Deterministic AI Surrogate Modeling for Fast Hydrodynamic Evolution of the Quark Gluon Plasma

Yeonju Go (BNL), S. Lee (BNL), B. Schenke (BNL), M. Chamizo-Llatas (BNL), J. Huang (BNL), Y. Ren (BNL), C. Shen (WSU)

Strangeness in Quark Matter (SQM2026)
UCLA, Los Angeles , CA, USA
Mar. 22 - 27, 2026

U.S. DEPARTMENT OF ENERGY

Brookhaven National Laboratory

SQM 2026

Heavy Ion Collision Simulation

arXiv: 2303.17254

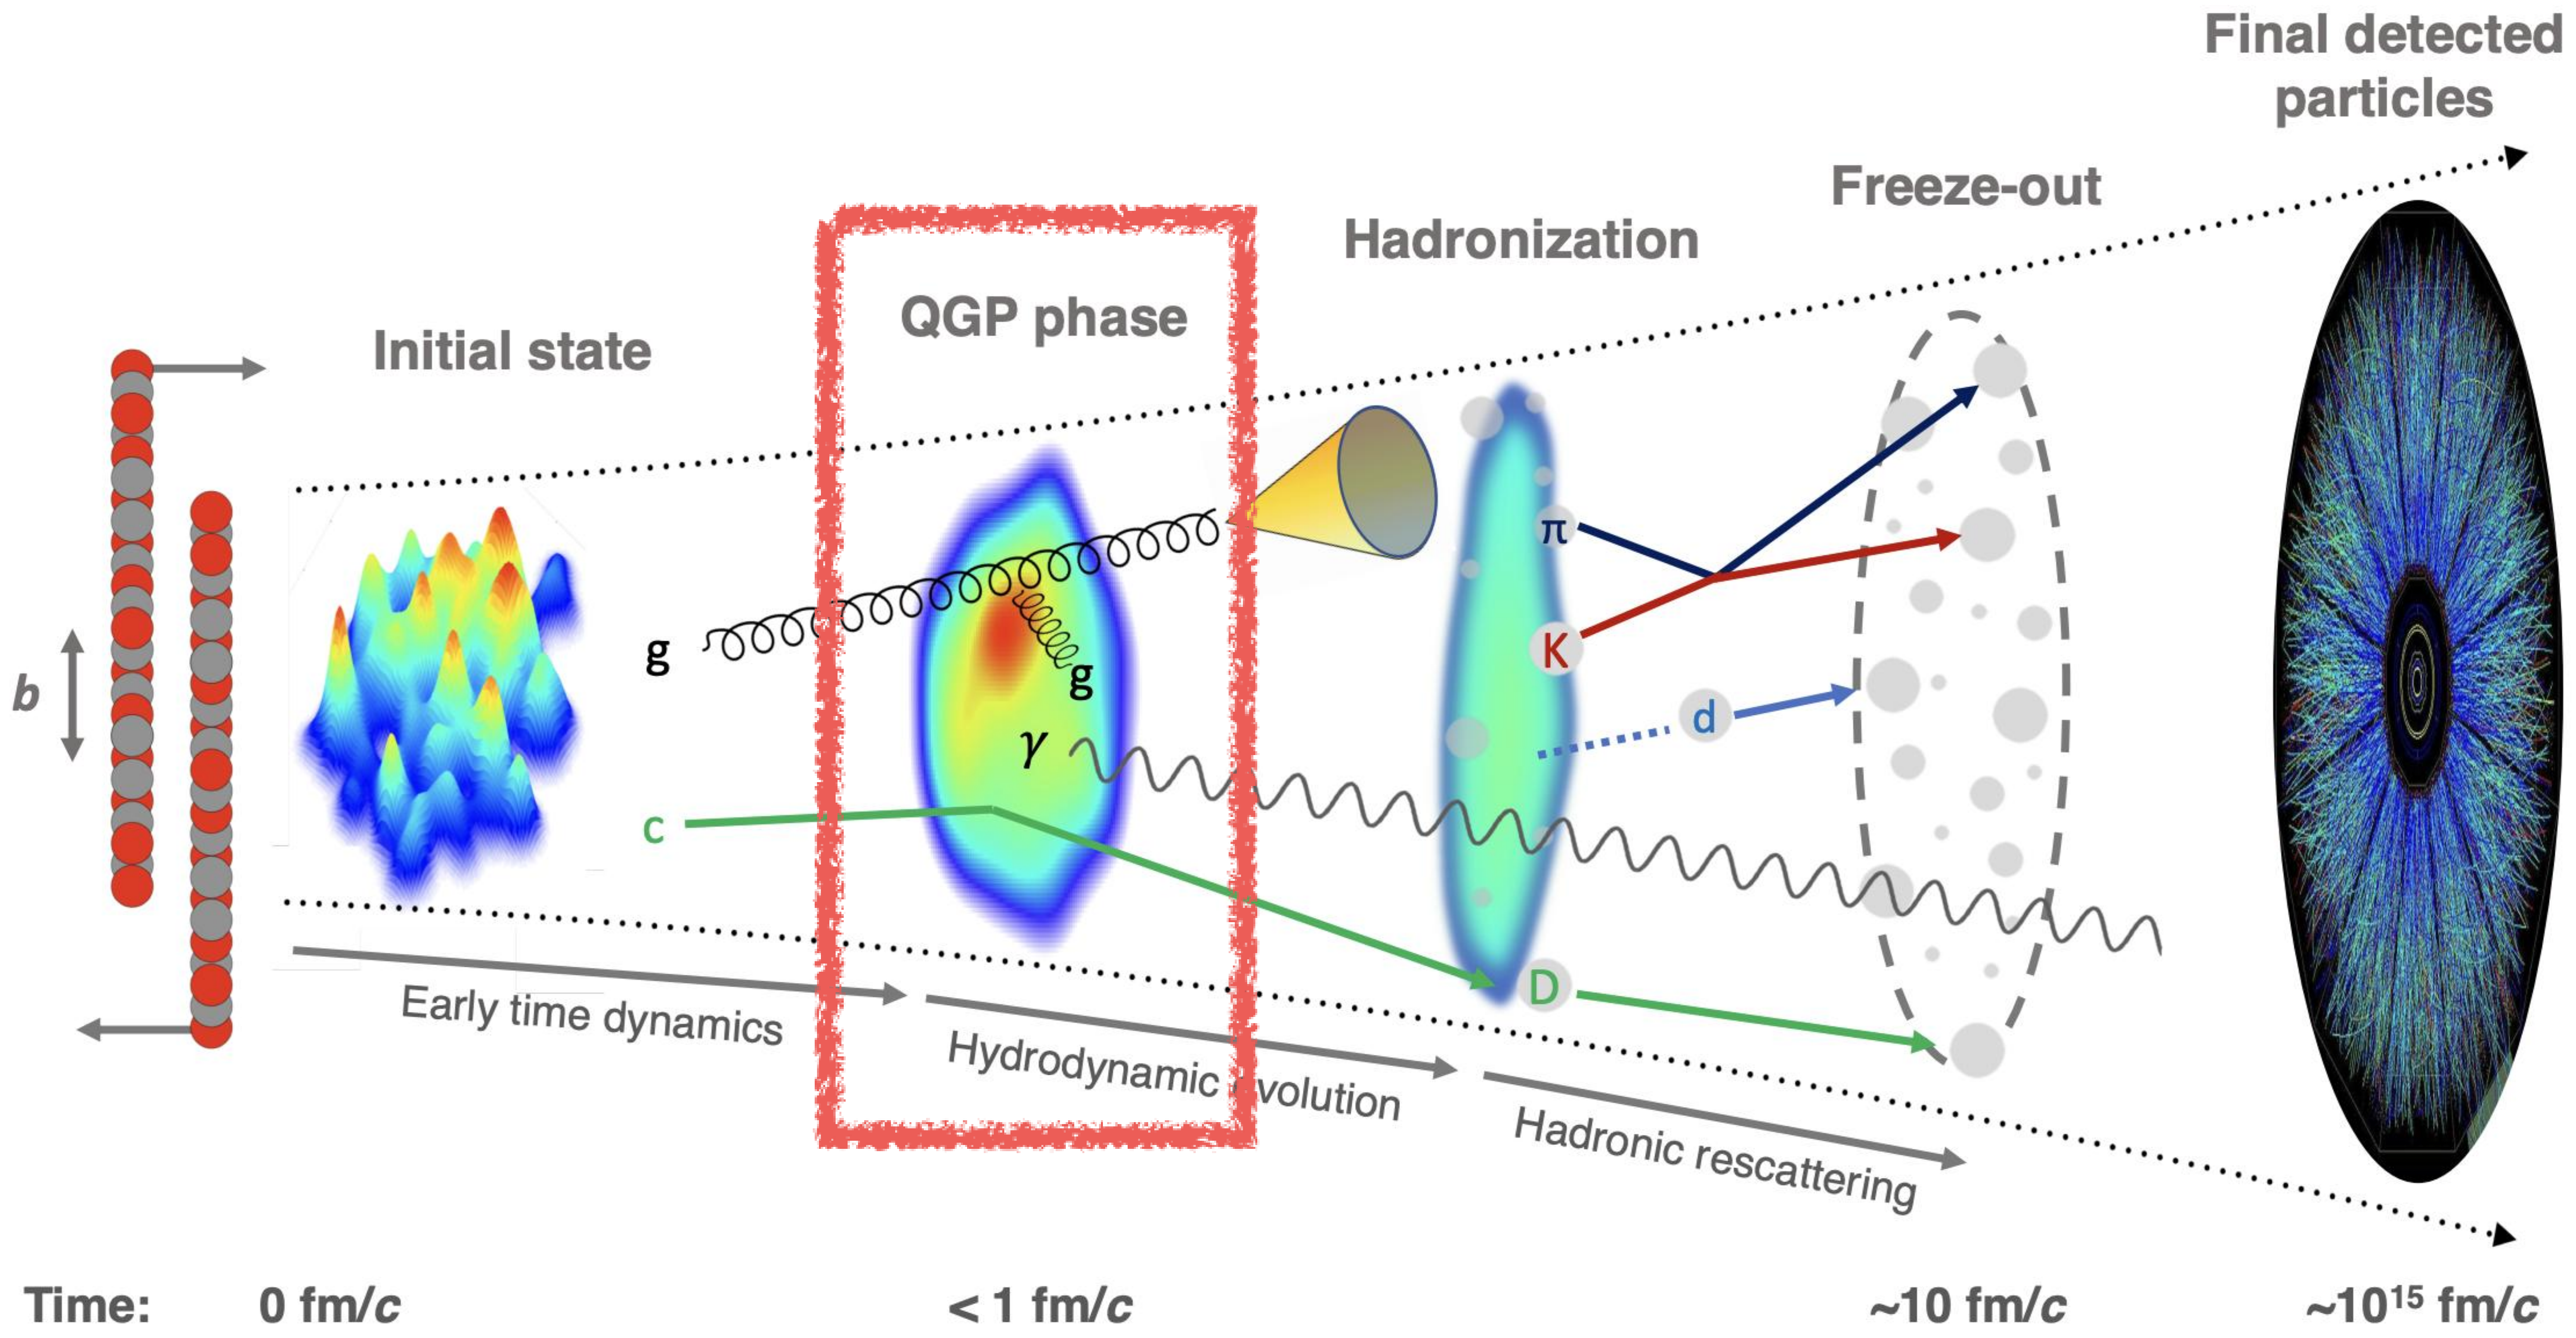

- Heavy ion collision simulations **connect QGP properties** to experimentally measured **final-state observables**
- A large number of simulated events is required for parameter scans, uncertainty quantification, and event-by-event inference

Hydrodynamic Medium Evolution

- Bayesian inference of transport coefficients (e.g. η/s and ζ/s) requires millions of events
- Event-by-event *hydrodynamic medium (QGP) evolution simulations* are computationally very demanding

	2+1D	3+1D
Simulation time per event with 1 CPU	~10 minutes	~20 hours

Quark Gluon Plasma 3+1D Energy Density Evolution *IP-Glasma+MUSIC*

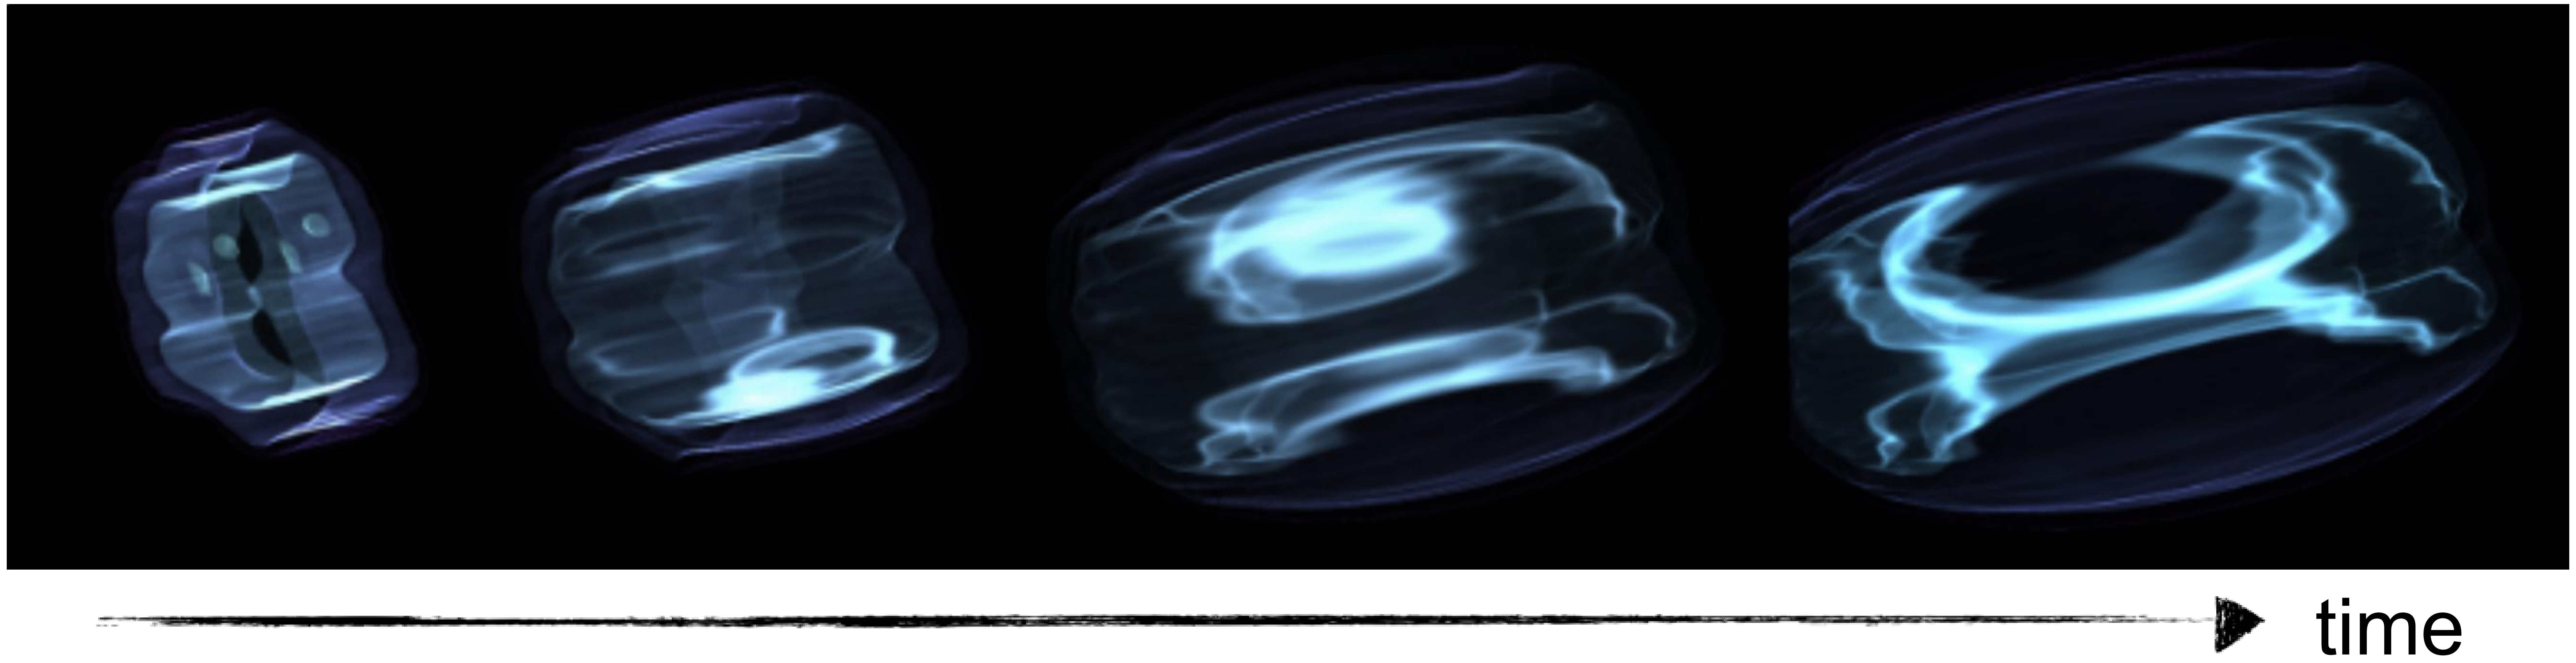

→ ***Use AI Surrogate Model for Fast Hydrodynamic Evolution***

What Machine Learning Model Should We Use?

- Generative Adversarial Network ?
- Diffusion Model ?

  → demonstrated to be good for fast heavy ion event generation with calorimeter detector simulation *PRC 110 (2024) 034912*
- …

Training sample (HIJING+GEANT4)

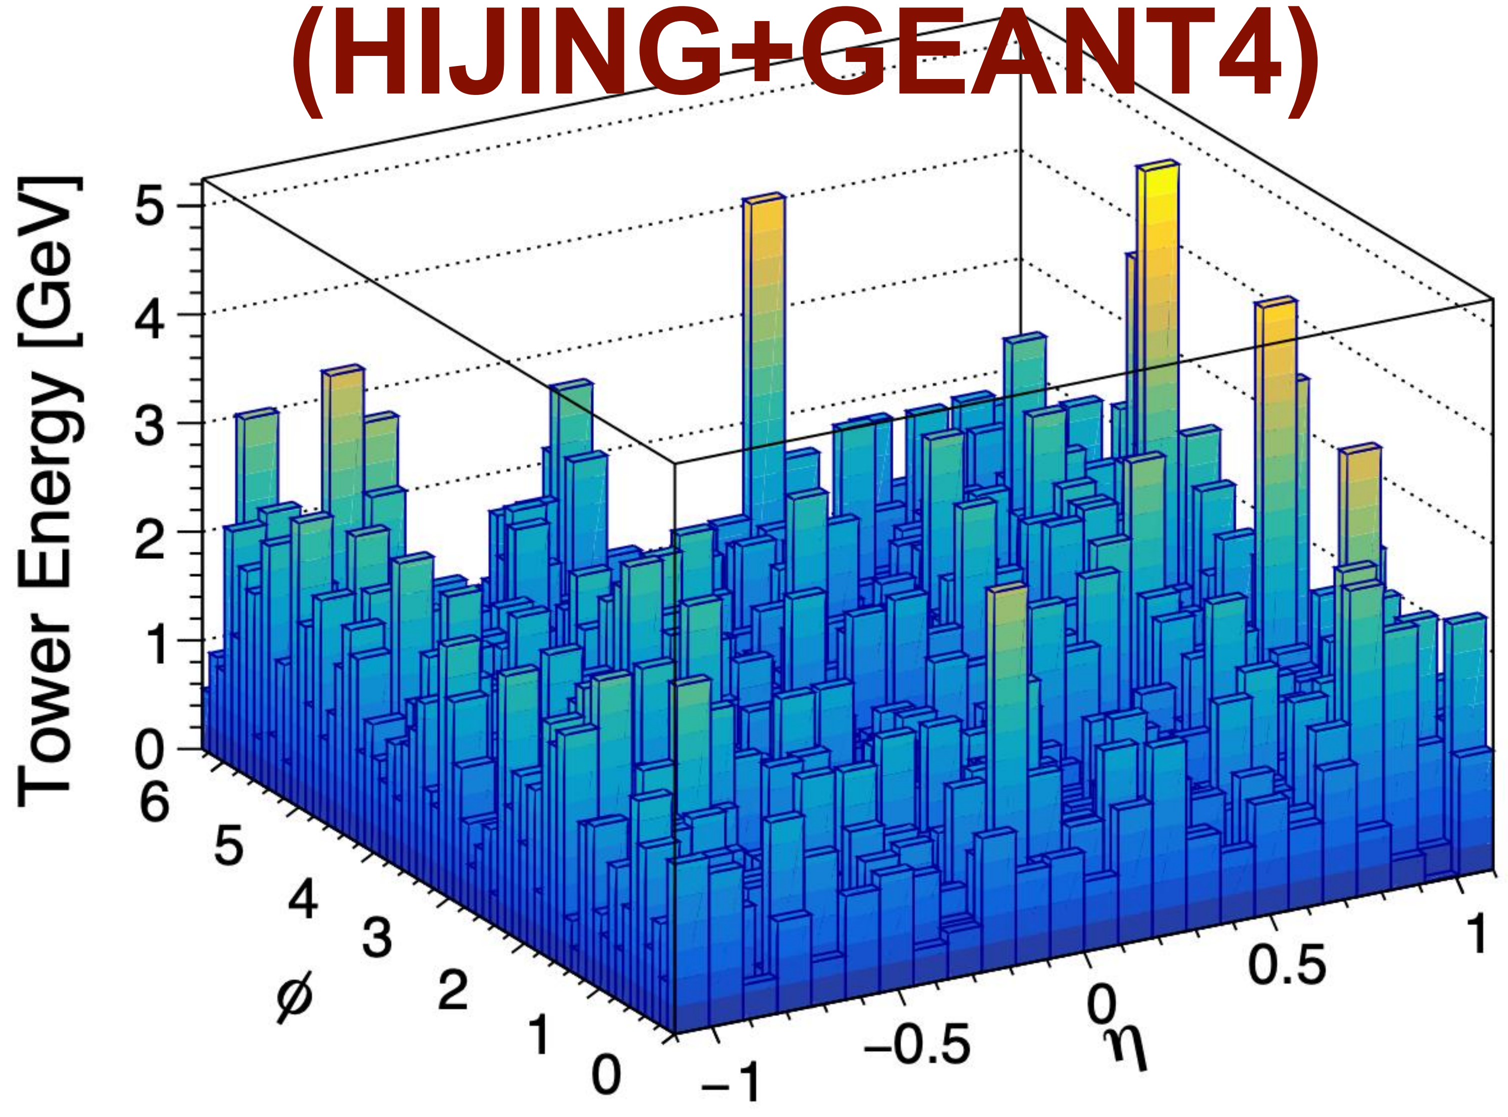

40 minutes per event *on CPU*

Generated (Denoising Diffusion Probabilistic Model)

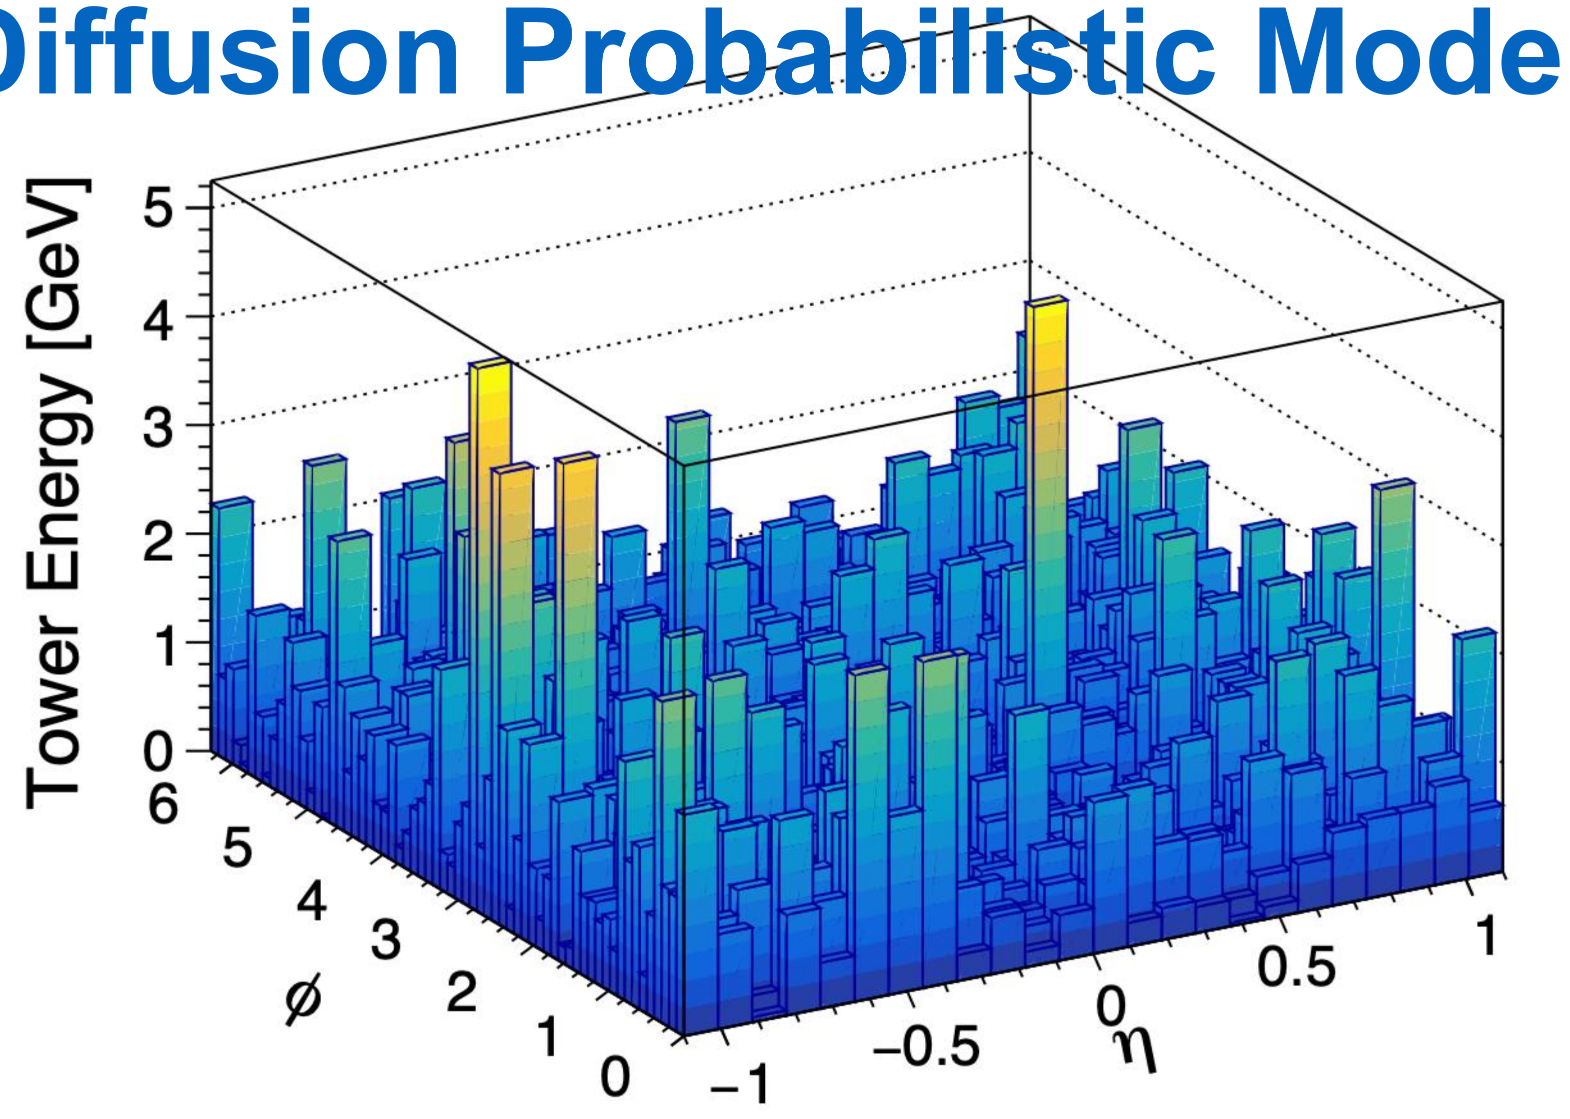

1.34 seconds per event *on GPU*

See Shuhang Li's talk on Tuesday

Fourier Neural Operator (FNO)

arXiv:2010.08895 [cs.LG]

- **FNO** learns mappings between **functions (operators)** instead of fixed inputs/outputs
  → ideal for Partial Differential Equations (PDE)-governed evolution
- Uses **Fourier Transform** (FFT) to model global patterns efficiently in frequency space
- Well matched to smooth but structured hydrodynamic fields across time

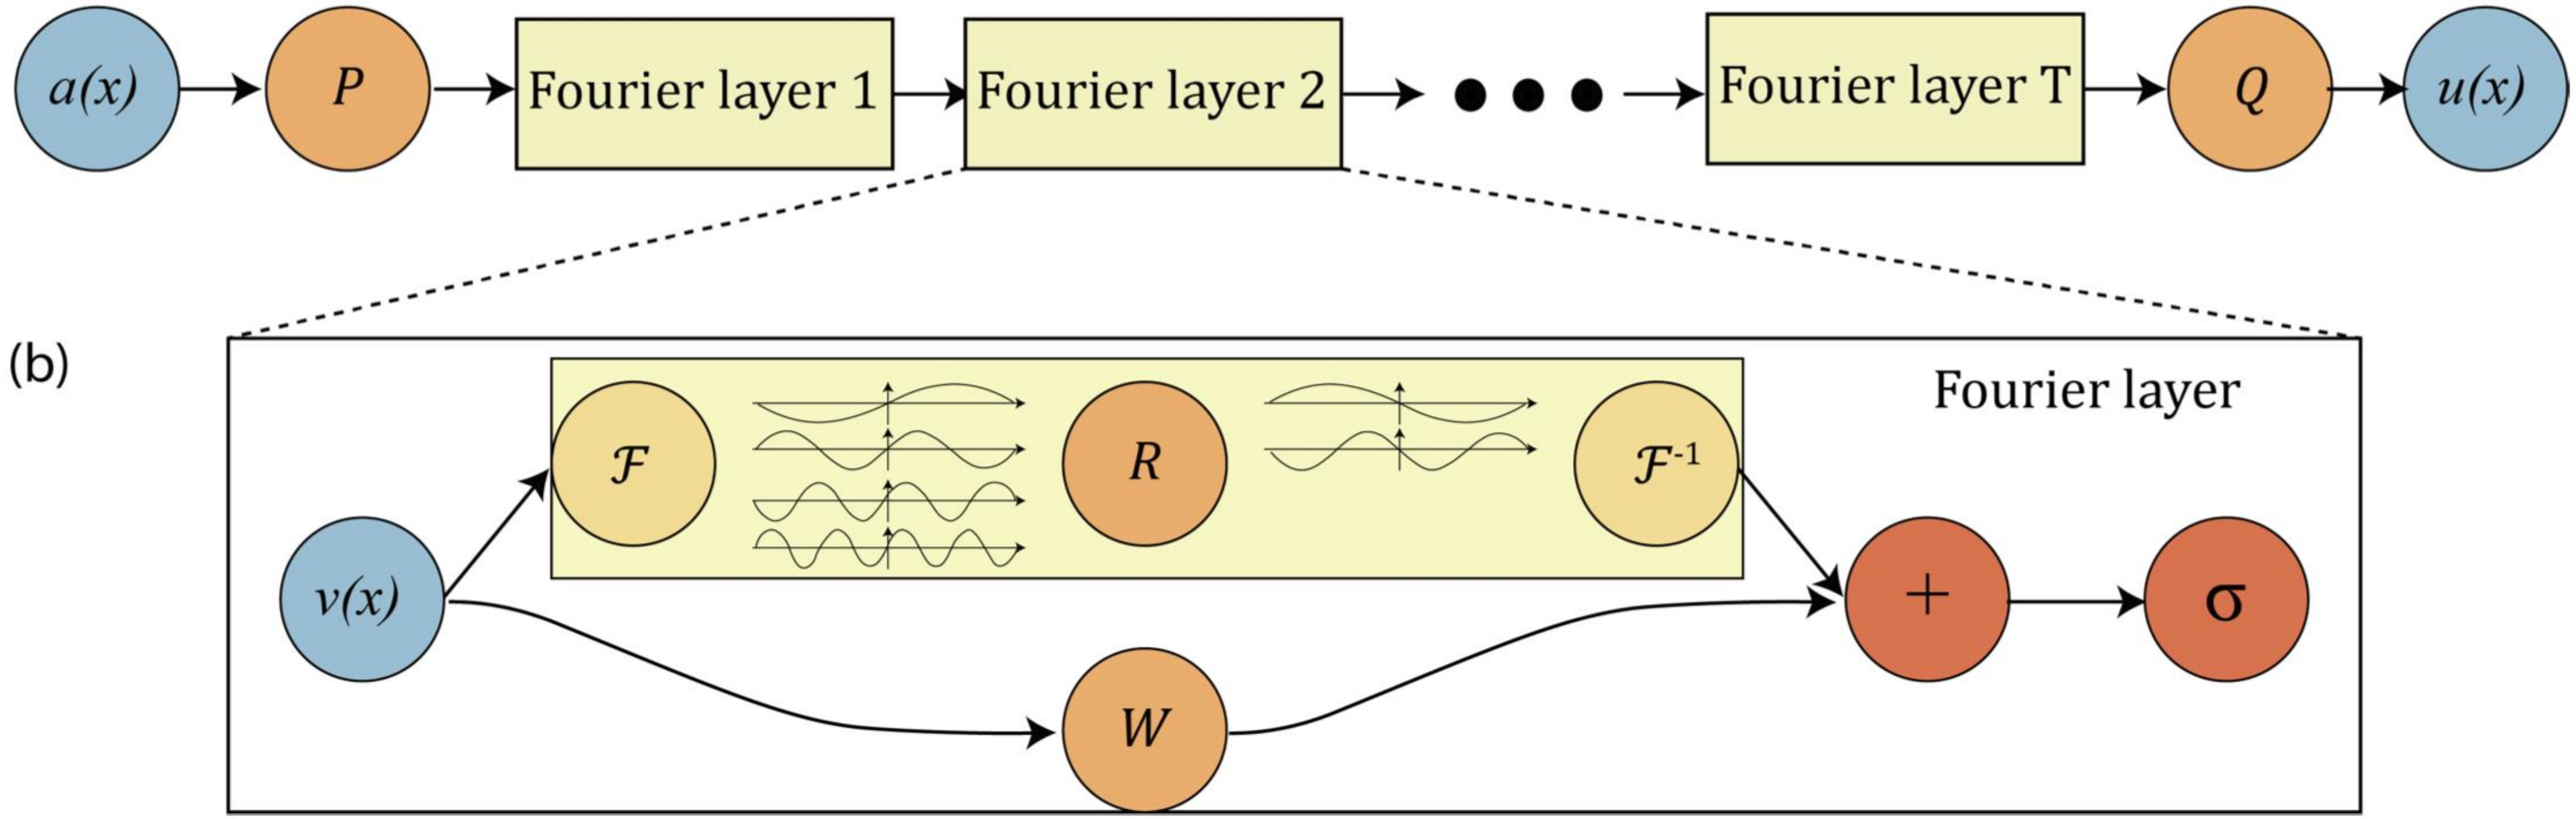

Fourier Neural Operator (FNO)

arXiv:2010.08895 [cs.LG]

- **FNO** learns mappings between **functions (operators)** instead of fixed inputs/outputs
  → ideal for Partial Differential Equations (PDE)-governed evolution
- Uses **Fourier Transform** (FFT) to model global patterns efficiently in frequency space
- Well matched to smooth but structured hydrodynamic fields across time
- up to **~1000× faster** than traditional PDE solvers
- **Resolution invariant**: it learns the underlying operator on the continuous function, not just a mapping between values on one fixed grid

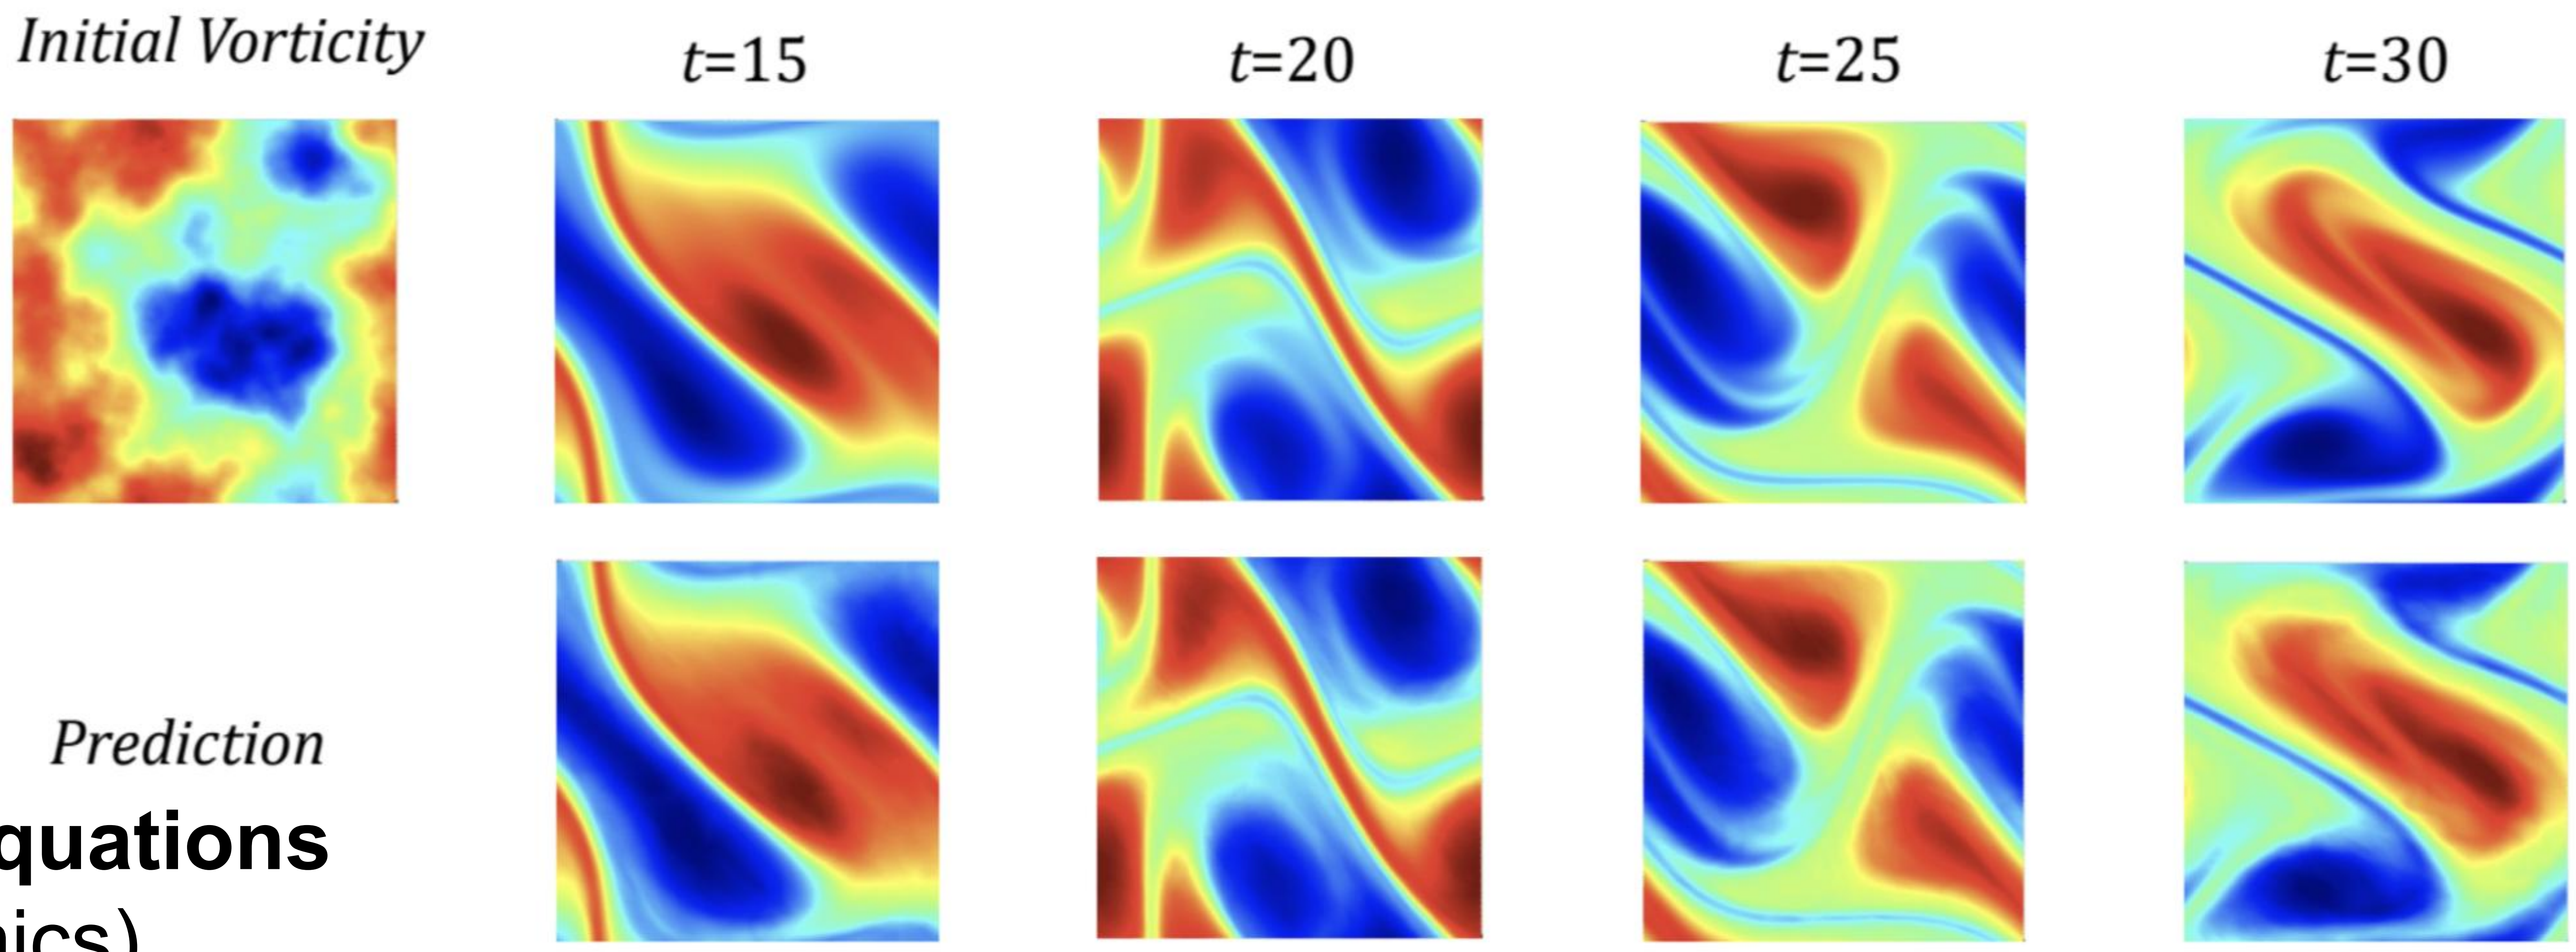

Navier Stokes equations
(fluid dynamics)

Previous Work

Slide added after presentation thanks to comments from J. Putschke et al.

- Fast developing field: stacked U-net accelerated 2+1D simulation in 2021 (*H. Huang et. al. Phys. Rev. Research* 3, 023256)
- Previous work (*D. Stewart, J. Putschke, Phys.Rev.C 113 (2026) 1, 014904*) has shown effectiveness of FNO in fast prediction of hydrodynamic QGP evolution
  - ➞ predicted energy density and velocity for different time step in 2+1D, experimental observables for flow and jet quenching

Stewart, Putschke, *Phys.Rev.C 113 (2026) 1, 014904*

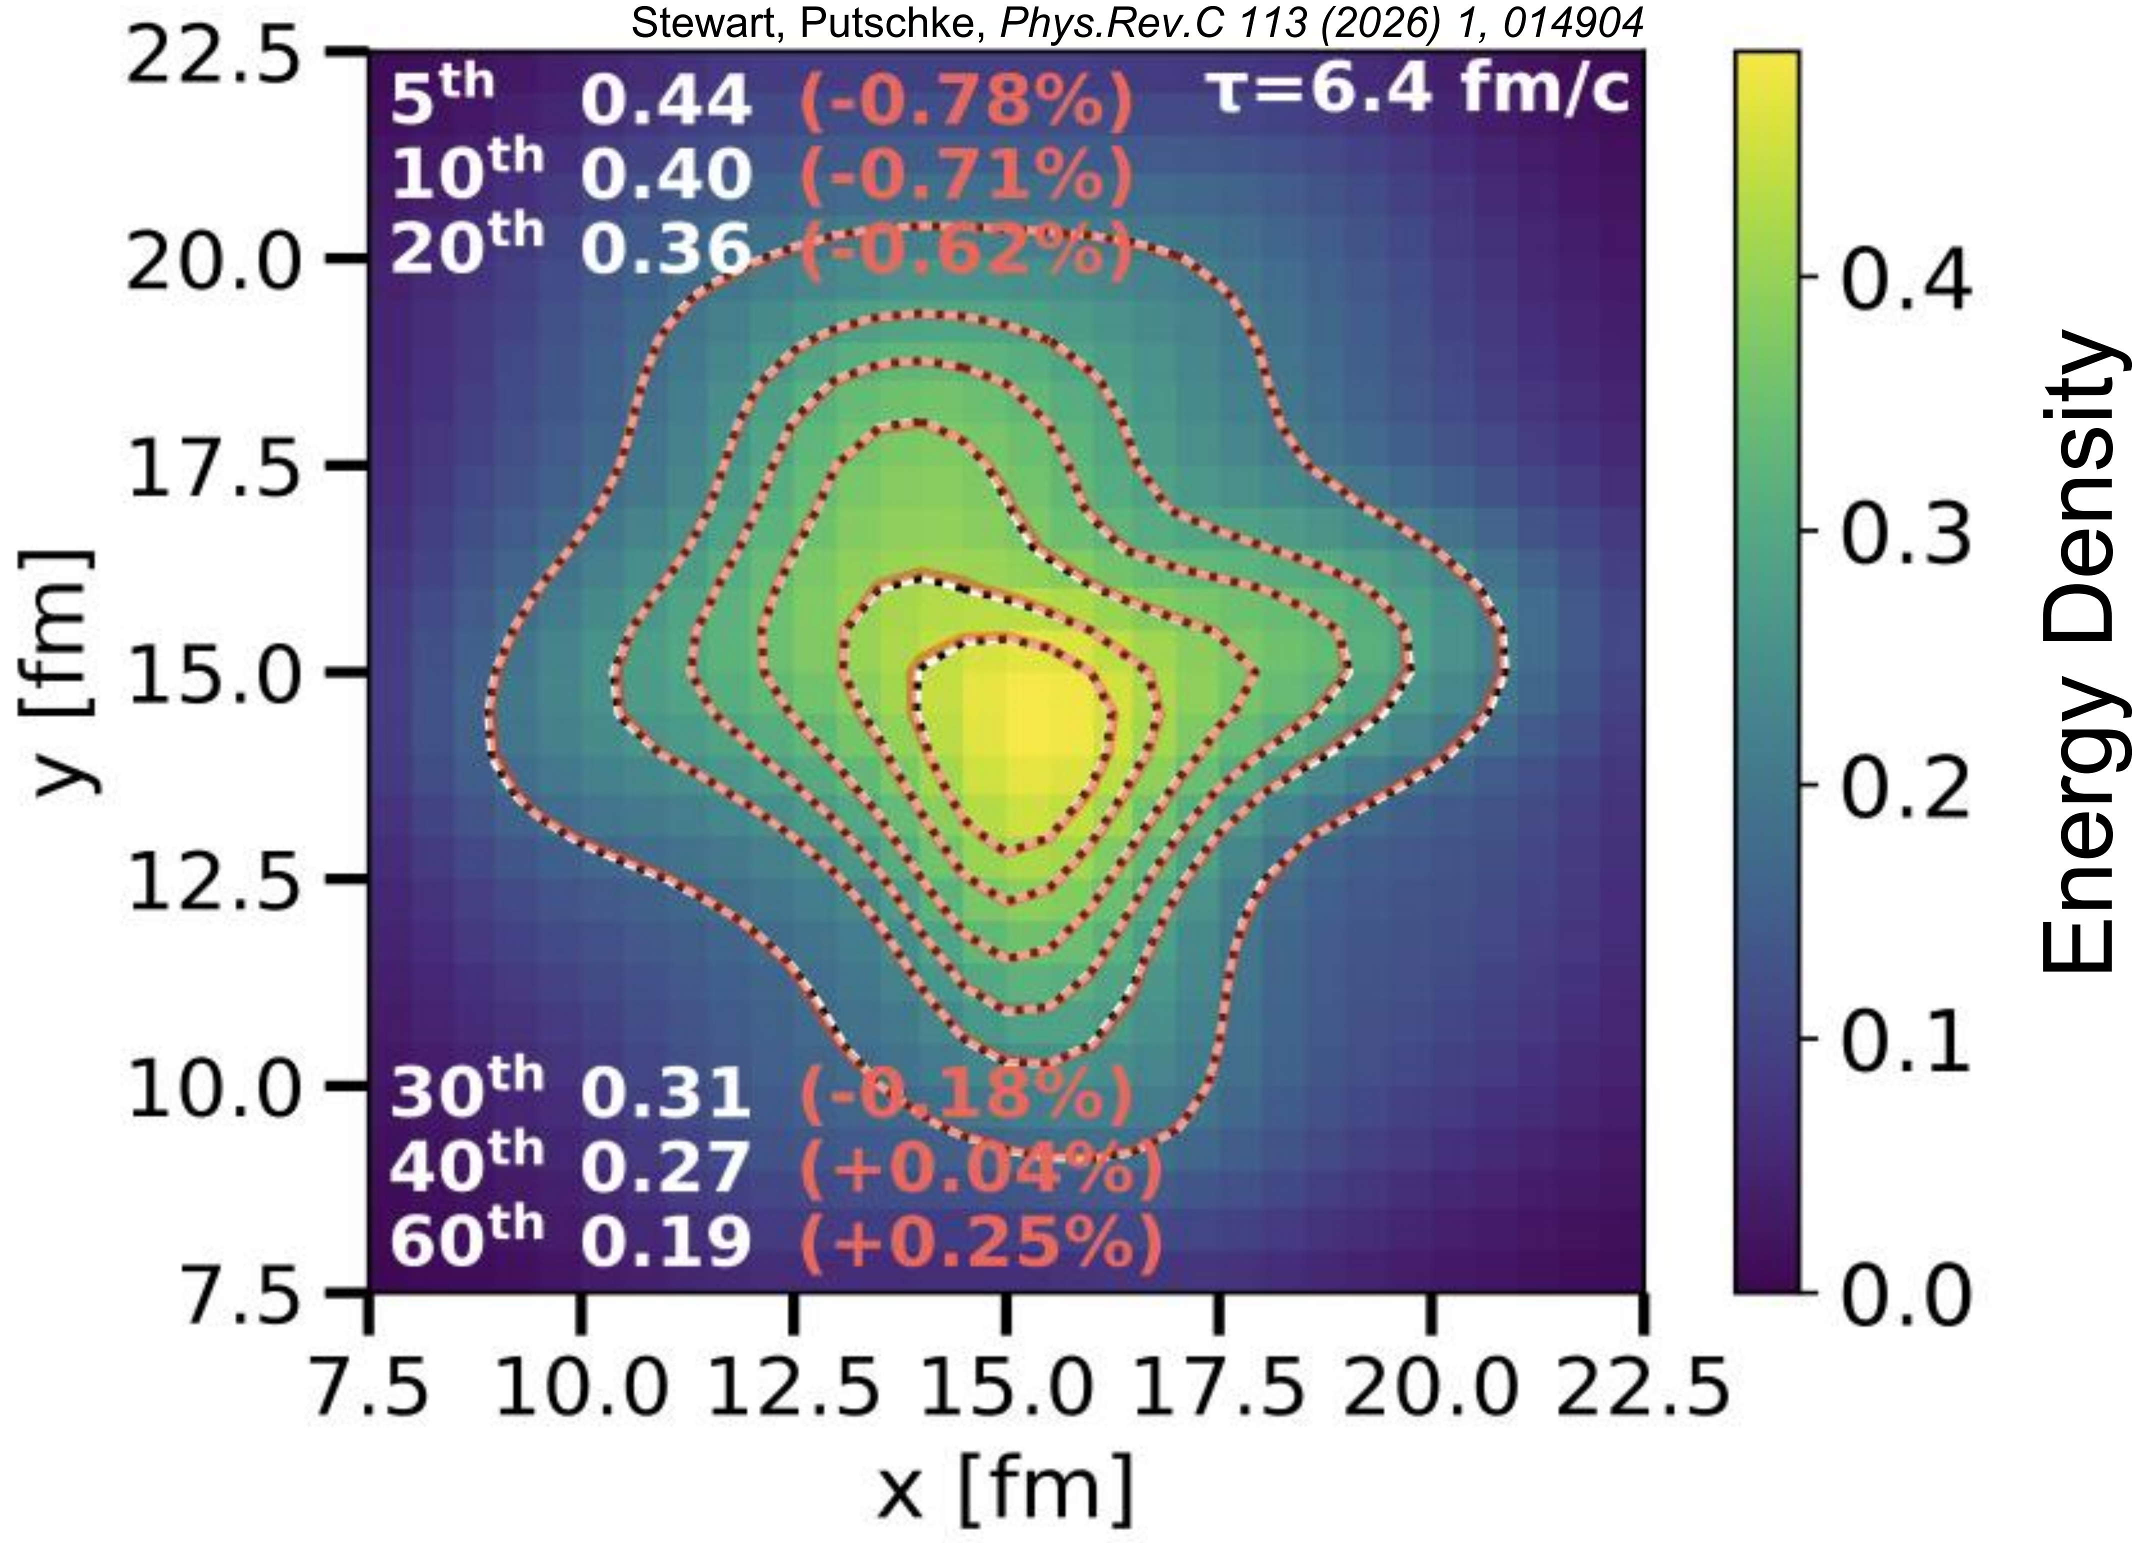

Stewart, Putschke, *Phys.Rev.C 113 (2026) 1, 014904*

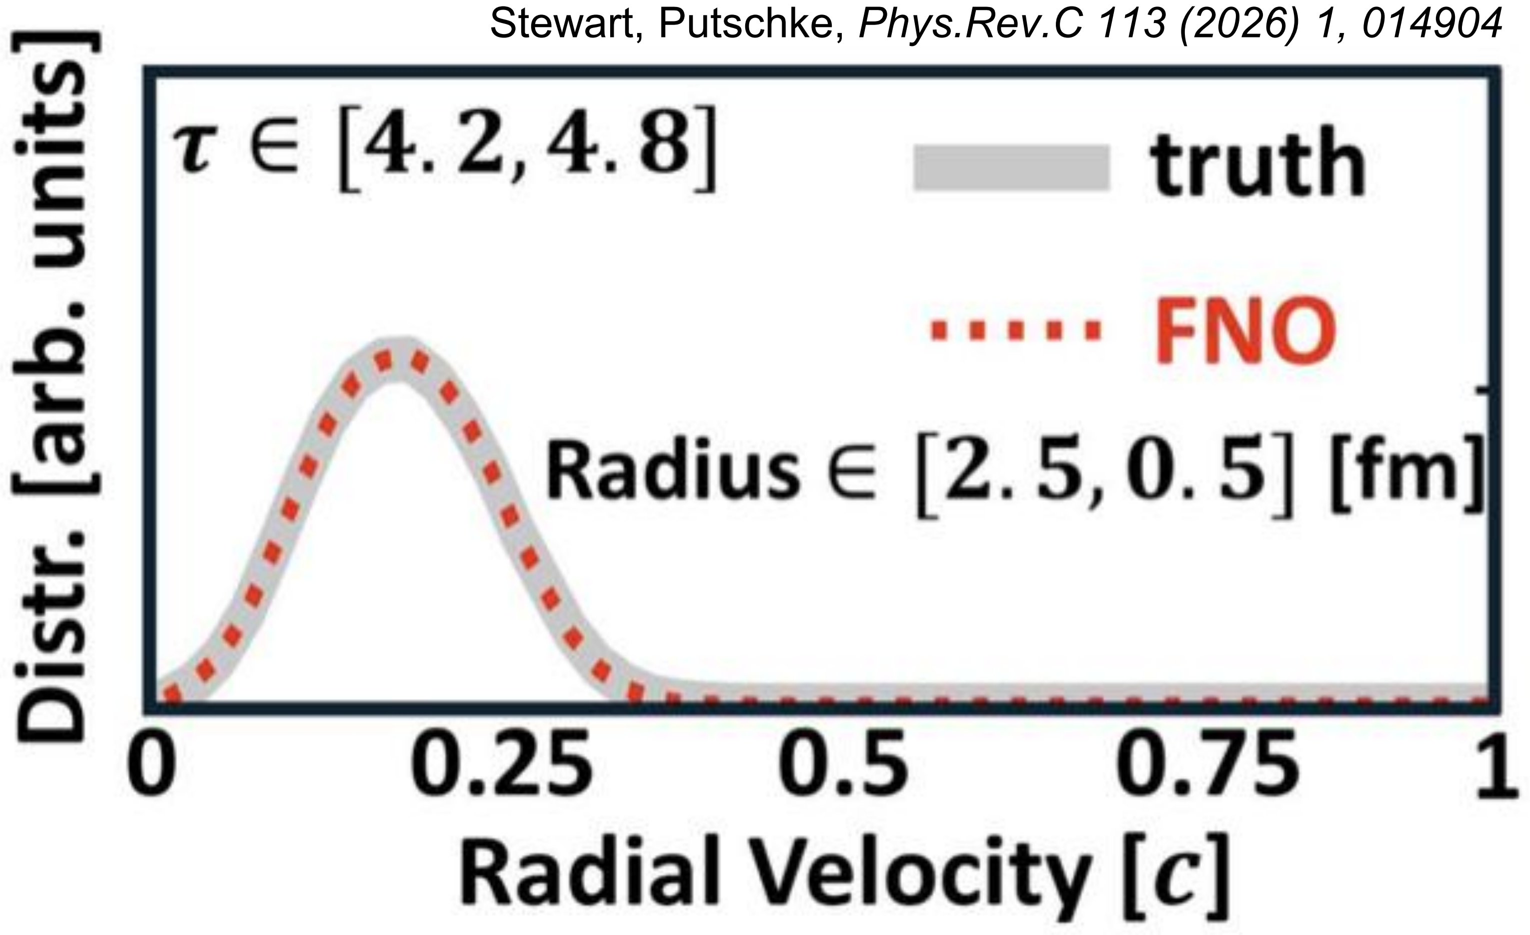

Our work extends this approach to additional hydrodynamic variables, including also temperature and pressure, aiming for the next-stage goal of full 3+1D evolution

Training FNO

- **Data**
  - → **IP-Glasma** + **MUSIC** for Au+Au collisions at $\sqrt{s_{\mathrm{NN}}}$=200 GeV
  - → 0-40% centrality trained in a single model
  - → Dimensions: **2+1D (x, y, time)**
  - → Total 612 events: 582 events for training, 30 events testing

- **Training Object**
  - → 5 fields (**energy density** ϵ, **Temperature** T, **Pressure** P, **Flow velocity** u_x, u_y)
  - → 60 time steps for **0.4 fm/c - 10 fm/c** ($\Delta\tau$ = 0.16 fm/c)
  - → No explicit normalization is applied. Instead, ϵ, T **and** P **are scaled with time** during training, which plays a role similar to normalization.
  - → **Relative loss** with Sobolev and L2 norms

Training Time and Inference Time

Initial condition (input)

FNO Prediction (output)

Initial condition (input)		FNO Prediction (output)			
ϵ, T, P, u_x, u_y at τ = 0 fm/c	→	ϵ, T, P, u_x, u_y at τ = 0.4 fm/c	ϵ, T, P, u_x, u_y at τ = 0.56 fm/c	. . .	ϵ, T, P, u_x, u_y at τ = 10.0 fm/c

	2+1D	3+1D
Simulation time per event with 1 CPU	~10 minutes	~20 hours
FNO prediction per event with 1 GPU	< 1 second	Target: < 10 seconds
Speed up	100-1000	TBD

- Training Epoch: 1000
- Training time: 15 hours for 1 GPU
- Memory usage: 5GB per inference

IP-Glasma+MUSIC Hydro 2+1D Evolution

Result: Energy Density

[GeV/fm^3]

	Prediction	Ground truth	Absolute Error	Relative Error
τ = **0.4** fm/c				
τ = **2.0** fm/c				
τ = **5.2** fm/c				

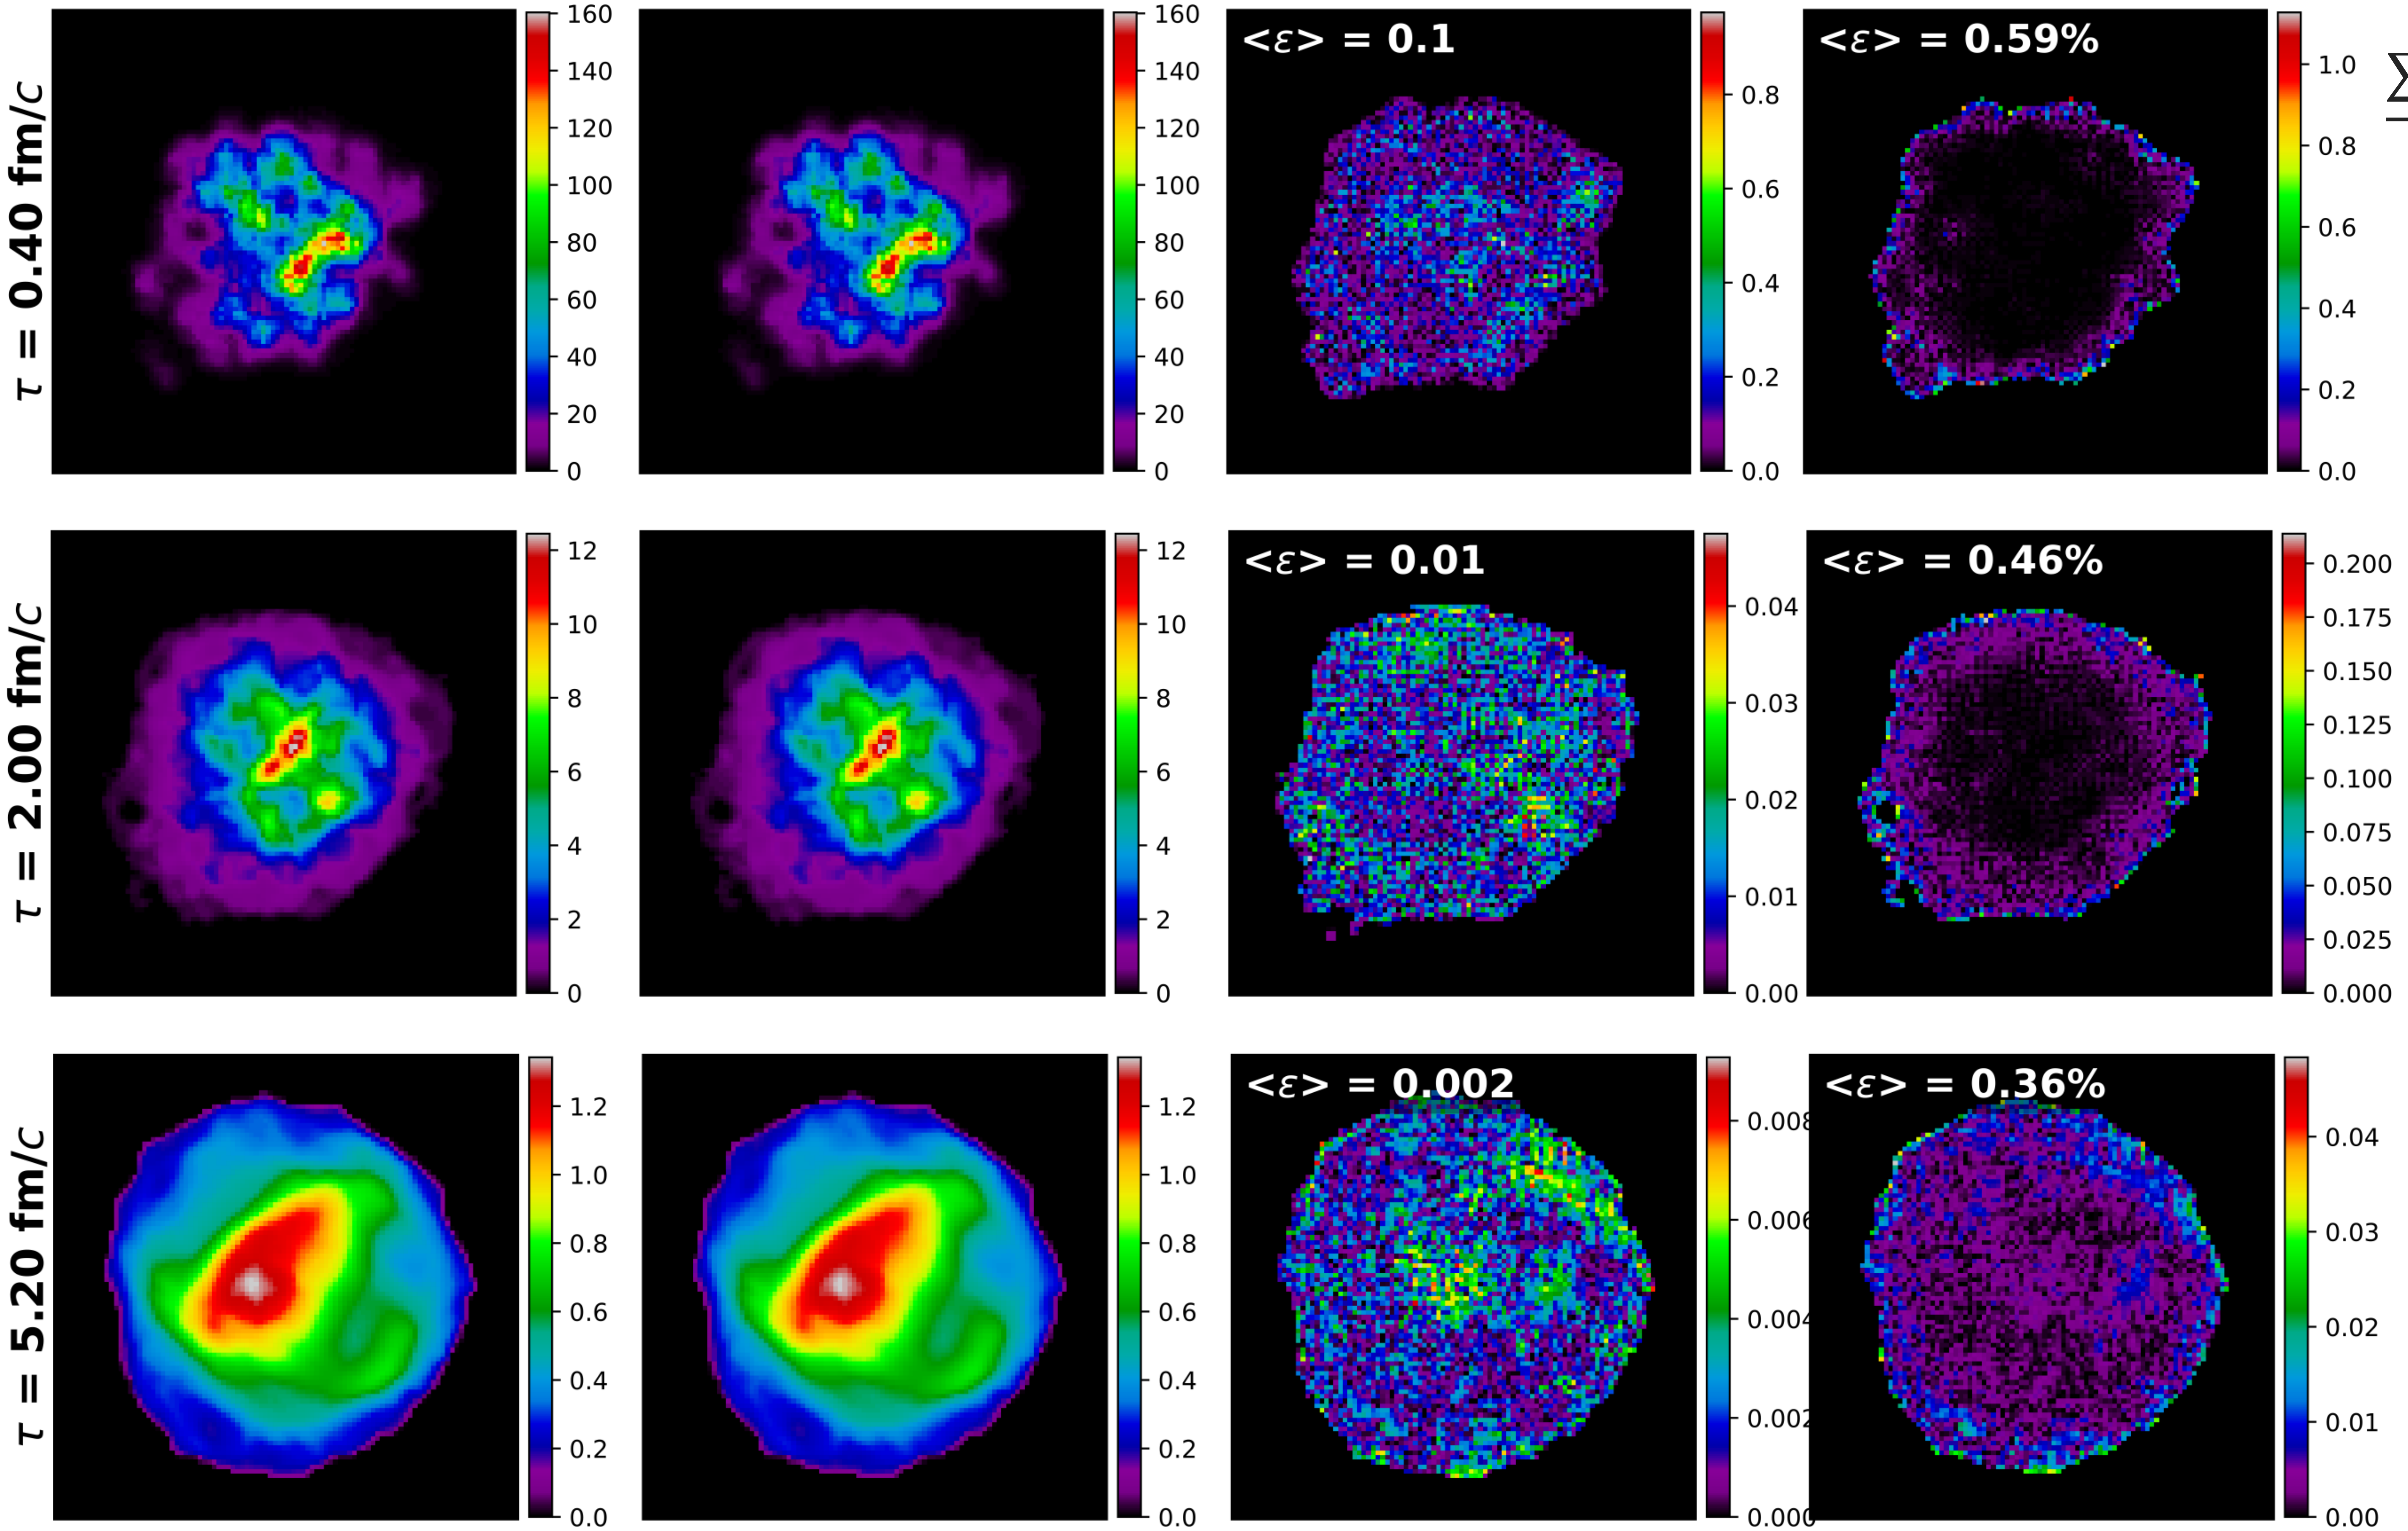

Mean Relative Error $<\epsilon> \equiv$

$$\frac{\sum_{x,y}|\mathrm{pred.}(x,y) - \mathrm{truth}(x,y)|}{\sum_{x,y}|\mathrm{truth}(x,y)|}$$

- Energy density prediction: relative error < **0.6** % in all time steps

Result: Pressure

[GeV/fm^3]

Prediction | **Ground truth** | **Absolute Error** | **Relative Error**

τ = **0.4** fm/c

τ = **2.0** fm/c

τ = **5.2** fm/c

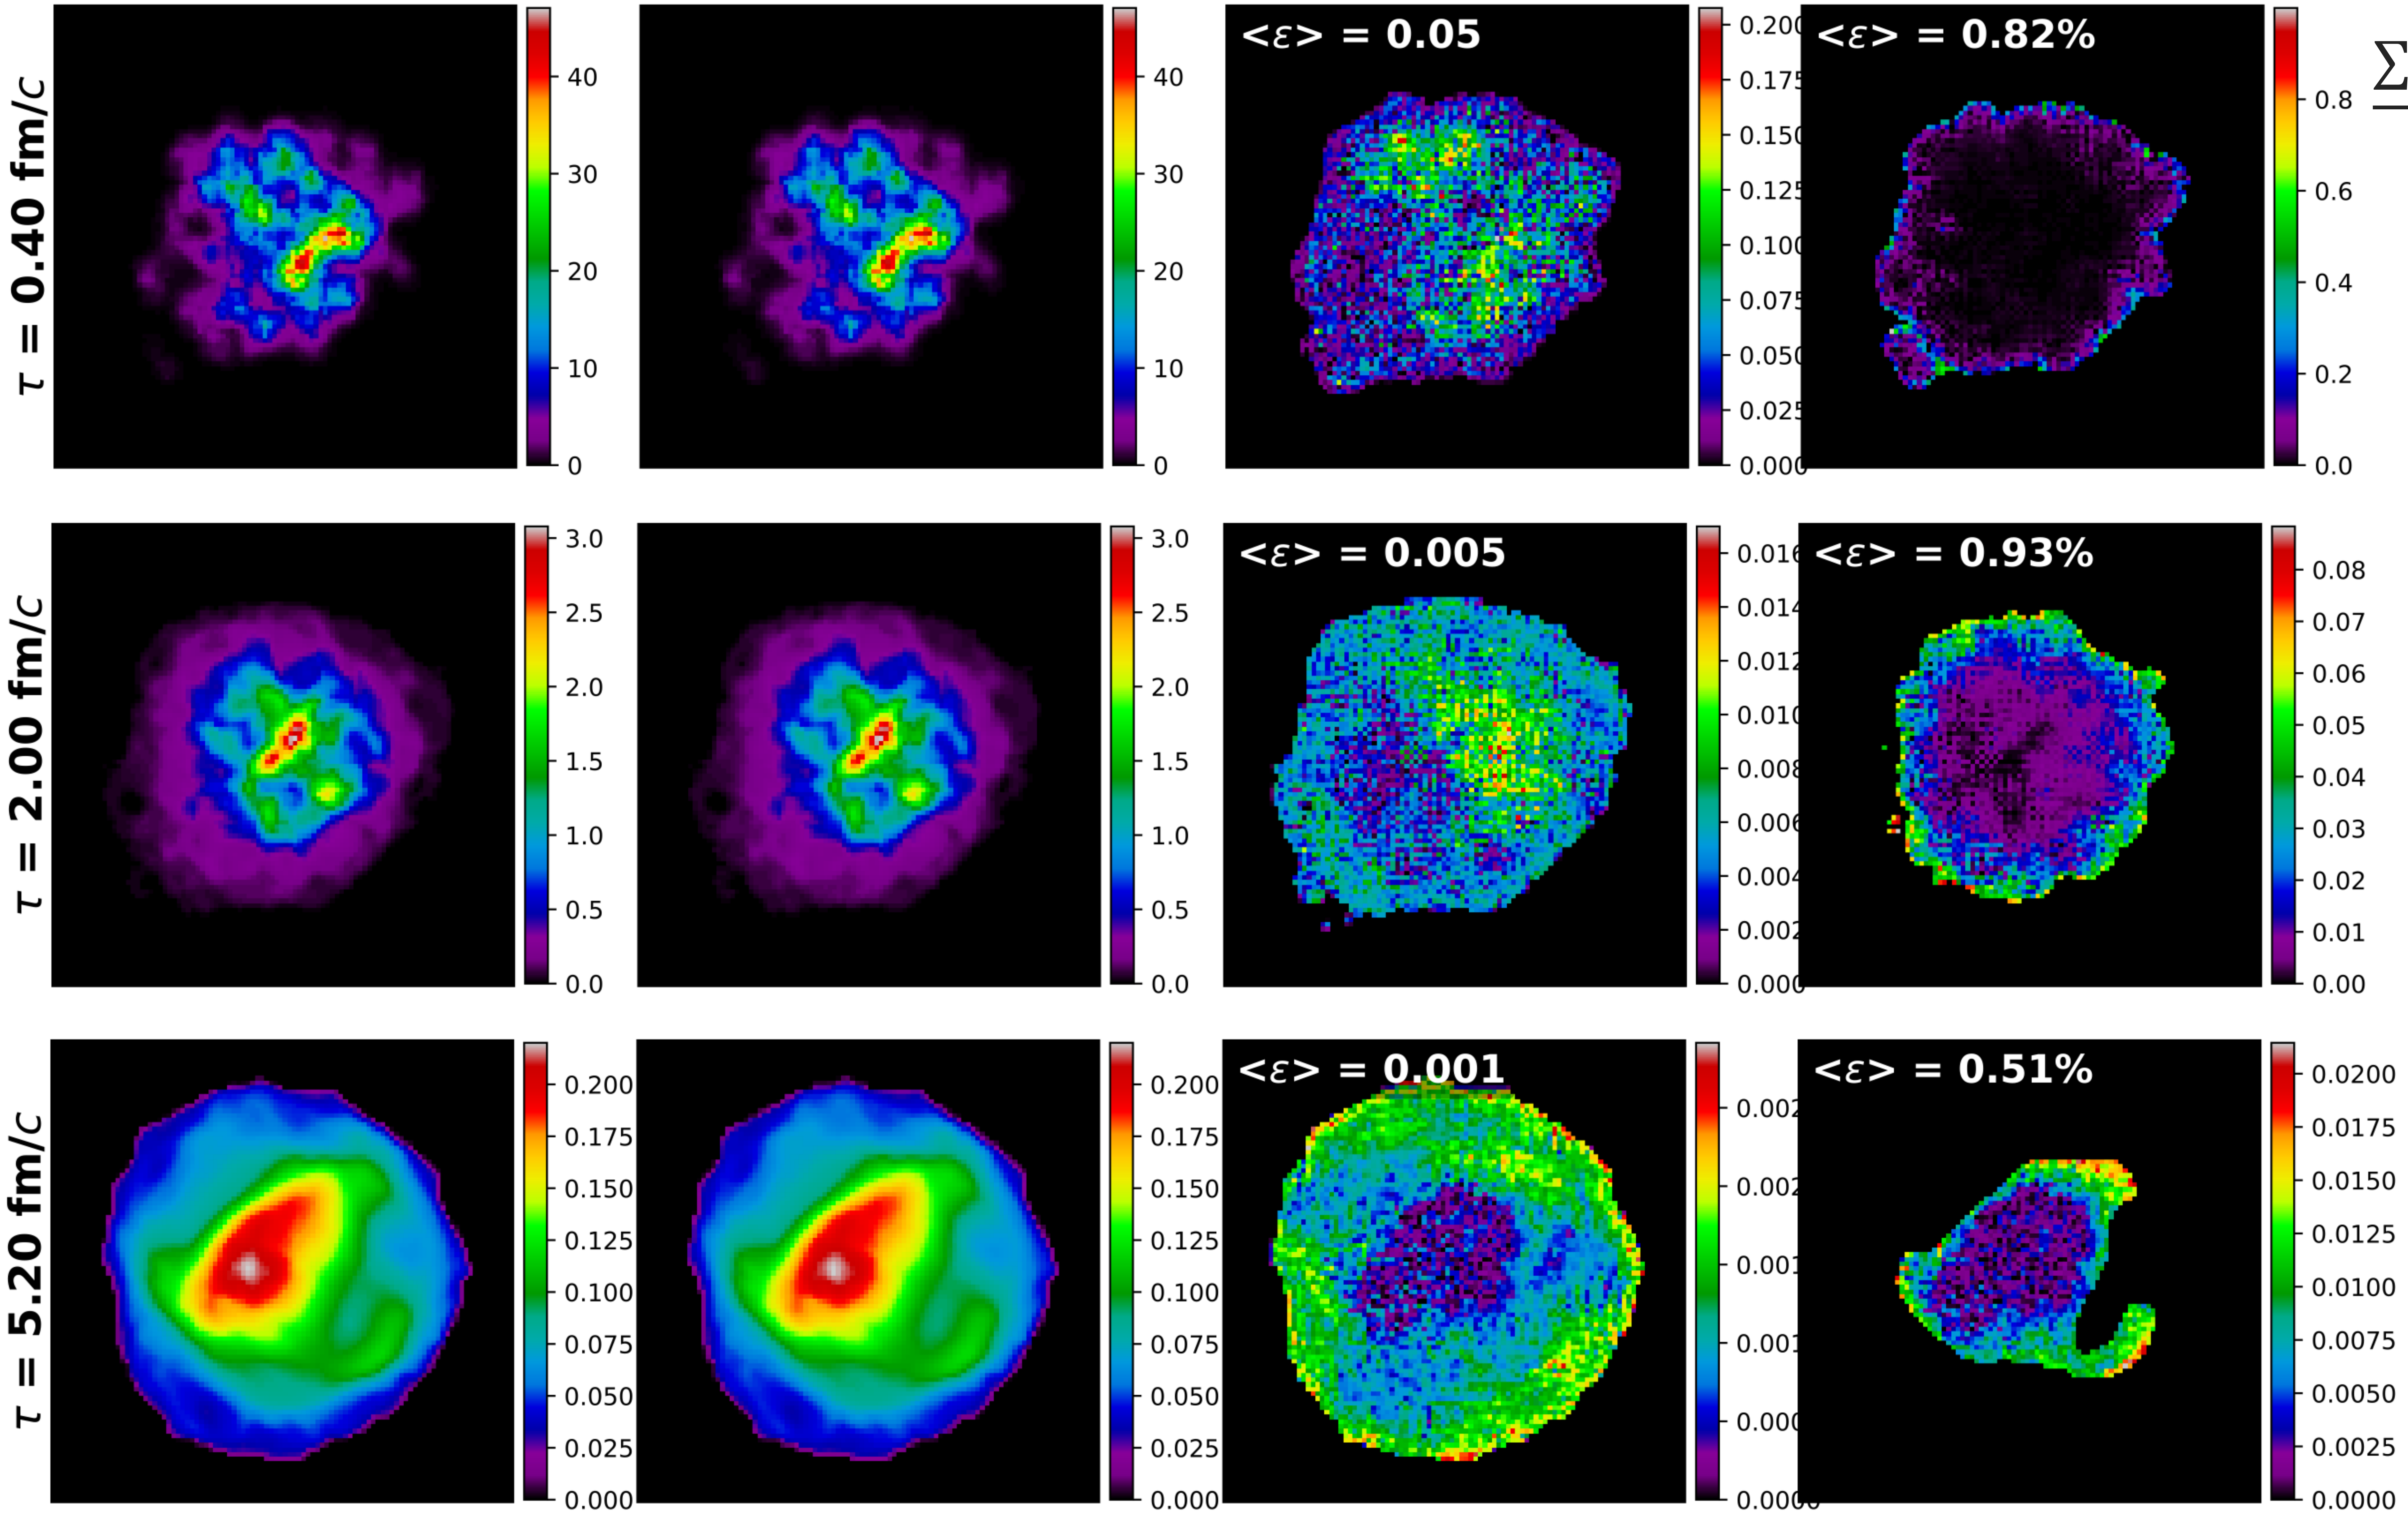

Mean Relative Error $<\epsilon> \equiv$

$$\frac{\sum_{x,y}|\mathrm{pred.}(x,y) - \mathrm{truth}(x,y)|}{\sum_{x,y}|\mathrm{truth}(x,y)|}$$

- Pressure prediction: relative error **< 1.0 %** in all time steps

Result: Temperature

[GeV]

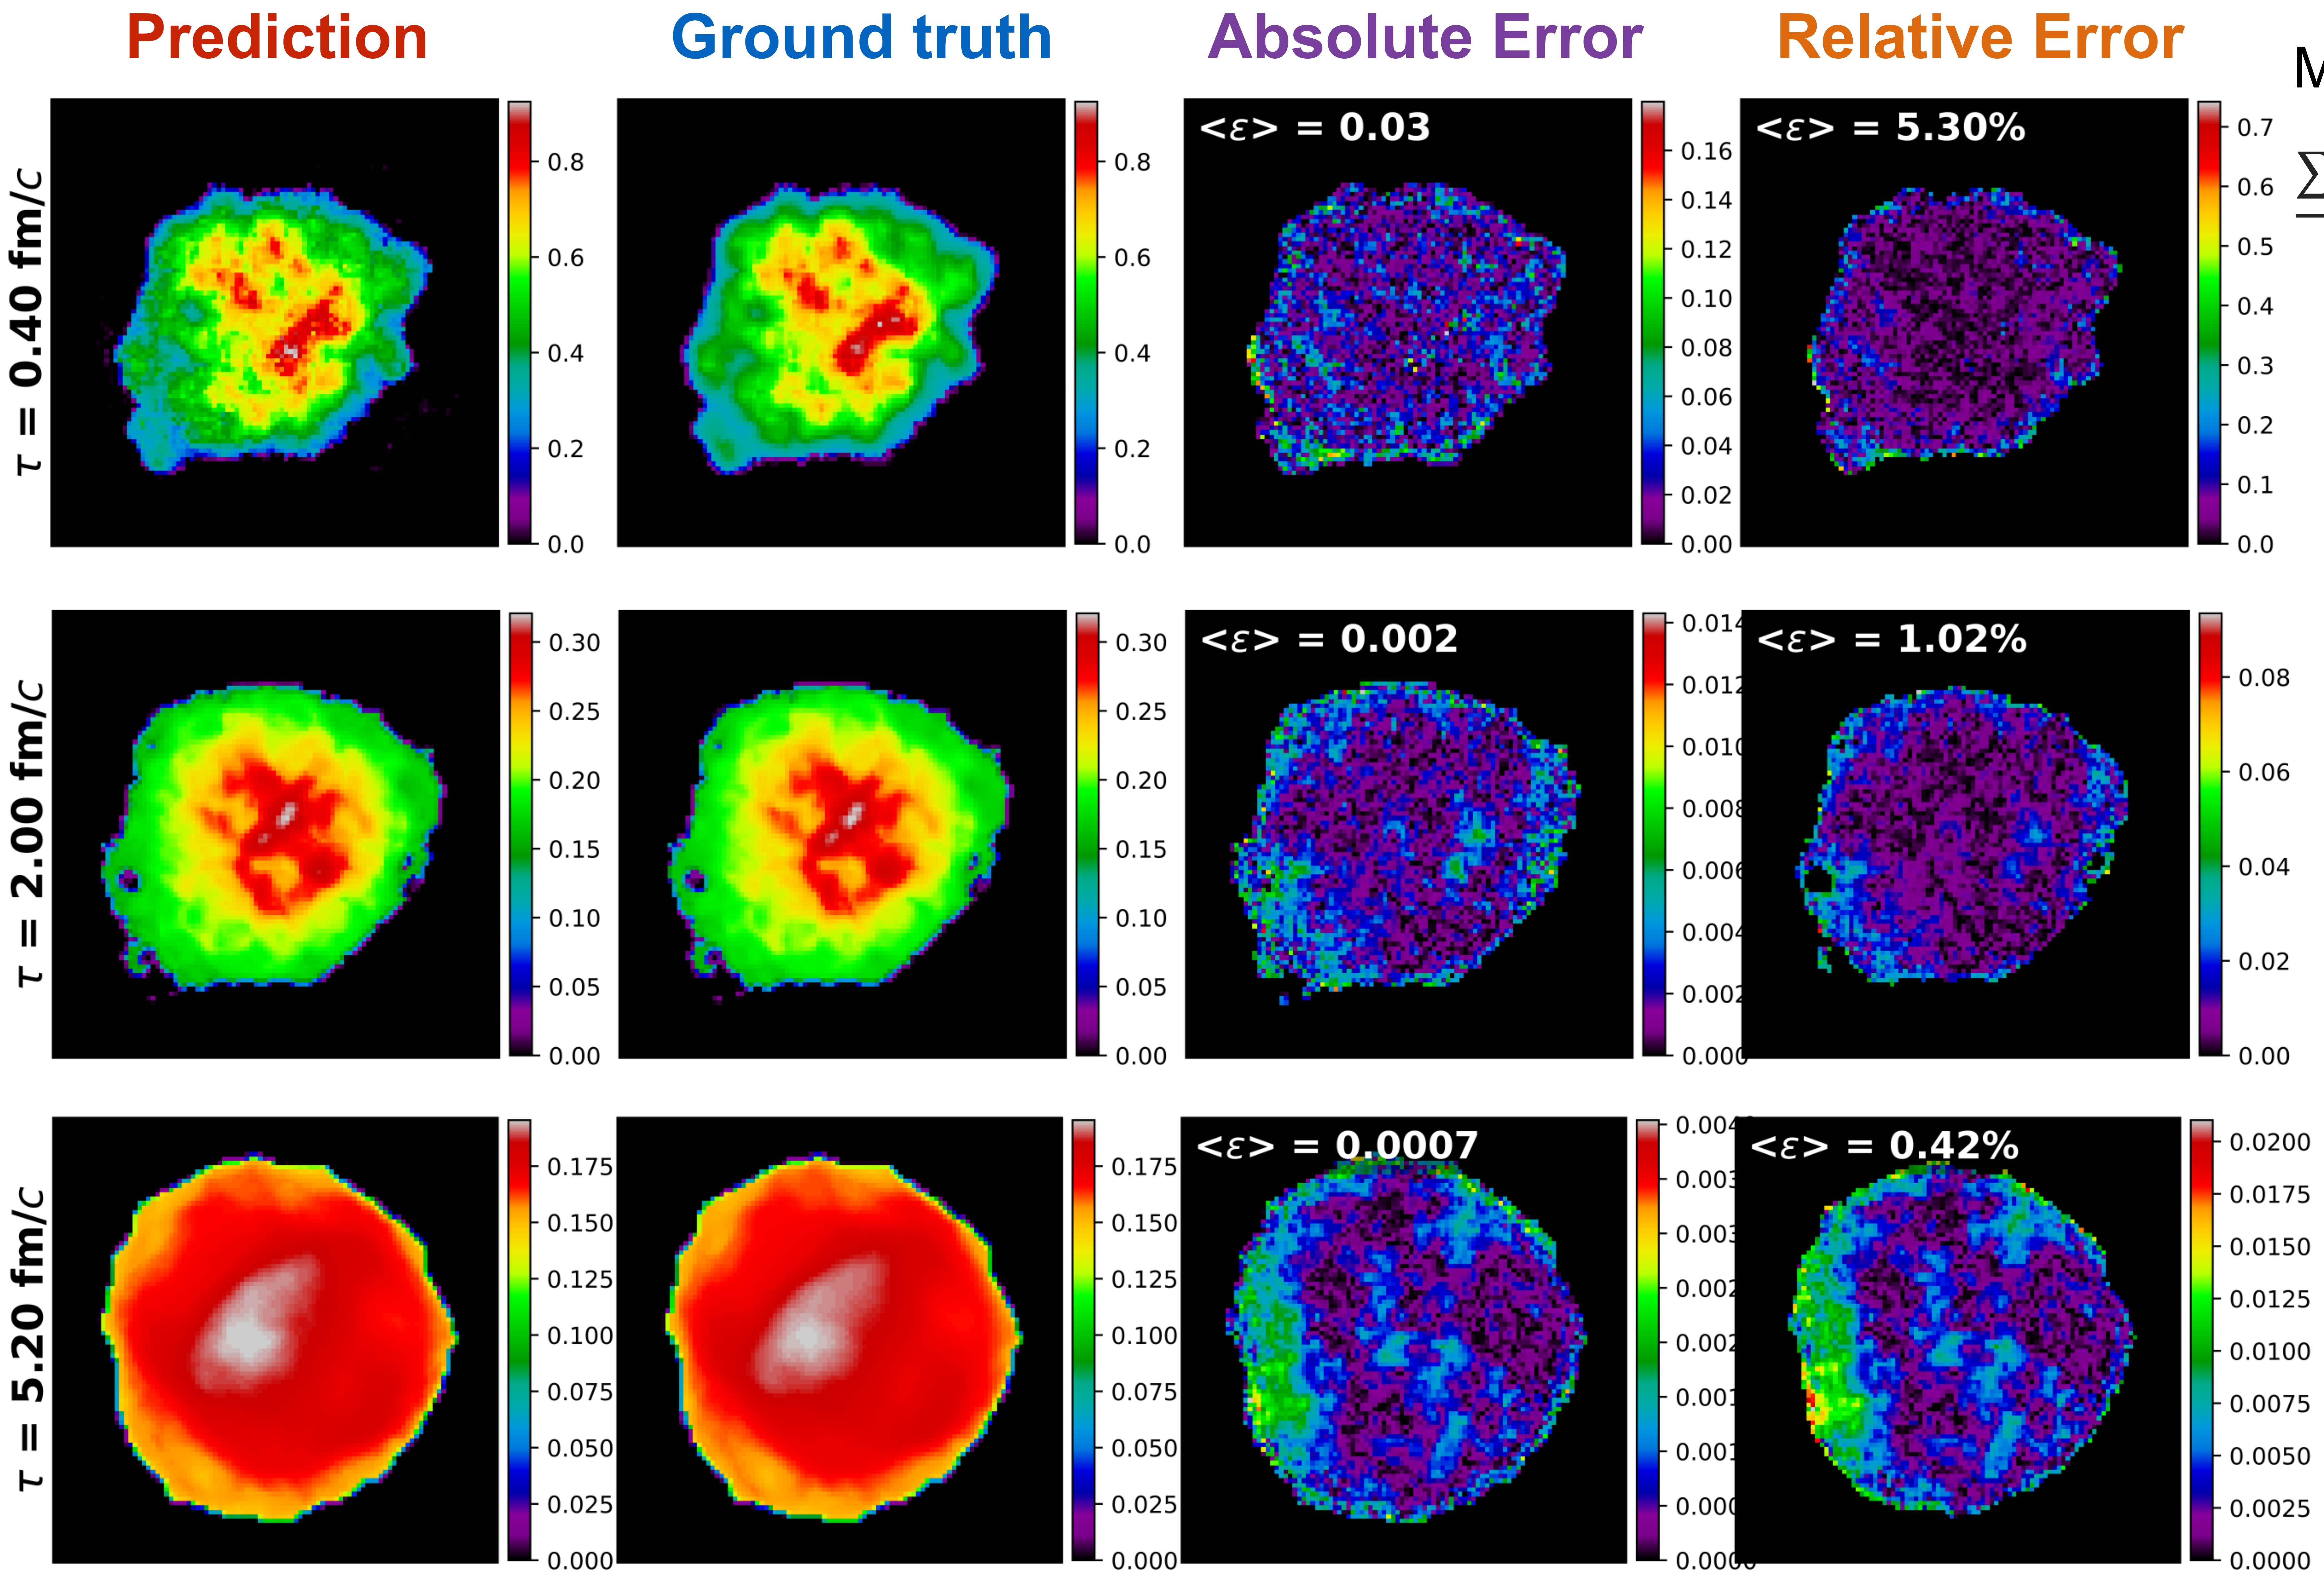

τ = **0.4** fm/c

τ = **2.0** fm/c

τ = **5.2** fm/c

Mean Relative Error $<\epsilon> \equiv \frac{\sum_{x,y}|\mathrm{pred.}(x,y) - \mathrm{truth}(x,y)|}{\sum_{x,y}|\mathrm{truth}(x,y)|}$

- Temperature prediction: relative error is larger at very early time step and rapidly decrease to < 1 %

Result: Flow Velocity u_x

Prediction | Ground truth | Absolute Error | Relative Error

τ = 0.4 fm/c

τ = 2.0 fm/c

τ = 5.2 fm/c

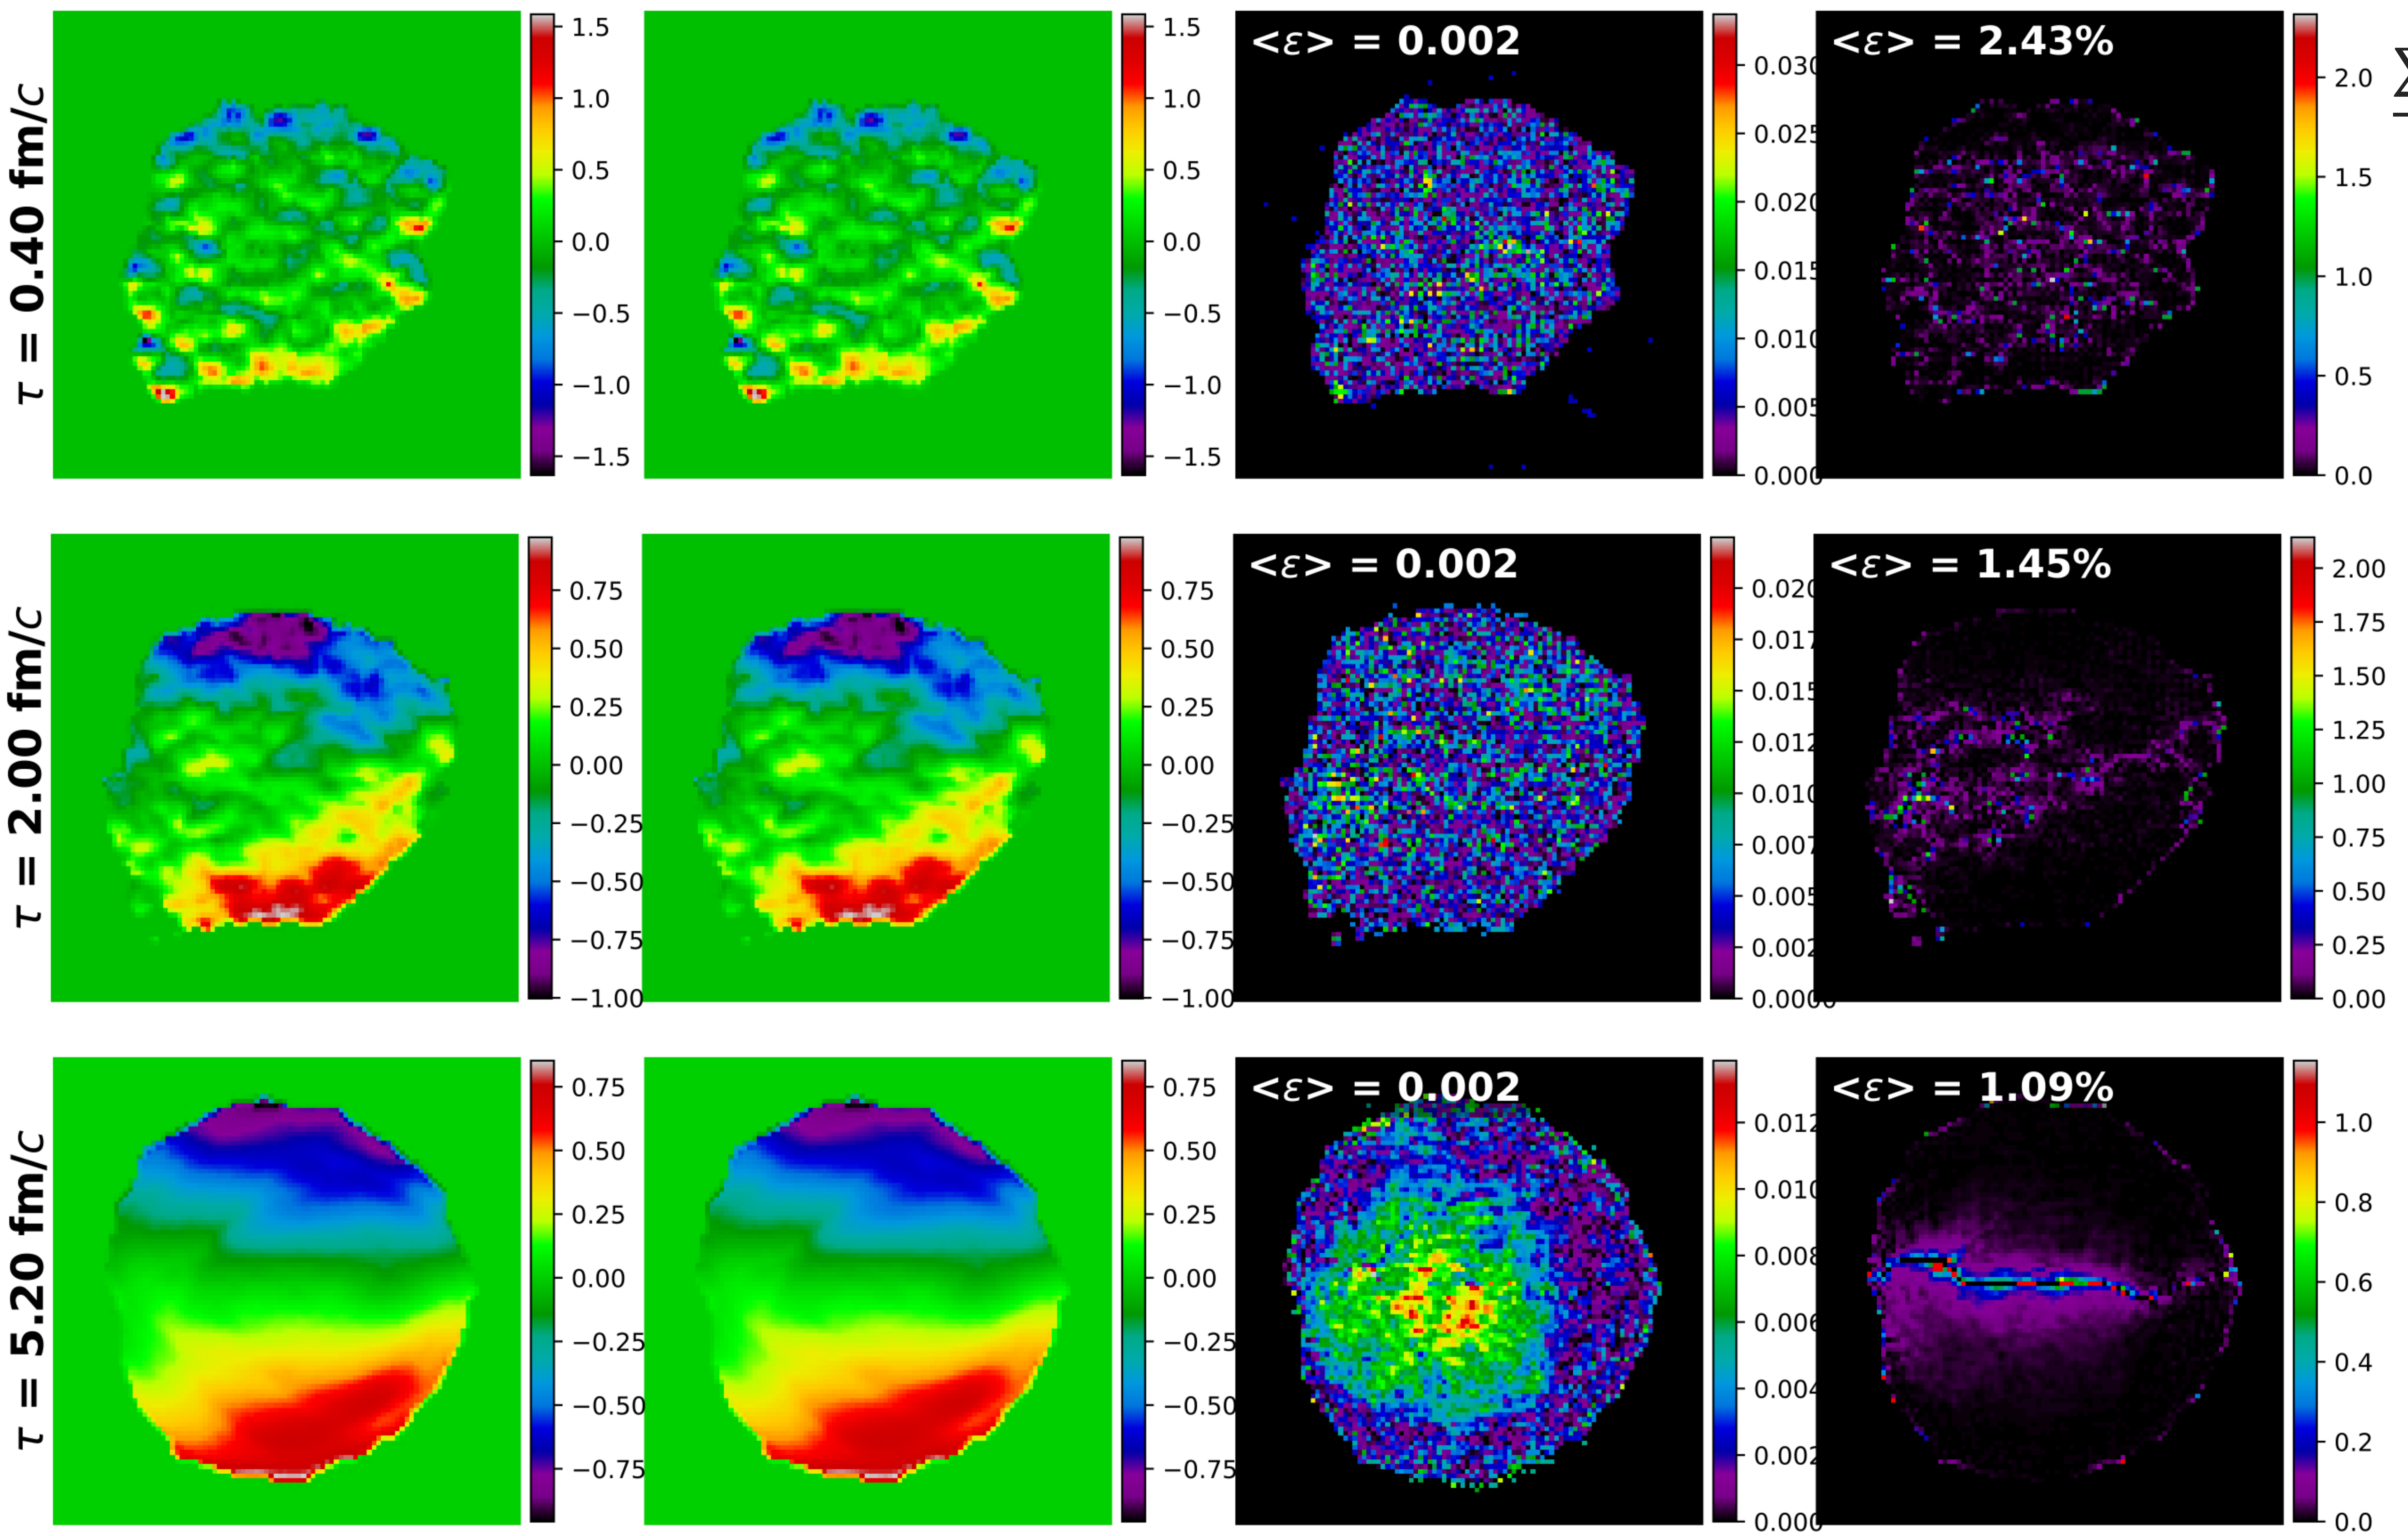

Mean Relative Error $<\epsilon> \equiv$

$$\frac{\sum_{x,y}|\text{pred.}(x,y) - \text{truth}(x,y)|}{\sum_{x,y}|\text{truth}(x,y)|}$$

- Signed fields
- Early-time flow fields are more challenging due to finer spatial structure

Result: Flow Velocity u_y

Prediction | Ground truth | Absolute Error | Relative Error

τ = **0.4** fm/c

τ = **2.0** fm/c

τ = **5.2** fm/c

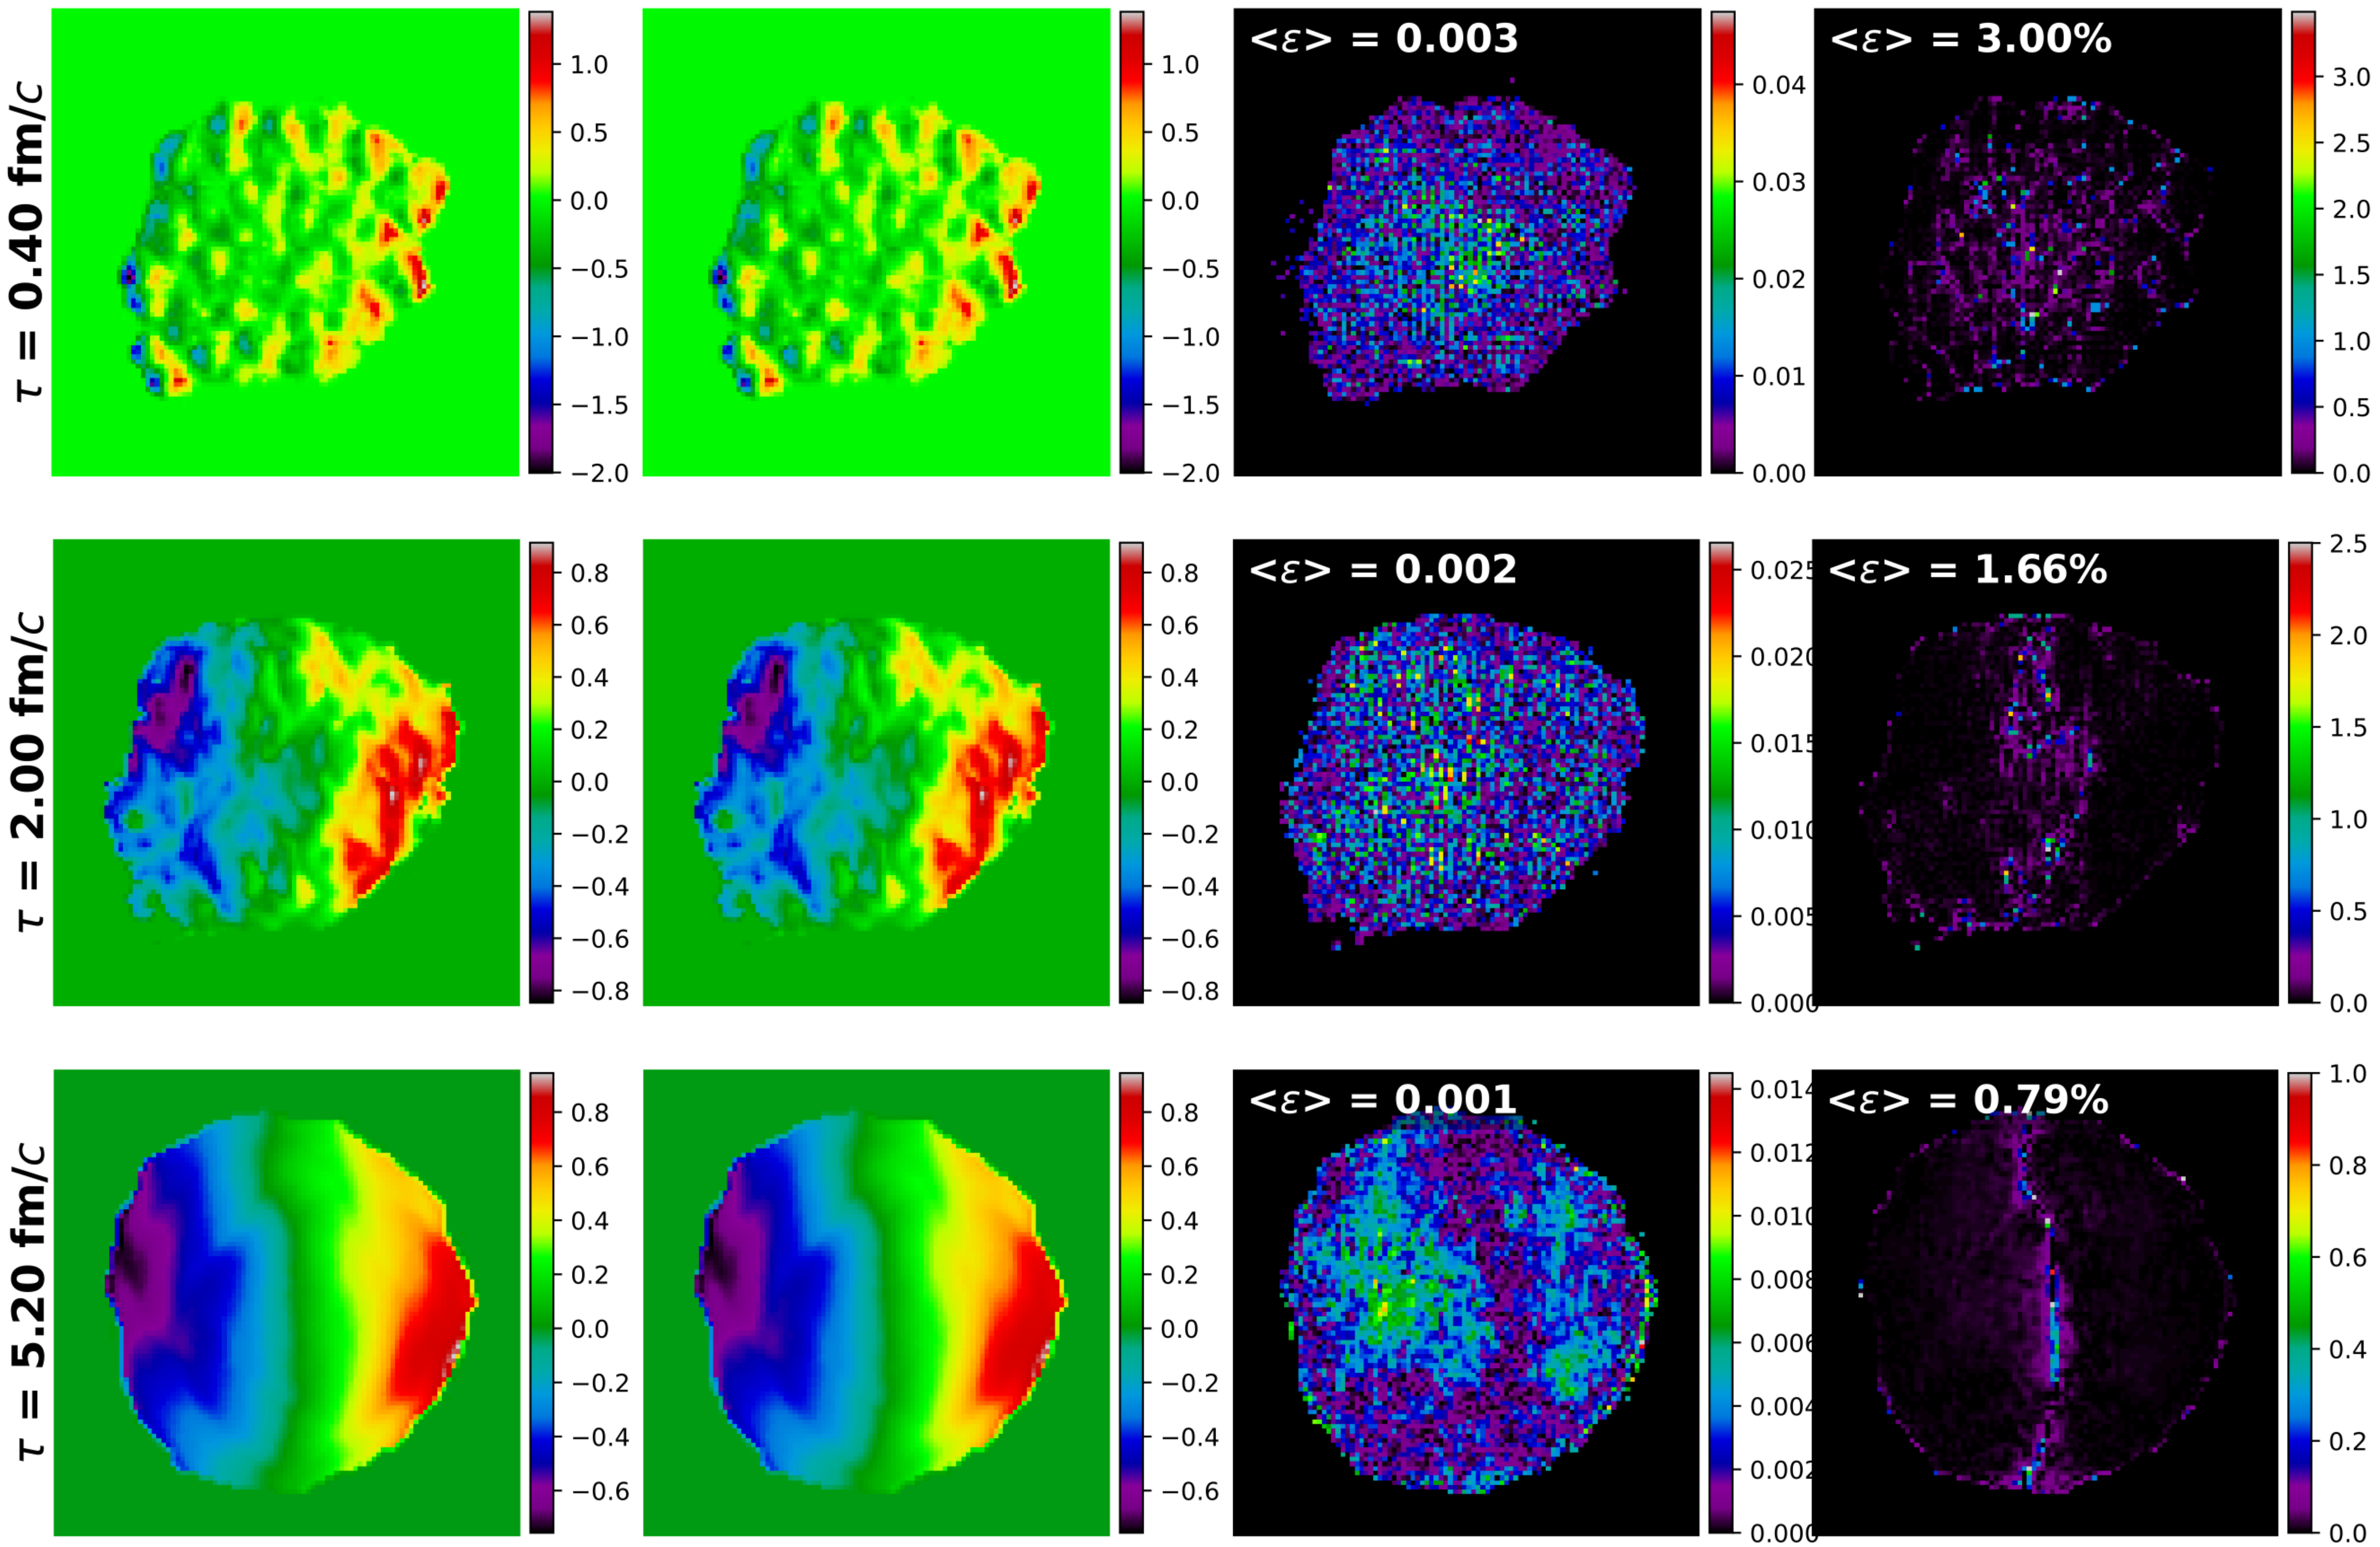

Mean Relative Error $<\epsilon> \equiv$

$$\frac{\sum_{x,y}|\mathrm{pred.}(x,y)-\mathrm{truth}(x,y)|}{\sum_{x,y}|\mathrm{truth}(x,y)|}$$

- Signed fields
- Early-time flow fields are more challenging due to finer spatial structure

Mean Relative Error

Mean relative error (over time and events)

Centrality	Temperature	Energy density	Pressure	Flow velocity (u_x)	Flow velocity (u_y)
0-5%	1.0%	0.7%	1.0%	1.4%	1.2%
5-10%	1.1%	0.5%	0.7%	1.2%	1.1%
10-20%	1.4%	0.5%	0.7%	1.2%	1.1%
20-30%	1.5%	0.5%	0.8%	1.1%	1.0%
30-40%	1.8%	0.6%	1.2%	1.1%	1.1%

→ Peripheral events have more fluctuations and asymmetry, which is harder for FNO

Temperature	Energy density	Pressure	Flow velocity (u_x)	Flow velocity (u_y)
1.3%	0.5%	0.9%	1.2%	1.1%

Overall, energy density is reproduced at the sub-percent level, while the other variables remain near the 1% level

Fourier Spectrum Validation: Energy Density

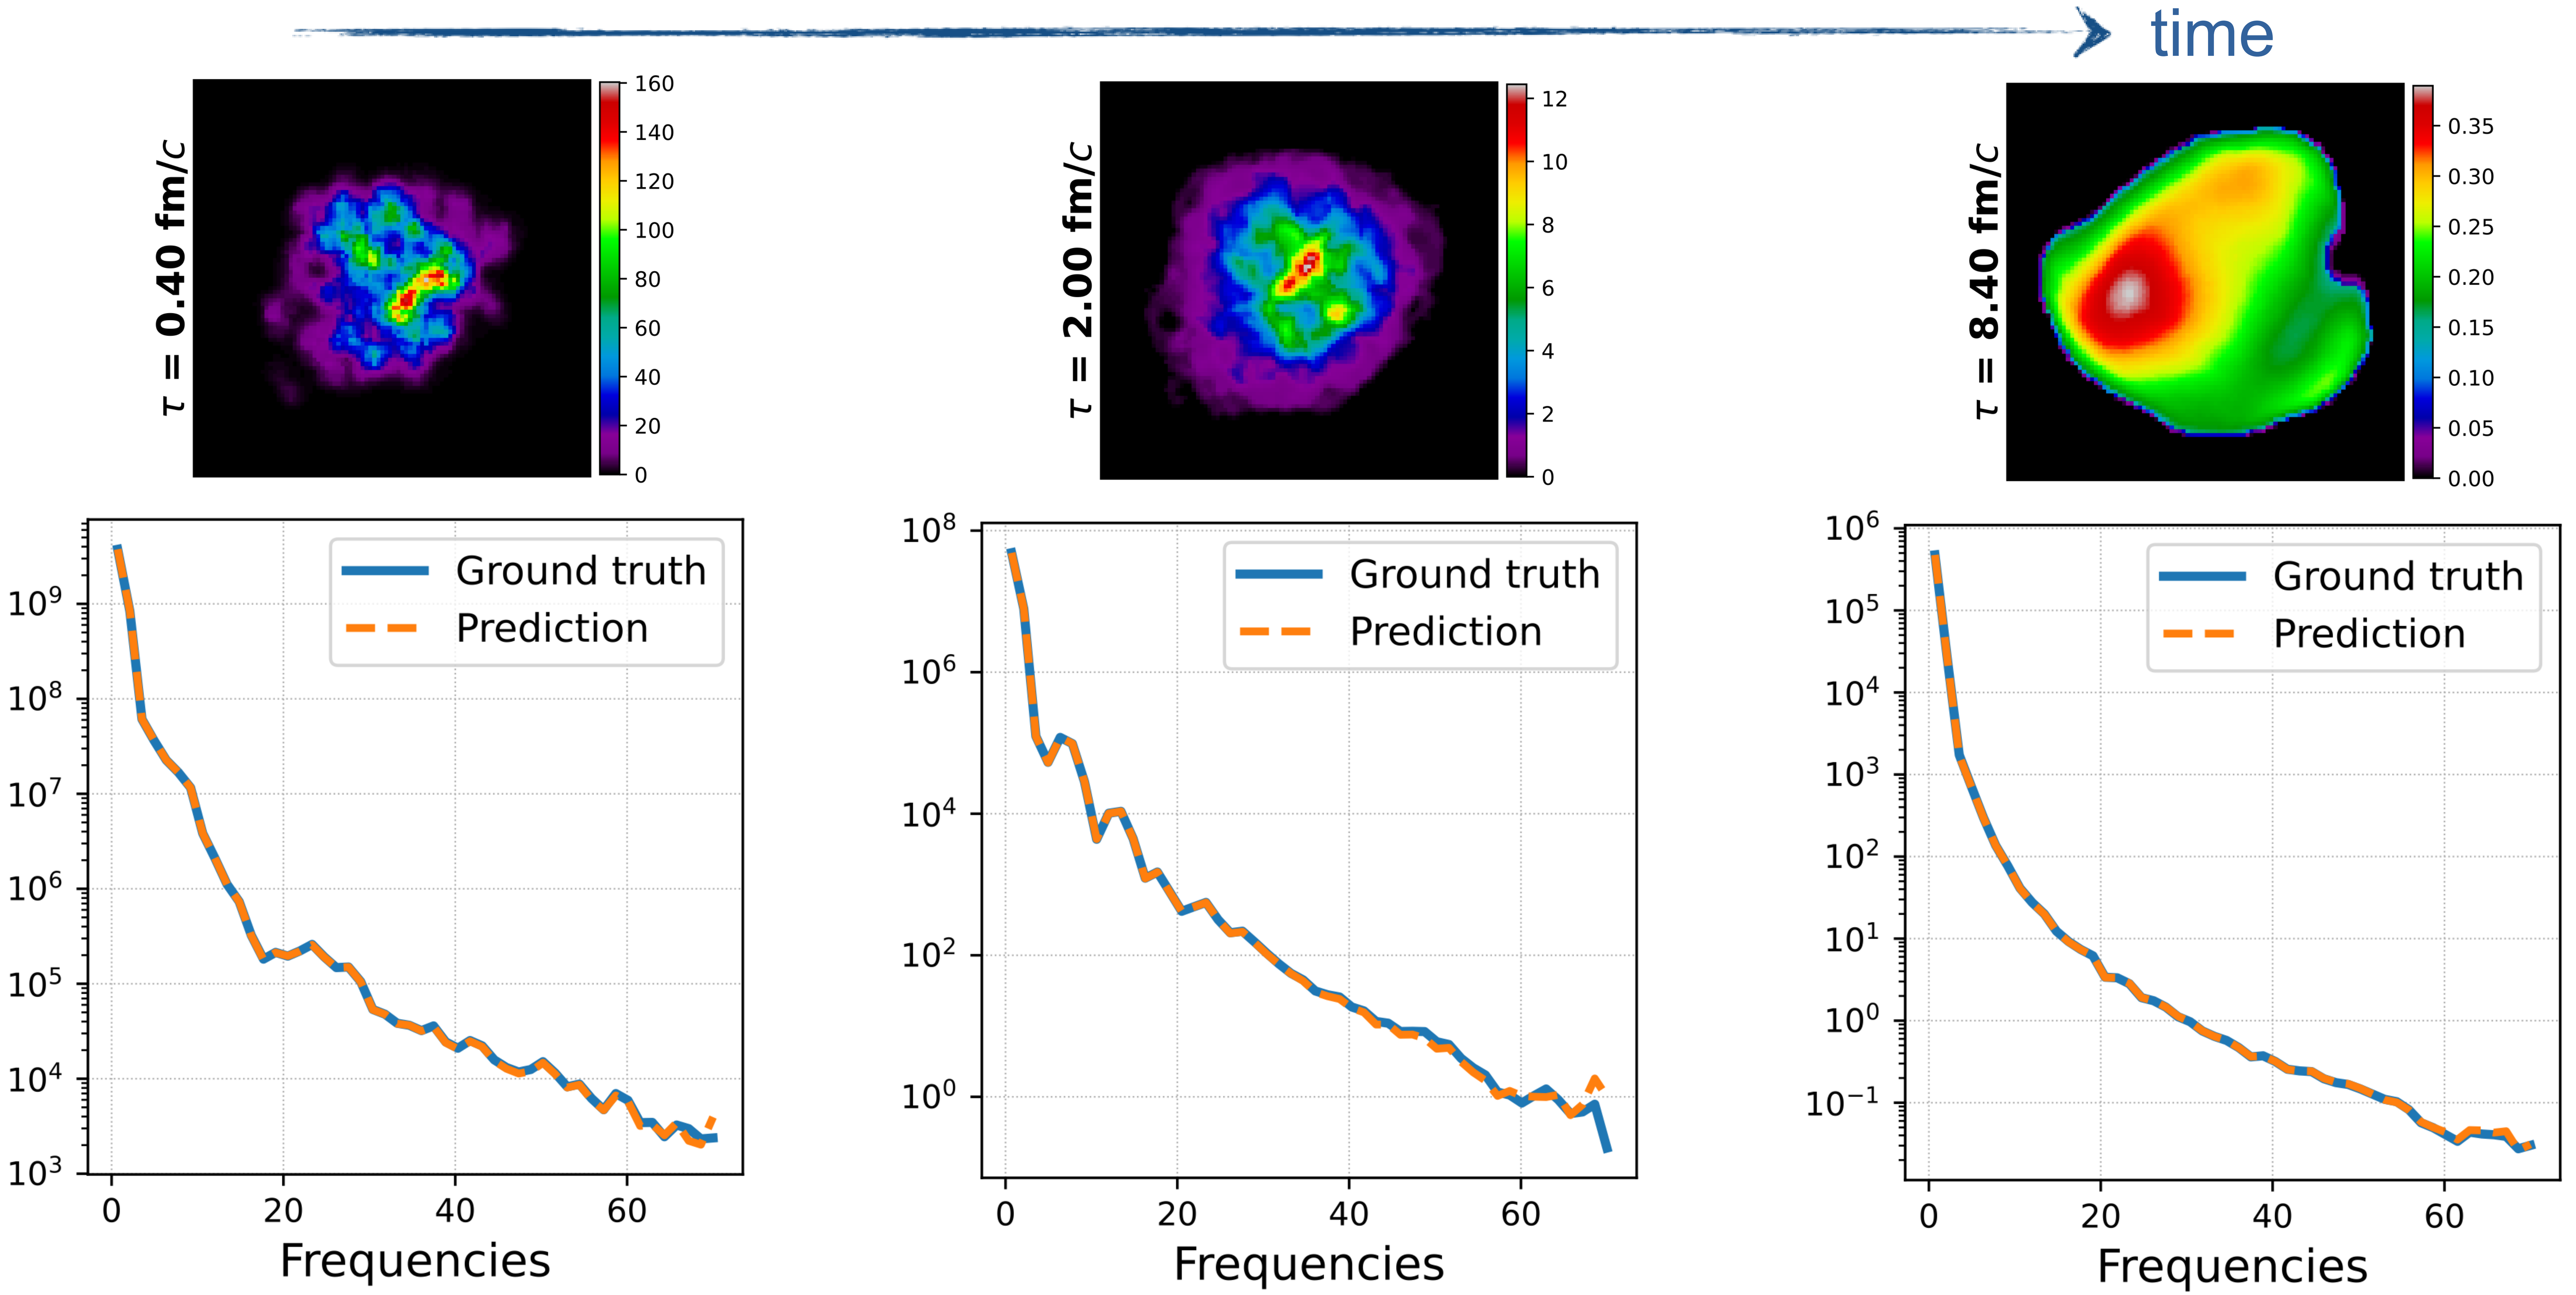

- Fourier frequency distributions for energy density are well reproduced in FNO

Fourier Spectrum Validation: Temperature

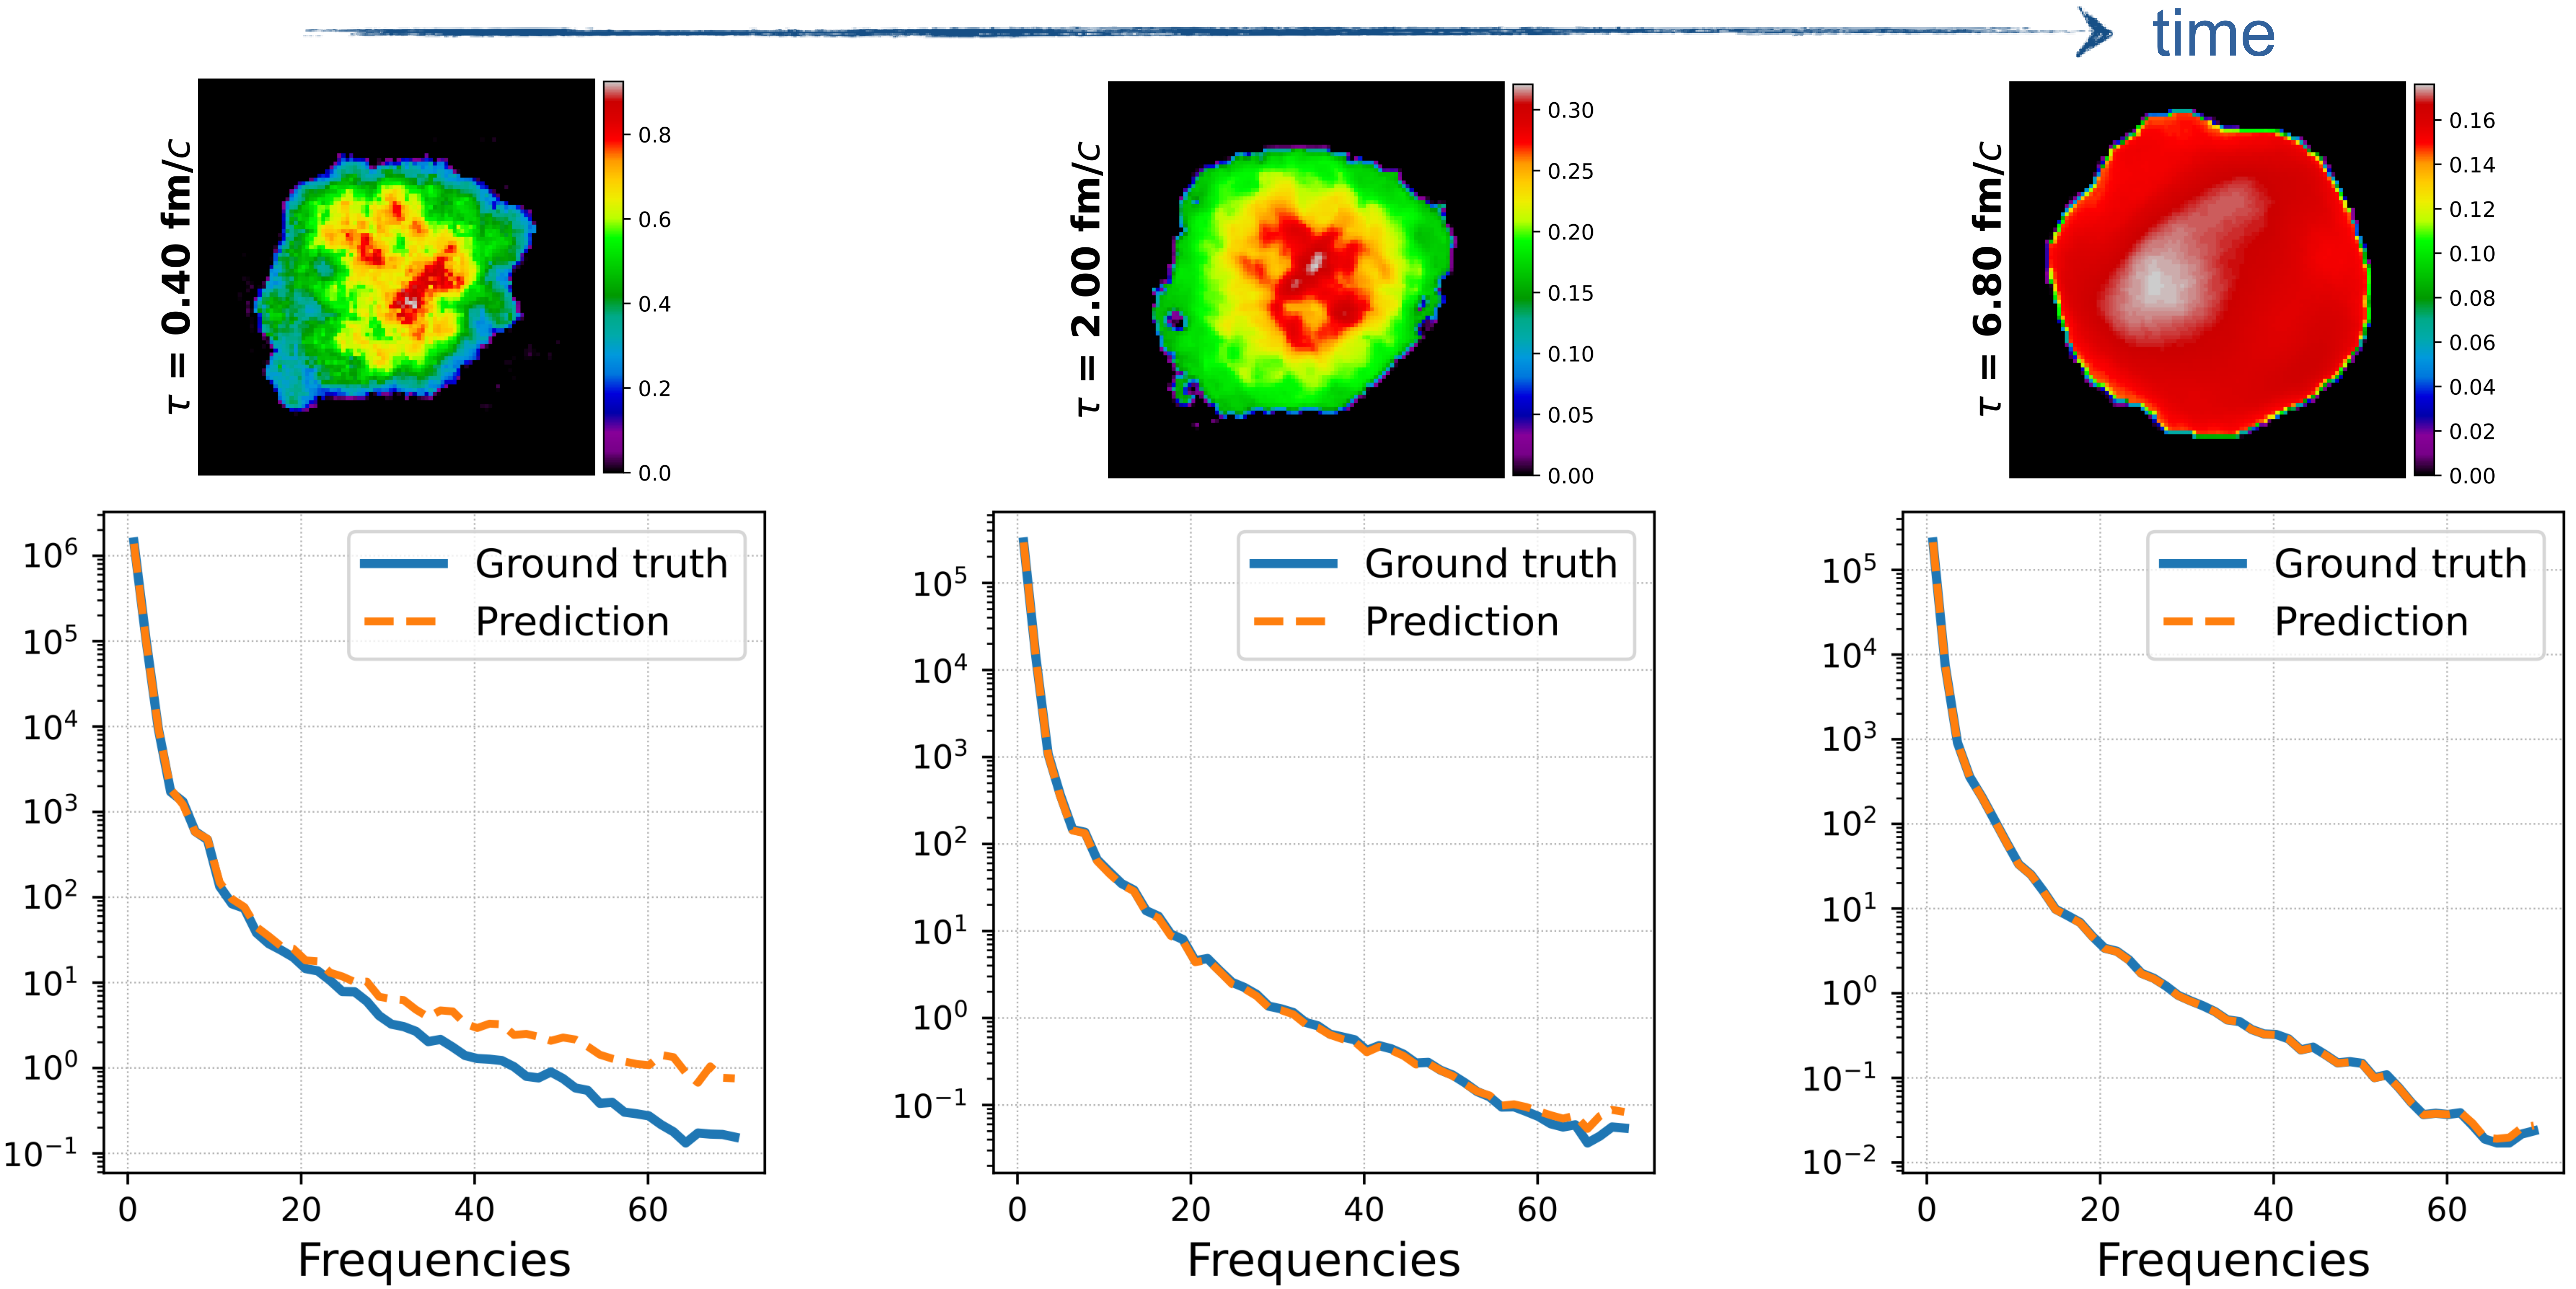

- Early-time high-frequency temperature modes remain more challenging

Out-of-distribution Handling

- Test on **different initial condition** i.e. smaller hotspot size (spiker) than the data used to train the model
  - ➞ training: hotspot size of ~0.2 fm, testing: hotspot size of ~0.1 fm
- The model still captures the global structure with increased error, as expected for unresolved higher-frequency features
  - ➞ Training on mixed hotspot sizes would reduce the error

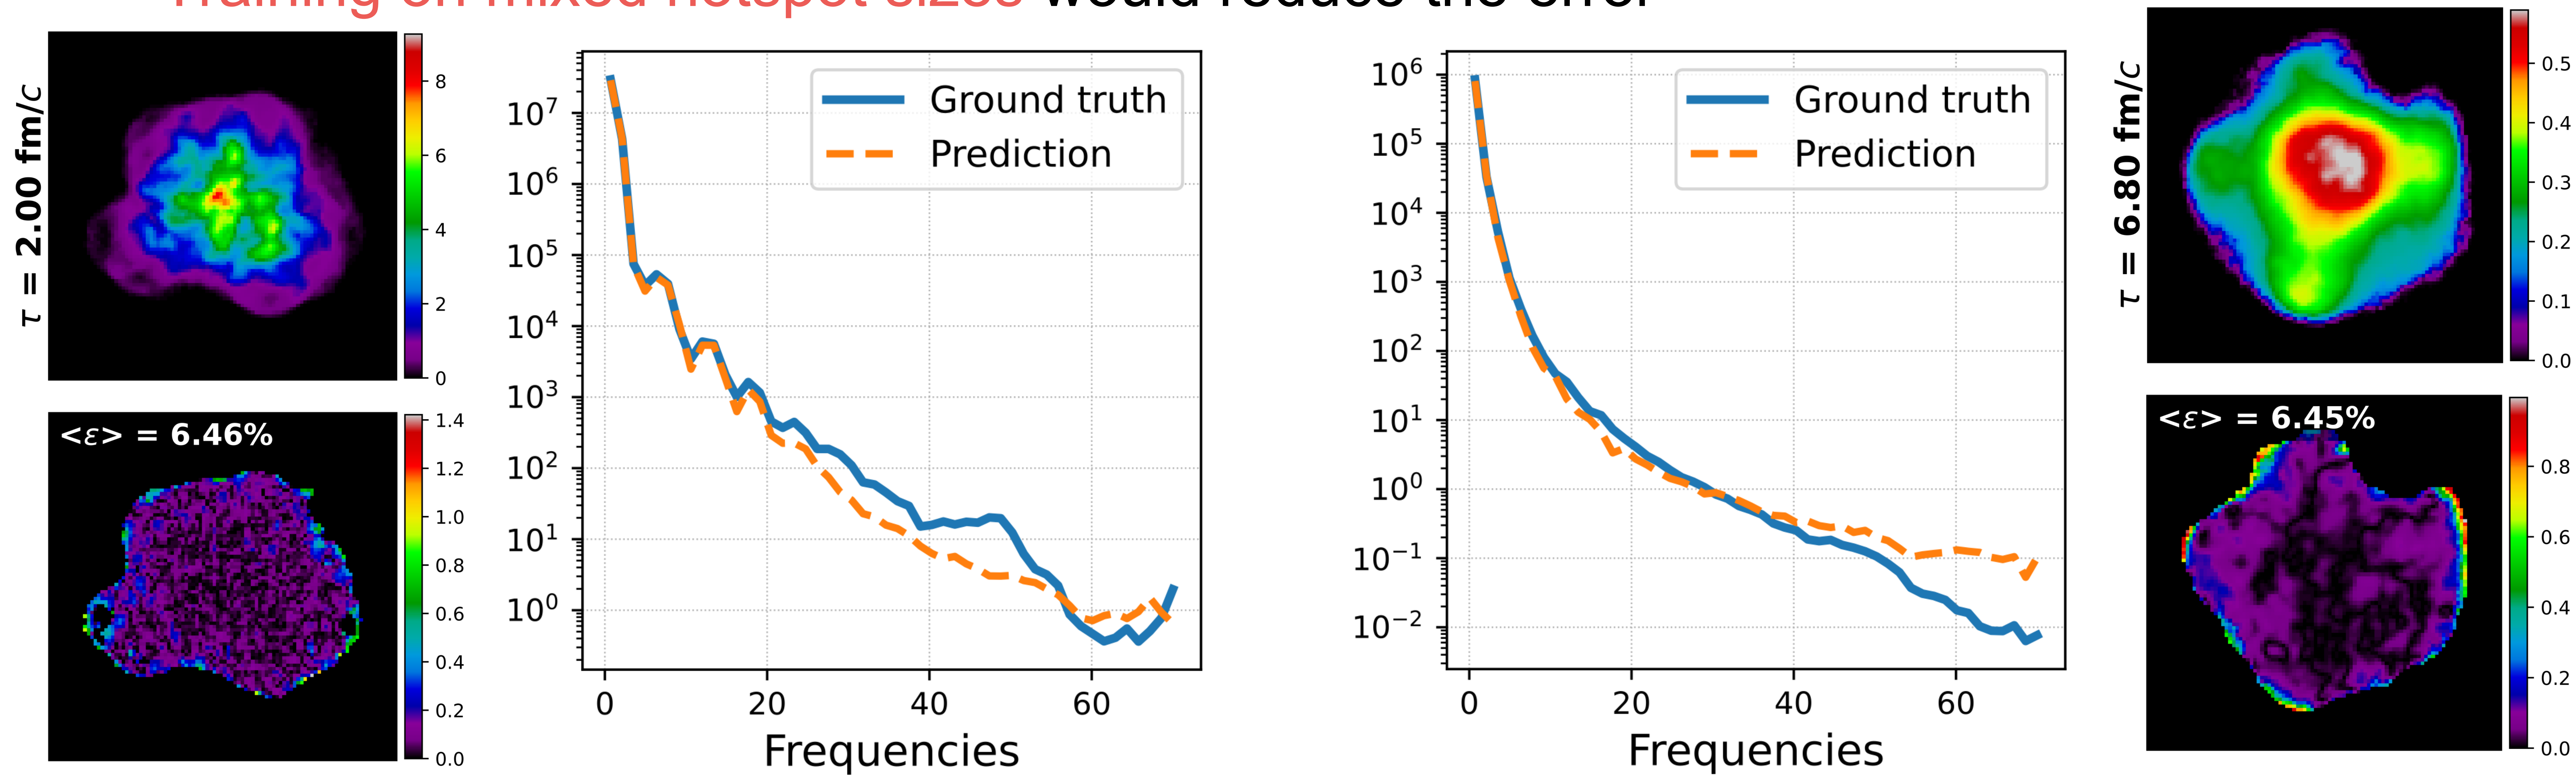

Outlook

- **Near term**
  - → Train on mixed initial conditions for robust generalization
  - → Further optimize performance for different physics use cases
    e.g. tighten accuracy for flow observables, relax where tolerable
  - → **3+1D FNO**
    - full traditional simulation ~20 hours → target 10 seconds
    - Key challenge: memory constraints
- **Longer term**
  - → *Initial condition → FNO hydro → hadronization (e.g. UrQMD afterburner) → experimental detector simulation*
  - → Make Bayesian inference and global multi-observable analyses computationally feasible

Summary

Slide modified after presentation thanks to comments from J. Putschke et al.

- **Fourier Neural Operator (FNO)** allows fast simulation of hydrodynamic evolution of Quark Gluon Plasma

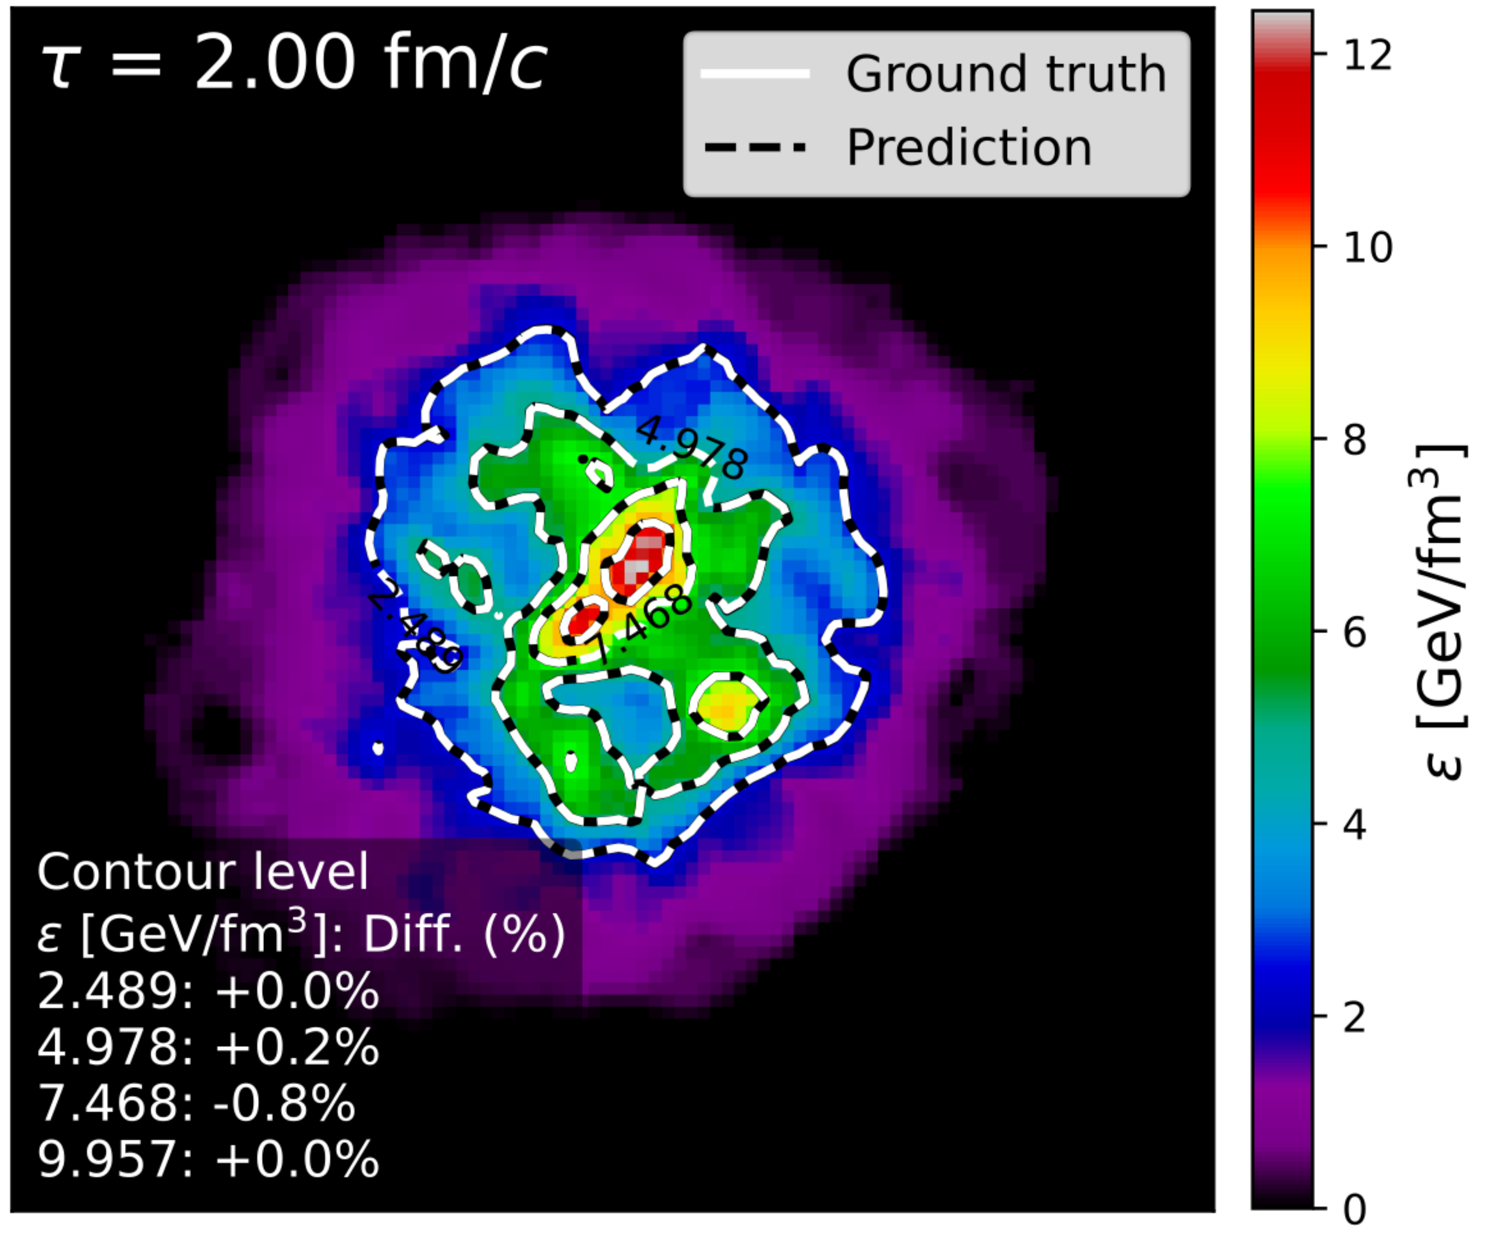

- FNO for **2+1D** QGP evolution achieved
  - → 2 to 3 orders of magnitude speedup
  - → < 2% relative error reproducing across $(\epsilon, T, P, u_x, u_y)$

- **3+1D** hydro simulation using FNO is in progress

This framework opens a path toward
large-scale event-by-event inference *and*
computationally intensive multi-observable studies

Please also see earlier work of this topic by D. Stewart, J. Putschke
Phys.Rev.C 113 (2026) 1, 014904

Backup