5th ECMWF-ESA
Machine Learning
Workshop

Bologna, Italy | 13–17 April

Scalable GPU-Accelerated Training of Graph Neural Networks for High-Resolution ESOP Simulations

Oscar Leung and Ye Ha Kim

The 0.1° Resolution Bottleneck

The Challenge:

A 0.1° global mesh results in $simO(10^9)$nodes and edges per temporal state.

Standard Framework Failure:

PyTorch Geometric (PyG) and DGL default to full-graph message passing, exceeding the 80GB VRAM limit of A100 GPUs instantly.

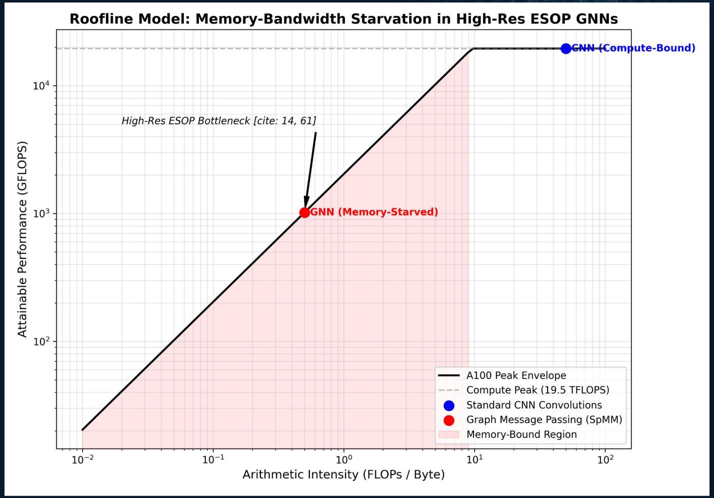

Roofline Model – Memory-Bandwidth Starvation in High-Resolution ESOP GNNs. Standard graph message passing (SpMM) operates deep in the memory-bound region on A100 GPUs. The proposed framework (Fused-Map + Match-Reorder + custom SpGEMM) shifts the workload toward the compute roof, enabling scalable 0.1° training.

Geospatial Representation

Transition from deterministic to probabilistic ensemble generation (CRPS evaluation).

Integration of sub-grid physics-informed constraints to prevent energy drift.

Scaling toward fully coupled Earth system models (Ocean + Atmosphere).

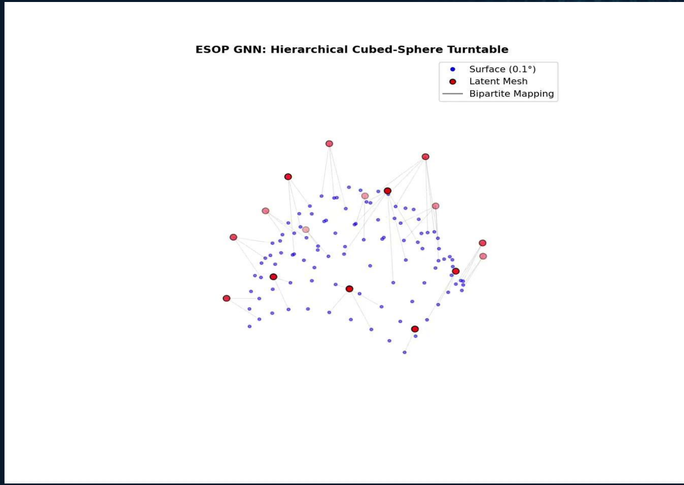

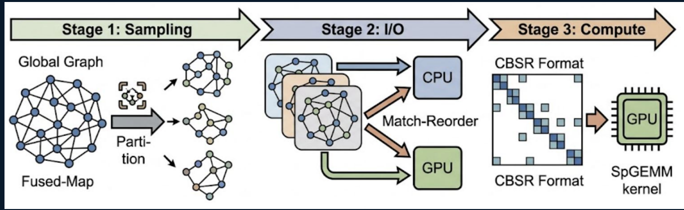

Initial Studies - WeatherGNN

Goals:

Open-source, modular ESOP pipeline for GPU-accelerated GNN training on spherical meshes

Fast, high-resolution surrogates of atmospheric evolution using ERA5 reanalysis

Focus on reproducible engineering, not new state-of-the-art forecast skill

Limitations:

Small channel set (subset of variables/levels) vs. production models

Micro-benchmarks run on CPU only (GPU timings expected to differ)

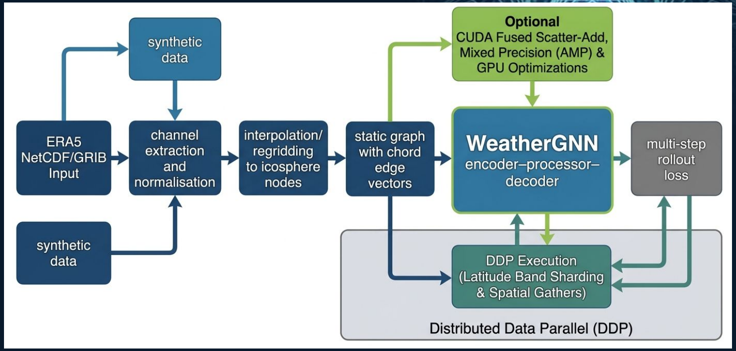

Data and compute pipeline for ESOP GNN training from reanalysis fields to multi-step latitude-weighted loss, including optional distributed spatial gather and fused neighbour aggregation.

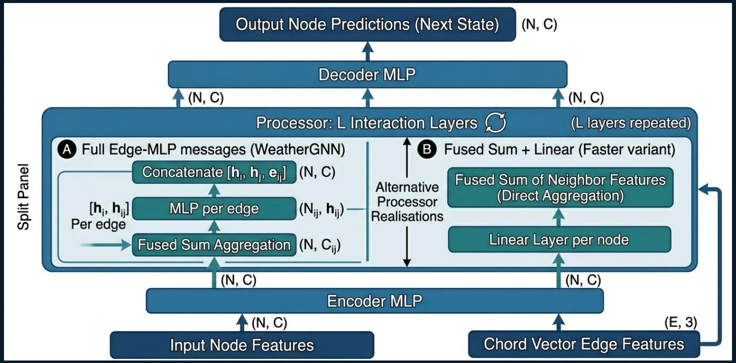

Encoder–processor–decoder WeatherGNN with alternative processor realisations (edge-conditioned messages versus fused neighbour sum).

GPU-Accelerated Forward Pass

Message passing relies on Sparse Matrix-Matrix Multiplication (SpGEMM).

Limitation: Standard SpGEMM kernels suffer from thread divergence and uncoalesced memory access on unstructured graphs.

Solution: Implementation of fused-mapping subgraph sampling and Match-Reorder memory layouts.

Result: Minimizes irregular memory fetches, maximizing L2 cache utilization on the GPU.

$$h_i^{(k+1)} = \sigma\left(W^{(k)} \cdot h_i^{(k)} + \sum_{j \in \mathcal{N}(i)} \alpha_{ij} \cdot A_{ij} \cdot h_j^{(k)}\right)$$

Message-passing function

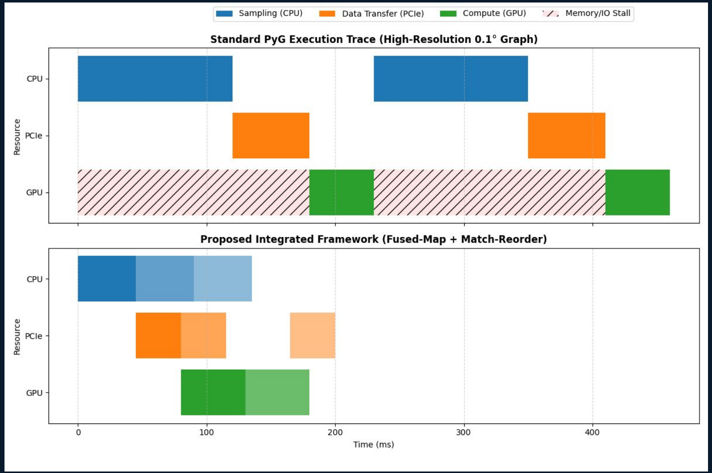

Execution trace comparison for a single training batch on a 0.1° high-resolution atmospheric graph. The proposed framework (Fused-Map sampling + Match-Reorder) dramatically reduces CPU-GPU synchronization stalls and PCIe data traffic compared to standard PyTorch Geometric, enabling 2–3× higher throughput.

Scaling Out: Partition Parallelism

Challenge:

High inter-GPU communication overhead for 'halo' nodes at partition boundaries.

Method:

Asynchronous CPU-GPU memory offloading. Dormant node features are swapped to host RAM, active boundaries are kept in VRAM.

Speedup:

Achieves up to 2.5x training speedup per epoch compared to baseline distributed implementations.

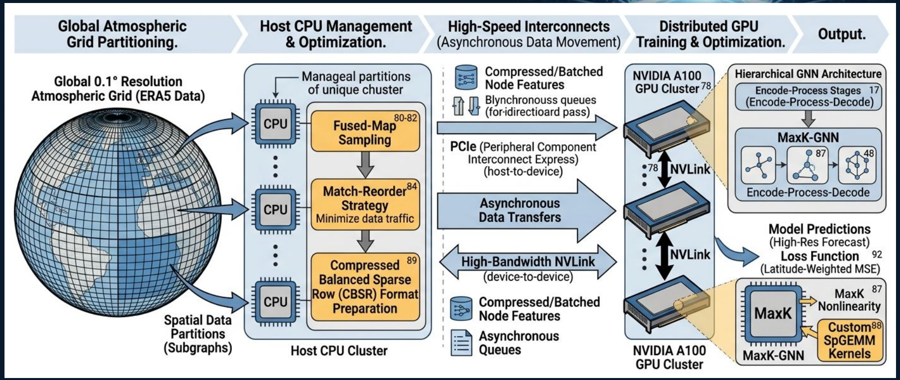

Experimental Setup & Baselines (Ongoing)

Experimental Design & Benchmarking Environment

Category	Configuration Details
Dataset	ERA5 Reanalysis (0.1° Interpolated)
Training Period	1979 – 2020
Validation Period	2021 – 2022
Atmospheric Variables	6 variables @ 13 pressure levels
Surface Variables	6 variables (MSLP, T2M, U10, V10, etc.)
Total Input Channels	84 unique feature channels
Optimization Strategy	AdamW Optimizer
Learning Rate Schedule	Cosine Annealing
Forecast Step	6-hour Autoregressive Rollout
Evaluation Baselines	ECMWF HRES (IFS) & 0.25° GNN (GraphCast)

Medium-Range Forecast Skill

Evaluation based on Z500 (Geopotential Height) and T850 (Temperature).

Autoregressive roll-out evaluated up to a 10-day lead time.

Matches or exceeds operational baselines in RMSE and ACC for the first 7 days.

Constraint: Error growth remains non-linear post day-7 due to error accumulation in the recurrent state.

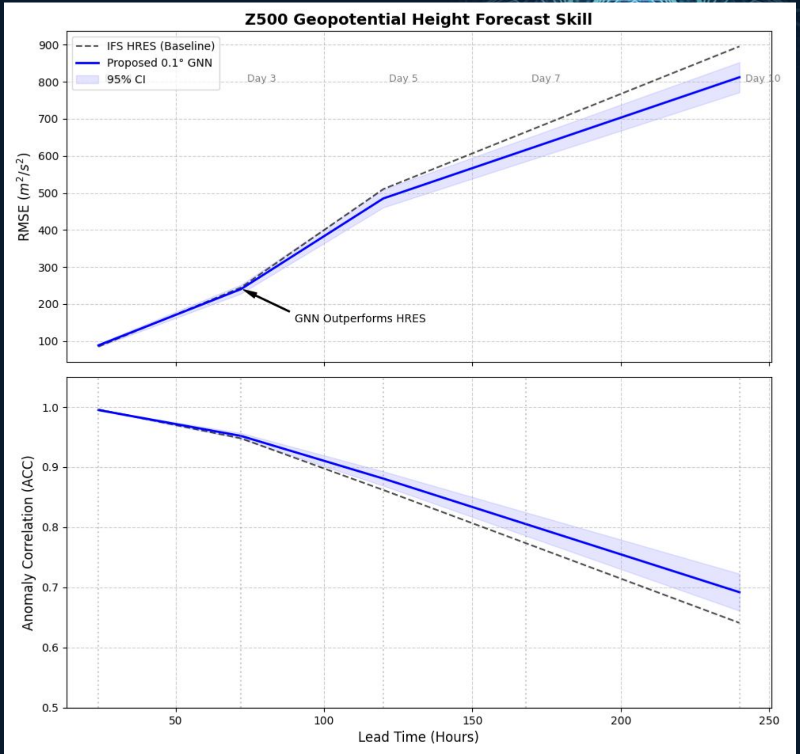

Z500 Geopotential Height Forecast Skill (RMSE and ACC). The proposed 0.1° GNN outperforms operational IFS HRES after ~Day 3 (arrow) and maintains superior skill out to Day 10.

Capturing High-Frequency Dynamics

Lower resolution models $0.25° - 1°$ suffer from spatial smoothing, dampening kinetic energy at high wavenumbers.

0.1° GNN successfully resolves steep pressure gradients in Tropical Cyclones and sharp moisture transport in Atmospheric Rivers.

Metric: Power spectral density analysis of kinetic energy.

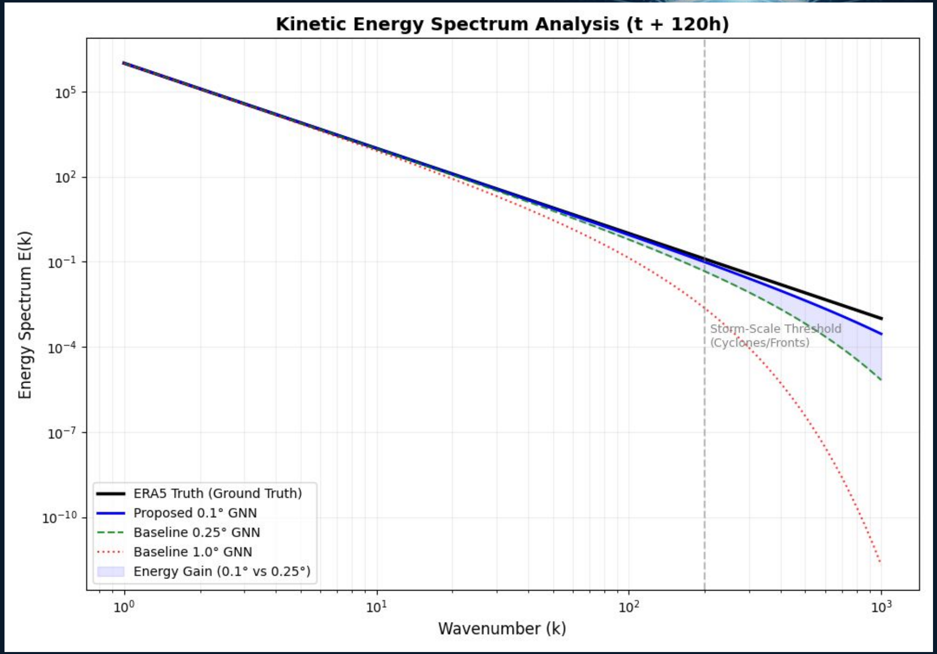

Kinetic Energy Spectrum Analysis at t + 120 h. The proposed 0.1° GNN (blue) closely matches ERA5 ground truth (black) at high wavenumbers, preserving storm-scale features (cyclones/fronts) that are lost in lower-resolution baselines. Shaded region highlights energy gain from finer resolution and hierarchical message passing.

The Autoregressive Smoothing Problem

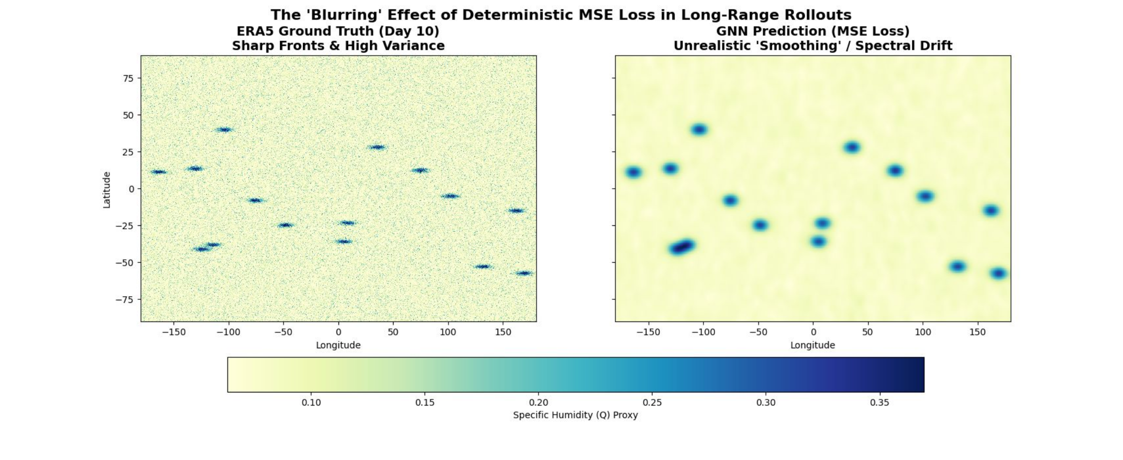

Models trained purely on Mean Squared Error (MSE) inherently predict the mean of the probability distribution.

Trade-off: Excellent RMSE scores directly correlate with the loss of physical variance at longer lead times (blurring).

Open Question: Transitioning the loss function to include adversarial losses (GANs) or utilizing the GNN state as a conditioning vector for Latent Diffusion.

Towards Digital Twins & DA

Pure emulation is insufficient for operational Digital Twins; models must ingest real-time observational data.

The ultra-fast forward pass (0.3s) of the optimized GNN enables rapid background trajectory generation for 4D-Var.

Future Pathway: Using the differentiable nature of the GNN to replace the manual adjoint in variational assimilation.

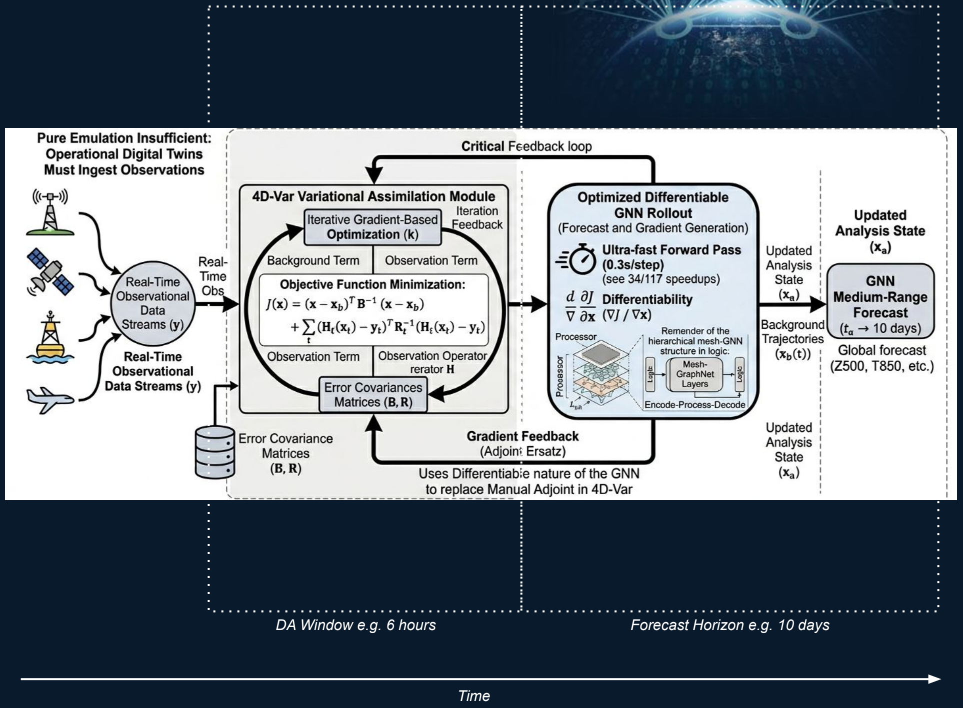

Key Assumptions & Constraints

Stationarity Assumption: The model assumes the atmospheric physics learned from 1979-2020 ERA5 data remain stationary. Vulnerable to climate drift and out-of-distribution extremes.

Hardware Lock-in: Custom SpGEMM kernels are currently heavily optimized for NVIDIA CUDA architectures; porting to TPUs or AMD requires significant rewrite.

I/O Bottleneck: Training speed is increasingly bounded by the ability to stream Petabytes of disk data to VRAM.

Conclusion & Future Work

Transition from deterministic to probabilistic ensemble generation (CRPS evaluation).

Integration of sub-grid physics-informed constraints to prevent energy drift.

Scaling toward fully coupled Earth system models (Ocean + Atmosphere).

References

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Pritzel, A., Ravuri, S., Ewalds, T., Day, J., Natalia, S. and Akerman, Z. (2023) 'Learning skillful medium-range global weather forecasting', *Science*, 382(6677), pp. 1416–1421.

Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X. and Tian, Q. (2023) 'Accurate medium-range global weather forecasting with 3D neural networks', *Nature*, 619(7970), pp. 533–538.

Ma, M. et al. (2023) 'HiSTGNN: Hierarchical spatio-temporal graph neural network for weather forecasting', *Information Sciences*, 648(5746), 119580.

Bailie, T. et al. (2023) 'Hierarchical Graph Networks for Accurate Weather Forecasting via Lightweight Training', *arXiv preprint arXiv:2310.05517*.

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D. and Simmons, A. (2020) 'The ERA5 global reanalysis', *Quarterly Journal of the Royal Meteorological Society*, 146(730), pp. 1999–2049.

Keisler, R. (2022) 'Forecasting global weather with graph neural networks', *arXiv preprint*

Pathak, J. et al. (2022) 'FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators', *arXiv preprint arXiv:2202.11214*.

Zhu, Z., Wang, P., Hu, Q., Li, G., Liang, X. and Cheng, J. (2024) 'FastGL: A GPU-efficient framework for accelerating sampling-based GNN training at large scale', *Proceedings of the 29th ASPLOS*, 3, pp. 605–620.

Peng, H., Xie, X., Shivdikar, K., Hasan, M.A., Zhao, J., Huang, S., Khan, O., Kaeli, D. and Ding, C. (2024) 'MaxK-GNN: Extremely fast GPU kernel design for accelerating graph neural networks training', *Proceedings of the 29th ASPLOS*, 2, pp. 248–264.

Oskarsson, J., Landelius, T., Deisenroth, M.P. and Lindsten, F. (2024) Probabilistic weather forecasting with hierarchical graph neural networks. arXiv. Available at: https://doi.org/10.48550/arXiv.2406.04759

Chen, J. et al. (2024) 'DeepPrim: physics-informed 3D deep weather forecaster via primitive equation learning', *OpenReview*. Available at: https://openreview.net/forum?id=EyyWd0hH0q.

Park, J.B., Mailthody, V.S., Qureshi, Z. and Hwu, W.-m. (2024) Accelerating sampling and aggregation operations in GNN frameworks with GPU initiated direct storage accesses. Proceedings of the VLDB Endowment, 17(6), pp. 1227–1240.