

PAPER • OPEN ACCESS

## Computational budget optimization for Bayesian parameter estimation in heavy-ion collisions

To cite this article: Brandon Weiss *et al* 2023 *J. Phys. G: Nucl. Part. Phys.* **50** 065104

View the [article online](#) for updates and enhancements.

You may also like

- [Dark Quest. I. Fast and Accurate Emulation of Halo Clustering Statistics and Its Application to Galaxy Clustering](#)  
Takahiro Nishimichi, Masahiro Takada, Ryuichi Takahashi *et al.*
- [The Aemulus Project. III. Emulation of the Galaxy Correlation Function](#)  
Zhongxu Zhai, Jeremy L. Tinker, Matthew R. Becker *et al.*
- [The Circumgalactic Medium from the CAMELS Simulations: Forecasting Constraints on Feedback Processes from Future Sunyaev–Zeldovich Observations](#)  
Emily Moser, Nicholas Battaglia, Daisuke Nagai *et al.*

# Computational budget optimization for Bayesian parameter estimation in heavy-ion collisions

Brandon Weiss<sup>1</sup> , Jean-François Paquet<sup>1,2</sup>  and Steffen A Bass<sup>1</sup> 

<sup>1</sup> Department of Physics, Duke University, Durham NC 27708, United States of America

<sup>2</sup> Department of Physics and Astronomy and Department of Mathematics, Vanderbilt University, Nashville TN 37235, United States of America

E-mail: [jean-francois.paquet@vanderbilt.edu](mailto:jean-francois.paquet@vanderbilt.edu)

Received 7 February 2023, revised 12 April 2023

Accepted for publication 27 April 2023

Published 16 May 2023



CrossMark

## Abstract

Bayesian parameter estimation provides a systematic approach to compare heavy-ion collision models with measurements, leading to constraints on the properties of nuclear matter with proper accounting of experimental and theoretical uncertainties. Aside from statistical and systematic model uncertainties, interpolation uncertainties can also play a role in Bayesian inference, if the model's predictions can only be calculated at a limited set of model parameters. This uncertainty originates from using an emulator to interpolate the model's prediction across a continuous space of parameters. In this work, we study the trade-offs between the emulator (interpolation) and statistical uncertainties. We perform the analysis using spatial eccentricities from the T<sub>R</sub>ENTo model of initial conditions for nuclear collisions. Given a fixed computational budget, we study the optimal compromise between the number of parameter samples and the number of collisions simulated per parameter sample. For the observables and parameters used in the present study, we find that the best constraints are achieved when the number of parameter samples is slightly smaller than the number of collisions simulated per parameter sample.

Keywords: Bayesian inference, Gaussian process emulator, heavy-ion collisions

(Some figures may appear in colour only in the online journal)



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

## 1. Introduction

One of the main goals of the heavy-ion program pursued at the relativistic heavy-ion collider (RHIC) and the large Hadron collider (LHC) is a quantitative understanding of the properties of nuclear matter under extreme conditions as described by Quantum Chromodynamics (QCD). At sufficiently high temperatures and densities accessible during these nuclear collisions, a transient state of matter, the quark-gluon plasma (QGP), is formed but decays again as the collision system cools down and disintegrates. Collisions of heavy ions leading to QGP formation are multistage processes: the initial impact of the nuclei, the formation and expansion of the QGP, reconfinement into hadrons, and subsequent hadronic interactions [1–4]. A wide range of properties of nuclear matter enters numerical simulations of heavy-ion collisions. Some of these properties are constrained by external means, for example, lattice QCD calculations of the equation of state of nuclear matter [5]. Other properties, such as transport coefficients or parameters entering the description of the early stage of the collision, are often parametrized and constrained by comparison with data.

Measurements from the RHIC and the LHC can constrain these physical parameters. Because of the heterogeneity and a large number of measurements and model parameters, it is beneficial to perform model-to-data comparisons using statistical techniques such as Bayesian parameter estimation: this allows for a systematic propagation of uncertainties from experimental measurements to physical parameters. A number of such studies have been performed over the past decade [6–20].

A key ingredient of Bayesian parameter estimation as used in heavy-ion collisions is emulation: given the model’s prediction at a *discrete* sample of parameters, an emulator will interpolate the model’s prediction over a *continuous* range of parameters. Furthermore, emulators such as Gaussian processes provide an estimate of their own interpolation uncertainty; in effect, this allows emulation to be used in model-to-data comparisons even if the emulator’s interpolation uncertainty is not negligible compared to the other uncertainties in the problem—a situation that is often unavoidable in simulations with a large number of parameters. In Bayesian parameter estimation, this emulator uncertainty is included into the physical parameter’s final uncertainty determination, alongside experimental, statistical and other theoretical uncertainties.

The most straightforward approach to reducing emulator uncertainty is to increase the number of parameter samples at which the model observables are evaluated; these parameter samples are generally referred to as the emulator’s ‘design points’. Evidently, this increase in the number of design points must be balanced with other demands on the available computational budget. In particular, simulations of heavy-ion collisions are stochastic: there are event-by-event fluctuations in the outcome of the collisions. Comparisons with measurements require simulating and averaging over a large number of collisions. At a fixed computational budget, one must determine the proper trade-off between reducing emulator interpolation uncertainties or statistical ones. In this work, we use the T<sub>R</sub>ENTo model of initial conditions for heavy-ion collisions [21] to study this trade-off between statistical and emulator uncertainties.

## 2. Methodology

We use the model T<sub>R</sub>ENTo as a proxy for observables in heavy-ion collisions; this is based on the known correlation [22] between (i) initial spatial anisotropies of models like T<sub>R</sub>ENTo, and (ii) measurable momentum anisotropies of final state hadrons. To interpolate the output of

the T<sub>R</sub>ENTo model, we use Gaussian process emulators, which have been the standard emulation technique used for model-to-data comparisons in heavy-ion physics [6–8, 10]. To quantify the optimal trade-off between statistical and emulator uncertainty, we use ‘closure tests’, which apply Bayesian parameter inference to model calculations. We summarize the overall approach and methods below.

### 2.1. T<sub>R</sub>ENTo

T<sub>R</sub>ENTo (‘Reduced Thickness Event-by-event Nuclear Topology’) is a model that parametrizes the energy or entropy deposition resulting from the impact of nuclei at high energy [21]. Its primary input is the distribution of nucleons inside the different species of nuclei. Given this distribution of nucleons, it uses the inelastic nucleon–nucleon cross-section to determine the nucleons that participate in the collision, and infers the energy deposition from it. Three parameters of the T<sub>R</sub>ENTo model, which we study in this work, are the effective nucleon size  $w$ , the fluctuation parameter  $k$  and the reduced thickness  $p$ . The effective nucleon size  $w$  is related to the area over which the energy is distributed for a given participant nucleon. The fluctuation parameter  $k$  allows for variation in the amount of energy carried by each nucleon, while the reduced thickness  $p$  parametrizes how the energy density is constructed from the profile of nucleons sampled from each nucleus. The output of T<sub>R</sub>ENTo is a density profile in the plane transverse to the collision axis. We interpret T<sub>R</sub>ENTo’s output as an energy density profile  $\epsilon(r, \phi)$  with arbitrary normalization. From this density profile  $\epsilon(r, \phi)$ , we compute single-event spatial anisotropies  $\varepsilon_n$  [22]:

$$\varepsilon_n e^{in\Phi_n} = \frac{\int_0^\infty dr r \int_0^{2\pi} d\phi r^n \epsilon(r, \phi) e^{in\phi}}{\int_0^\infty dr r \int_0^{2\pi} d\phi r^n \epsilon(r, \phi)}. \quad (1)$$

Although these anisotropies can be related to the momentum anisotropy of hadrons measured at the end of the collision [23–27], in this work, we focus on the initial spatial eccentricities themselves.

We compute the average of an ensemble of T<sub>R</sub>ENTo events with an arithmetic average:

$$\langle \varepsilon_n \rangle = \frac{1}{M_{\text{ev}}} \sum_{j=1}^{M_{\text{ev}}} \varepsilon_n \{\text{event } j\}, \quad (2)$$

where  $M_{\text{ev}}$  is the number of T<sub>R</sub>ENTo events used to evaluate the initial spatial eccentricity. We used up to four harmonics in this work,  $n = 2$  to 5. The variable sensitivity of the spatial eccentricities to the model parameters as well as the dependence on  $n$  of the statistical uncertainty of  $\langle \varepsilon_n \rangle$  makes the spatial eccentricities a good choice of observable to study the interplay of interpolation and statistical uncertainties.

We studied minimum bias Pb–Pb collisions with Woods–Saxon nucleon distributions. We used an inelastic nucleon–nucleon cross-section of 64 mb, corresponding approximately to  $\sqrt{s_{NN}} = 2.76$  TeV collisions.

### 2.2. Bayesian parameter estimation

We note  $y_{\text{th},j}(\mathbf{p})$  the model’s prediction for the ‘ $j$ ’th observables of interest, evaluated when the model’s parameters are set to  $\mathbf{p}$ . At  $\mathbf{p}$ , we compute  $y_{\text{th},j}(\mathbf{p}) = \langle \varepsilon_{n_j} \rangle(\mathbf{p})$ , where the index  $j$  runs over all the observables of interest, which in our case are the different ‘ $n$ ’ in equation (2).

In general, given this information, we want to find the probability that parameters  $\mathbf{p}$  are consistent with the data  $\{y_{\text{exp},j}\}$  and with the experimental and theoretical uncertainties  $\{\sigma_{\text{exp},j}\}$

and  $\{\sigma_{\text{th},j}(\mathbf{p})\}$ . The first step is to define a metric to quantify the level of model-to-data agreement: the likelihood function. We make the typical choice of a Gaussian likelihood function:

$$\mathcal{L}(\{y_{\text{exp},j}\}|\mathbf{p}) = \frac{\exp\left(-\frac{1}{2}\sum_j \frac{[y_{\text{th},j}(\mathbf{p}) - y_{\text{exp},j}]^2}{\sigma_{\text{th},j}(\mathbf{p})^2 + \sigma_{\text{exp},j}^2}\right)}{\sqrt{(2\pi)^n \prod_j (\sigma_{\text{th},j}(\mathbf{p})^2 + \sigma_{\text{exp},j}^2)}}. \quad (3)$$

We assumed a diagonal covariance matrix. In this work, instead of comparing the model with data, we will use closure tests: we will replace  $y_{\text{exp},j}$  by model calculations, as discussed later. These model calculations do stand for experimental data, and consequently we keep the label ‘exp’ to denote these quantities.

Given our model of the collision, the probability that the parameter set  $\mathbf{p}$  is consistent with the data and the uncertainties is given by the posterior distribution of model parameters,  $\mathcal{P}(\mathbf{p}|\{y_{\text{exp}}\})$ :

$$\mathcal{P}(\mathbf{p}|\{y_{\text{exp}}\}) = \frac{\mathcal{L}(\{y_{\text{exp}}\}|\mathbf{p})\text{Prior}(\mathbf{p})}{\int d\mathbf{p}\mathcal{L}(\{y_{\text{exp}}\}|\mathbf{p})\text{Prior}(\mathbf{p})}, \quad (4)$$

where  $\text{Prior}(\mathbf{p})$  is the prior distribution. In this work, we take the prior to be constant over a finite parameter range, for all the parameters.

The posterior distribution  $\mathcal{P}(\mathbf{p}|\{y_{\text{exp}}\})$  is a probability distribution whose dimensionality is equal to the number of model parameters. Different projections of the posterior distribution summarize the constraints on the model parameters. As long as the number of parameters is small and the model is relatively fast computationally, it is straightforward to sample and marginalize the posterior  $\mathcal{P}(\mathbf{p}|\{y_{\text{exp}}\})$ . When the number of parameters increases or if the model is expensive, it can become overly burdensome to evaluate a model’s prediction at a large number of values of the parameter  $\mathbf{p}$  to compute  $\mathcal{P}(\mathbf{p}|\{y_{\text{exp}}\})$ . A solution is to prepare a fast proxy, such as a Gaussian process [28], to emulate the model’s prediction. We review Gaussian process emulators briefly in the next subsection. As far as the Bayesian inference is concerned, the consequence of using Gaussian process emulators is substituting the model’s prediction by the emulator’s prediction

$$y_{\text{th},j}(\mathbf{p}) \rightarrow y_{\text{emul},j}(\mathbf{p}), \quad (5)$$

as well as adding the emulator interpolation uncertainty into the sum of uncertainties


$$\sigma_{\text{th},j}(\mathbf{p})^2 \rightarrow \sigma_{\text{th},j}(\mathbf{p})^2 + \sigma_{\text{emul},j}(\mathbf{p})^2. \quad (6)$$

### 2.3. Emulation with Gaussian processes

The idea behind emulation is to first sample the parameter space of the model,  $\{\mathbf{p}_k\}$ ; the parameter samples are the design points. The model’s predictions are then calculated for all design points:  $\{y_{\text{th},j}(\mathbf{p}_k)\}$ . The emulator then interpolates the model’s predictions over a continuous range of model parameters  $\mathbf{p}$ .

To sample the parameter space, we use a low discrepancy sequence generator [29, 30]. Figure 1 shows an example for 16 design points within a two-dimensional parameter space.


For each observable ( $\langle\epsilon_n\rangle$ ), a Gaussian process is trained on the set of values of  $\langle\epsilon_n\rangle(\mathbf{p})$  calculated at the design points  $\{\mathbf{p}_k\}$ . Gaussian process emulators are probabilistic



**Figure 1.** 16 design points in a two-dimensional parameter space (reduced thickness on the  $x$ -axis and nucleon width on the  $y$ -axis), as sampled with a low discrepancy sequence generator [29, 30].

interpolators, which assume that each point in parameter space is a probability distribution. At the design points, the width of the probability distribution should include the statistical uncertainty of the observables. Between the design points, the width of the probability distribution should increase to account for the interpolation uncertainty. To handle stochastic simulations, Gaussian process emulators must be set up such that variations between parameter samples are divided into (i) statistical uncertainty, and (ii) the actual variation from the parameter dependence of the model. A mathematical description of the approach can be found in [15], section V-B-2. In short, a white noise kernel accounts for the uncorrelated (‘short range’) fluctuations originating from the statistical uncertainty, and a squared-exponential kernel accounts for the longer-range parameter dependence. The parameters of these kernels, for example, the relative size of the kernels, are determined by numerical optimization. This approach is used in most applications of Bayesian parameter inference in heavy-ion physics [6–8, 10–12, 14–20, 31–37]. The result of this implementation is that the emulator uncertainty accounts for both the statistical and interpolation uncertainties.

Figure 2 shows the mean-value predictions of the emulator for  $\langle \epsilon_2 \rangle$  and  $\langle \epsilon_3 \rangle$  compared to the actual value of these observables. The points all lie close to the line  $y_{\text{emul},j}(\mathbf{p}) = y_{\text{th},j}(\mathbf{p})$ , which suggests that the emulator does indeed interpolate well, although this should not be seen as strong validation: one would need to compare the emulator with calculations not used to train the emulator to perform proper validation. Nevertheless, the use of design points rather than validation points in figure 2 allows us to illustrate clearly that emulator uncertainty originates both from the statistical uncertainty of each design point and the interpolation uncertainty across the parameter space. The observable shown on the right,  $\langle \epsilon_3 \rangle$ , intrinsically has a larger statistical uncertainty than  $\langle \epsilon_2 \rangle$  at an equal number of averaged events  $M_{\text{ev}}$ . The left panel of figure 2 shows a scenario where the spread around the  $y_{\text{emul},j}(\mathbf{p}) = y_{\text{th},j}(\mathbf{p})$  line is small, but with a small number of samples in the parameter space: this illustrates a scenario where the total emulator uncertainty is likely dominated by the interpolation uncertainty, rather than the statistical uncertainty at each design point. On the other hand, the right panel of figure 2 illustrates the case where statistical uncertainties would dominate over the



**Figure 2.** (Left) Value for  $\langle \varepsilon_2 \rangle$  as predicted by a Gaussian process emulator trained with 16 design points (y-axis), compared to the actual value of the model at the design points (x-axis); (right) same as left, for  $\langle \varepsilon_3 \rangle$  with 256 design points. The red line corresponding to perfect emulation,  $y_{\text{emul},j}(\mathbf{p}) = y_{\text{th},j}(\mathbf{p})$ , is shown for reference. Statistical uncertainties on each sample are not shown for clarity.

interpolation uncertainties: the spread around  $y_{\text{emul},j}(\mathbf{p}) = y_{\text{th},j}(\mathbf{p})$  is large, and cannot be reduced by increasing the number of design points. The focus of the present work is to study this trade-off and find how to optimize interpolation and statistical uncertainties so as to avoid these two substandard scenarios.


#### 2.4. Closure tests

Closure tests are a straightforward application of Bayesian parameter inference, with experimental data replaced by calculations from the model itself. It is a self-consistency check made non-trivial by (i) the presence of uncertainties, as well as (ii) loss of information from the model to the observables. A longer discussion of closure tests, in the context of heavy-ion collisions, can be found in section VI of [15]. We describe it briefly here, with a focus on metrics to quantify the success of closure tests.

In closure tests, we first select a set of parameters that are considered the reference parameter for the closure test; we refer to these parameters as the ‘reference’ or ‘truth’ parameters,  $\mathbf{p}^{\text{truth}}$ . We then compute event-averaged eccentricities  $\langle \varepsilon_n^{\text{truth}} \rangle = \langle \varepsilon_n \rangle(\mathbf{p}^{\text{truth}})$ , which have a statistical uncertainty  $\Delta \langle \varepsilon_n^{\text{truth}} \rangle$ . While these uncertainties are purely statistical, they are meant to mimic statistical and systematic uncertainties found in measurements. This uncertainty can be dialed by changing the number of T<sub>R</sub>ENTo events  $M_{\text{ev}}^{\text{truth}}$  used to compute  $\langle \varepsilon_n^{\text{truth}} \rangle$ .

Separately, we train Gaussian process emulators for  $\langle \varepsilon_n \rangle(\mathbf{p})$  over a chosen range of model parameters, as described in the previous section. The two sources of uncertainties in the emulator are the number of design points (parameter samples)  $N_d$ , which controls the interpolation uncertainty, and the number of T<sub>R</sub>ENTo events per design point,  $M_{\text{ev}}$ , which controls the statistical uncertainty of  $\langle \varepsilon_n \rangle(\mathbf{p})$  at each design point.

Using the Gaussian process emulators as a proxy for T<sub>R</sub>ENTo, we then perform Bayesian parameter estimation to attempt to recover the parameters  $\mathbf{p}^{\text{truth}}$  from the values of the observables  $\langle \varepsilon_n^{\text{truth}} \rangle \pm \Delta \langle \varepsilon_n^{\text{truth}} \rangle$ . The result of the Bayesian inference is the posterior  $\mathcal{P}(\mathbf{p} | \{ \langle \varepsilon_n^{\text{truth}} \rangle \})$  (equation (4)). This posterior distribution depends on three sources of uncertainties: the uncertainty on the reference observables ( $\Delta \langle \varepsilon_n^{\text{truth}} \rangle$ ), which mimics experimental uncertainty, and the statistical and interpolation uncertainty in the emulator. If all uncertainties were small, one would expect the posterior to be a narrow peak at the parameter truth  $\mathbf{p}^{\text{truth}}$ . In practice, some of these uncertainties are significant, and the posterior



**Figure 3.** Schematic comparison of 1-parameter posterior distributions (red solid, green dashed and blue dotted Gaussians) with a single fixed ‘truth’ value of the parameter (vertical teal line). The green dashed line shows how the posterior increases as *accuracy* increases when compared to the blue dotted line, and the red solid line shows how the posterior increases as *precision* increases when compared to the blue dotted line.

distribution thus has some finite width in parameter space. If closure is successful, the posterior should enclose the parameter set  $\mathbf{p}^{\text{truth}}$ .

To quantify the degree of agreement of the posterior— $\mathcal{P}(\mathbf{p}|\{\{\epsilon_n^{\text{truth}}\}\})$ —with the known true value  $\mathbf{p}^{\text{truth}}$  of the parameters used in the closure test, we use two different metrics. The first one is the value of the posterior at the parameter truth,  $\mathcal{P}(\mathbf{p}^{\text{truth}}|\{\{\epsilon_n^{\text{truth}}\}\})$ . The posterior at the parameter truth balances accuracy and precision. As illustrated in figure 3, an increase in accuracy (shift towards the true value of the parameter) will increase the posterior value. An increase in precision (narrower posterior) will also increase the posterior at the parameter truth, but only to a certain extent: the posterior at the truth will actually start to decrease if the results get too confident about the incorrect value. The posterior at the truth has the benefit of being simple to interpret and inexpensive to compute.

For completeness, we studied a second metric, the Akaike information criterion (AIC) [38], given by:

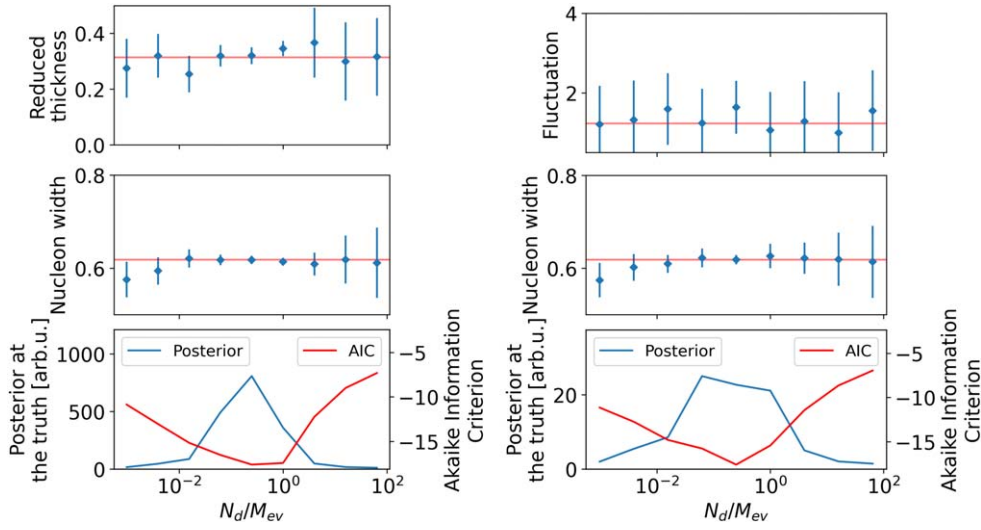
$$\text{AIC} \equiv -2 \ln(\mathcal{L}_{\text{max}}) + 2k, \quad (7)$$

where  $\mathcal{L}_{\text{max}}$  is the maximum likelihood in the parameter space and  $k$  is the number of parameters. The AIC is a metric used for model selection, to help determine how successful a model is at describing measurements given its number of parameters [39, 40]. In the case of closure tests, while we know that the emulator and the reference (truth) calculations originate from the same model, the uncertainties in the problem can evidently obscure this fact. We will see in this work the AIC’s success at quantifying closure. Note that the AIC is more expensive to compute, as it requires identifying the maximum of likelihood.

### 3. Results

As discussed in section 2, the result of Bayesian parameter inference generally depends on the experimental uncertainty as well as on the statistical and interpolation uncertainty of the emulator. Additional theoretical uncertainties can also play a role [15].





**Figure 4.** Closure tests with  $N_d M_{ev} = 2^{16}$  total events,  $M_{ev}^{\text{truth}} = 2^{16}$  events to compute the reference observables, two parameters (nucleon width and reduced thickness on the left, and nucleon width and reduced fluctuation parameter on the right), and two observables ( $\langle \varepsilon_2 \rangle$  and  $\langle \varepsilon_3 \rangle$ ) in Pb–Pb collisions at  $\sqrt{s_{NN}} = 2.76$  TeV. Top two panels show the mode and width of the one-parameter marginalized likelihood, while the bottom panel shows the value of the posterior at the true value of the parameter as well as the Akaike information criterion (AIC), all as a function of the ratio of the number of design points  $N_d$  to the number of T<sub>RENT</sub>o simulations per design points  $M_{ev}$ .

Because we use closure tests in this manuscript, the experimental uncertainty is replaced by the statistical uncertainty of the reference (‘truth’) calculations, which is controlled by changing the number  $M_{ev}^{\text{truth}}$  of T<sub>RENT</sub>o events that are averaged over to compute the reference observables. As for the emulator, its uncertainty is determined by the number of samples of the model parameters  $N_d$  and the number of collisions simulated per parameter sample  $M_{ev}$ .

Tests are performed at a fixed computation budget: the product  $N_{\text{tot}} = N_d M_{ev}$  is kept fixed. The set of  $(N_d, M_{ev})$  pairs is chosen to vary by a factor of 2. We use nine different sets of  $(N_d, M_{ev})$ , resulting in nine different groups of emulators, each group having one Gaussian process emulator per observable.

As a base case, we first used  $N_{\text{tot}} = N_d M_{ev} = 2^{16}$  total events for a Pb–Pb collision. We used  $M_{ev}^{\text{truth}} = 2^{16}$  collisions to compute the reference ‘truth’ observables. We begin with only two T<sub>RENT</sub>o parameters at the time—nucleon width and reduced thickness in one case, and nucleon width and fluctuation in the other. To test the robustness of the methodology, we later vary the data uncertainty, and the number of parameters and of observables.

After performing Bayesian parameter inference in the framework described above, we obtain a 2-parameter posterior distribution. By marginalizing over each parameter, we obtain 1-parameter marginalized posterior distributions, whose mode and width (variance) we calculate and plot as a function of the ratio of the number of design points  $N_d$  to the number of T<sub>RENT</sub>o simulations per design points  $M_{ev}$ ; this is shown in the left top two panels of figure 4. The true (reference) value of the parameters used in the tests is shown as a red horizontal line. We first see that, for any number of design points, the constraints from the

posterior are consistent with the true value of the parameters. This is not a trivial observation: it indicates that the emulator is quantifying properly the interpolation and statistical uncertainties, however large or small these uncertainties are. In this sense, we see that Bayesian parameter inference provides valid constraints for any set of design points  $N_d$  and number of TRENTo events  $M_{ev}$ . On the other hand, one can optimize the results of the Bayesian inference by a careful compromise between emulator and statistical uncertainty. With large numbers of design points (large  $N_d/M_{ev}$ ), the emulator uncertainty is expected to be small, but large statistical uncertainties on the  $\langle \epsilon_n \rangle$  of each design point prevent accurate constraints on the parameter. Using a large number of TRENTo events with a small number of design points (small  $N_d/M_{ev}$ ) leads to similar sub-optimal results. To quantify this tension, we use the two metrics discussed in section 2.4: the value of the posterior at the true value of the closure test parameters, and the AIC. Both are plotted on the bottom left panel of figure 4, still as a function of the ratio  $N_d/M_{ev}$ . The maximum value of the posterior at the truth, and the minimum value of the AIC, are observed around  $N_d/M_{ev} = 0.25$ , which corresponds to  $N_d = 2^7$  design points. Overall, the different panels all indicate that certain apportionment of the number of design points does result in significantly better constraints on the model parameter *for the same computational expense*. The optimal number is found to be between  $N_d/M_{ev} = 0.1$  and 1. This suggests that the optimal value of  $N_d$  is of the order of  $\sqrt{N_{tot}}$  or slightly smaller.

Repeating this exercise by swapping one of the TRENTo parameters (reduced thickness) with another (fluctuation parameter) yields the right panels of figure 4. Because the fluctuation parameter is more challenging to constrain from  $\langle \epsilon_2 \rangle$  and  $\langle \epsilon_3 \rangle$ , the dependence on the number of design points  $N_d$  of the top panel is modest. Yet this does not prevent the other parameter to be better constrained (middle right panel), and the optimal number of design points remains in the vicinity of  $N_d/M_{ev} = 0.25$ .

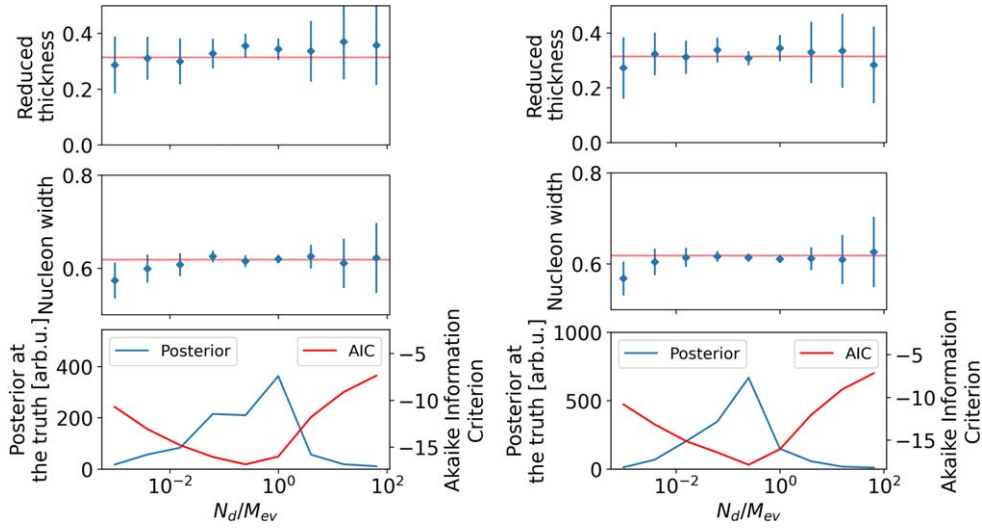
To ensure the robustness of our conclusions, we repeat this same analysis varying the uncertainty on the reference calculations (which are the proxied for data uncertainty) as well as varying the number of parameters and observables.

### 3.1. Dependence on the uncertainty of the reference calculation ('data')

To compute the reference ('truth') observables,  $M_{ev}^{truth}$  TRENTo events are averaged. Changing  $M_{ev}^{truth}$  controls the uncertainty on the reference observables,  $\Delta \langle \epsilon_n^{truth} \rangle$ . In this section, we vary  $M_{ev}^{truth}$  to verify the effect on the optimal number of design points.

In the previous example, because  $N_{tot} = N_d M_{ev}$  and  $M_{ev}^{truth}$  were both equal to  $2^{16}$ , the interpolation and statistical uncertainty in the emulator (controlled by  $N_{tot}$ ) were always large compared to the uncertainty of the reference calculations that are stand-ins for measurements (controlled by  $M_{ev}^{truth}$ ): at best, the statistical uncertainties for the calculations used in the emulator could be equal to the statistical uncertainty from the reference calculations (uncertainty  $\propto 1/\sqrt{2^{16}}$ ) when  $N_d = 1$ , which would evidently be a case with extreme interpolation uncertainty.

We first repeat the previous example with  $M_{ev}^{truth} = 2^{14}$  events while keeping  $N_{tot} = N_d M_{ev} = 2^{16}$ . This means that when  $N_d = 2^2 = 4$ , the statistical uncertainty of the reference calculation is similar to that in the calculations used for the emulator; when  $N_d = 2^4 = 16$ , the statistical uncertainty in the emulator is approximately a factor 2 larger than in that of the reference calculations, etc. The results of the closure tests are shown on the left panel of figure 5. Overall, the results are similar: the constraints on the model parameters are consistent with the true value of the parameters at any ratio  $N_d/M_{ev}$ , but the range  $N_d/M_{ev} \approx 0.1-1$  is found to be optimal.



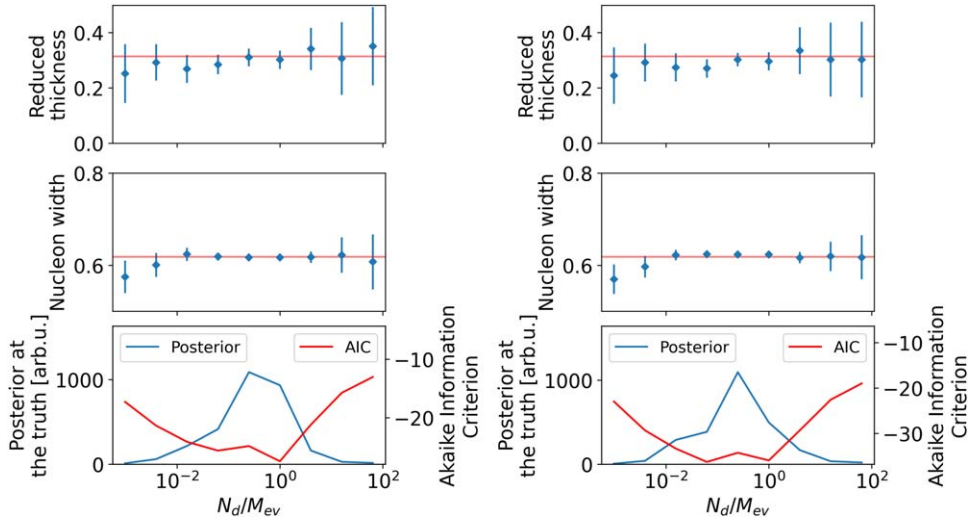
**Figure 5.** Same as left side of figure 4, with (left)  $M_{\text{ev}}^{\text{truth}} = 2^{14}$  events and (right)  $M_{\text{ev}}^{\text{truth}} = 2^{18}$  events.

We repeat this test in the opposite direction, using  $M_{\text{ev}}^{\text{truth}} = 2^{18}$  events to reduce the uncertainty on the reference calculations while keeping the emulator budget the same ( $N_{\text{tot}} = N_d M_{\text{ev}} = 2^{16}$ ). This is equivalent to a scenario where the emulator uncertainties are always much larger than the data uncertainties. This yields the right panel of figure 5. Again, the optimal range of design points remains  $N_d/M_{\text{ev}} \approx 0.1-1$ .

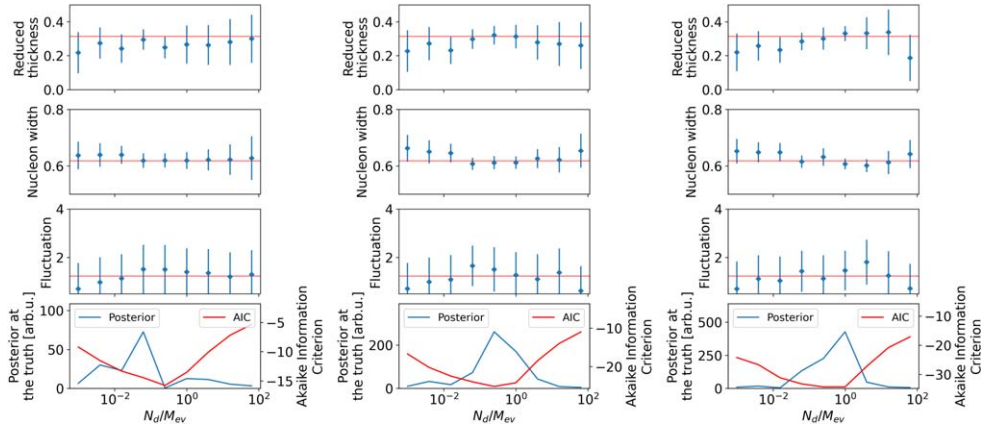
### 3.2. Dependence on the observables and the parameters

Emulator performance depends on multiple factors, in particular the number of parameters, and the complexity of the parameter dependence of the observables. The results of Bayesian parameter inference evidently also depend on the factors just listed that impact emulator performance, and further depend on the information carried by the different observables. We repeat the closure test shown on figure 4 by adding new observables: we first add  $\langle \varepsilon_4 \rangle$ , and then  $\langle \varepsilon_5 \rangle$ . The results are shown respectively on the left and right panels of figure 6. Again, the results remain similar, with the optimal number of design points in the same range as found in previous tests.

We repeat this test, this time with three T<sub>R</sub>ENTO parameters instead of two, starting with two observables ( $\langle \varepsilon_n \rangle$ ,  $n = 2, 3$ ), then three ( $n = 2, 3, 4$ ) and four ( $n = 2, 3, 4, 5$ ) observables. The results are shown respectively on the left, center and right panels of figure 7. The results are consistent with all previous ones: closure is successful since the constraints on all parameters are consistent with the truth value of the parameter for any number of design points. Moreover, constraints on the parameters are still best when the ratio of the number of design points  $N_d$  to the number of T<sub>R</sub>ENTO simulations per design points  $M_{\text{ev}}$  is in the range  $N_d/M_{\text{ev}} \approx 0.1-1$ .



**Figure 6.** Same as left side of figure 4, with (left) three observables ( $\langle \varepsilon_n \rangle$ ,  $n = 2, 3, 4$ ), and (right) four observables ( $\langle \varepsilon_n \rangle$ ,  $n = 2, 3, 4, 5$ ).



**Figure 7.** Same as figure 4, with three TR-ENTo parameters instead of two, and between two (left), three (center) or four (right) observables, with the observables being  $\langle \varepsilon_n \rangle$ ,  $n = 2, 3$ ,  $n = 2, 3, 4$  and  $n = 2, 3, 4, 5$  respectively.

#### 4. Discussion

There is a consistent trend among the results shown in figures 4–7: closure was best when the number of design points was slightly smaller than the square root of the total number of events ( $\sqrt{N_{tot}}$ ), which is equivalent to  $N_d/M_{ev} \approx 0.1-1$ . This trend holds true for all the variations tested in this study.

A rule of thumb that is used at times to determine the number of parameter samples (design points) for a Gaussian process emulator is 10 times the number of parameters ( $10N_{params}$ ) [41]. This guidance is evidently not meant to be precise, if only because it does not take into

account the unavoidable trade-offs between statistical and emulator uncertainty. In the two-parameter case,  $N_d = 10N_{\text{params}} \approx 2^4$ , which is considerably smaller than the optimal  $N_d \approx 2^6 - 2^8$  found in the previous section. We also note that the optimal number of design points did not change significantly when three model parameters were used instead of two (figure 7), which suggests a modest dependence of this optimum on the number of parameters.

While statistical uncertainties generally converge at the rate  $1/\sqrt{M_{\text{ev}}}$ , interpolation uncertainties depend on the details of the model's parameter dependence, and the range of parameters (the prior) over which an emulator is being trained. Rather than attempting to provide general guidance, we focus on application to heavy-ion physics, where our results guidance will apply most directly.

## 5. Implications for Bayesian inference in heavy-ion collisions

As seen in the results from section 3, emulator uncertainty is propagated properly in the Bayesian parameter estimation: even if one uses an emulator with large interpolation or statistical uncertainties, the parameter constraints remain consistent with the true value of the parameter used in the closure test. In this sense, we expect the Gaussian process emulators used in Bayesian inference procedure of heavy-ion data to be robust. On the other hand, we saw that optimizing the interpolation and statistical uncertainty of the emulator can lead to considerably improved constraints on the model parameters at the same computational cost. In our study, we found an optimal ratio of design points to TRENTo events per design point,  $N_d/M_{\text{ev}}$ , around 0.1–1.

The optimum  $N_d/M_{\text{ev}} \approx 1$  is expected if the interpolation uncertainty in the emulator scales approximately as  $1/\sqrt{N_d}$ : that is, if both interpolation and statistical uncertainties scale the same, one should aim for a roughly even number for  $N_d$  and  $M_{\text{ev}}$ . Tests published in [17] found some evidence of  $1/\sqrt{N_d}$  convergence of observables in realistic models of heavy-ion collisions. This convergence is likely to depend on the details of the model and on the choice of observables. The exact implementation of the emulator is expected to impact this convergence as well; it will be valuable to quantify the dependence on  $N_d$  of the interpolation uncertainty if significant changes are made to the emulation method.

To apply our work's conclusions to heavy-ion collisions, one must remember that there are different sources of stochasticity in the collisions, only one of which was considered in this work: fluctuations in the initial impact geometry of the nuclei. Fluctuations in the number of particles produced will also introduce statistical uncertainty, which was not studied in this work. It is common, though not universal, that the two statistical uncertainties are reduced by oversampling the particles produced for each initial geometry of the collision. The factor  $M_{\text{ev}}$  in the ratio  $N_d/M_{\text{ev}} \approx 0.1-1$  corresponds to the number of initial geometries samples and *not* the total number of oversamples.

Finally, we note that the choice of  $M_{\text{ev}}$  might be guided by the need for accuracy for certain statistics-hungry observables, rather than the benefits of the emulator as a whole [19, 20]. It will be useful to repeat the current study with a broader range of observables.

## 6. Conclusion

We studied the interplay between interpolation and statistical uncertainties for Bayesian parameter inference using Gaussian process emulators. We used the TRENTo model of initial conditions in heavy-ion collisions, with observables given by the event-averaged spatial

anisotropies  $\langle \varepsilon_n \rangle$  (equation (2)), which are analogous to momentum anisotropy observables encountered in heavy-ion collisions. Using  $N_d$  samples of the T<sub>R</sub>ENTo parameter space and  $M_{ev}$  T<sub>R</sub>ENTo events at each parameter point, we performed closure tests to determine the impact of the emulator uncertainty on constraining model parameters. We found that, for a budget of  $N_{tot}$  total T<sub>R</sub>ENTo events, the optimal number of design points was around  $N_d \approx (0.25-1.0)\sqrt{N_{tot}}$ , which corresponds to  $N_d/M_{ev} \approx 0.1-1$ . We found this optimum by looking at multiple different metrics: the maximum value and width of the marginalized posterior, the value of the posterior at the true value of the parameter, and the AIC. We found the posterior at the truth to be very effective at quantifying closure.

In view of our results, we believe it would be beneficial for Bayesian inference applications in heavy-ion physics to experiment with different ratios  $N_d/M_{ev}$ . Simulating a large number of collision events, which reduce all sources of stochastic uncertainties, can give access to statistics-hungry observables [19, 20]; yet using a large number  $M_{ev}$  of events at each parameter sample will not provide significant benefits for a large number of observables whose emulation is limited by interpolation uncertainty. An improved balance between interpolation and statistical uncertainties can help reduce the emulator's uncertainty below current experimental uncertainties, allowing us to make full use of heavy-ion measurements from the RHIC and the LHC.

## Acknowledgments

The authors thank Simon Mak for insightful discussions, and Ron Solz and Matthew Luzum for valuable feedback. BW thanks Duke University and Wayne State University for financial support. SAB and J-FP have been supported by the US Department of Energy Grant No. DE-FG02-05ER41367. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the US Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award NP-ERCAP0022230. Calculations were also performed on the Duke Compute Cluster.

## Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: [https://github.com/Duke-QCD/emulator\\_uncertainty\\_optimization\\_arXiv\\_2301.08385](https://github.com/Duke-QCD/emulator_uncertainty_optimization_arXiv_2301.08385).

## ORCID iDs

Brandon Weiss  <https://orcid.org/0000-0003-1606-1089>  
Jean-François Paquet  <https://orcid.org/0000-0001-8736-8171>  
Steffen A Bass  <https://orcid.org/0000-0002-9451-0954>

## References

- [1] Bass S A 2002 Microscopic reaction dynamics at SPS and RHIC *Nucl. Phys. A* **698** 164–70
- [2] Gale C, Jeon S and Schenke B 2013 Hydrodynamic modeling of heavy-ion collisions *Int. J. Mod. Phys. A* **28** 1340011
- [3] Heinz U and Snellings R 2013 Collective flow and viscosity in relativistic heavy-ion collisions *Ann. Rev. Nucl. Part. Sci.* **63** 123–51

- [4] Derradi de Souza R, Koide T and Kodama T 2016 Hydrodynamic approaches in relativistic heavy ion reactions *Prog. Part. Nucl. Phys.* **86** 35–85
- [5] Philipsen O 2013 The QCD equation of state from the lattice *Prog. Part. Nucl. Phys.* **70** 55–107
- [6] Petersen H, Coleman-Smith C, Bass S A and Wolpert R 2011 Constraining the initial state granularity with bulk observables in Au+Au collisions at  $\sqrt{s_{NN}} = 200$  GeV *J. Phys. G* **38** 045102
- [7] Novak J, Novak K, Pratt S, Vredevoogd J, Coleman-Smith C and Wolpert R 2014 Determining fundamental properties of matter created in ultrarelativistic heavy-ion collisions *Phys. Rev. C* **89** 034917
- [8] Sangaline E and Pratt S 2016 Toward a deeper understanding of how experiments constrain the underlying physics of heavy-ion collisions *Phys. Rev. C* **93** 024908
- [9] Pratt S, Sangaline E, Sorensen P and Wang H 2015 Constraining the eq. of state of super-Hadronic matter from heavy-ion collisions *Phys. Rev. Lett.* **114** 202301
- [10] Bernhard J E, Marcy P W, Coleman-Smith C E, Huzurbazar S, Wolpert R L and Bass S A 2015 Quantifying properties of hot and dense QCD matter through systematic model-to-data comparison *Phys. Rev. C* **91** 054910
- [11] Bernhard J E, Moreland J S, Bass S A, Liu J and Heinz U 2016 Applying Bayesian parameter estimation to relativistic heavy-ion collisions: simultaneous characterization of the initial state and quark-gluon plasma medium *Phys. Rev. C* **94** 024907
- [12] Moreland J S, Bernhard J E and Bass S A 2020 Bayesian calibration of a hybrid nuclear collision model using p-Pb and Pb–Pb data at energies available at the CERN Large Hadron Collider *Phys. Rev. C* **101** 024911
- [13] Bernhard J E 2018 Bayesian parameter estimation for relativistic heavy-ion collisions *PhD Thesis* Duke U. 4
- [14] Bernhard J E, Moreland J S and Bass S A 2019 Bayesian estimation of the specific shear and bulk viscosity of quark–gluon plasma *Nat. Phys.* **15** 1113–7
- [15] Everett D *et al* 2021 Multisystem Bayesian constraints on the transport coefficients of QCD matter *Phys. Rev. C* **103** 054904
- [16] Nijs G, van der Schee W, Gürsoy U and Snellings R 2021 Transverse momentum differential global analysis of heavy-ion collisions *Phys. Rev. Lett.* **126** 202301
- [17] Nijs G, van der Schee W, Gürsoy U and Snellings R 2021 Bayesian analysis of heavy ion collisions with the heavy ion computational framework Trajectum *Phys. Rev. C* **103** 054909
- [18] Everett D *et al* 2021 Phenomenological constraints on the transport properties of QCD matter with data-driven model averaging *Phys. Rev. Lett.* **126** 242301
- [19] Parkkila J E, Onnerstad A and Kim D J 2021 Bayesian estimation of the specific shear and bulk viscosity of the quark-gluon plasma with additional flow harmonic observables *Phys. Rev. C* **104** 054904
- [20] Parkkila J E, Onnerstad A, Taghavi S F, Mordasini C, Bilandzic A, Virta M and Kim D J 2022 New constraints for QCD matter from improved Bayesian parameter estimation in heavy-ion collisions at LHC *Phys. Lett. B* **835** 137485
- [21] Moreland J S, Bernhard J E and Bass S A 2015 Alternative ansatz to wounded nucleon and binary collision scaling in high-energy nuclear collisions *Phys. Rev. C* **92** 011901
- [22] Luzum M and Petersen H 2014 Initial state fluctuations and final state correlations in relativistic heavy-ion collisions *J. Phys. G* **41** 063102
- [23] Ollitrault J-Y 1992 Anisotropy as a signature of transverse collective flow *Phys. Rev. D* **46** 229–45
- [24] Miller M and Snellings R 2003 Eccentricity fluctuations and its possible effect on elliptic flow measurements arXiv:nucl-ex/0312008
- [25] Alver B *et al* 2007 System size, energy, pseudorapidity, and centrality dependence of elliptic flow *Phys. Rev. Lett.* **98** 242302
- [26] Alver B and Roland G 2010 Collision geometry fluctuations and triangular flow in heavy-ion collisions *Phys. Rev. C* **81** 054905
- [27] Gardim F G, Grassi F, Luzum M and Ollitrault J-Y 2012 Mapping the hydrodynamic response to the initial geometry in heavy-ion collisions *Phys. Rev. C* **85** 024908
- [28] Williams C K and Rasmussen C E 2006 *Gaussian Processes for Machine Learning* vol 2 (Cambridge, MA: MIT Press)
- [29] Roberts M [https://github.com/j-f-paquet/bayesian\\_parameter\\_estimation\\_tutorial/blob/master/parameter\\_space\\_sampling/param\\_sampling.ipynb](https://github.com/j-f-paquet/bayesian_parameter_estimation_tutorial/blob/master/parameter_space_sampling/param_sampling.ipynb)

- [30] Niederreiter H 1992 *Random Number Generation and Quasi-Monte Carlo Methods* (Philadelphia, PA: SIAM)
- [31] Paquet J-F *et al* 2021 Revisiting Bayesian constraints on the transport coefficients of QCD *Nucl. Phys. A* **1005** 121749
- [32] Cao S *et al* 2021 Determining the jet transport coefficient  $\hat{q}$  from inclusive hadron suppression measurements using Bayesian parameter estimation *Phys. Rev. C* **104** 024905
- [33] Nijs G and van der Schee W 2022 Predictions and postdictions for relativistic lead and oxygen collisions with the computational simulation code Trajectum *Phys. Rev. C* **106** 044903
- [34] Xie M, Ke W, Zhang H and Wang X-N 2022 Information field based global Bayesian inference of the jet transport coefficient arXiv:2206.01340 [hep-ph]
- [35] Xie M, Ke W, Zhang H and Wang X-N 2022 Global constraint on the jet transport coefficient from single hadron, dihadron and  $\gamma$ -hadron spectra in high-energy heavy-ion collisions arXiv:2208.14419 [hep-ph]
- [36] Liyanage D, Sürer O, Plumlee M, Wild S M and Heinz U 2023 Bayesian calibration of viscous anisotropic hydrodynamic simulations of heavy-ion collisions arXiv:2302.14184 [nucl-th]
- [37] Heffernan M R, Gale C, Jeon S and Paquet J-F 2023 Bayesian quantification of strongly-interacting matter with color glass condensate initial conditions arXiv:2302.09478 [nucl-th]
- [38] Akaike H 1974 A new look at the statistical model identification *IEEE Trans. Autom. Control* **19** 716–23
- [39] Liddle A R 2007 Information criteria for astrophysical model selection *Mon. Not. R. Astron. Soc.* **377** L74–8
- [40] Liddle A R 2004 How many cosmological parameters *Mon. Not. R. Astron. Soc.* **351** L49–53
- [41] Loeppky J L, Sacks J and Welch W J 2009 Choosing the sample size of a computer experiment: a practical guide, *Technometrics* **51** 366–76