Contribution ID: **67** Type: **not specified**

DeepCosmoNet

Thursday 29 May 2025 10:45 (15 minutes)

Understanding the large-scale structure of the universe, known as the Cosmic Web, relies heavily on the analysis of numerical 3D N-body simulations.
Analyzing the output of those simulations remains computationally intensive due to large 3D point clouds and complex clustering tasks.
Traditional methods—such as FoF, SUBFIND, and ZOBOV often operate as separate, CPU-bound tools with limited parallel scalability and require costly post-processing.

We introduce a unified DeepCosmoNet framework that segments both high-density (halos and subhalos) and low-density (voids) structures in one GPU-optimized pass. For halos and subhalos, we apply particle-level Graph Neural Networks accelerated by tree-based nearest-neighbor searches, avoiding the inefficiencies of voxelization for sparse objects. For voids, we voxelize the density field and employ 3D CNNs framing void identification as a bounding-box detection problem, with a parallelized 3D Non-Maximum Suppression.

Our pipeline achieves end-to-end speedups of two orders of magnitude over CPU benchmarks, leveraging hybrid parallelism to maximize throughput.
The approach exhibits near-linear scaling on multi-GPU clusters, demonstrating exceptional throughput and scalability for large simulation volumes.

Furthermore, the quality of the detections is notably high; we find that the detections produced by our method are cleaner than the ground truth labels themselves.
By shifting the computational bottleneck from data processing to scientific interpretation, DeepCosmoNet enables accelerated and scalable exploration of the Cosmic Web.

Presenter: DEL ZOPPO, Vincenzo (KOEXAI)

Session Classification: Bandi a Cascata