

Direct Constraints on the Nuclear Hamiltonian from Relativistic Light-Ion Collisions

Matthew Luzum

University of São Paulo (USP)

Special thanks to **Shihang Shen** for providing NLEFT pinhole configurations.

EMMI Workshop: Light-Ion Collisions 2026 · September 21–25, 2026

Motivation: Model-independent Probe of Nuclear Structure with Collisions

The classical rotor picture vs. quantum many-body reality

1. Classical Rotor Picture

- ▶ Standard heavy-ion analyses probe nuclear structure by sampling nucleons from deformed Woods–Saxon densities:

$$\rho(r, \theta, \phi) = \frac{\rho_0}{1 + \exp\left[\frac{r - R_0(1 + \beta_2 Y_{20} + \beta_3 Y_{30})}{a}\right]}$$

- ▶ Classical rotor: independent nucleons + probabilistic orientation
- ▶ Can incorporate refinements:
 - Fluctuating deformation parameters $P(\beta_2, \gamma)$ to mimic shape softness
 - Minimum exclusion distance ($d_{\min}$) to account for short-range repulsion

2. Quantum Reality

▶ Limitations of the classical rotor:

- Doesn't capture all correlations of a realistic nuclear wave function:
- Real nuclei are quantum $J^\pi = 0^+$ ground states; apparent deformation arises from collective correlations.
- Multi-nucleon clustering, Fermi-Dirac statistics, non-collective correlations

▶ Wanted:

- A **model-independent** extraction directly linked to fundamental nuclear theory.

The Ab Initio Tool: NLEFT Pinholes & Reweighting

Lattice Hamiltonian, pinhole projection, and complex weights

1. Lattice Hamiltonian & Parameters

- ▶ $^{16}\text{O}, ^{20}\text{Ne}$ from PRL 2025 [Giacalone et al.].
- ▶ **Discrete Coordinate Pinholes:** Samples nucleon positions $\vec{r} = (\vec{r}_1, \dots, \vec{r}_A)$ in 0^+ ground state

- ▶ **Leading-Order (LO) Hamiltonian:**

$$H_{\text{LO}} = H_{\text{kin}} + \frac{c_2}{2} \sum_{\mathbf{r}} \tilde{\rho}(\mathbf{r})^2 + \frac{c_3}{6} \sum_{\mathbf{r}} \tilde{\rho}(\mathbf{r})^3$$

SU(4) symmetry suppresses the MC sign problem.

- ▶ **Smearing Parameters (s_L, s_{NL}):**

$$\tilde{\rho}(\mathbf{r}) = \sum_i \tilde{a}_i^\dagger \tilde{a}_i(\mathbf{r}) + s_L \sum_{\langle \mathbf{r}' | \mathbf{r} \rangle} \sum_i \tilde{a}_i^\dagger \tilde{a}_i(\mathbf{r}')$$

$$\tilde{a}_i(\mathbf{r}) = a_i(\mathbf{r}) + s_{NL} \sum_{\langle \mathbf{r}' | \mathbf{r} \rangle} a_i(\mathbf{r}')$$

s_L : local, s_{NL} : non-local; $\langle \mathbf{r}' | \mathbf{r} \rangle$ nearest neighbors.

2. Complex Amplitudes and Perturbations

- ▶ Each configuration carries:
- ▶ **Baseline Amplitude $A(\vec{r})$ (per snapshot):**

$$A = \langle \psi | M^{L_t/2} | \vec{r} \rangle \langle \vec{r} | M^{L_t/2} | \psi \rangle$$

Fluctuates in phase/sign even at baseline ($\vec{\alpha} = 0$).

- ▶ **Perturbations $B_j(\vec{r})$ (Lu et al., PLB 2019):**

$$B_j = -a_t \sum_k \langle \psi | M \dots H_1^{(j)} \dots M | \psi \rangle$$

- ▶ **Reweighted Configuration Weight:**

$$W(\vec{\alpha}) = A + \sum_{j=1}^5 \Delta \alpha_j B_j$$

- ▶ **5 Explored Shifts:** $\pm 10\% c_2, c_3$; $\pm 5\% s_L, s_{NL}$; Coulomb.
- ▶ **Key Advantage:** Enables continuously probing the Hamiltonian parameter space with a single set of collision simulations!

The Statistical Challenge: Amplitude fluctuation/cancellation

Effective sample size, independent nuclei, and exploding event multipliers

Statistical Efficiency

► Effective Sample Size (N_{eff}):

$$N_{\text{eff}} \equiv \frac{|\sum_i w_i|^2}{\sum_i |w_i|^2}, \quad f_{\text{eff}} \equiv \frac{N_{\text{eff}}}{N}$$

Extra events for equal precision: $\text{Penalty} \equiv 1/f_{\text{eff}}$.

► Collision Squaring Rule (A, B independent):


$$W_{\text{coll}} = W_A W_B \implies f_{\text{coll}} = f_{\text{single}}^2$$

- ^{16}O base: $3.3\times \rightarrow 10.9\times \approx 3.3^2$
- ^{20}Ne base: $10.6\times \rightarrow 113.2\times \approx 10.6^2$

► Why Brute Force Fails → Strategy:

- Parameter shifts explode penalties, requiring huge numbers of hydro events.
- **Roadmap:** Fast pre-collision & Glauber exploration first; ML surrogate for full hydro.

Event Multiplier for Equal Precision



$^{20}\text{Ne} + ^{20}\text{Ne}$ parameter shifts: Incur little to no extra penalty above baseline ($\approx 110\times - 115\times$ for all channels, unlike $^{16}\text{O} + ^{16}\text{O}$).

Warmup: Probing Parameter Space via Pre-Collision Structure

Part 1: Nuclear size response (Δr_{rms}) and universal contraction

Nuclear Size Response (Δr_{rms})

► Physical Ground-State Scale:


- Sets the overall nuclear wave function scale:
- $r_{\text{rms}}(^{16}\text{O}) = 2.57 \text{ fm}$, $r_{\text{rms}}(^{20}\text{Ne}) = 3.02 \text{ fm}$.

► Contraction with Attractive LECs:

- Contact attraction (+10% c_2): shrinks r_{rms} by -2.9% (^{16}O) and -4.0% (^{20}Ne).
- Non-local force (+5% s_{NL}): contracts even more: -5.4% (^{16}O) and -7.3% (^{20}Ne).
- Coulomb repulsion expands both (+1.3% and +2.1%).

► Universal Response Ratio:

$$\frac{\Delta r_{\text{rms}}(s_{NL})}{\Delta r_{\text{rms}}(c_2)} \approx 1.85 \quad \text{identical in both nuclei!}$$




Warmup: Probing Parameter Space via Pre-Collision Structure

Part 2: Multi-nucleon α -clustering (C_4) and shape competition in ^{20}Ne

Multi-Nucleon α -Clustering (C_4)

- ▶ **Connected 4-Body Cumulant (C_4):** Subtracts disconnected 2-body contractions to isolate genuine 4-nucleon α -clustering.
- ▶ **Intrinsic Clustering Hierarchy:** ^{16}O has $2\times$ **stronger** intrinsic clustering than ^{20}Ne ($R_4 = 0.145$ vs 0.074).
- ▶ **LEC Sensitivity: ^{16}O vs ^{20}Ne :**
 - ^{16}O (**uniform**): Attractive LECs (c_2, s_{NL}) enhance both clusters equally ($\sim +25\%$ and $\sim +45\%$).
 - ^{20}Ne (**shape competition**): Non-local s_{NL} causes compact tetrahedral to **surge** $+94\%$, while linear chain **collapses** -47% ! (c_2 : $+51\%$ vs -33%).

4-Body α -Clustering Response (C_4)



Warmup: Probing Parameter Space via Pre-Collision Structure

Part 3: Ground-state multipoles B_2^2, B_3^2 and independent sensitivities

Nuclear eccentricities B_2^2 and B_3^2

► Intrinsic Orientation-Averaged Multipoles:

$$B_n^2 \equiv \frac{2\pi}{3} \frac{\langle E_n \rangle}{\langle r_{\perp}^n \rangle^2}, \quad E_n \equiv \frac{|Q_n|^2 - \sum_i r_i^{2n}}{A(A-1)}$$


Formulated to directly dictate collision eccentricities $\varepsilon_n \{2\}$.

► Independent Sensitivities (No Collinearity):

- +10% c_2 selectively drives B_2^2 ($3\times$ larger than B_3^2).
- +10% c_3 selectively drives B_3^2 ($5\times$ larger than B_2^2 !).

► Question: How directly can we actually access these quantities?

Ground-State Multipoles (B_2^2, B_3^2)



Testing the Approximate Relation: $\langle \varepsilon_n^2 \rangle$ vs B_n^2

Relating collision eccentricity to intrinsic multipoles via participant Glauber


The Analytic Relation (Eq. 12)

- **Analytic formula** (Duguet et al., PRL 2025):

$$\langle \varepsilon_n^2 \rangle = \frac{1}{2A} \left[M_n^{(\perp)} + (A-1) \frac{3}{2\pi} B_n^2 \right]$$

- **Participant Glauber Assumptions:** Entropy $S \propto T_A + T_B$ (TRENTo $p = 1$), complete participation ($N_{\text{part}} = 2A$), central ($b = 0$), linear fluctuation expansion δs .
- **Two Competing Contributions:**
 - **Radial kurtosis** $M_n^{(\perp)}$: Single-nucleon shot noise.
 - **Intrinsic multipole** B_n^2 : Two-body correlation.
- **^{16}O is Shot-Noise Dominated:** 1-body term is 88.3% of $\langle \varepsilon_2^2 \rangle$! 2-body B_2^2 is only 11.7%.

Eq. 12 Baseline: Shot-Noise vs. Intrinsic Shape




Testing the Approximate Relation: $\langle \varepsilon_n^2 \rangle$ vs B_n^2

Hamiltonian parameter sensitivities and internal competition in ^{16}O

Parameter Sensitivities & Sign Flips

- ▶ **Core Question:** Does collision eccentricity $\varepsilon_n\{2\}$ follow intrinsic ground-state deformation B_n^2 ?
- ▶ **Radial Response:** $M_n^{(\perp)}$ is **not constant**—stronger binding reduces 1-body ($\Delta M_n^{(\perp)} < 0$), directly opposing multipole correlation growth ($\Delta B_n^2 > 0$).
- ▶ **Internal Sign Flips (^{16}O , $n = 2$):**
 - +10% c_3 : Radial -5.8 beats shape $+2.1 \Rightarrow \Delta \langle \varepsilon_2^2 \rangle = -3.8 \times 10^{-4}$
 - -5% s_L : Radial -13.0 beats shape $+11.3 \Rightarrow \Delta \langle \varepsilon_2^2 \rangle = -1.7 \times 10^{-4}$
- ▶ **Takeaway:** Collision eccentricity moves **opposite** to nuclear eccentricity ($\langle \varepsilon_2^2 \rangle \downarrow$ while $B_2^2 \uparrow$)!

Eq. 12 Parameter Sensitivity (^{16}O , $n = 2$)




Testing the Approximate Relation: $\langle \varepsilon_n^2 \rangle$ vs B_n^2

Collision realities: Non-zero impact parameter and spectator protons

Collision Realities ($b \neq 0$ & Spectators)

- ▶ **Analytic Formula Assumptions:** Formula assumes central collisions ($b = 0$) and complete participation ($N_{\text{part}} = 2A$).
- ▶ **Non-Zero Impact Parameter ($b \neq 0$):** Even selecting complete participation ($N_{\text{part}} = 32$), collisions peak at $\langle b \rangle = 0.64 \text{ fm} \neq 0$ (0.88 fm with ZDC veto).
- ▶ **ZDC Proton Blindness:** ZDCs detect only neutrons Leaves **1.83 spectator protons** in collision volume ($\langle N_{\text{part}} \rangle = 30.2$).
- ▶ **Takeaway:** Ideal selection not realizable in light-ion collisions



Testing the Approximate Relation: $\langle \varepsilon_n^2 \rangle$ vs B_n^2

Breakdown in practice: Scale inflation and centrality-driven sign flips

Breakdown: Scale Inflation & Sign Flips

► Scale Inflation by Spectators:


- Formula reference ($b = 0$): $\langle \varepsilon_2^2 \rangle = 0.0619$.
- Spectator fluctuations inflate scale by +27% (ZDC veto) to +34% (Top 1% N_{part}); tails masquerade as deformation!

► Centrality-Driven Sign Flips:


- Formula & MB: negative slope ($\partial \langle \varepsilon_2^2 \rangle / \partial c_2 < 0$).
- **Top 1% N_{part} selection: response LEAPS TO POSITIVE** ($\partial \langle \varepsilon_2^2 \rangle / \partial c_2 = +1.92 \times 10^{-3}$)!

- **Conclusion:** Analytic formula fails in realistic collisions. Simulations are necessary for precision extraction

(a) Scale Inflation by Spectators



(b) Sensitivity : The Striking Sign Flips



Full Simulations: 3D Dynamics & Neural Surrogates


3D initial conditions, experimental subevents, and $^{16}\text{O} + ^{16}\text{O}$ validation

The 3D Model & Neural Surrogate

- ▶ **3D Simulation Chain:** 3D MC-Glauber (strings) $\rightarrow$ 3D MUSIC viscous hydro $\rightarrow$ UrQMD cascade.
- ▶ **Surrogate Strategy (Hirvonen et al.):** Full hydro too slow; **neural network** evaluates in ms.
- ▶ MC-Glauber $\rightarrow$ NN $\rightarrow Q_n(\eta, p_T)$
- ▶ **Event-Level Targets** $[M, \text{Re}Q_n, \text{Im}Q_n]$: Single-event vectors give cumulants under reweighting **analytically for any shift** $\vec{\lambda}$:

$$\langle Q \dots \rangle_W = \sum_i w_i(\vec{\lambda}) Q_i / \sum_i w_i(\vec{\lambda})$$

- ▶ Trained on 100 events $\times$ 1000 design points



$^{16}\text{O} + ^{16}\text{O}$ **Anchor Validation:** Surrogate (blue) matches 3D hydro (black) across ALICE, ATLAS, CMS, all centralities, and cumulants.

Full Simulations: 3D Dynamics & Neural Surrogates


Predicting event-level flow vectors and $^{20}\text{Ne} + ^{20}\text{Ne}$ validation

The 3D Model & Neural Surrogate

- ▶ **3D Simulation Chain:** 3D MC-Glauber (strings) $\rightarrow$ 3D MUSIC viscous hydro $\rightarrow$ UrQMD cascade.
- ▶ **Surrogate Strategy (Hirvonen et al.):** Full hydro too slow; **neural network** evaluates in ms.
- ▶ MC-Glauber $\rightarrow$ NN $\rightarrow Q_n(\eta, p_T)$
- ▶ **Event-Level Targets** $[M, \text{Re}Q_n, \text{Im}Q_n]$: Single-event vectors give cumulants under reweighting **analytically for any shift $\vec{\lambda}$** :

$$\langle Q \dots \rangle_W = \sum_i W_i(\vec{\lambda}) Q_i / \sum_i W_i(\vec{\lambda})$$

- ▶ Trained on 100 events $\times$ 1000 design points




$^{20}\text{Ne} + ^{20}\text{Ne}$ **Anchor Validation:** Accurate surrogate reproduction of flow harmonics and cumulant ratios.

Synthetic Closure Test: Validating Idealized Parameter Extraction

Recovering injected truth, precision limits

Validation & Precision Limits

- ▶ **Setup:** 892 observables (O-O + Ne-Ne) at injected $\vec{\lambda}_{\text{true}} = [3, 3, -1, 0.5]\%$ with realistic 1.5% LHC errors ($\sigma = 0.015$).
- ▶ **Inference Validation** ($\leq 1\sigma$): MCMC cleanly recovers all 4 injected values (blue crosshairs match posterior modes).
- ▶ **Precision Lower Bound:** At 1.5% data error, statistical sensitivity achieves **sub-percent precision:** $\sigma(c_i) \approx \pm 0.4\%$, $\sigma(s_i) \approx \pm 0.14\%$.
- ▶ Still need model parameter variation, folding in of NN uncertainty




Global Bayesian Framework & LHC Light-Ion Dataset

5-stage simulation pipeline and 892 experimental calibration points

Machine Learning Simulation Pipeline

Bridging ab initio lattice EFT nuclear structure with LHC light-ion collisions via fast neural surrogates




Pipeline Synthesis: Conditioned NN surrogates evaluate 1M events/design in ms $\rightarrow$ Multi-output GP emulator interpolates 26-dim parameter space $\rightarrow$ Global 32-dim MCMC simultaneously fits **892 light-ion data points** across ALICE, ATLAS, and CMS ($v_2\{2\}, v_3\{2\}, v_2\{4\}, v_3\{4\}, dN_{\text{ch}}/d\eta$).

Confronting LHC Data & Status of Nuclear Constraints

Simultaneous joint calibration: $^{16}\text{O} + ^{16}\text{O}$ flow observables (ALICE, ATLAS, CMS)

Model Calibration & LEC Status

- ▶ **Global Fit to 892 Points:** Simultaneous fit in O–O and Ne–Ne. Prior (grey) → Posterior (blue 95% CI) captures all flow harmonics and $dN_{\text{ch}}/d\eta$.
- ▶ **Status of Nuclear LEC Posteriors:** Extracting definitive LEC constraints is an active work in progress:
- ▶ **Roadmap:** Validated posteriors with full uncertainty quantification still to come.




$^{16}\text{O} + ^{16}\text{O}$ Flow Observables (150 bins: ALICE, ATLAS, CMS)

Confronting LHC Data & Status of Nuclear Constraints

Simultaneous joint calibration: $^{20}\text{Ne} + ^{20}\text{Ne}$ flow observables (ALICE, ATLAS, CMS)

Model Calibration & LEC Status

- ▶ **Global Fit to 892 Points:** Simultaneous fit in O–O and Ne–Ne. Prior (grey) → Posterior (blue 95% CI) captures all flow harmonics and $dN_{\text{ch}}/d\eta$.
- ▶ **Status of Nuclear LEC Posteriors:** Extracting definitive LEC constraints is an active work in progress:
- ▶ **Roadmap:** Validated posteriors with full uncertainty quantification still to come.




$^{20}\text{Ne} + ^{20}\text{Ne}$ Flow Observables (150 bins: ALICE, ATLAS, CMS)

Confronting LHC Data & Status of Nuclear Constraints

Longitudinal profiles: CMS $^{16}\text{O} + ^{16}\text{O}$ $dN_{\text{ch}}/d\eta$ (8 centrality classes)

Model Calibration & LEC Status

- **Global Fit to 892 Points:** Simultaneous fit in O–O and Ne–Ne. Prior (grey) → Posterior (blue 95% CI) captures all flow harmonics and $dN_{\text{ch}}/d\eta$.
- **Status of Nuclear LEC Posteriors:** Extracting definitive LEC constraints is an **active work in progress**:
- **Roadmap:** Validated posteriors with full uncertainty quantification still to come.




CMS $^{16}\text{O} + ^{16}\text{O}$ $dN_{\text{ch}}/d\eta$ (8 centrality classes)

Confronting LHC Data & Status of Nuclear Constraints

Longitudinal profiles: ATLAS $^{20}\text{Ne} + ^{20}\text{Ne}$ $dN_{\text{ch}}/d\eta$ (20 centrality classes)

Model Calibration & LEC Status

- ▶ **Global Fit to 892 Points:** Simultaneous fit in O–O and Ne–Ne. Prior (grey) → Posterior (blue 95% CI) captures all flow harmonics and $dN_{\text{ch}}/d\eta$.
- ▶ **Status of Nuclear LEC Posteriors:** Extracting definitive LEC constraints is an **active work in progress**:
- ▶ **Roadmap:** Validated posteriors with full uncertainty quantification still to come.



ATLAS $^{20}\text{Ne} + ^{20}\text{Ne}$ $dN_{\text{ch}}/d\eta$ (20 centrality classes)

Summary & Outlook

Connecting ab initio nuclear theory to relativistic light-ion collisions

Ab Initio Nuclear Theory Meets Relativistic Collisions

- ▶ **Ab Initio Reweighted Nuclear Configurations:** NLEFT coordinate pinholes enable continuous exploration of the nuclear Hamiltonian with a single collision set.
- ▶ **Statistical Bottleneck:** Weight fluctuation/cancelation leads to **effective sample size collapse** and **exploding event multipliers**. Brute-force hydro across parameter space is impossible.
- ▶ **Analytic Formulas Break Down in Practice for Light Ions:**
 - $v_n\{2\}$ doesn't track nuclear eccentricity
 - Realistic centrality selections break analytic connection between nuclear eccentricity and measured flow
 - Full numerical simulations are necessary

Surrogates & Future Roadmap

- ▶ **3D Dynamics & Neural Surrogates:**
 - 3D model matches experimental subevents & centrality selection for each measurement.
 - Event-level surrogate neural network unlocks huge speedup and exact analytical reweighting.
- ▶ **Global Data Confrontation:** Joint Bayesian framework calibrates 892 LHC light-ion data points across ALICE, ATLAS, and CMS in O-O and Ne-Ne.
- ▶ **Roadmap:**
 - Fine tune neural network
 - Fold in uncertainty from neural network into likelihood
 - Validation of convergence of reweighted perturbations across full prior
 - External: new configurations with improved convergence and interactions