HEIDi

A generative AI for heavy-ion reactions and cosmic rays

Lina Jeritslev

ITP, Goethe-Universität Frankfurt; FIAS

2nd GSI/FAIR AI Workshop, 31st August 2026

Simulation of particle collisions

- UrQMD: Ultra-relativistic Quantum Molecular Dynamics
- Microscopic transport model: tracks every hadron, every collision, every decay
- Well established over a big energy range (SIS, AGS, SPS, RHIC)
- But computationally expensive → generating the large event samples many analyses need becomes impractical

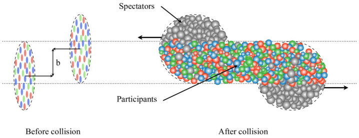

Figure: Simulation of particle collision with UrQMD

Generative AI

- Idea: instead of simulating every collision, learn the statistical structure of UrQMD output directly
- Each collision event = a point cloud of final-state particles in momentum space
- HEIDi (Heavy-ion Events through Intelligent Diffusion) is based on a diffusion model built for 3D point clouds[1]
- Diffusion models: add noise to real data, then learn to reverse it

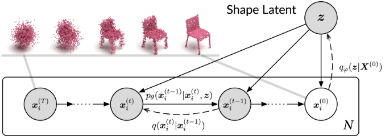

Figure: Noise → structured point cloud.

[1] S. Luo and W. Hu. *Diffusion probabilistic models for 3D point cloud generation*. 2021.

HEIDi

- Three components work together:
  1. **PointNet encoder**: compresses an event into a compact context vector z
  2. **Flow-based decoder**: makes the distribution of z flexible
  3. **Diffusion model**: generates new point clouds conditioned on z

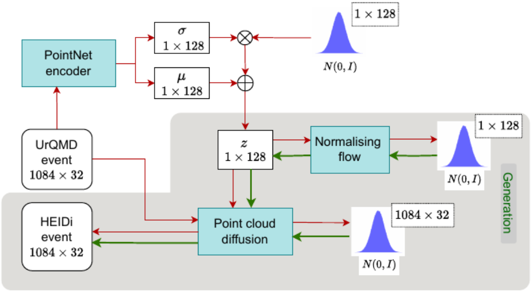

Figure: HEIDi: encoder + flow-based decoder + diffusion model.

The diffusion process

- **Forward process:** gradually add Gaussian noise to a real event over T steps (training only)
- **Reverse process:** start from pure noise, remove it step by step, guided by the context vector z (generation)

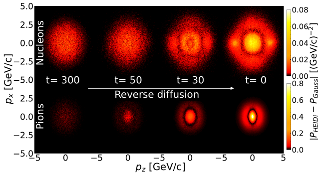

Figure: HEIDi transforms random Gaussian noise into realistic particle momenta and IDs.[2]

A neural network learns "how much noise was just added" at every step, and how to undo it.

[2] M. O. Kuttan et al. *Towards a foundation model for heavy-ion collision experiments through point cloud diffusion*. 2024

Training objective

Training maximizes the probability of reproducing real events, equivalent to minimizing:

$$\mathcal{L}(\theta,\phi) = \underbrace{\mathbb{E}_q\left[\sum_{t=2}^{T}\sum_{i=1}^{N} D_{\mathrm{KL}}\left(q(\mathbf{x}_i^{t-1}|\mathbf{x}_i^{t},\mathbf{x}_i^{0}) \parallel \tilde{q}_\theta(\mathbf{x}_i^{t-1}|\mathbf{x}_i^{t},\mathbf{z})\right)\right]}_{\mathrm{I}}$$

$$\underbrace{-\sum_{i=1}^{N}\log\tilde{q}_\theta(\mathbf{x}_i^{0}|\mathbf{x}_i^{1},\mathbf{z})}_{\mathrm{II}} + \underbrace{D_{\mathrm{KL}}\left(\tilde{q}_\phi(\mathbf{z}|\mathbf{X}^{0}) \parallel p(\mathbf{z})\right)}_{\mathrm{III}}$$

- **I** Does the reverse (denoising) step match the true noise removal, at every timestep and every particle?
- **II** Can the model reconstruct the exact final particle from the last noisy step?
- **III** Does the learned context vector z stay close to a simple prior?

Results: multiplicity distributions

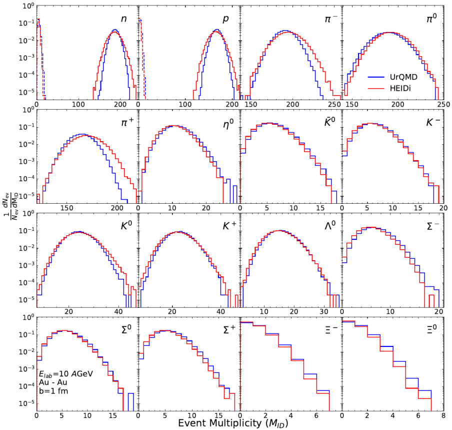

Figure: Multiplicity distribution of various hadron species for Au-Au collisions with $b = 1$ fm at 10 AGeV.[4]

[4] M. O. Kuttan et al. *Ultra fast, event-by-event heavy-ion simulations for next generation experiments.* 2025

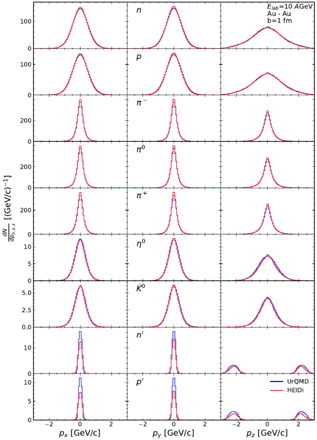

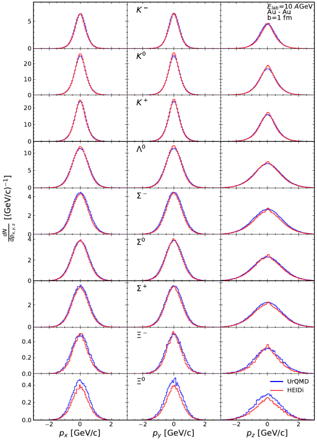

Momentum distributions of various hadron types generated by HEIDi for Au-Au collisions with $b = 1\,\mathrm{fm}$ at $10\,\mathrm{AGeV}$.[5]

[5] M. O. Kuttan et al. *Ultra fast, event-by-event heavy-ion simulations for next generation experiments.* 2025.

Application: cosmic-ray air showers

- Air showers: a chain reaction of secondary collisions
- UrQMD simulates a single, elementary collision. Shower frameworks like SENECA call it millions of times per shower
- $\Rightarrow$ speeding up this elementary step speeds up the entire shower simulation
- We use p–O collisions as this elementary building block

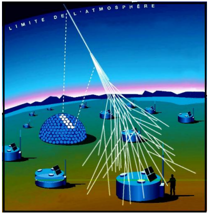

Figure: Pierre Auger Observatory: detecting air showers via water tanks and fluorescence telescopes.

Results: multiplicity distributions

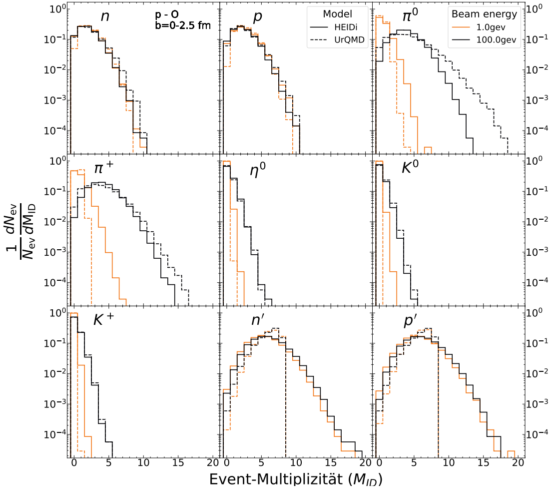

Figure: Multiplicity distribution of various hadron species for p-O collisions with $b = 0 - 2.5\,\mathrm{fm}$ at 1 and 100 GeV.

Results: momentum distributions

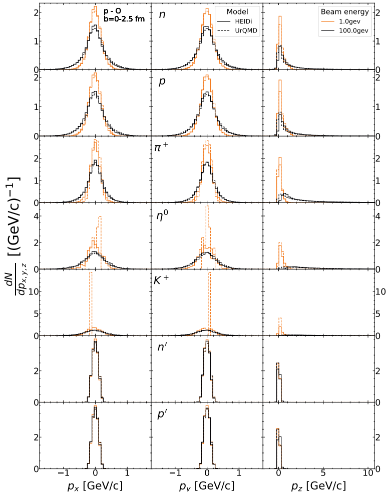

Figure: Momentum distributions of various hadron types generated by HEIDi for p-O collisions with $b = 0 - 2.5\,\text{fm}$ at 1 and 100 AGeV.

Speed: orders of magnitude faster

- UrQMD (cascade mode): $\approx 3\,\mathrm{s}$ per event
- HEIDi (after hyperparameter optimization, on an Nvidia A100 GPU): $\approx 1.6\,\mathrm{ms}$ per event

Speedup

$$\frac{3000\,\mathrm{ms}}{1.6\,\mathrm{ms}} \approx 1875 \times \text{ faster}$$

Conclusion

- HEIDi reproduces multiplicity and momentum distributions of UrQMD across many hadron species
- $\sim$1000$\times$ faster than UrQMD, while preserving intra-event correlations
- HEIDi is highly flexible and adaptable, as demonstrated by its application to cosmic-ray air showers