

UFlex: A Flexible & Efficient Multiscale Neural Physics Simulator

Wednesday, 11 February 2026 14:50 (25 minutes)

U-Net-style architectures are widely adopted for modeling physical systems, as their multiscale structure enables efficient processing of high-resolution data while reflecting the hierarchical structure of many physical phenomena. The most successful U-Net-based neural physics simulators typically combine convolutions with transformers; however, when applied to regular grids, they depend on highly structured components that limit adaptability across different spatial and spatiotemporal dimensions. In this work, we revisit transformer-based U-Nets with the goal of maximizing flexibility without sacrificing multiscale efficiency. We introduce a U-Net composed entirely of transformer blocks operating on a one-dimensional latent sequence, making it easy to extend across spatial and spatiotemporal dimensions. Evaluated on seven challenging benchmarks (four 2D and three 3D), our model scales to resolutions of up to 512×512 in 2D and $256 \times 128 \times 256$ in 3D, while reducing training memory and accelerating training compared to state-of-the-art transformer baselines, all while achieving competitive or state-of-the-art predictive accuracy.

Presenter: SITTONI, Pietro (Gran Sasso Science Institute)