning [21,22]. Bayesian analysis uses global fitting to simultaneously determine multiple model parameters, utilizing all available experimental data. On the other hand, deep learning identifies observables that are sensitive to specific physical properties, enabling the extraction of relevant information. This work improves upon our previous analysis of Ref. [23] by using Bayesian inference as the optimality criterion. The previous approach of minimizing the χ^2 value had several limitations, including the neglect of experimental data correlations, interpolation uncertainties, and parameter correlations. In addition, we improved our theoretical description with respect to Ref. [23] by separating

the chemical and kinematic freeze-outs, and incorporating the partial chemical equilibrium to describe the later stages of the evolution, as well as by exploiting a parametrization from Yang-Mills theory for the shear viscosity to entropy ratio of the QGP [24,25] instead of using a constant value.

Our model for simulating high-energy nuclear collisions combines three distinct components. The TRENTO model [26] was utilized for the initial conditions, while the FLUIDuM model [27], featuring a mode splitting technique for very fast computations, was used for the relativistic fluid dynamic expansion with viscosity. Additionally, the FASTRESO code [28] was used to take resonance decays into account. Despite FLUIDuM being significantly faster than event-by-event hydrodynamic codes, it is still beneficial to develop an emulator that can function as a quick substitute for the full model. Where previous Bayesian analyses [9–18] utilized Gaussian process regression, we for the first time exploit neural network emulation. It has been demonstrated in Ref. [29] that infinitely wide neural networks can approximate Gaussian process regression. Furthermore, the use of neural networks has several computational benefits, including significantly reduced training time, lower memory usage, and the ability to handle any number of inputs and outputs. This newly developed framework will have additional applications in the precise fitting of charm and beauty observables contributing to further constrain the heavy-quark spatial diffusion coefficient, widening our knowledge of QGP transport properties [30,31].

In the present work, we determine the specific shear and bulk viscosity to entropy ratios of the QGP, as well as the freeze-out temperatures $T_{\rm kin}$ and $T_{\rm chem}$, the starting time of a fluid description τ_0, and the normalization of the initial entropy profile. We employ our newly developed fast framework to comprehensively explore systematic uncertainties in the extracted model parameters by systematically varying critical components of our analysis. We compare against experimental measurements of transverse momentum spectra of identified charged hadrons (π, K, p) and strange hyperons (Λ) with $p_{\rm T} < 2$ GeV/c in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV and $\sqrt{s_{NN}} = 5.02$ TeV, and Xe-Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV from the ALICE Collaboration [32–35]. Given that this paper primarily serves as a proof of concept for our new Bayesian inference framework, our study will be restricted to the 0–5% centrality range. We specifically chose a very central bin because we expect the background-fluctuation splitting ansatz that underlies FLUIDuM to work best for central collisions, where the profiles tend to be the most symmetric. Furthermore, the restriction to only one centrality class was imposed such that we are able to keep the initial-state parameters, which play an important role in the centrality dependence of the observables, fixed. In a forthcoming publication, we will explore the use of anisotropic flow observables, as well as experimental data from different centrality classes, to gain further insight into the key parameters of the QGP. Nevertheless, it is important to underscore that through the analysis of transverse momentum spectra exclusively, we can derive significant constraints on essential aspects such as the bulk viscosity, the freeze-out temperatures, τ_0, and the normalization. This work highlights the potential of such analyses to provide valuable insights into the properties of the QGP.

This paper is organized as follows. We summarize the details of the initial conditions, the hydrodynamic evolution, and the hadronization procedures in Sec. II. The procedure of the Bayesian analysis is discussed in Sec. III. We then discuss the extraction of the model parameters in Sec. IV, and end with a summary and future directions in Sec. V.

II. MODELING OF THE DIFFERENT STAGES OF A HEAVY-ION COLLISION

The following section provides a brief overview of the various components of our theoretical model. We first examine the initial conditions, which are determined using the TRENTO model [26]. Subsequently, we turn to the time evolution as implemented in FLUIDuM [27], which solves the equations of relativistic fluid dynamics with shear and bulk viscosity and corresponding relaxation times. The newly introduced partial chemical equilibrium will be discussed next, together with the kinetic freeze-out and the implementation of strong resonance decays performed with FASTRESO [28].

A. Initial conditions: TRENTO

As in our previous work [23], we use the TRENTO model parametrization [26] for the initial conditions. This is an effective model, intended to generate realistic Monte Carlo initial transverse entropy (or energy) profiles without assuming specific physical mechanisms. It involves positioning nucleons with a Gaussian width w using a fluctuating Glauber model, while ensuring a minimum distance d between them. Each nucleon contains m randomly placed constituents with a Gaussian width of v. TRENTO uses an entropy deposition parameter p that interpolates among qualitatively different physical mechanisms for entropy production [26]. Furthermore, additional multiplicity fluctuations are introduced by multiplying the density of each nucleon by random weights sampled from a gamma distribution with unit mean and shape parameter k.

For this study, the TRENTO parameters are not estimated via the Bayesian analysis. Instead, they are set based on the current state of knowledge in literature. As reviewed extensively in Ref. [36], we set $w = 0.5$ fm, $m = 4$, $v = 0.4$ fm, $p = 0$, and use the TRENTO output as entropy density. In addition, we set $k = 1$ and $d = 0.75$ fm, based on the outcome of the Bayesian analysis of Ref. [13]. For the nucleon-nucleon cross section, the measurements by the ALICE Collaboration [37] are used, i.e., $x = 61.8$, 67.6, and 68.4 mb for Pb-Pb at $\sqrt{s_{NN}} = 2.76$ TeV, Pb-Pb at $\sqrt{s_{NN}} = 5.02$ TeV, and Xe-Xe at $\sqrt{s_{NN}} = 5.44$ TeV, respectively. The Pb ion is sampled from a spherically symmetric Woods-Saxon distribution with radius $R = 6.65$ fm and surface thickness $a = 0.54$ fm, while the Xe ion comes from a deformed spheroidal Woods-Saxon distribution with $R = 5.60$ fm, $a = 0.49$ fm, and deformation parameters $\beta_2 = 0.21$ and $\beta_4 = 0.0$ [38]. Using this set of parameters (which we call the "central" configuration in the remaining) we have generated the transverse density $T_{\rm R}(x, y)$ for 1.5×10^6 minimum-bias events with impact parameter sampled from the range $b \in [0 \, {\rm fm}, 20 \, {\rm fm}]$.

To classify the TRENTO events in centrality classes, we utilize the integrated transverse density $\int d^2x\, T_R(\vec{x})$, which is expected to have a linear monotonic relationship with multiplicity [39]. This allows us to divide the events into narrow multiplicity classes of 1%, each of which can be treated as an ensemble of events with random orientation in the transverse plane as in experiment. For each centrality class, we determine the averaged or expected entropy density profile as

$$s(r) = \frac{\text{Norm}}{\tau_0}\langle T_R(r)\rangle, \quad (1)$$

where the average $\langle\cdots\rangle$ is taken over all the events in the class with a random reaction plane angle. Consequently, the average is independent of the azimuthal angle ϕ by construction. We have introduced a centrality-class-dependent normalization constant Norm to account for possible variation of the fits from the TRENTO multiplicity scaling. We account for the longitudinal expansion effect (Bjorken flow) at early times by scaling the Norm by the initialization time τ_0. To enlarge the statistic of the average entropy profile $\langle T_R(r)\rangle$ (as we will see in the next section, it can be taken independently of ϕ for this work), we perform the event average over all events in one centrality class:

$$\langle T_R(r)\rangle = \frac{1}{2\pi}\int_0^{2\pi} d\phi\,\langle T_R(r,\phi)\rangle. \quad (2)$$

As in Ref. [23], we produce the averaged entropy densities for the larger experimental centrality classes by averaging the corresponding distributions from the more narrow classes and propagating those.

Note that a free-streaming phase that evolves the energy density profile at $\tau = 0$ for a short timescale to finite τ_0 is currently not included in our framework. To mimic the effect of a free-streaming phase, which can be seen as a suppression/dilution of the "spikiness" of the initial entropy density profiles, we also test another set of TRENTO parameters (called the "free-steaming" configuration later on) in which the width of the nucleons is increased (i.e., $w = 1.0$ fm, $v = 0.9$ fm, and $d = 1.25$ fm). A larger nucleon width in TRENTO is observed to suppress the spikiness in the initial entropy density. Systematic variations of the other TRENTO parameters are not considered.

B. Hydrodynamic evolution: FLUID*u*M

The software package FLUID*u*M [27], which utilizes a theoretical framework based on relativistic fluid dynamics with mode expansion [40–42], is used to solve the equations of motion for relativistic fluids. This involves decomposing the fluid fields into background and fluctuation components. Specifically, the fluid fields are represented as $\Phi(\tau, r, \phi, \eta) = \Phi_0(\tau, r) + \epsilon\,\Phi_1(\tau, r, \phi, \eta)$, where Φ_0 is the background solution, Φ_1 is the perturbation around it, and ϵ is the formal expansion parameter (set to $\epsilon \to 1$ at the end). The background and perturbations can be solved accurately and efficiently using numerical algorithms [27]. The background-fluctuation splitting is justified due to two statistical symmetries that are approximately observed in high-energy nuclear collisions. First, the collision energy is

sufficiently high, leading to an approximate boost invariance of the system at midrapidity. This implies that observables at midrapidity exhibit only a mild dependence on the rapidity, and consequently the dependence from it can be a perturbation. Second, there is a statistical symmetry related to the random orientation of the reaction plane angle. In each collision event, the two ions collide with a particular relative orientation in the reaction plane, which is uncorrelated with other events. When calculating observables averaged over multiple events, the azimuthal angle's dependence is effectively removed, unless the reaction plane angle is explicitly reconstructed and corrected for each event. Moreover, the average will always be independent of the azimuthal angle, and every dependence form is a perturbation.

For our study, we are interested in examining the azimuthally averaged transverse momentum spectra of identified particles at midrapidity. Therefore, we do not consider azimuthally and rapidity-dependent perturbations and only require the background solution to the fluid evolution equations, neglecting terms of quadratic or higher order in perturbation amplitudes. The causal equations of motion are obtained from second-order Israel-Stewart hydrodynamics [43].

We introduce a novel aspect here with respect to Ref. [23], by assuming the shear viscosity to entropy ratio η/s to be temperature dependent. In particular, we exploit the calculation in Yang-Mills theory of Ref. [24] with recently updated parameters [25]. The calculation provides an analytic fit formula describing the temperature dependence of η/s as a direct sum of a glueball resonance gas contribution with an algebraic decay at small temperatures and a high-temperature contribution consistent with hard-thermal-loop (HTL) resummed perturbation theory. The fit function in SU(3) Landau gauge Yang-Mills theory is

$$\frac{\eta}{s}(T)_{\text{YM}} = a\left(\frac{T}{T_c} - d\right)^2 + \frac{b}{(T/T_c)^\delta}. \quad (3)$$

The first term has been changed with respect to Ref. [24] for simplicity since the differences do not play a role for hydrodynamic applications [25]. The best fit to the full Yang-Mills results is given by the parameters $a = 0.0613$, $b = 0.00588$, $d = -0.709$, and $\delta = 40.3$. In the low-temperature regime, the pure glueball resonance gas is not replaced by a hadron resonance gas; the b and δ parameters are adjusted to 0.02 and 6.0 to capture this regime better. Finally, an overall correction factor of 4/3 is used to take the differences in scales and the running couplings in Yang-Mills theory and QCD into account [24]. On top of this, we add a global scaling $(\eta/s)_{\text{scale}}$ as a parameter to be estimated with the Bayesian analysis

$$\frac{\eta}{s}(T)_{\text{QCD}} = (\eta/s)_{\text{scale}}\frac{4}{3}\left[a\left(\frac{T}{T_c} - d\right)^2 + \frac{0.02}{(T/T_c)^6}\right]. \quad (4)$$

The Yang-Mills and QCD η/s distributions are presented in the left panel of Fig. 1. We would like to emphasise that our η/s parametrization is consistently applied throughout both the partonic and hadronic phases. In contrast, other Bayesian analyses [9–20] commonly adopt a different approach, switching to the URQMD [44] or SMASH [45]
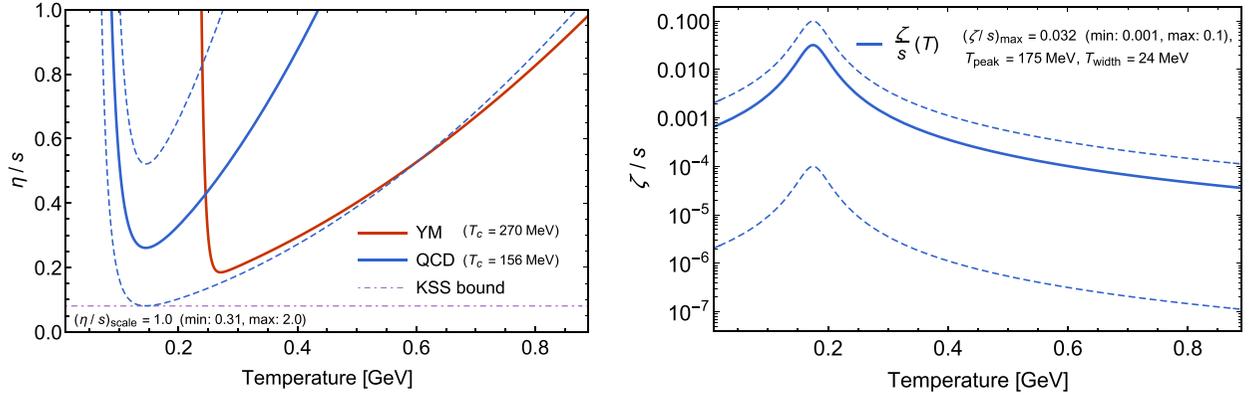
FIG. 1. Temperature dependence of the shear (left) and bulk (right) viscosity to entropy ratios as defined in Eqs. (4), and (5), respectively. Both the Yang-Mills and QCD η/s distributions are shown, as well as the Kovtun-Son-Starinet (KSS) bound from anti–de Sitter and conformal field theory (AdS-CFT) calculations, $\eta/s = 1/4\pi$. The dashed blue lines indicate the minimum and maximum values used in the Bayesian analysis for $(\eta/s)_{\mathrm{scale}}$ and $(\zeta/s)_{\mathrm{max}}$.

transport models at temperature T_{switch} (typically around 145–155 MeV), where $(\eta/s)(T_{\mathrm{switch}}) \approx 1$ is employed.

The bulk viscosity to entropy ratio ζ/s is also considered temperature dependent. We assume it to be of the Lorentzian form (shown in the right panel of Fig. 1)

$$\frac{\zeta}{s}(T) = \frac{(\zeta/s)_{\mathrm{max}}}{1 + \left(\frac{T - T_{\mathrm{peak}}}{T_{\mathrm{width}}}\right)^2}. \tag{5}$$

For now, the peak temperature T_{peak} and width T_{width} are fixed to 175 MeV and 24 MeV, respectively, based on Ref. [9]. The maximum temperature $(\zeta/s)_{\mathrm{max}}$ is a free parameter for the Bayesian analysis. We would like to remark that we are aware of the calculations of Refs. [46,47] (and its parametrization in Ref. [48]), where the bulk viscosity coefficient shows a maximum near the QCD phase transition temperature T_c and starts to decrease almost exponentially while approaching the hadron gas model, as in our assumption. On the other hand, the temperature dependence of the bulk viscosity in SU(3) gluodynamics is also studied within lattice QCD simulations by Ref. [49], which shows that, below T_c, ζ/s continues to rise. The same behavior is observed for a pion gas in Ref. [50] where the bulk viscosity to entropy ratio rapidly increases at low temperatures. This rise might be assigned to the rapid decrease of the entropy density below the transition in SU(3) gluodynamics, where the (de)confinement phase transition is of the first order, while in QCD it is a crossover. A bulk viscosity that keeps increasing as the temperature drops will lead to a large bulk viscous pressure on the freeze-out surface and might result in some nonphysical results during particlization via the Cooper-Frye procedure. Since a consensus on the shape of the ζ/s as a function of the T is not yet available in the literature, we choose to keep using the bulk parametrization as in Eq. (5) already used in our previous work [23].[1] On the freeze-out surface we take the particle distribution function to be given by the equilibrium

Bose-Einstein or Fermi-Dirac distribution (depending on the species), modified by additional corrections due to bulk and shear viscous dissipation,

$$f = f_{\mathrm{eq}} + \delta f^{\mathrm{bulk}} + \delta f^{\mathrm{shear}}. \tag{6}$$

For the viscous corrections we use the commonly employed parametrizations [23,51,52]

$$\delta f^{\mathrm{bulk}} = f_{\mathrm{eq}}(1 \pm f_{\mathrm{eq}})\left[\frac{\bar{E}_p}{T}\left(\frac{1}{3} - c_s^2\right) - \frac{m^2}{3T\bar{E}_p}\right]\frac{\pi_{\mathrm{bulk}}}{\zeta/\tau_{\mathrm{bulk}}}, \tag{7}$$

$$\delta f^{\mathrm{shear}} = f_{\mathrm{eq}}(1 \pm f_{\mathrm{eq}})\frac{\pi_{\rho\nu}p^\rho p^\nu}{2(\epsilon + p)T^2}. \tag{8}$$

Here m is the mass of the primary resonance.

The bulk and shear relaxation times are taken as derived in Ref. [53]:

$$\frac{\tau_{\mathrm{bulk}}}{\zeta/(\epsilon + p)} = \frac{1}{15\left(\frac{1}{3} - c_s^2\right)^2} + \frac{a_{\mathrm{offset}}}{\zeta/(\epsilon + p)}, \tag{9}$$

$$\frac{\tau_{\mathrm{shear}}}{\eta/(\epsilon + p)} = \begin{cases} 5 & \text{if } T \geqslant T_{\mathrm{chem}}, \\ 5 + (T - T_{\mathrm{chem}})a_{\mathrm{slope}} & \text{if } T < T_{\mathrm{chem}}, \end{cases} \tag{10}$$

where ϵ is the energy density, p is the pressure, c_s is the (temperature dependent) velocity of sound, and $a_{\mathrm{offset}} = 0.1$ fm/c is a small offset such that a causal evolution of the radial expansion is ensured [43]. For the same reason, we adjusted the relation between η and τ_{shear} for temperatures $T < T_{\mathrm{chem}}$, incorporating an additional $(T - T_{\mathrm{chem}})a_{\mathrm{slope}}$ term with respect to Ref. [53]. This was necessary to ensure that the relaxation time is much larger than the characteristic scale of the hadron resonance gas, for which we expect the scatterings to become more sparse. A value of $a_{\mathrm{slope}} = 3$ MeV^{-1} was considered for the central analysis after stability checks on the high-p_{T} region ($p_{\mathrm{T}} > 2$ GeV/c) of the transverse momentum spectra.

---

[1]It is at present also unclear how bulk viscous effects interfere with the chemical freeze-out and a phase of partial chemical equilibrium

between chemical and kinetic freeze-outs. We plan to revisit this topic in future work.

TABLE I. The predefined parameter intervals for the six model parameters across the three collision systems. The Norm and τ_0 are considered to be system-dependent parameters. Note that the intervals for the τ_0 parameters for Pb-Pb $\sqrt{s_{NN}} = 5.02$ TeV and Xe-Xe $\sqrt{s_{NN}} = 5.44$ TeV were adjusted *a posteriori*.

	$(\zeta/s)_{\mathrm{max}}$	$(\eta/s)_{\mathrm{min}}$	T_{chem} (MeV)	T_{kin} (MeV)	Norm	τ_0 (fm/c)
Pb-Pb (2.76 TeV)					5–80	0.01–3.0
Pb-Pb (5.02 TeV)	10^{-4}–0.3	0.08–0.52	130–155	110–140	80–140	2.0–7.0
Xe-Xe (5.44 TeV)					70–150	2.0–7.0

C. Freeze-out and resonance decays: FASTRESO

As the system cools down and dilutes, it changes from a QGP to a hadronic gas. Because particle scatterings are no longer efficient in maintaining equilibrium, the fluid dynamic description breaks down, necessitating the conversion of fluid fields to the distribution of hadronic degrees of freedom. The Cooper-Frye procedure is used to convert fluid fields to the spectrum of hadron species on a freeze-out surface, which in our work is assumed to be a surface of constant temperature [54].

In our previous work [23], a single freeze-out was considered, where both the chemical and kinetic distributions freeze out simultaneously at temperature T_{freeze}. Here, we improve on this description by introducing a partial chemical equilibrium (PCE) after the chemical freeze-out and before the kinetic freeze-out, i.e., $T_{\mathrm{kin}} < T < T_{\mathrm{chem}}$. During the PCE, the mean free time for elastic collisions is still smaller than the characteristic expansion time of the expanding fireball, thereby keeping the gas in a state of local kinetic equilibrium. The chemical equilibrium is not maintained since the mean free path of the inelastic collisions exceeds this threshold.

Our description follows the pioneering work of Refs. [55,56], in which the different particle species in a hadronic gas are treated as being in chemical equilibrium with each other, while the overall gas is not. The effective number of long living hadrons, taken in our case as the ones with a characteristic lifetime longer than 10 fm/c (the approximate lifetime of the system [57]), are fixed at the values they had at the chemical freeze-out. The term "effective" includes the hadrons which would be produced when all unstable resonances decay, i.e., $\bar{N}_i = N_i + \sum_j b_j^i N_j$. Here, N_i is the number of hadrons i, N_j the number of resonances j, and b_j^i the number of hadrons i formed in the decay of resonance j (including the branching ratio). The corresponding effective chemical potential is given in a similar way. We then assume the isentropic evolution of ideal hydrodynamics (i.e., entropy is conserved), meaning that the ratio of effective particle number density ($\bar{n}_i = \bar{N}_i/V$) and entropy density is constant until the kinetic freeze-out. From the relation

$$\frac{\bar{n}_i(T, \mu)}{s(T, \mu)} = \frac{\bar{n}_i(T_{\mathrm{chem}}, 0)}{s(T_{\mathrm{chem}}, 0)}, \quad (11)$$

one can then obtain $\mu_i(T)$.

On the freeze-out surface we take the particle distribution function to be given by the equilibrium Bose-Einstein or Fermi-Dirac distribution (depending on the species), modified by additional corrections due to bulk and shear viscous dissipation and decays of unstable resonances, as explained in detail in our previous work [23]. For the viscous corrections, we use commonly employed parametrizations [51,52], while the resonance decays are efficiently calculated with the FASTRESO package [28]. We use a list of approximately 700 resonances from Refs. [58–60].

III. PROCEDURE OF THE BAYESIAN ANALYSIS

As outlined in Sec. II, our central framework revolves around certain free parameters: the overall normalization constant Norm, $(\eta/s)_{\mathrm{min}}$ and $(\zeta/s)_{\mathrm{max}}$ values[2] in the shear and bulk viscosity to entropy ratio parametrizations, the initial fluid time τ_0, and the two freeze-out temperatures T_{kin} and T_{chem}. While Norm and τ_0 are considered system-dependent parameters, the others are assumed to converge to the same values for different collision systems and energies. The primary objective of our Bayesian analysis is to determine these six model parameters simultaneously, allowing them to vary within predefined intervals (see Table I). These intervals are based on physical considerations and knowledge from previous studies [23,25,32,61,62]. It is worth mentioning that we have confirmed *a posteriori* that the optimal values fall within these intervals rather than on their boundaries, and, in cases where no clear convergence was obtained, larger intervals were employed.

In contrast to previous Bayesian analyses in the field of heavy-ion physics that employed Gaussian process regression [9–18], we introduce a novel approach by utilizing neural network emulation. This paper marks the first application of neural network emulation in this context. The subsequent two subsections provide a concise introduction to our newly developed framework, which combines neural networks and Markov-chain Monte Carlo simulations. For a more comprehensive understanding of both components, we refer the reader to Ref. [63], where a detailed overview is presented.

A. Neural network implementation

Although FLUIDuM is recognized for its very fast execution speeds, the extensive parameter exploration involved in Bayesian analyses necessitates an approach to speed up the simulations. One common strategy is to train machine

---

[2]For convenience, the $(\eta/s)_{\mathrm{scale}}$ parameter will be converted from now on to $(\eta/s)_{\mathrm{min}}$, representing the minimum of the shear viscosity-to-entropy parametrization, always located at $T_{\mathrm{min}} = 145$ MeV [24,25].

learning models to emulate the complete model and utilize these fast emulators within the Markov-chain Monte Carlo simulations. In our study, we employ an ensemble of artificial neural networks (ANNs) to emulate the TRENTO + FLUIDuM + FASTRESO model. A neural network is a computational algorithm in the field of supervised learning inspired by the functioning of the brain. When appropriately constructed and trained, NNs can effectively identify and model complex relationships between inputs and outputs, making them highly promising tools for our research.

In addition to speed, the emulator also needs to exhibit a high level of accuracy. Achieving this requires careful consideration of the neural network's architecture and the training process. The training necessitates large datasets to achieve the required accuracy for replacing the simulation outputs. For each collision system, we use the outputs of ten thousand complete FLUIDuM + FASTRESO simulations, with parameters distributed within the ranges presented in Table I. The parameter values are generated using Latin hypercube sampling, which ensures a uniform density in an efficient way. In order to enhance the training performance of the neural networks, we apply a normalization technique to both the input parameters and the output values, scaling them to the range of $[-1, 1]$. This procedure leads to a significant increase in the convergence rate and speed of the training process. The emulator model incorporates this normalization, automatically converting the input parameters to the appropriate scale before feeding them through the networks. Similarly, the resulting outputs are scaled back to their original ranges.

Simple neural networks provide point predictions without any measure of uncertainty or confidence, which can arise from, e.g., insufficient model complexity or missing information due to unknown data. Ensemble methods offer an alternative approach where predictions are not solely reliant on a single model; instead, they combine predictions from multiple diverse models within an ensemble. The predictions, denoted as f_i, of the ensemble members $i \in 1, 2, \ldots, M$ are averaged for the model input $\boldsymbol{x}$,

$$f_{\text{emu}}(\boldsymbol{x}) = \frac{1}{M} \sum_{i=1}^{M} f_i(\boldsymbol{x}), \tag{12}$$

while the spread among the different ensemble members[3] and their mean correlation ρ are used to estimate the error on the ensemble prediction via

$$\sigma_{\text{emu}}(\boldsymbol{x}) = \sqrt{\frac{\frac{1}{M} + \frac{M-1}{M}\rho}{1 - \rho}} \sqrt{\frac{1}{M} \sum_{i=1}^{M} [f_i(\boldsymbol{x}) - f_{\text{emu}}(\boldsymbol{x})]}$$

$$= c\hat{\sigma}_{\text{emu}}(\boldsymbol{x}). \tag{13}$$

The correlation among the different neural networks effectively introduces a correction factor, denoted as c, to the standard deviation of the neural network predictions $\hat{\sigma}_{\text{emu}}(\boldsymbol{x})$.

———
[3]It was tested with a Shapiro-Wilk test that the distribution of the individual neural network predictions can be assumed everywhere to come from a normal distribution [63].
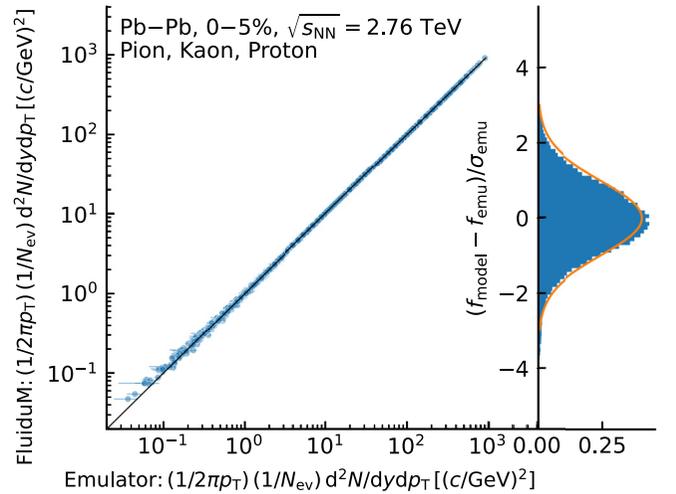

FIG. 2. Correlation between the NN ensemble prediction and the original FLUIDuM + FASTRESO output for ten different parameter configurations $\boldsymbol{x}$. Each data point represents the $dN^2/dy\,dp_T$ value in one experimental p_T interval for either pions, kaons, or protons measured in the 0–5% most central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [32]. On the right, the residual distribution is shown in units of the corrected (see text for more details) neural network ensemble uncertainty. A normal distribution with mean 0 and standard deviation 1 is shown in orange to guide the eye.

This correction factor is determined by fitting a t distribution to $[f_{\text{model}}(\boldsymbol{x}) - f_{\text{emu}}(\boldsymbol{x})]/\hat{\sigma}_{\text{emu}}(\boldsymbol{x})$ (here f_{model} represents the original FLUIDuM + FASTRESO output), which should ideally follow a standard normal distribution if the prediction uncertainty accurately captures the prediction error. We assume that the correlation is independent of the position in parameter space and the considered p_T intervals, allowing the fit to be performed once, considering all configurations $\boldsymbol{x}$ and data points.

For each collision system, we train 100 neural networks using the input data, splitting the data sample into 80% for training and 20% for testing and validation purposes. The PYTORCH Python library [64] is employed to construct the neural networks, while the hyperparameters of the neural networks are optimized using a grid search approach facilitated by the TUNE Python library [65]. The nature of our problem favours shallow neural networks [1–3 hidden layers with $\mathcal{O}(100)$ nodes], a learning rate around 0.01, and the use of a leaky rectified linear unit activation function.

The neural network ensemble effectively captures the true output of the model simulations, as demonstrated in Fig. 2. The figure presents the correlation between the $(1/2\pi p_T)(1/N_{\text{ev}})\,d^2N/dy\,dp_T$ values within experimental p_T intervals for pions, kaons, and protons, as estimated by the emulator model and originally simulated for ten different model parameter configurations $\boldsymbol{x}$. It is evident that the emulator accurately reproduces the FLUIDuM + FASTRESO output, and the emulator uncertainties, which have been appropriately corrected for the leftover correlation among the individual neural networks (as discussed above), effectively capture the residual spread. Furthermore, our ensemble consisting

of 100 neural networks (NNs) exhibits a prediction error, measured in units of the experimental data uncertainty, of $[f_{\mathrm{emu}}(\boldsymbol{x}) - f_{\mathrm{model}}(\boldsymbol{x})]/\sigma_{\mathrm{exp}} = 2.5 \times 10^{-3}$. This level of accuracy is considered sufficient for the purposes of this study. Additionally, our emulator significantly reduces the computation time by a factor 10^4.

B. Markov-chain Monte Carlo

To obtain the posterior probability densities of the N model parameters, we employ Bayesian inference. The posterior density is inferred from a probabilistic model, and $(N - 1)$-dimensional integrals are performed to obtain marginal probability densities for each parameter. Due to the complexity of our case, an analytic treatment is not feasible. Therefore, we employ the numerical Markov-chain Monte Carlo (MCMC) method, which is the most efficient approach for exploring the probability space. In this method, samples are drawn randomly but not independently by constructing a so-called Markov chain. Each element of the chain is sampled in dependence on its preceding element and only its preceding element.

For our MCMC sampling, we utilize the EMCEE Python library [66], which implements the affine-invariant ensemble sampler with the so-called "stretch move." The logarithmic probability is quantified as the log of the posterior probability,

$$\log[P(\boldsymbol{f}|\boldsymbol{x})] \propto \begin{cases} -\frac{1}{2}[\boldsymbol{f}_{\mathrm{emu}}(\boldsymbol{x}) - \boldsymbol{f}_{\mathrm{exp}}]^{\intercal} \boldsymbol{\Sigma}^{-1}[\boldsymbol{f}_{\mathrm{emu}}(\boldsymbol{x}) - \boldsymbol{f}_{\mathrm{exp}}] & \text{if } x_i^{\min} \leqslant x_i \leqslant x_i^{\max} \; \forall i, \\ -\infty & \text{otherwise.} \end{cases} \tag{14}$$

Here, x_i represents the ith model input parameter, $x_i^{\min}$ and $x_i^{\max}$ denote their limits (as defined in Table I), $\boldsymbol{f}_{\mathrm{exp}}$ represents a vector of all the experimental data, $\boldsymbol{f}_{\mathrm{emu}}(\boldsymbol{x})$ is the corresponding vector for the emulator model for input $\boldsymbol{x}$, and $\boldsymbol{\Sigma}$ represents the covariance matrix, which consists of the data and model covariance matrices ($\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_{\mathrm{exp}} + \boldsymbol{\Sigma}_{\mathrm{emu}}$). Since no information about the correlations between the experimental uncertainties exists, $\boldsymbol{\Sigma}_{\mathrm{exp}}$ is diagonal with its elements given by the squared sum of the statistical and systematic uncertainties. The model covariance matrix is computed from the ensemble output as

$$\Sigma_{\mathrm{emu}}^{j,k}(\boldsymbol{x}) = c \frac{1}{M - 1}$$
$$\times \sum_{i=1}^{M} \left[\left(f_i^j(\boldsymbol{x}) - f_{\mathrm{emu}}^j(\boldsymbol{x}) \right) \left(f_i^k(\boldsymbol{x}) - f_{\mathrm{emu}}^k(\boldsymbol{x}) \right) \right], \tag{15}$$

where the iterators j and k denote specific output values of the vectors. All entries have to be scaled with the correction factor c, as discussed in Sec. III A, to account for the correlation among the different neural networks in the ensemble. In our MCMC sampling, the prior probability distribution is given by a uniform distribution,

$$p(\boldsymbol{x}) \propto \begin{cases} 1 & \text{if } x_i^{\min} \leqslant x_i \leqslant x_i^{\max} \; \forall i, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

To ensure that the Markov chain has converged sufficiently, we require the sampling error of the MCMC method to be smaller than 1%. The sampling error decreases by $\sqrt{\tau_f/N_{\mathrm{samples}}}$, where τ_f is the integrated autocorrelation time [63]. This implies that the product of the number of walkers (chosen to be 300 in our case) and the length of the chains must be greater than $10\,000\tau_f$. To prevent too early stopping, we also require that the change of τ_f (calculated every 100 MCMC steps) is smaller than 1%.

In order to provide the reader with an initial demonstration of the power of our emulator-MCMC approach, Fig. 3

presents a comparison between the simulated and emulated model predictions and the corresponding experimental data for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV [32]. The top row displays the output of the ten thousand FLUIDuM + FASTRESO simulations used for the training of the NNs, and the bottom row presents the predictions generated by the neural-network ensemble emulator using $n = 100$ random parameter configurations sampled from the Bayesian posterior distribution obtained through the MCMC chain. While the training points exhibit a significant spread, which is expected as the parameters are varied over wide ranges (see Table I), the emulator predictions sampled from the posterior distributions are significantly more constrained and closely follow the experimental data points.

IV. RESULTS

In this study, we focus on comparing against the experimental measurements of transverse momentum spectra of identified charged hadrons and strange hyperons in the 0–5% most central Pb-Pb and Xe-Xe collisions. In contrast to traditional Bayesian inference analyses, which often focus on deriving a single "best-fit" scenario, our approach centres on using our framework to assess additional sources of systematic uncertainties, such as potential physics that is not included/parametrized in our model, the selection of experimental data, and the treatment of possible unknown correlated uncertainties in the experimental data. Through this extensive exploration of the parameter space, we aim to better understand the underlying dynamics of the system. We emphasise that, in our analysis, all configurations are equally valid, and we do not aim to single out one central result as the definitive outcome.

The soft particle momentum range $p_{\mathrm{T}} < 2$ GeV/c (with variations up to $p_{\mathrm{T}} < 3$ GeV/c) always serves as our primary window of investigation, as it is believed to be governed by fluid dynamic approximations of QCD dynamics. This momentum range is sensitive to important factors such as radial flow, viscous transport coefficients, and the initial conditions
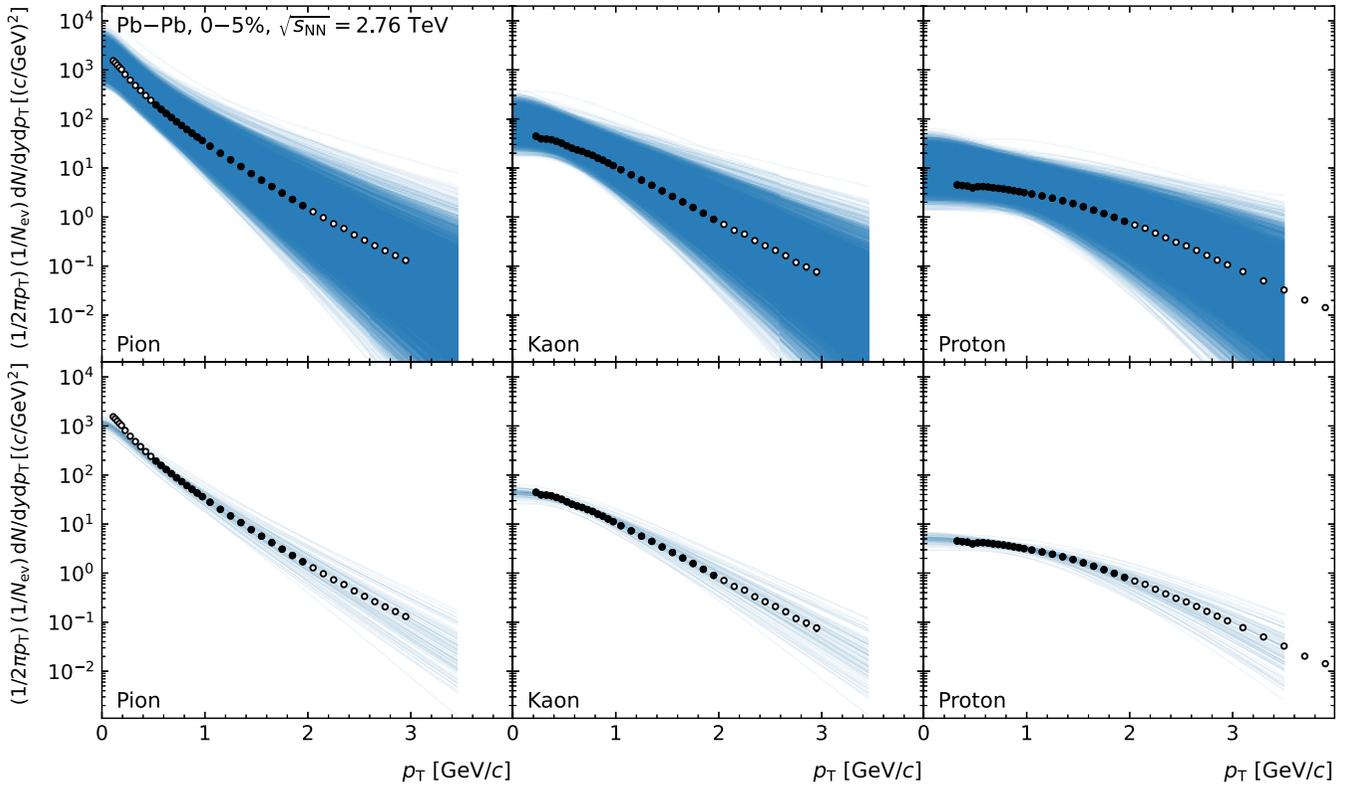
FIG. 3. Comparison between simulated (top row; $n = 10\,000$) or emulated (bottom row; $n = 100$) FLUID*u*M + FASTRESO predictions with experimental data for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV in the 0–5% centrality interval from the ALICE Collaboration [32]. Note that the solid markers show the experimental data used in the MCMC procedure, while the open markers show the full experimental measurement. See text for more details.

of the QGP [67–70]. Similarly to our previous work [23], our default setup excludes pions in the $p_T < 0.5$ GeV/c range due to the well-known enhancement relative to hydrodynamic simulations [14,71]. This low-p_T pion excess is believed to arise from physics features not fully accounted for in hydrodynamic simulations, such as (i) Bose-Einstein condensation [72,73], (ii) increased population of resonances [74], (iii) correct treatment of the finite width of ρ meson [75], and (iv) effects of critical chiral fluctuations [76].

A. Comparison of different collision system configurations

We start by exploring the impact of different collision systems on the constraints of our model parameter. The Bayesian posterior distributions are presented in Fig. 4. The diagonal panels display the marginalized distributions of each individual model parameters, while the off-diagonal panels illustrate the joint distributions for pairs of these parameters marginalized over all others. The contours represent the $(0.5, 1, 1.5, 2)\sigma$ equivalent regions, encompassing 11.8%, 39.3%, 67.5%, and 86.4% of the samples. In addition, the median values and 68% confidence intervals of the individual model parameters are reported in Table II. Specifically, we present three configurations: (i) the combination of the Pb-Pb at $\sqrt{s_{NN}} = 2.76$ TeV, Pb-Pb at $\sqrt{s_{NN}} = 5.02$ TeV, and Xe-Xe at $\sqrt{s_{NN}} = 5.44$ TeV collision systems in blue; (ii) the Pb-Pb at $\sqrt{s_{NN}} = 2.76$ TeV and Xe-Xe case in red; and

(iii) only the $\sqrt{s_{NN}} = 2.76$ TeV data in green. The experimental data incorporate the pion ($0.5 < p_T < 2$ GeV/c), kaon ($0.2 < p_T < 2$ GeV/c), and proton ($0.3 < p_T < 2$ GeV/c) $(1/2\pi p_T)(1/N_{ev})\,d^2N/dy\,dp_T$ spectra in the 0–5% centrality class. The central TRENTO configuration is used.

Remarkably, even with a limited number of p_T spectra observables, we find that most model parameters are well constrained, underscoring the potential of Bayesian analyses in providing valuable insights into the properties of the QGP. The one parameter that appears to prefer higher values beyond the upper bound of the allowed interval is the minimum value of the shear viscosity to entropy ratio parametrization. We attribute this behavior to the limited sensitivity of the current observables to the shear viscosity of the system. Future work incorporating experimental flow data is expected to further address this point. However, we emphasise that our η/s parametrization is consistently applied throughout both the partonic and hadronic phases, possibly leading to higher minimum values compared to other Bayesian analyses, where an afterburner is used for the hadronic phase [9,10,12–20].

When comparing the three scenarios, most parameters are compatible within their uncertainties (which naturally decrease when more experimental data is incorporated). However, noticeable differences emerge, particularly in the maximum temperature for the bulk viscosity to entropy ratio and the chemical freeze-out temperature, which appear to be influenced by the inclusion of the data for Pb-Pb at $\sqrt{s_{NN}} =$
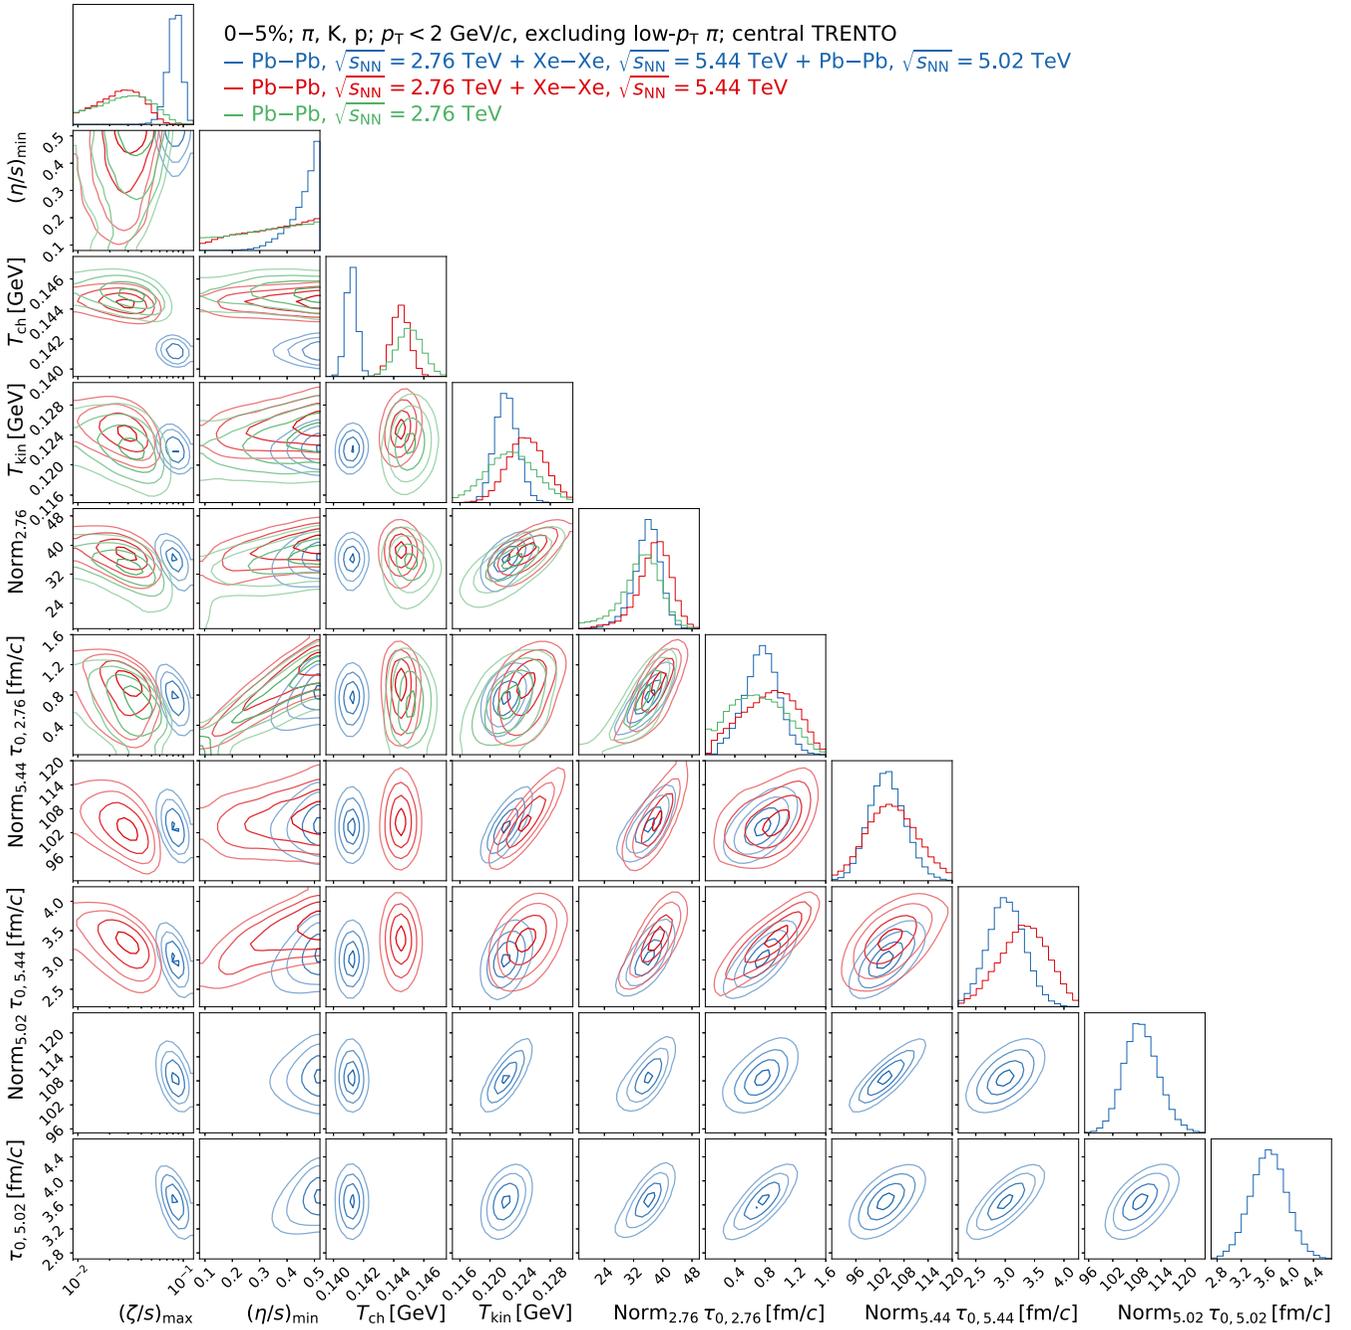
FIG. 4. Bayesian posterior distribution of the model input parameters for three different combinations of the collision systems. The diagonal panels show the marginalized distributions of individual model parameters, while the off-diagonal panels present the correlations among pairs of model parameters. See text for more details.

5.02 TeV, the experimental data which we observed to be the most difficult to fit. Since no immediate physical explanation is available, we recommend further investigation in a future analysis, especially as additional experimental observables are included. Another notable distinction is observed in the Norm and τ_0 parameters for the various collision systems, with a substantial increase between the Pb-Pb at $\sqrt{s_{NN}} = 2.76$ TeV and the other two systems.

B. Effect of modified initial conditions and inclusion of strange hyperons

Shifting the focus to two other crucial aspects of our analysis, we transition to Fig. 5, where we explore the implications of the following scenarios: (i) a modification of the initial TRENTO configuration to mimic a free-streaming phase, as discussed in Sec. II A, and (ii) the inclusion of the experimental $(1/2\pi p_T)(1/N_{ev})\,\mathrm{d}^2N/\mathrm{d}y\,\mathrm{d}p_T$ data of the strange hyperon

TABLE II. Posterior parameter estimates corresponding to Fig. 4. The reported values correspond to the median values and 68% confidence intervals.

	2.76–5.44–5.02 TeV	2.76–5.44 TeV	2.76 TeV
$(\zeta/s)_{\max}$	$0.083^{+0.012}_{-0.012}$	$0.022^{+0.016}_{-0.014}$	$0.023^{+0.022}_{-0.016}$
$(\eta/s)_{\min}$	unconstr.	unconstr.	unconstr.
T_{chem} (MeV)	141^{+0}_{-0}	144^{+1}_{-1}	145^{+1}_{-1}
T_{kin} (MeV)	122^{+2}_{-1}	125^{+3}_{-2}	123^{+3}_{-3}
Norm$_{2.76}$	$36.0^{+3.1}_{-3.6}$	$38.2^{+3.8}_{-4.3}$	$34.6^{+4.5}_{-5.3}$
$\tau_{0,2.76}$ (fm/c)	$0.76^{+0.21}_{-0.23}$	$0.87^{+0.34}_{-0.40}$	$0.69^{+0.39}_{-0.37}$
Norm$_{5.44}$	$103.7^{+4.6}_{-4.4}$	$104.9^{+6.7}_{-6.2}$	n/a
$\tau_{0,5.44}$ (fm/c)	$3.01^{+0.31}_{-0.31}$	$3.32^{+0.40}_{-0.44}$	n/a
Norm$_{5.02}$	$109.0^{+4.5}_{-4.1}$	n/a	n/a
$\tau_{0,5.02}$ (fm/c)	$3.66^{+0.30}_{-0.30}$	n/a	n/a

Λ ($0.6 < p_T < 2$ GeV/c). We compare the resulting posterior distributions with the single Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV case discussed earlier.

When employing the modified TRENTO settings to mimic the free-streaming phase, we observe a notable decrease in the values of the Norm and τ_0 parameters. This decrease can be attributed to the larger integrated transverse density $\int d^2x \, T_R(\vec{x})$ in this configuration. Because of the relationship expressed by Eq. (1), there exists a correlation between the magnitudes of the Norm and τ_0 parameters. Consequently, interpreting the decrease of τ_0 in a physical context becomes challenging, and the smaller value indicates only partly the (artificial) smoothening out of the initial entropy density profile by reducing its "spikiness."

Furthermore, the inclusion of the Λ hyperon data leads to an increase in the chemical freeze-out temperature. This finding is consistent with previous observations using hydrodynamic simulations [23,61], suggesting that strange and
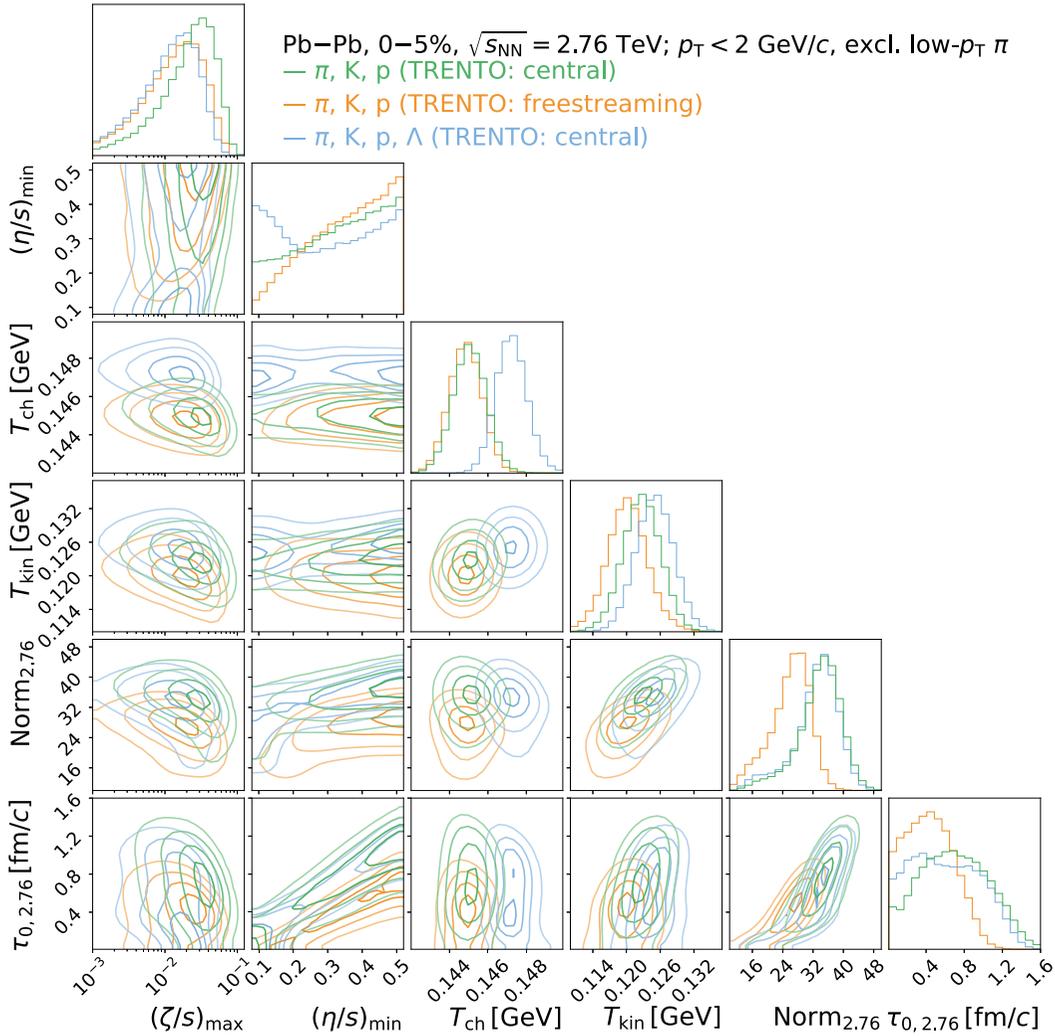

FIG. 5. Bayesian posterior distribution of the model input parameters for three different configurations exploiting the $\sqrt{s_{NN}} = 2.76$ TeV Pb-Pb experimental data. The diagonal panels show the marginalized distributions of individual model parameters, while the off-diagonal panels present the correlations among pairs of model parameters. See text for more details.
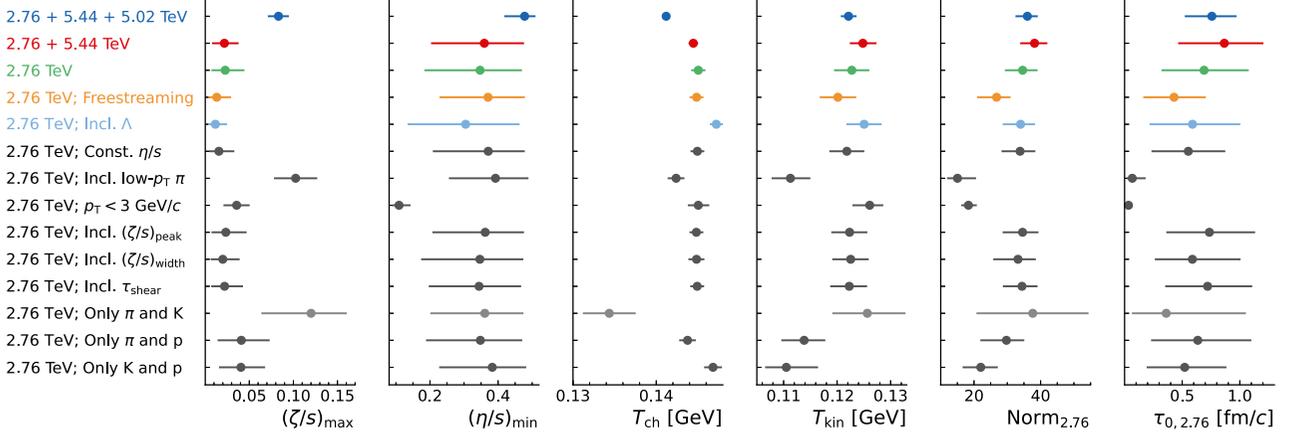
FIG. 6. The median values and 68% confidence intervals of the marginalized Bayesian posterior distributions for each model parameter across all analyzed configurations. Note that, due to the collision system dependency, the Norm and τ_0 parameters for Xe-Xe and Pb-Pb at $\sqrt{s_{NN}} = 5.02$ TeV are reported only in Table II. The configuration involving only pions and kaons is depicted in a lighter color, reflecting the lack of convergence to a single minimum in the MCMC procedure for this specific setup.

multistrange baryons are more sensitive to changes in the transition temperature between the fluid evolution and the hadronic transport phases. This aligns with proposals in the literature indicating that strange hadrons may undergo chemical freeze-out earlier than nonstrange particles [77–80].

C. Exploring variations in the analysis setup

Figure 6 summarizes the previously discussed variations by reporting the median values and 68% confidence intervals of the marginalized Bayesian posterior distributions for each model parameter. Several smaller variations, exclusively applied to the Pb-Pb collision system at $\sqrt{s_{NN}} = 2.76$ TeV, are included as well. We assessed the effect of (i) using a constant shear viscosity to entropy ratio instead of the temperature-dependent version; (ii) including pions with transverse momentum below 500 MeV/c; (iii) increasing the transverse momentum limit to 3 GeV/c for all particle species; (iv)–(vi) including $(\zeta/s)_{\text{peak}}$ (150–200 MeV), $(\zeta/s)_{\text{width}}$ (10–100 MeV), or τ_{shear} ($a_{\text{slope}} = 0$–10 MeV^{-1}) as the seventh model variable in the MCMC procedure; and (vii)–(ix) considering only pions and kaons, pions and protons, and kaons and protons.

Overall, the values of the extracted model parameters demonstrate a reasonable stability across all configurations. When considering a constant shear viscosity to entropy ratio, the results remain consistent with those obtained from temperature-dependent formulation [Eq. (3)]. The observation that the constant value for η/s closely resembles $(\eta/s)_{\text{min}}$ in the temperature-dependent parametrization suggests that, in the latter case, we are primarily sensitive to the region close to the phase transition. The inclusion of low transverse momentum pions impacts several model variables. Specifically, it leads to a substantially smaller τ_0, marginally lower freeze-out temperatures, and a slightly larger $(\zeta/s)_{\text{max}}$. Given the expected nonhydrodynamic origin of the enhancement in the low-p_T pion spectra [72–76], this underscores the importance of restricting the analysis within the range where

hydrodynamics is expected to be applicable. Increasing the upper p_T limit, instead, yields a notable impact on the $(\eta/s)_{\text{min}}$ value. This parameter now tends to favor lower values, nearing the theoretical lower bound of $1/4\pi$ derived from AdS-CFT. This behavior likely arises from the shear correction in the computation of the final particle spectrum that scales with p_T^2 [28,81], i.e., expanding the analyzed p_T interval includes points that are much more sensitive to the shear corrections and we would thus expect a change in the $(\eta/s)_{\text{min}}$ parameter.

The introduction of $(\zeta/s)_{\text{peak}}$, $(\zeta/s)_{\text{width}}$, or a_{slope} as the seventh model variable in the MCMC procedure has a negligible impact on the values of the original six parameters. These parameters remain consistent with those obtained from the central Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV configuration. However, it is important to note that none of the posterior distributions for these seventh model parameters converge; all three distributions touch the upper limit of their respective intervals. This emphasizes the need for first-principle calculations of the bulk viscosity which could replace the currently assumed Lorentzian form. Finally, excluding one of the three particle species (pions, kaons, or protons) results in small changes in the Norm and freeze-out temperatures, but they are still compatible within (1–2)σ. When considering only pions and kaons, the posterior distributions exhibit a dual structure, where one component aligns well with the central Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV configuration, while the other shows significantly different values. Due to this configuration's lack of convergence, we choose to exclude it from further analysis.

D. Discussion and comparison to experimental data

We start the discussion by comparing our extracted model parameters in Fig. 6 with those obtained from analogous Bayesian inference analyses. The values for $(\zeta/s)_{\text{max}}$ are in agreement with analyses employing the same functional form [9,10,15,16,19,20], and are compatible [18] or lower by (2–3)σ [11,12,17] when compared to Bayesian analyses that introduce a slight adjustment to allow for an asymmetry in
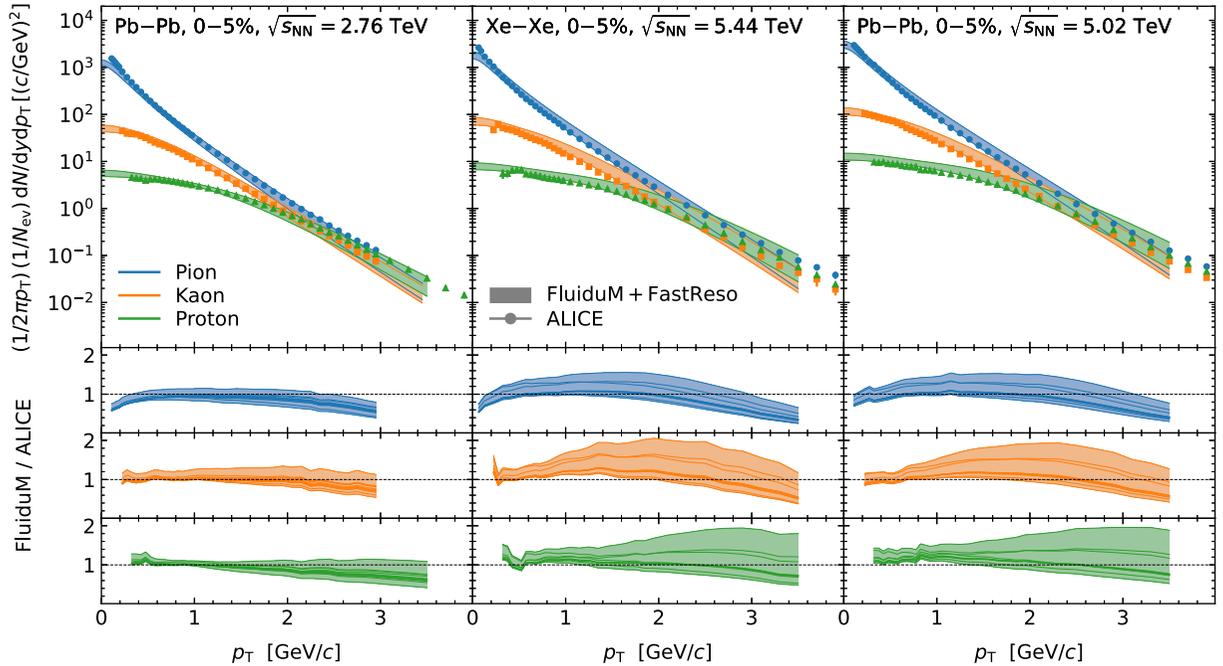
FIG. 7. The top panels display the transverse momentum spectra for pions, kaons, and protons in the 0–5% centrality class for the three collision systems. The spectra as simulated with the FLUID*u*M + FASTRESO framework, using the extracted model parameters from our Bayesian analysis (see Fig. 6), are compared with experimental data from the ALICE Collaboration [32,34,35]. The bottom panels present the respective ratios between theory and experimental data for each hadron species. Note the experimental uncertainties are not taken into account in the ratio.

temperature around the peak. Regarding the $(\eta/s)_{\mathrm{min}}$ parameter, all other Bayesian analyses typically find values around 0.10 [9–20], while our posterior distributions hint at values beyond the upper bound of 0.52. We remind the reader that we attribute this behavior to two factors: (i) a limited sensitivity of the current observables to the shear viscosity of the system and (ii) a different strategy regarding the hadronic phase of the system. On account of reason (ii), a direct comparison of our extracted freeze-out temperatures with those from other Bayesian analyses seems unfeasible. These analyses typically employ a single freeze-out temperature T_{switch} (converging to values between 130 and 160 MeV), while their hadronic afterburner continues to evolve until yields and momentum distributions cease changing. However, since most of the hadronic yields vary by less than 20% as a consequence of inelastic collisions in the afterburner phase, the switching and chemical freeze-out temperature are typically associated with each other [82]. In this context, it is noteworthy that we obtain values that are compatible with the switching temperatures found in Refs. [17,18], 5–15 MeV lower with respect to Refs. [13–16,19,20], and slightly higher than values in Refs. [11,12]. Instead, the statistical hadronization model of Ref. [62] suggests chemical freeze-out temperatures of approximately 10 MeV higher, while blast-wave fits to the p_{T} distributions of identified hadrons in central collisions estimate kinematic freeze-out temperatures about 20 MeV lower than ours [32]. The extracted τ_0 values for the Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV system align roughly with the time of the prehydrodynamic phase in the other Bayesian analyses, typically spanning 0.3–1.0 fm/c [9–17,19,20]. However, due to

the correlation with the Norm parameter [as given in Eq. (1)], our consideration of this parameter as system dependent (resulting in significantly larger values for the Xe-Xe and Pb-Pb $\sqrt{s_{NN}} = 5.02$ TeV collision systems), and the absence of a free-streaming phase in our setup, we would refrain from interpreting them as being the same parameter in a physical sense.

In Fig. 7, we translate the values of our model parameters from Fig. 6 into a final FLUID*u*M + FASTRESO prediction for the pion, kaon, and proton $(1/2\pi p_{\mathrm{T}})(1/N_{\mathrm{ev}}) d^2N/dy\, dp_{\mathrm{T}}$ spectra in the 0–5% centrality class for the three collision systems. For the analysis configurations for which the MCMC procedure was only run for the Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV system, we use the Norm and τ_0 from Table II for the calculations of Pb-Pb at $\sqrt{s_{NN}} = 5.02$ TeV and Xe-Xe at $\sqrt{s_{NN}} = 5.44$ TeV. The theoretical uncertainty is estimated as the full envelope of the various trials.[4] The simulations exhibit good quantitative agreement with experimental measurements. The theoretical uncertainties are approximately 20% for the Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV system and about 40% for the Xe-Xe and $\sqrt{s_{NN}} = 5.02$ TeV Pb-Pb systems. This difference can be attributed to the fact that we primarily explored variations of the analysis configuration using the Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV system. In future work, where we will include anisotropic flow observables and explore additional centrality classes, we intend to conduct systematic variations for all three systems

---

[4]Excluding the "$p_{\mathrm{T}} < 3$ GeV/c" and "only π and K" cases as previously discussed.

simultaneously. We expect that the resulting theoretical uncertainties will be of comparable magnitude for each system.

V. SUMMARY

In summary, we presented a Bayesian analysis employing our TRENTO + FLUIDuM + FASTRESO framework to determine key parameters of the QGP. These parameters included the shear and bulk viscosity to entropy ratios, the initialization time, the initial entropy density, and the freeze-out temperatures.

Modern machine learning tools were utilized to perform a global search in multidimensional space to extract the posterior distributions of the model parameters. In contrast to previous Bayesian analyses employing Gaussian process regression, this study pioneered the use of neural network ensemble emulation. This innovation offered several computational advantages, including significantly reduced training time, lower memory usage, the ability to handle any number of inputs and outputs, and a rigorous uncertainty determination.

Another notable improvement in this work concerns our theoretical framework. Instead of a single freeze-out, we now separate the chemical and kinematic freeze-outs, incorporating the partial chemical equilibrium to describe the later stages of the evolution. In addition, we replaced the use of a constant value for the shear viscosity to entropy ratio with a Yang-Mills theory based parametrization.

Our theoretical model is compared against experimental measurements of transverse momentum spectra for identified charged hadrons (π, K, p) and strange hyperons (Λ) in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV and $\sqrt{s_{NN}} = 5.02$ TeV, as well as Xe-Xe collisions at $\sqrt{s_{NN}} = 5.44$ TeV, from the ALICE Collaboration. Our focus lays on the 0–5% centrality class, where the FLUIDuM background-fluctuation splitting ansatz works best. Furthermore, the restriction to only one centrality class was imposed such that we are able to keep the initial-

state parameters, which play an important role in the centrality dependence of the observables, fixed. In future work, we will include anisotropic flow observables and explore additional centrality classes.

Thanks to the computational efficiency of our framework, we could employ it to assess systematic uncertainties by varying key components of the analysis. In other words, we did not focus on a single "best-fit" scenario but conducted an extensive exploration of the parameter space to better understand the underlying dynamics of the system. Multiple variations—including data from various collision systems, modified initial conditions, inclusion and exclusion of hadron species, variations in the shear and bulk viscosity to entropy ratios, and the use of different $p_{\rm T}$ ranges—were performed. Overall, the extracted model parameters exhibit a reasonable stability across the configurations. Furthermore, our results are in agreement with previous Bayesian analyses, except for the $(\eta/s)_{\rm min}$, which we attributed to limited sensitivity of the observables and a different theoretical strategy concerning the hadronic phase.

Finally, we translated our extracted parameters into FLUIDuM +FASTRESO predictions for the pion, kaon, and proton $(1/2\pi p_{\rm T})(1/N_{\rm ev})\, {\rm d}^2N/{\rm d}y\, {\rm d}p_{\rm T}$ spectra in the 0–5% centrality class for the three collision systems. They demonstrate strong quantitative agreement with experimental measurements from the ALICE Collaboration.

[1] W. Busza, K. Rajagopal, and W. van der Schee, Annu. Rev. Nucl. Part. Sci. **68**, 339 (2018).

[2] K. Aamodt *et al.* (ALICE Collaboration), Phys. Rev. Lett. **105**, 252302 (2010).

[3] J. Adams Jr. *et al.* (STAR Collaboration), Nucl. Phys. A **757**, 102 (2005).

[4] K. Adcox *et al.* (PHENIX Collaboration), Nucl. Phys. A **757**, 184 (2005).

[5] S. Acharya *et al.* (ALICE Collaboration), Phys. Rev. Lett. **123**, 142301 (2019).

[6] C. Aidala *et al.* (PHENIX Collaboration), Nat. Phys. **15**, 214 (2019).

[7] V. Khachatryan *et al.* (CMS Collaboration), Phys. Rev. Lett. **116**, 172302 (2016).

[8] G. Aad *et al.* (ATLAS Collaboration), Phys. Rev. Lett. **116**, 172301 (2016).

[9] J. S. Moreland, J. E. Bernhard, and S. A. Bass, Phys. Rev. C **101**, 024911 (2020).

[10] J. E. Bernhard, J. S. Moreland, and S. A. Bass, Nat. Phys. **15**, 1113 (2019).

[11] D. Everett *et al.* (JETSCAPE Collaboration), Phys. Rev. Lett. **126**, 242301 (2021).

[12] D. Everett *et al.* (JETSCAPE Collaboration), Phys. Rev. C **103**, 054904 (2021).

[13] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, Phys. Rev. Lett. **126**, 202301 (2021).

[14] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, Phys. Rev. C **103**, 054909 (2021).

[15] G. Nijs and W. van der Schee, Phys. Rev. C **106**, 044903 (2022).

[16] G. Nijs and W. van der Schee, arXiv:2304.06191.

[17] M. R. Heffernan, C. Gale, S. Jeon, and J.-F. Paquet, arXiv:2302.09478.

[18] D. Liyanage, O. Sürer, M. Plumlee, S. M. Wild, and U. Heinz, arXiv:2302.14184.

[19] J. E. Parkkila, A. Onnerstad, and D. J. Kim, Phys. Rev. C **104**, 054904 (2021).

[20] J. E. Parkkila, A. Onnerstad, S. F. Taghavi, C. Mordasini, A. Bilandzic, M. Virta, and D. J. Kim, Phys. Lett. B **835**, 137485 (2022).

[21] L.-G. Pang, K. Zhou, N. Su, H. Petersen, H. Stöcker, and X.-N. Wang, Nat. Commun. **9**, 210 (2018).

[22] J. Steinheimer, L. Pang, K. Zhou, V. Koch, J. Randrup, and H. Stoecker, J. High Energy Phys. 12 (2019) 122.

[23] D. Devetak, A. Dubla, S. Floerchinger, E. Grossi, S. Masciocchi, A. Mazeliauskas, and I. Selyuzhenkov, J. High Energy Phys. 06 (2020) 044.

[24] N. Christiansen, M. Haas, J. M. Pawlowski, and N. Strodthoff, Phys. Rev. Lett. **115**, 112002 (2015).

[25] J. M. Pawlowski (private communication).

[26] J. S. Moreland, J. E. Bernhard, and S. A. Bass, Phys. Rev. C **92**, 011901(R) (2015).

[27] S. Floerchinger, E. Grossi, and J. Lion, Phys. Rev. C **100**, 014905 (2019).

[28] A. Mazeliauskas, S. Floerchinger, E. Grossi, and D. Teaney, Eur. Phys. J. C **79**, 284 (2019).

[29] R. M. Neal, Priors for infinite networks, in *Bayesian Learning for Neural Networks* (Springer, New York, 1996), pp. 29–53.

[30] F. Capellino, A. Dubla, S. Floerchinger, E. Grossi, A. Kirchner, and S. Masciocchi, Phys. Rev. D **108**, 116011 (2023).

[31] F. Capellino, A. Beraudo, A. Dubla, S. Floerchinger, S. Masciocchi, J. Pawlowski, and I. Selyuzhenkov, Phys. Rev. D **106**, 034021 (2022).

[32] B. Abelev *et al.* (ALICE Collaboration), Phys. Rev. C **88**, 044910 (2013).

[33] B. B. Abelev *et al.*, Phys. Rev. Lett. **111**, 222301 (2013).

[34] S. Acharya *et al.* (ALICE Collaboration), Phys. Rev. C **101**, 044907 (2020).

[35] S. Acharya *et al.* (ALICE Collaboration), Eur. Phys. J. C **81**, 584 (2021).

[36] G. Giacalone, arXiv:2208.06839.

[37] S. Acharya *et al.* (ALICE Collaboration), CERN Report No. ALICE-PUBLIC-2018-011, 2018 (unpublished), https://cds.cern.ch/record/2636623.

[38] B. Bally, M. Bender, G. Giacalone, and V. Somà, Phys. Rev. Lett. **128**, 082301 (2022).

[39] G. Giacalone, A matter of shape: seeing the deformation of atomic nuclei at high-energy colliders, Ph.D. thesis, Université Paris-Saclay, 2020 (unpublished), arXiv:2101.00168.

[40] S. Floerchinger and U. A. Wiedemann, Phys. Lett. B **728**, 407 (2014).

[41] S. Floerchinger and U. A. Wiedemann, Phys. Rev. C **89**, 034914 (2014).

[42] S. Floerchinger and U. A. Wiedemann, J. High Energy Phys. 08 (2014) 005.

[43] S. Floerchinger and E. Grossi, J. High Energy Phys. 08 (2018) 186.

[44] N. Demir and S. A. Bass, Phys. Rev. Lett. **102**, 172302 (2009).

[45] J. B. Rose, J. M. Torres-Rincon, A. Schäfer, D. R. Oliinychenko, and H. Petersen, Phys. Rev. C **97**, 055204 (2018).

[46] F. Karsch, D. Kharzeev, and K. Tuchin, Phys. Lett. B **663**, 217 (2008).

[47] J. Noronha-Hostler, J. Noronha, and C. Greiner, Phys. Rev. Lett. **103**, 172302 (2009).

[48] G. S. Denicol, T. Kodama, T. Koide, and P. Mota, Phys. Rev. C **80**, 064901 (2009).

[49] N. Y. Astrakhantsev, V. V. Braguta, and A. Y. Kotov, Phys. Rev. D **98**, 054515 (2018).

[50] E. Lu and G. D. Moore, Phys. Rev. C **83**, 044901 (2011).

[51] D. Teaney, Phys. Rev. C **68**, 034913 (2003).

[52] J.-F. Paquet, C. Shen, G. S. Denicol, M. Luzum, B. Schenke, S. Jeon, and C. Gale, Phys. Rev. C **93**, 044906 (2016).

[53] G. S. Denicol, S. Jeon, and C. Gale, Phys. Rev. C **90**, 024912 (2014).

[54] F. Cooper and G. Frye, Phys. Rev. D **10**, 186 (1974).

[55] H. Bebie, P. Gerber, J. L. Goity, and H. Leutwyler, Nucl. Phys. B **378**, 95 (1992).

[56] P. Huovinen, Eur. Phys. J. A **37**, 121 (2008).

[57] K. Aamodt *et al.* (ALICE Collaboration), Phys. Lett. B **696**, 328 (2011).

[58] P. Alba *et al.*, Phys. Rev. D **96**, 034517 (2017).

[59] P. Alba, V. Mantovani Sarti, J. Noronha, J. Noronha-Hostler, P. Parotto, I. Portillo Vazquez, and C. Ratti, Phys. Rev. C **98**, 034909 (2018).

[60] P. Parotto (private communication).

[61] S. Ryu, J. F. Paquet, C. Shen, G. Denicol, B. Schenke, S. Jeon, and C. Gale, Phys. Rev. C **97**, 034910 (2018).

[62] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel, Nature (London) **561**, 321 (2018).

[63] Y. Seemann, Multidimensional parameter optimization in fluid dynamic simulations of heavy-ion collisions, Master's thesis, Universität Heidelberg, 2022 (unpublished).

[64] A. Paszke *et al.*, in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, edited by H. Wallach *et al.* (IEEE, Piscataway, NJ, 2019).

[65] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, arXiv:1807.05118.

[66] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, Publ. Astron. Soc. Pac. **125**, 306 (2013).

[67] C. Gale, S. Jeon, and B. Schenke, Int. J. Mod. Phys. A **28**, 1340011 (2013).

[68] J. E. Bernhard, J. S. Moreland, S. A. Bass, J. Liu, and U. Heinz, Phys. Rev. C **94**, 024907 (2016).

[69] D. A. Teaney, in *Quark-Gluon Plasma 4* (World Scientific, Singapore, 2010), pp. 207–266.

[70] A. Dubla, S. Masciocchi, J. M. Pawlowski, B. Schenke, C. Shen, and J. Stachel, Nucl. Phys. A **979**, 251 (2018).

[71] C. Gale, S. Jeon, B. Schenke, P. Tribedy, and R. Venugopalan, Phys. Rev. Lett. **110**, 012302 (2013).

[72] V. Begun and W. Florkowski, Phys. Rev. C **91**, 054909 (2015).

[73] V. Begun, Phys. Rev. C **94**, 054904 (2016).

[74] E. Schnedermann, J. Sollfrank, and U. W. Heinz, Phys. Rev. C **48**, 2462 (1993).

[75] P. Huovinen, P. M. Lo, M. Marczenko, K. Morita, K. Redlich, and C. Sasaki, Phys. Lett. B **769**, 509 (2017).

[76] E. Grossi, A. Soloviev, D. Teaney, and F. Yan, Phys. Rev. D **104**, 034025 (2021).

[77] H. van Hecke, H. Sorge, and N. Xu, Nucl. Phys. A **661**, 493 (1999).

[78] N. Arbex, F. Grassi, Y. Hama, and O. Socolowski, Phys. Rev. C **64**, 064906 (2001).

[79] S. Chatterjee, R. M. Godbole, and S. Gupta, Phys. Lett. B **727**, 554 (2013).

[80] R. Bellwied, S. Borsanyi, Z. Fodor, S. D. Katz, and C. Ratti, Phys. Rev. Lett. **111**, 202302 (2013).

[81] A. Kirchner, E. Grossi, and S. Floerchinger, arXiv:2308.10616.

[82] H. Song, S. A. Bass, and U. Heinz, Phys. Rev. C **83**, 024912 (2011).