

Accuracy-cost trade-offs & Efficiency in Neural Physics Simulation

Monday, 26 January 2026 16:00 (30 minutes)

Accuracy-cost trade-offs and data and parameter efficiency are two fundamental aspects of machine learning for scientific computing. In the first part of this talk, we address test-time control of model performance. We introduce the Recurrent-Depth Simulator (RecurrSim), an architecture-agnostic framework that enables explicit control over accuracy-cost trade-offs in neural simulators without retraining or architectural redesign. By adjusting the number of recurrent iterations, users can flexibly trade computational cost for accuracy at inference time. ReccurrSim achieves physically plausible long-horizon simulations across standard fluid dynamics benchmarks and large-scale 3D compressible Navier-Stokes problems, where a 0.8B-parameter ReccurrFNO outperforms 1.6B-parameter baselines while using 13.5% less training memory. The framework generalizes across diverse architectures, including transformers and operator-learning models.

In the second part of the talk, we focus on data and parameter efficiency. We introduce Neural-HSS, a novel architecture, inspired by the structure of Green's functions for elliptic PDEs, based on Hierarchical Semi-Separable matrices. Neural-HSS is provably data-efficient, satisfies exactness properties in low-data regimes for a broad class of PDEs, and empirically demonstrates superior performance on large-scale elliptic and multi-physics PDEs across diverse scientific domains.

Primary author: SITTONI, Pietro (Gran Sasso Science Institute- GSSI)

Co-authors: TUDISCO, Francesco (GSSI); GUGLIELMI, Nicola (Gran Sasso Science Institute); ZANGRANDO, Emanuele (Gran Sasso Science Institute); CASULLI, Angelo Alberto (Scuola Normale Superiore di Pisa); Mr MAJID, Harris Abdul (The University of Edinburgh)

Presenter: SITTONI, Pietro (Gran Sasso Science Institute- GSSI)

Session Classification: Session 2