
Leveraging machine learning for improved numerical weather prediction through large-scale spectral nudging

Syed Zahid Husain^{*1}, Leo Separovic¹, Jean-François Caron¹, Christopher Subich¹, Rabah Aider¹, Mark Buehner¹, Stéphane Chamberland¹, Ervig Lapalme¹, Ron McTaggart-Cowan¹, Ayrton Zadra Moraes¹, Paul Vaillancourt¹, and Jing Yang¹

¹Environment and Climate Change Canada – Canada

Abstract

The recent emergence of machine learning-based weather prediction (MLWP) models has begun to disrupt the existing landscape of operational meteorological forecasting, which has long been dominated by physics-based numerical weather prediction (NWP) models. Various MLWP models have been shown to outperform state-of-the-art NWP models in terms of accuracy and computational speed. However, most deterministic MLWP models for medium-range forecasting suffer from major limitations, including low effective resolution and a narrow range of predicted variables, which prevent these models from completely replacing their NWP counterparts for operational weather forecasting.

This study examines the relative strengths and weaknesses of these two competing forecasting paradigms using two models from Environment and Climate Change Canada (ECCC): GEM (Global Environmental Multiscale) and ECCC-GraphCast, which represent the physics-based and machine learning-based approaches, respectively. The dynamical core of the GEM model solves the elastic Euler equations over a global grid based on the Yin-Yang system, whereas ECCC-GraphCast employs a graph neural network with an encoder-processor-decoder architecture and is based on open-source code from Google DeepMind. ECCC-GraphCast is trained by ECCC using ERA5 reanalyses covering 1979–2015, and is later fine-tuned with more recent analyses (2016–2021) from the European Centre for Medium-Range Weather Forecasts.

By analyzing global predictions from GEM and ECCC-GraphCast against observations and analyses in both physical and spectral spaces, this study demonstrates that large-scale predictions from ECCC-GraphCast outperform those from GEM, particularly for longer lead times. Building on this insight, a hybrid NWP-MLWP system is developed (Husain et al. 2024), wherein GEM-predicted large-scale state variables (temperature and horizontal wind) are spectrally nudged toward GraphCast inferences, while GEM is allowed to freely generate fine-scale details, which are often critical for weather extremes.

Results from different verifications reveal that this hybrid approach effectively leverages the strengths of ECCC-GraphCast to improve the prediction skill of the GEM model. For example, the root-mean-square error (RMSE) of the 500-hPa geopotential height is reduced by 5–10%, with an overall predictability improvement of 12–18 hours, peaking at day 7 of a

^{*}Speaker

10-day forecast. Notably, the trajectories of tropical cyclones are predicted with increased accuracy without any significant reduction in intensity, addressing a well-documented weakness of deterministic MLWP models. In addition to improving forecasting skill, this new hybrid system ensures that meteorologists have access to all forecast variables, including those relevant for high-impact weather.

At present, this hybrid system—currently in the process of becoming operational at ECCO—is limited to nudging scales 2,500km and larger. Nudging only large scales provides limited benefits at shorter lead times (e.g. up to day 3) because there is little divergence between predictions, but nudging fine scales is unfeasible because the GraphCast predictions are excessively smooth at these scales. This fine-scale smoothing issue, common in most deterministic MLWP models, primarily results from the use of mean squared error (MSE) as the loss function, which leads these models to learn to attenuate fine-scale features to avoid the so-called “double penalty” effect. To address this limitation, a parameter-free alternative loss function, namely adjusted MSE (AMSE), has recently been devised at ECCO (Subich et al. 2025). AMSE is designed to enable a machine-learning model to retain fine-scale spectral variance while maintaining predictive skill comparable to the MSE-based version of the model. As a result, the use of AMSE substantially improves the effective resolution of ECCO-GraphCast and, among other improvements pertaining to weather extremes, fixes the weak intensity bias in tropical cyclone predictions. In the future, ECCO-GraphCast, trained and fine-tuned to minimize AMSE, is expected to allow for nudging at scales as small as 500 km within the hybrid NWP-MLWP system, potentially enhancing its predictive accuracy for shorter lead times (the first 2–3 days).

References:

Husain, S.Z., Separovic, L., Caron, J.-F., Aider, R., Buehner, M., Chamberland, S., Lapalme, E., McTaggart-Cowan, R., Subich, C., Vaillancourt, P., Yang, J., and Moraes, A.Z. (2024): Leveraging data-driven weather models for improving numerical weather prediction skill through large-scale spectral nudging. arXiv:2407.06100.

Subich, C., Husain, S.Z., Separovic, L., and Yang, J. (2025): Fixing the double penalty in data-driven weather forecasting through a modified spherical harmonic loss function. arXiv:2501.19374.

Keywords: Spectral nudging, Numerical weather prediction, Machine learning, based weather prediction