

Bayesian analysis in heavy-ion collisions

Govert Nijs

Sep 23, 2026

Based on:

- Giacalone, GN, van der Schee, 2305.00015, 2606.03993
- ALICE, ALI-PREL-601142




Utrecht
University

Heavy ion collisions in the Bayesian era

- Over the past decade, Bayesian analysis has transformed heavy ion theory.
- We can now quantitatively compare with data, and our fits get ever more precise.
- On the one hand, this allows for more precise extractions of parameters. . .
- . . . and on the other hand we become sensitive to details which were previously unimportant.
 - This is a burden if we want to achieve our precision extractions.
 - But also an opportunity, as these details are often interesting in their own right.
- One of these details is the connection with nuclear structure.
- We have used this new connection in two directions:
 - Using measured heavy ion data to infer the shapes of nuclei, which are interesting for related fields.
 - Using the known difference in shape between ^{16}O and ^{20}Ne to learn something about the smallest droplet of QGP.




Some simple intuition



- Hydrodynamics can be intuitively understood:
 - Pressure gradients drive expansion.
 - Hotter systems expand faster, resulting in more transverse momentum.
 - Spatially anisotropic systems expand preferentially along the short axis, resulting in momentum anisotropy in the final state.
- The nuclear structure of the colliding nuclei influences the initial shape of the QGP that is formed.
 - More compact (2D projections of) nuclei result in more transverse momentum.
 - Deformed nuclei generate a larger v_2 .


Bayesian analysis workflow

- In principle, Bayesian analysis is simply a fit to data.
- In practice the process is more complicated:
 - Generate a large number of randomly chosen parameter sets called *design points*.
 - Run the model for each one to obtain the prior.
 - Train the emulator.
 - Run the MCMC to obtain the posterior.
- The posterior then is a list of likely parameter sets.




Bayesian analysis workflow

- In principle, Bayesian analysis is simply a fit to data.
- In practice the process is more complicated:
 - Generate a large number of randomly chosen parameter sets called *design points*.
 - Run the model for each one to obtain the prior.
 - Train the emulator.
 - Run the MCMC to obtain the posterior.
- The posterior then is a list of likely parameter sets.



Bayesian analysis workflow


- In principle, Bayesian analysis is simply a fit to data.
- In practice the process is more complicated:
 - Generate a large number of randomly chosen parameter sets called *design points*.
 - Run the model for each one to obtain the prior.
 - Train the emulator.
 - Run the MCMC to obtain the posterior.
- The posterior then is a list of likely parameter sets.




Govert Nijs

Uses of Bayesian analysis: viscosities

- We know the QGP phase is described by viscous hydrodynamics.
 - We know exactly what the free parameters are, i.e. η/s , ζ/s , ...
- We can use Bayesian analysis to find data-preferred values for these parameters.
- The values of the parameters provide an interface with microscopic theories of the QGP.




Uses of Bayesian analysis: parameterized phenomenology



- For the initial state, there is no single widely accepted model.
- With a phenomenological model such as T_RENTO, aspects of microscopic models can be tested, such as the scaling shown here, parameterized by p .
 - IP-Glasma and EKRT are ruled in.
 - KLN and wounded nucleon are ruled out.


Using the posterior parameter values to make predictions

- The posterior parameter values can be used to make predictions for new observables.
 - When using multiple samples from the posterior, this includes systematic uncertainty from the parameter estimation.
- Here shown is the prediction for ultracentral $\langle p_T \rangle$.



Using the posterior parameter values to make predictions

- The posterior parameter values can be used to make predictions for new observables.
 - When using multiple samples from the posterior, this includes systematic uncertainty from the parameter estimation.
- Here shown is the prediction for ultracentral $\langle p_T \rangle$.
- Precise agreement between theory and experiment.



Neutron skin


- In a ^{208}Pb nucleus, neutrons sit further from the center than protons.

- This is quantified by the *neutron skin*:


$$\Delta r_{np} = \langle r^2 \rangle_n^{1/2} - \langle r^2 \rangle_p^{1/2},$$

i.e. the *difference* in RMS radii of the neutron and proton distributions.

- Heavy nuclei and neutron stars are sensitive to the same nuclear interactions.
 - A constraint on Δr_{np} translates directly into a constraint on the radius of a $1.4M_\odot$ neutron star.
 - We can learn something about the low T , high μ_B region even at LHC energies!



How to measure neutron skin?



- To measure the neutron skin, we need the distributions of protons and neutrons inside the nucleus.
 - The proton distribution distribution is well-known from electron scattering.
- Several different methods are in use for the neutron distribution:
 - Polarized electron scattering off ^{208}Pb (PREX).
 - Photon tomography of ^{197}Au (STAR).
- Heavy ion collisions provide a completely orthogonal method.
 - Sensitive to the total matter distribution inside the nucleus.
 - Purely gluonic measurement.


The Woods-Saxon distribution

- Nucleon positions are drawn from a Woods-Saxon distribution:

$$\rho_{\text{WS}}(r) \propto \frac{1}{1 + \exp\left(\frac{r-R}{a}\right)}.$$

- We fix R for both protons and neutrons.
- We fix a for protons, while varying a_n as a parameter.
- Neutron skin $\Delta r_{np} = \langle r^2 \rangle_n^{1/2} - \langle r^2 \rangle_p^{1/2}$ strongly depends on a_n :

$$\langle r^2 \rangle_{\text{WS}} = \frac{12a^2 \text{Li}_5(-e^{R/a})}{\text{Li}_3(-e^{R/a})}.$$



	proton	neutron
R [fm]	6.68	6.69
a [fm]	0.447	a_n


The Woods-Saxon distribution

- Nucleon positions are drawn from a Woods-Saxon distribution:

$$\rho_{\text{WS}}(r) \propto \frac{1}{1 + \exp\left(\frac{r-R}{a}\right)}.$$

- We fix R for both protons and neutrons.
- We fix a for protons, while varying a_n as a parameter.
- Neutron skin $\Delta r_{\text{np}} = \langle r^2 \rangle_n^{1/2} - \langle r^2 \rangle_p^{1/2}$ strongly depends on a_n :

$$\langle r^2 \rangle_{\text{WS}} = \frac{12a^2 \text{Li}_5(-e^{R/a})}{\text{Li}_3(-e^{R/a})}.$$






	proton	neutron
R [fm]	6.68	6.69
a [fm]	0.447	a_n




Do we have observables sensitive to a_n ?

- Initial geometry is sensitive to a_n .
Larger nuclei lead to:
 - Larger hadronic PbPb cross-section,
 - Larger initial QGP size,
 - Smaller initial QGP eccentricity.
- Final state observables are in turn sensitive to initial geometry. Larger Δr_{np} leads to:
 - Larger hadronic PbPb cross-section,
 - Smaller charged particle yield,
 - Smaller mean transverse momentum,
 - Smaller elliptic flow.




Do we have observables sensitive to a_n ?

- Initial geometry is sensitive to a_n .
Larger nuclei lead to:
 - Larger hadronic PbPb cross-section,
 - Larger initial QGP size,
 - Smaller initial QGP eccentricity.
- Final state observables are in turn sensitive to initial geometry. Larger Δr_{np} leads to:
 - Larger hadronic PbPb cross-section,
 - Smaller charged particle yield,
 - Smaller mean transverse momentum,
 - Smaller elliptic flow.




Do we have observables sensitive to a_n ?

- Initial geometry is sensitive to a_n .
Larger nuclei lead to:
 - Larger hadronic PbPb cross-section,
 - Larger initial QGP size,
 - Smaller initial QGP eccentricity.
- Final state observables are in turn sensitive to initial geometry. Larger Δr_{np} leads to:
 - Larger hadronic PbPb cross-section,
 - Smaller charged particle yield,
 - Smaller mean transverse momentum,
 - Smaller elliptic flow.




Do we have observables sensitive to a_n ?

- Initial geometry is sensitive to a_n .
Larger nuclei lead to:
 - Larger hadronic PbPb cross-section,
 - Larger initial QGP size,
 - Smaller initial QGP eccentricity.
- Final state observables are in turn sensitive to initial geometry. Larger Δr_{np} leads to:
 - Larger hadronic PbPb cross-section,
 - Smaller charged particle yield,
 - Smaller mean transverse momentum,
 - Smaller elliptic flow.




What does a_n correlate with?

- a_n is not the only parameter affecting the initial geometry, leading to correlations. a_n :
 - anticorrelates with p ,
 - mildly anticorrelates with both w and q .
- Correlations highlight the importance of global analysis.
- Parameters are not degenerate, allowing us to extract a_n , and with it, Δr_{np} .




Bayesian analysis result using LHC data

- Resulting posterior for Δr_{np} is compatible with PREX II and *ab initio* nuclear theory.
- Slightly stronger constraint than PREX II ($\Delta r_{np} = 0.283 \pm 0.071$).
- Result is in principle improvable with better Bayesian analyses.
 - May be hard to do in practice.
 - The current analysis already took 2M CPUh.
- We vary a Woods-Saxon parameter, but Δr_{np} is defined purely in terms of nucleon positions.
 - The observable we report is well-defined.



The neutron skin of ^{48}Ca

- JLab measured not only the neutron skin of ^{208}Pb , but also of ^{48}Ca .
- Nuclear structure models have trouble describing both of these measurements simultaneously.
- Given collisions of ^{48}Ca at the LHC, we could probably extract the neutron skin.
 - Including the (isospin symmetric) ^{40}Ca as a baseline, we could possibly get even more precision.



The deformation of ^{129}Xe


- We describe ^{129}Xe with a *deformed* Woods-Saxon distribution:

$$\rho_{\text{WS}}(r) \propto \frac{1}{1 + \exp\left(\frac{r - R(\theta, \phi)}{a}\right)},$$

$$R(\theta, \phi) = R \left(1 + \beta_2 \cos \gamma Y_2^0 + \beta_2 \sin \gamma Y_2^2 \right),$$

with Y_l^m spherical harmonics.

- We do not distinguish protons and neutrons.
- Observables in the central region are sensitive to β_2 and γ :
 - v_2 is sensitive to β_2 .
 - ρ_2 is sensitive to β_2 and γ .



The deformation of ^{129}Xe


- We describe ^{129}Xe with a *deformed* Woods-Saxon distribution:

$$\rho_{WS}(r) \propto \frac{1}{1 + \exp\left(\frac{r - R(\theta, \phi)}{a}\right)},$$


$$R(\theta, \phi) = R \left(1 + \beta_2 \cos \gamma Y_2^0 + \beta_2 \sin \gamma Y_2^2 \right),$$

with Y_l^m spherical harmonics.

- We do not distinguish protons and neutrons.
- Observables in the central region are sensitive to β_2 and γ :
 - v_2 is sensitive to β_2 .
 - ρ_2 is sensitive to β_2 and γ .



A selection of posterior observables




- Precise description of $v_n\{k\}$ and ρ_2 .
 - We fit PbPb as well as XeXe data.
 - For ρ_2 we also fit to the difference of XeXe compared to PbPb, for $v_n\{k\}$ we fit to the ratio.




Details are important!

- Some observables, such as ρ_2 , depend sensitively on the measurement details.
- Bin size matters:
 - Measurement of ρ is done in small centrality bins.
 - Should compute in small bins, then average.
- Definition of centrality:
 - Different definitions give different answers.
 - In this case, many definitions are available.
 - Should compare to a measurement which uses a definition that our model can reliably compute.
- Ignoring these details can result in biased answers!




Details are important!

- Some observables, such as ρ_2 , depend sensitively on the measurement details.
- Bin size matters:
 - Measurement of ρ is done in small centrality bins.
 - Should compute in small bins, then average.
- Definition of centrality:
 - Different definitions give different answers.
 - In this case, many definitions are available.
 - Should compare to a measurement which uses a definition that our model can reliably compute.
- Ignoring these details can result in biased answers!



Bayesian analysis result using LHC data

- The deformation parameters β_2 and γ are well constrained.
- Without ρ_2 , only β_2 is constrained.
- ^{129}Xe is truly triaxial, shaped like a kiwi.
- Extracted values agree well with ab initio nuclear structure computations.




1-, 2- and 3-nucleon correlators

		^{208}Pb	^{129}Xe
1-body	$\langle R_2 \rangle_1 [\text{fm}^2]$	$20.92(14)_{\text{syst}}$	$15.47(42)_{\text{syst}}$
	$\langle R_3 \rangle_1 [\text{fm}^3]$	$115.0(13)_{\text{syst}}$	$74.0(33)_{\text{syst}}$
	$\langle R_4 \rangle_1 [\text{fm}^4]$	$681(12)_{\text{syst}}$	$384(26)_{\text{syst}}$
	$\langle R_6 \rangle_1 [\text{fm}^6]$	$28.4(9)_{\text{syst}} \times 10^3$	$12.6(15)_{\text{syst}} \times 10^3$
2-body	$\langle \mathcal{E}_2 \mathcal{E}_{-2} \rangle_2 / \langle R_2 \rangle_1^2$	$0.29(3)_{\text{stat}}(56)_{\text{syst}} \times 10^{-3}$	$13.1(1)_{\text{stat}}(19)_{\text{syst}} \times 10^{-3}$
	$\langle \mathcal{E}_3 \mathcal{E}_{-3} \rangle_2 / \langle R_3 \rangle_1^2$	$-29(38)_{\text{stat}} \times 10^{-6}$	$4.3(1)_{\text{stat}}(26)_{\text{syst}} \times 10^{-3}$
3-body	$(\langle R_2 \mathcal{E}_2 \mathcal{E}_{-2} \rangle_3 - \langle R_2 \rangle_1 \langle \mathcal{E}_2 \mathcal{E}_{-2} \rangle_2) / \langle R_2 \rangle_1^3$	$1.3(14)_{\text{stat}}(28)_{\text{syst}} \times 10^{-6}$	$0.01(1)_{\text{stat}}(30)_{\text{syst}} \times 10^{-3}$
	$(\langle R_2 \mathcal{E}_3 \mathcal{E}_{-3} \rangle_3 - \langle R_2 \rangle_1 \langle \mathcal{E}_3 \mathcal{E}_{-3} \rangle_2) / \langle R_2 \rangle_1 \langle R_3 \rangle_1^2$	$-1.4(17)_{\text{stat}} \times 10^{-6}$	$0.139(7)_{\text{stat}}(75)_{\text{syst}} \times 10^{-3}$


- Woods-Saxon parameters are not well-defined observables.
- We can rephrase in terms of nucleon position correlators in the transverse plane:
 - β_n carry the same information as the 2-body correlators shown above.
 - γ carries the same information as the first 3-body correlator.
- The correlators are extracted with a systematic uncertainty propagated from the Bayesian fit.

Nuclear fluctuations




- Instead of taking Woods-Saxon parameters fixed, one can draw them from a distribution.
- ALICE did this for β_2 and γ .
- In principle, fitting to higher moments can constrain this.
- In practice, only a mild constraint is obtained.

Nuclear fluctuations



- Instead of taking Woods-Saxon parameters fixed, one can draw them from a distribution.
- ALICE did this for β_2 and γ .
- In principle, fitting to higher moments can constrain this.
- In practice, only a mild constraint is obtained.

Nuclear fluctuations




- Instead of taking Woods-Saxon parameters fixed, one can draw them from a distribution.
- ALICE did this for β_2 and γ .
- In principle, fitting to higher moments can constrain this.
- In practice, only a mild constraint is obtained.

Neutrinoless double beta decay

- For certain nuclei, single beta decay is kinematically forbidden, but double beta decay is allowed.
- If the neutrino is a Majorana particle, then the two neutrinos can annihilate: neutrinoless double beta decay.
 - In the spectrum for the electrons, this would show up as a peak at the end.
- The decay rate is given by:

$$\Gamma \propto |M|^2 m_{\beta\beta}^2.$$




- $m_{\beta\beta}$ is the *effective Majorana mass* of the neutrinos.
- $|M|$ is a *nuclear matrix element* for this decay.
- $|M|$ is very complicated to compute, and cannot be measured by a different process.

Neutrinoless double beta decay

- For certain nuclei, single beta decay is kinematically forbidden, but double beta decay is allowed.
- If the neutrino is a Majorana particle, then the two neutrinos can annihilate: neutrinoless double beta decay.
 - In the spectrum for the electrons, this would show up as a peak at the end.
- The decay rate is given by:

$$\Gamma \propto |M|^2 m_{\beta\beta}^2.$$

- $m_{\beta\beta}$ is the *effective Majorana mass* of the neutrinos.
- $|M|$ is a *nuclear matrix element* for this decay.
- $|M|$ is very complicated to compute, and cannot be measured by a different process.




Neutrinoless double beta decay

- For certain nuclei, single beta decay is kinematically forbidden, but double beta decay is allowed.
- If the neutrino is a Majorana particle, then the two neutrinos can annihilate: neutrinoless double beta decay.
 - In the spectrum for the electrons, this would show up as a peak at the end.
- The decay rate is given by:

$$\Gamma \propto |M|^2 m_{\beta\beta}^2.$$

- $m_{\beta\beta}$ is the *effective Majorana mass* of the neutrinos.
- $|M|$ is a *nuclear matrix element* for this decay.
- $|M|$ is very complicated to compute, and cannot be measured by a different process.




Neutrinoless double beta decay


- For certain nuclei, single beta decay is kinematically forbidden, but double beta decay is allowed.
- If the neutrino is a Majorana particle, then the two neutrinos can annihilate: neutrinoless double beta decay.
 - In the spectrum for the electrons, this would show up as a peak at the end.
- The decay rate is given by:

$$\Gamma \propto |M|^2 m_{\beta\beta}^2.$$

- $m_{\beta\beta}$ is the *effective Majorana mass* of the neutrinos.
- $|M|$ is a *nuclear matrix element* for this decay.
- $|M|$ is very complicated to compute, and cannot be measured by a different process.



Nuclear matrix elements from heavy ion collisions



- Nuclear matrix elements are very correlated with the *difference* in shape of the nuclei.
 - This leaves imprints in heavy ion collisions of these nuclei.
 - The nuclei involved are isobar pairs, making extracting the difference in shape easier.
 - The 'shape' can be made well-defined using multinucleon correlators.
- We are currently expanding on this, and will conduct a sensitivity study for $^{76}\text{Ge} \mapsto ^{76}\text{Se}$.

Conclusions


- Bayesian analysis allows us to quantitatively describe a large amount of experimental data simultaneously.
- We can gain quantitative insight into the parameters of our model, *including the nuclear structure of the colliding nuclei*.
 - LHC based determinations of the neutron skin of ^{208}Pb are competitive with other experimental methods.
 - The deformation of ^{129}Xe can be determined in a way that directly probes the nucleon positions.
 - These quantities can be reported in terms of genuine observables in terms of correlators.
- We still only probe a small subset of available data in our fits.

A concrete example: light ion data

- A motivation for light ion running was that it is sensitive to different things compared to PbPb:
 - A shorter lifetime means more sensitivity to the pre-hydrodynamic stage and initial state.
 - A smaller, but controlled, system allows to study the path length dependence of energy loss.
- By now enough soft sector quantities have been measured that we started a Bayesian analysis including OO, NeNe and ratio data.
 - Still work in progress, so no showable results yet. . .
 - Stay tuned!

The LHC as a finite resource

- Nuclear structure has uncovered opportunities for injecting new species into the LHC.
 - More such opportunities are expected in the future.
- The LHC will shut down in 2041, and...
 - Proposals need to be well-motivated and worked out in detail.
 - Getting proposals approved takes a long time.
- We should urgently identify the most promising candidates.
- In parallel, we should try to describe as much LHC data as possible, to identify gaps in our knowledge where further experimental input is required.



Backup


Backup




Bayesian analysis details

- 3000 design points.
- 18k events per design point.
- Every 15th design point has $10\times$ more statistics, enabling to emulate 'hard' observables such as $SC(n, m)$ and $\rho(v_2\{2\}^2, \langle p_T \rangle)$.


Error budget




Posterior observables (1/3)



Posterior observables (2/3)



Posterior observables (3/3)



T_{RENT}o initial conditions

- Nucleons A and B become *wounded* with probability

$$P_{\text{wounded}} = 1 - \exp\left(-\sigma_{gg} \int d\mathbf{x} \rho_A(\mathbf{x}) \rho_B(\mathbf{x})\right), \quad \rho_A \propto \exp\left(\frac{-|\mathbf{x} - \mathbf{x}_A|^2}{2w^2}\right).$$

- Each wounded nucleon desposits energy into its nucleus's *thickness function* $\mathcal{T}_{A/B}$:

$$\mathcal{T}_{A/B} = \sum_{i \in \text{wounded } A/B} \gamma \exp(-|\mathbf{x} - \mathbf{x}_i|^2/2w^2),$$

with γ drawn from a gamma distribution with mean 1 and standard deviation σ_{fluct} .

- Actual formulas slightly modified because each nucleon has n_c constituents.

The T_RENTo phenomenological ansatz


- The standard T_RENTo formula combines thickness functions of the two nuclei $\mathcal{T}_A$ and $\mathcal{T}_B$ into a *reduced thickness* $\mathcal{T}$, interpreted as an energy density:

$$\mathcal{T} \propto \left(\frac{\mathcal{T}_A^p + \mathcal{T}_B^p}{2} \right)^{1/p},$$

with p a parameter.

- Some useful limits:

p	-1	0	1
$\mathcal{T}$	$\frac{2}{\frac{1}{\mathcal{T}_A} + \frac{1}{\mathcal{T}_B}}$	$\sqrt{\mathcal{T}_A \mathcal{T}_B}$	$\frac{\mathcal{T}_A + \mathcal{T}_B}{2}$



Free streaming pre-hydrodynamic stage

- T_RENTo creates matter at proper time $\tau = 0^+$.
- Propagate the matter using free streaming:

$$T^{\mu\nu}(x, y, \tau_{\text{hyd}}) = \frac{1}{\tau_{\text{hyd}}} \int d\phi \hat{p}^\mu \hat{p}^\nu \mathcal{T}(x - \tau_{\text{hyd}} \cos \phi, y - \tau_{\text{hyd}} \sin \phi),$$

with

$$\hat{p}^\mu = \begin{pmatrix} 1 & \cos \phi & \sin \phi \end{pmatrix},$$

giving us the stress tensor $T^{\mu\nu}$ at proper time $\tau = \tau_{\text{hyd}}$.

- Here τ_{hyd} is the time at which hydrodynamics is started.
- The factor $1/\tau_{\text{hyd}}$ is due to longitudinal expansion.



Basics of hydrodynamics

- Hydrodynamics is the ultimate effective theory. Knowledge of the underlying microscopic theory is completely summarized in transport coefficients.
- Only conservation laws survive: equation of motion is simply

$$\partial_\mu T^{\mu\nu} = 0.$$

- Not enough equations to close the system. Need additional assumption of *local thermal equilibrium*.
- We write $T^{\mu\nu}$ in terms of building blocks T , u^μ , $g^{\mu\nu}$ and ∂_μ .



Hydrodynamics in the 14-moment approximation

- Define ($g^{\mu\nu} = \text{diag}(1, -1, -1, -1)$):

$$\Delta^{\mu\nu} = g^{\mu\nu} - u^\mu u^\nu, \quad \nabla^\mu = \Delta^{\mu\nu} \partial_\nu, \quad D = u^\mu \nabla_\mu, \quad \sigma^{\mu\nu} = \nabla^{\langle\mu} u^{\nu\rangle},$$

with $\langle \rangle$ symmetrizing and removing the trace.

- We solve viscous hydrodynamics without currents, i.e.

$$\partial_\mu T^{\mu\nu} = 0, \quad T^{\mu\nu} = e u^\mu u^\nu - (P + \Pi) \Delta^{\mu\nu} + \pi^{\mu\nu},$$

- $\pi^{\mu\nu}$ and Π follow the 14-moment approximation:

$$\begin{aligned} -\tau_\pi \Delta_\alpha^\mu \Delta_\beta^\nu D \pi^{\alpha\beta} &= \pi^{\mu\nu} - 2\eta \sigma^{\mu\nu} + \delta_{\pi\pi} \pi^{\mu\nu} \nabla \cdot u \\ &\quad - \phi_7 \pi_\alpha^{\langle\mu} \pi^{\nu\rangle\alpha} + \tau_{\pi\pi} \pi_\alpha^{\langle\mu} \sigma^{\nu\rangle\alpha} - \lambda_{\pi\pi} \Pi \sigma^{\mu\nu}, \\ -\tau_\Pi D \Pi &= \Pi + \zeta \nabla \cdot u + \delta_{\Pi\Pi} \nabla \cdot u \Pi - \lambda_{\Pi\pi} \pi^{\mu\nu} \sigma_{\mu\nu}. \end{aligned}$$




Particlization

- At the freeze-out temperature T_{SW} , we turn the fluid back into particles.
- Particles are sampled thermally, and boosted with the fluid velocity u^μ .
- We use the PTB prescription to match $\pi^{\mu\nu}$ and Π across the transition, so that $T^{\mu\nu}$ is smooth.
- After particlization, we use SMASH as a hadronic afterburner.



Fitting to the p Pb and PbPb cross sections

- In the T_{REN}To model, the nucleon size is described by the Gaussian radius w .
- Previous analyses favored $w \approx 1$ fm.
 - This leads to a 3σ discrepancy in σ_{PbPb} .
- Fitting to the p Pb and PbPb cross sections lowers w to 0.6 fm.
 - σ_{PbPb} discrepancy is reduced to 1σ .
 - Many other observables fit slightly worse.
- Smaller width is now compatible with our knowledge of the gluonic structure of the proton at low x .







	$\sigma_{\text{PbPb}}[\text{b}]$	$\sigma_{p\text{Pb}}[\text{b}]$
with σ_{AA}	8.02 ± 0.19	2.20 ± 0.06
without σ_{AA}	8.95 ± 0.36	2.48 ± 0.10
ALICE/CMS	7.67 ± 0.24	2.06 ± 0.08





Implication for $\rho_n \equiv \rho(v_n\{2\}^2, \langle p_T \rangle)$

- Pearson correlation coefficient $\rho(v_n\{2\}^2, \langle p_T \rangle)$ between $v_n\{2\}^2$ and $\langle p_T \rangle$ is sensitive to the nucleon size.
- Postdiction without fitting to σ_{PbPb} and σ_{pPb} is qualitatively wrong:
 - $\rho(v_2\{2\}^2, \langle p_T \rangle)$ goes negative already at 30% centrality.
 - $\rho(v_3\{2\}^2, \langle p_T \rangle)$ has the wrong sign.
- Fitting to σ_{PbPb} and σ_{pPb} results in a much improved agreement.



Nucleon width and $\rho(v_n\{2\}^2, \langle p_T \rangle)$

- Previous study shows that $\rho(v_n\{2\}^2, \langle p_T \rangle)$ depends strongly on the nucleon size w .




Energy deposition in the initial state

- Nuclear thickness functions $\mathcal{T}_{A/B}$ deposit matter into the initial state energy density $\mathcal{T}$ as follows:


$$\mathcal{T} \propto \left(\frac{\mathcal{T}_A^p + \mathcal{T}_B^p}{2} \right)^{q/p} \stackrel{p \rightarrow 0}{=} (\mathcal{T}_A \mathcal{T}_B)^{q/2}.$$

- Previous analyses implicitly set $q = 1$.
- The fit to experimental data favors $q \approx 4/3$.
 - Previous default $q = 1$ is disfavored.
 - Binary scaling $q = 2$ is ruled out.
 - $q = 4/3$ indicates that $\sqrt{\mathcal{T}_A \mathcal{T}_B}$ behaves like an entropy density.



Future improvements

- We kept R_n fixed in the present analysis.
 - Bayesian analysis increases in difficulty with more parameters.
 - A priori it was not clear that this approach would work.
 - Decided to include only a_n in the first analysis.
- What can be expected from varying R_n in a future Bayesian analysis?
 - When varying R_n , as R_n grows, σ_{PbPb} increases and $\langle p_T \rangle$ decreases.
 - Smallness of σ_{PbPb} prefers smaller R_n , possibly leading to a smaller estimate of Δr_{np} .
 - In this case bulk viscosity would need to increase to compensate for $\langle p_T \rangle$.



Nuclear Hamiltonian parameters

- NLEFT uses an $SU(4)$ invariant Hamiltonian with four parameters: C_2 , C_3 , s_L and s_{NL} .
- We use configurations sampled according to this Hamiltonian.
 - This is very computationally expensive!
 - We cannot repeat this every time we want to vary the Hamiltonian parameters.
- Solution: use a weighted ensemble of configurations, where the weights depend on the parameters.
- Example:

$$w_{\text{config}} = 0.467 + 47.3 \frac{C_2 + 3.41 \cdot 10^{-7} \text{ MeV}^{-2}}{-3.41 \cdot 10^{-7} \text{ MeV}^{-2}} + 5.71 \frac{C_3 + 1.40 \cdot 10^{-14} \text{ MeV}^{-5}}{-1.40 \cdot 10^{-14} \text{ MeV}^{-5}} \\ + 26.8 \frac{s_L - 0.061}{0.061} + 160 \frac{s_{NL} - 0.50}{0.50}.$$

- We can use this to compute systematic errors, as was done for $^{20}\text{Ne}/^{16}\text{O}$.
- We can even vary these parameters directly in future Bayesian analyses!
- Eventually we could also redo the ^{208}Pb neutron skin with these *ab initio* calculations instead of a Woods-Saxon.

