Fixing the Double Penalty in Data-Driven Weather Forecasting Through a Modified Spherical Harmonic Loss Function

Christopher Subich¹ Syed Zahid Husain¹ Leo Separovic¹ Jing Yang¹

Abstract

Recent advancements in data-driven weather forecasting models have delivered deterministic models that outperform the leading operational forecast systems based on traditional, physics-based models. However, these data-driven models are typically trained with a mean squared error loss function, which causes smoothing of fine scales through a "double penalty" effect. We develop a simple, parameter-free modification to this loss function that avoids this problem by separating the loss attributable to decorrelation from the loss attributable to spectral amplitude errors. Finetuning the GraphCast model with this new loss function results in sharp deterministic weather forecasts, an increase of the model's effective resolution from 1,250 km to 160 km, improvements to ensemble spread, and improvements to predictions of tropical cyclone strength and surface wind extremes.

1. Introduction

The models developed in Weyn et al. (2020) and Keisler (2022) suggested that deep neural networks might "solve" the problem of medium-range weather forecasting with data-driven machine learning models. In 2023, the release of GraphCast (Lam et al., 2023), FourCastNet (Kurth et al., 2023), and Pangu-Weather (Bi et al., 2023) demonstrated forecast skill that met or surpassed that of the high-resolution forecast system (IFS) of the European Centre for Medium Range Weather Forecasts (ECMWF) at lead times (forecast lengths) up to 10 days, and some commenters (Bauer, 2024) anticipated that data-driven forecasting would soon supplant traditional numerical weather prediction (NWP) in all operational contexts. Since the

publication of these models, the field has been joined by many others, including the Artificial Intelligence Forecasting System (AIFS) developed by ECMWF itself (Lang et al., 2024a).

From the standpoint of machine learning, atmospheric forecasting is a large-scale generative problem comparable to predicting the next frame of a video. As a typical example, the version of the GraphCast model deployed experimentally by the National Oceanic and Atmospheric Administration (NOAA) (Sadeghi Tabas et al., 2025; NOAA, 2024) predicts the 6-hour forecast for six atmospheric variables at each of 13 vertical levels plus five surface variables, on a ¹/4° latitude/longitude grid, for about 86 million output degrees of freedom in aggregate. GraphCast takes two time-levels as input, so the input for this model has about 170 million degrees of freedom.

These first-generation data-driven weather models generally act as deterministic forecast systems, where each unique initial condition is mapped to a single forecast and verified against a "ground truth" from a data analysis system. The ERA5 atmospheric reanalysis (Hersbach et al., 2020) of ECMWF is most often used as the source of initial and verifying data for these forecast systems owing to its high quality and consistent behaviour from 1979 to present.

1.1. The Problem of Forecast Smoothing

Despite their overall forecast skill, deterministic data-driven forecast systems are universally understood to produce overly-smooth forecasts. A typical example of this behaviour is shown in figure 1 where a 3.5-day prediction of winter storm Eunice by the 13-level, ¹/₄° GraphCast model is too weak and overly smooth. This smoothing results in an under-prediction of localized extreme events, and it makes the model less suitable for downstream tasks such as spectral nudging (Husain et al., 2024) and data assimilation (Slivinski et al., 2025).

This smoothing is most-discussed in relation to the prediction of gridded, global weather fields, but it is still present in models that have radically different architectures. Allen et al. (2025) develops a model that operates directly in observation space without an underlying grid that still produces

¹Meteorological Research Division, Environment and Climate Change Canada, Dorval, Quebec, Canada. Correspondence to: Christopher Subich <christopher.subich@ec.gc.ca>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).



Figure 1. 10 m wind speed and mean sea level pressure for winter storm Eunice, 18 Feb 2022 at 0 h UTC. Top: HRES data at $\frac{1}{4}^{\circ}$ (ground truth), middle: 3.5d forecast produced by GraphCast, bottom: this work. This work produces an overall sharper forecast, with a better prediction of the winter storm's strength.

smooth forecasts of the future, and Han et al. (2024) shows diminished forecast activity (a bulk measure related to blurring) at longer lead times for a local-area model despite a nominal kilometer-scale resolution.

The conventional wisdom is that this smoothing is something that can be fixed in the context of an ensemble forecasting system, which produces realizations from the space of potential future forecasts. GenCast (Price et al., 2025) and AIFS-CRPS (Lang et al., 2024b) directly produce a stochastic ¹⁄4° forecast given initial values and a source of random noise. SEEDS (Li et al., 2024) and ArchesWeatherGen (Couairon et al., 2024) are examples of models that predict variations around an ensemble mean, using the generative step to "fill in the blanks" around a smooth baseline. Lippe et al. (2023) approaches this problem from a more general partial differential equation framework, and it develops a diffusion method that iteratively refines finer scales.

Of these examples, all but AIFS-CRPS use a diffusion technique with mean squared error (MSE) used as the de-noising loss function, while AIFS-CRPS instead uses the continuous ranked probability score (CRPS, Gneiting & Raftery (2007)) as its loss function to directly optimize the spread/error relationship of its produced ensemble.

However, we think that the problem of generating a good ensemble is distinct from the problem of forecast sharpness and effective resolution. Traditional NWP systems try to directly model the physics of the atmosphere, such that the system's forecasts are always plausible atmospheric states without excessive smoothing. Turning such a system into an ensemble prediction system involves supplying it with perturbed initial conditions and possibly stochastically perturbing the model's sub-grid parameterizations (Palmer, 2001; Berner et al., 2015).

In the machine learning space, Mahesh et al. (2024) develops a well-calibrated large ensemble using 29 independentlytrained instantiations of the Bonev et al. (2023) architecture. When combined with initial-condition perturbations, the result was a well-calibrated large ensemble, despite each individual ensemble member suffering from the smoothness typical of deterministic data-driven forecast systems.

Lagerquist & Ebert-Uphoff (2022) also develops a variety of loss functions based on the same spatial methods (such as filtering and max-pooling) to verify forecasts of convective events like thunderstorms in evaluation of high-resolution, limited-area models.

NEURALGCM

NeuralGCM (Kochkov et al., 2024) is one of the few global data-driven models that has addressed the problem of smoothing even in deterministic (non-ensemble) configurations. However, this model is difficult to compare with its peers. It has a hybrid architecture, combining a classical dynamical core with a learned network for sub-grid parameterizations that acts independently at each vertical column, and the classical dynamical core should cause fine-scale features to develop naturally. In addition, the model was trained using a weighted sum of several loss functions, one of which uses MSE only on a coarsened (smoothed) version of the forecast and verifying analysis while another matches the spherical harmonic power spectrum (but not phase) only at high wavenumbers (short scales). It is not clear which of these properties are necessary or sufficient to reduce the smoothing of deterministic NeuralGCM forecasts, and the use of several loss functions adds many degrees of freedom in their weighting and internal filtering.

1.2. This Work

The purpose of this work is to tackle the problem of smoothing in a purely deterministic, data-driven setting: can we produce a sharp forecast of the atmosphere without directly modelling ensemble uncertainty? Our answer is "yes." By modifying the MSE loss function to smoothly interpolate between amplitude-preservation and classical MSE, we can efficiently fine-tune a version of the GraphCast model to fix its smoothing problem and reproduce sharp forecasts. This greatly increases the model's effective resolution, producing better predictions of tropical cyclone intensity and surface wind speed. Section 2 describes the modified loss function, its theory of operation, and the fine-tuning procedure used for this work. Section 3 presents verification results of the finetuned GraphCast model, and section 4 concludes with discussion of the method's limitations and potential extensions. Appendix A discusses the loss function in the context of maximum likelihood estimation, and appendix B presents more detailed verification statistics.

2. Method

2.1. Smoothing Is Optimal Under Mean Squared Error

In the NWP community, model evaluation using the mean squared error is widely understood to suffer from a so-called "double penalty" (Hoffman et al., 1995; Ebert et al., 2013). Under MSE, a good forecast that correctly predicts a feature such as a storm but misses its location is penalized twice compared to a perfect forecast, once for missing the storm at its correct location and again for predicting a storm at an incorrect location. In traditional NWP, this double penalty makes model verification more difficult, particularly when studying the impact of improvements to forecast resolution that create more opportunities for misplaced predictions.

When MSE is used as the loss function to train a data-driven model, the double penalty problem is more than annoyance: it encourages the model to generate unrealistically smooth predictions by reducing the amplitude of unpredictable scales. To show this quantitatively, consider the case of predicting a single variable. Let $Y = \mathcal{N}(0, 1)$ be the target, and let X be the imperfect prediction of that target, modelled as a normal random variable with a standard deviation of $\sigma_X = \sqrt{\mathbb{E}(X^2)}$ and correlation coefficient of $\rho = \mathbb{E}(XY)/\sigma_X$, where $\mathbb{E}(\cdot)$ is the expectation operator. Writing X in terms of a correlated and an uncorrelated component gives:

$$X = \sigma_X(\rho Y + \sqrt{1 - \rho^2} \mathcal{N}(0, 1)),$$
(1)

and the corresponding expected MSE is:

$$\mathbb{E}\left(\mathrm{MSE}(X,Y)\right) = \mathbb{E}((X-Y)^2)$$
$$= \mathbb{E}(X^2) + \mathbb{E}(Y^2) - 2\mathbb{E}(XY)$$
$$= \sigma_X^2 + 1 - 2\sigma_X\rho. \tag{2}$$

For fixed Y, this MSE is optimized with a perfect prediction, when $\sigma_X = 1$ and $\rho = 1$. However, if $0 < \rho < 1$ because the process is only partially predictable, the MSE is optimized with respect to σ_X when $\sigma_X = \rho < 1$, leading to an underprediction of the process's natural variability.

2.2. Spectral Separation of the Mean Squared Error

Predictions of global weather are high-dimensional, but equations (1) and (2) can be extended to any decomposition (partition of unity) of the prediction and target fields that obeys Parseval's theorem. Taking this decomposition point-by-point, extending the analysis to include a nonzero mean, and taking the expectation over an ensemble of predictions gives rise to skill/spread evaluations. However, this decomposition is not possible at training time for a deterministic data-driven weather forecast, and instead we turn to a spherical harmonic decomposition.

Let $Y_k^l(i, j)$ be the complex-valued spherical harmonic mode with total wavenumber k and zonal wavenumber l at the (i, j) grid point on a latitude/longitude grid, normalized such that $\int Y_k^l (Y_m^n)^* = \delta_{km} \delta_{ln}$, where $(\cdot)^*$ is the complex conjugate¹. A scalar field x(i, j) defined on the latitude/longitude grid can be written in terms of spherical harmonics as:

$$x(i,j) = \sum_{k} \sum_{l=-k}^{k} \alpha_x(k,l) \, \mathbf{Y}_k^l(i,j),$$

with $\alpha_x(k, l)$ the corresponding spectral coefficient. For two fields x and y the latitude-weighted MSE is:

$$MSE(x, y) = \sum_{i} \sum_{j} dA(i, j) (x(i, j) - y(i, j))^{2}$$
$$= \sum_{k} \sum_{l=-k}^{k} |\alpha_{x}(k, l) - \alpha_{y}(k, l)|^{2}, \qquad (3)$$

where the dA term is incorporated into the normalization of Y_k^l . Importantly, α_x and α_y are independent with respect to zonal and total wavenumber, but the double summation here now allows us to group these terms in a physically meaningful way. Grouping terms in the inner (zonal) sum together gives rise to the power spectral density $PSD_k(x) = \sum_l |\alpha_x(k,l)|^2$ and coherence $Coh_k(x,y) =$ $\sum_l \Re(\alpha_x(k,l)\alpha_y^*(k,l))/\sqrt{PSD_k(x)PSD_k(y)}$ (where $\Re(\cdot)$ takes the real part) as scale-dependent analogs to variance and correlation respectively. Performing the appropriate substitutions:

$$MSE(x, y) = \sum_{k} PSD_{k}(x) + PSD_{k}(y) - 2\sqrt{PSD_{k}(x)PSD_{k}(y)} \operatorname{Coh}_{k}(x, y).$$
(4)

If x is taken to be a forecast field and y is the ground-truth analysis, as in (2) this is minimized when $\sqrt{PSD_k(x)PSD_k(y)^{-1}} = \operatorname{Coh}_k(x, y)$

This optimum leads to the observed smoothing in datadriven models through two factors:

• Fine scales (large k, short wavelengths) are generally less predictable than coarse scales (small k, large wavelengths), particularly at longer lead times, and

¹In practice, this work takes advantage of the property that $Y_{k}^{-l} = (Y_{k}^{l})^{*}$ to work with only non-negative wavenumbers.



Figure 2. Amplitude ratio (solid) and coherence (dashed) for the spherical harmonic mode with total wavenumber 100 for temperature at 850hPa during the training of a 1° version of the GraphCast model with an MSE loss function. At top, values for 6h lead time during the single-step pre-training phase and at bottom, values for 6h–72h during the forecast rollout (batches 300,000–311,000, incrementing one step every 1,000 batches).

 Data-driven models with conventional architectures learn to smooth fine scales (reducing the power spectral density) more quickly than they learn to predict them (increasing coherence).

This is illustrated in figure 2, which shows the amplitude ratio (square root of power spectral density ratio) and coherence for total wavenumber 100 (wavelength about 400 km) between predictions of the temperature field at the 850hPa level and the ground truth, for a 1° version of GraphCast during training with the curriculum of Lam et al. (2023). After a rapid adjustment from initially random outputs, the amplitude ratio and coherence closely track each other, with an initial smoothing followed by a gradual but partial sharpening as the model learns to predict this scale (with increasing coherence). When training is extended autoregressively to 12 steps (72h forecasts), smoothing increases at longer lead times as the forecast length increases.

2.3. Spectrally Adjusted Mean Squared Error

This smoothing is undesirable. It makes the produced forecasts less realistic, and it complicates model comparisons. Lam et al. (2023) performs extensive verification under an "optimal blurring" model to show that the purported forecast power of GraphCast is not just an artifact of its smoothing, and more straightforward verification methodologies such as that of Rasp et al. (2024) may conflate the effects of more-optimal smoothing with forecast skill even when evaluating at reduced resolution. It would instead be far more desirable if the loss function reflected our true goal, encouraging forecasts to correlate well to the ground-truth and retain realistic variation at finer scales.

Fortunately, beginning with MSE written in terms of its spectral decomposition, this is a simple modification. First, we write (4) in terms of a perfectly-correlated loss (with $Coh_k(x, y) = 1$) and a residual:

$$MSE(x,y) = \sum_{k} (\sqrt{PSD_k(x)} - \sqrt{PSD_k(y)})^2 + 2\sqrt{PSD_k(x)PSD_k(y)}(1 - Coh_k(x,y)).$$
(5)

Then, we seek to break the interaction between the spectral amplitudes and coherence contained in the second term of (5). One option would be to fix the role of x as a trial prediction and y as the verifying analysis and replace $\sqrt{\text{PSD}_k(x) \text{PSD}_k(y)}$ by $\text{PSD}_k(y)$, but the symmetry of the loss function can be retained by writing:

$$AMSE(x, y) = \sum_{k} (\sqrt{PSD_k(x)} - \sqrt{PSD_k(y)})^2 + 2\max(PSD_k(x), PSD_k(y))(1 - Coh_k(x, y)).$$
(6)

AMSE is now an adjusted mean squared error, which can act as a drop-in replacement during model training. Like its unmodified counterpart, AMSE is zero if and only if x = y, and it has the same Taylor expansion (in x) about x = y. The gradients of -AMSE(x, y) with respect to x (that is, minimizing AMSE) will always point in the direction of increased coherence ($Coh_k(x, y) \rightarrow 1$) and a correct spectral magnitude ($PSD_k(x) \rightarrow PSD_k(y)$), even if physical limits to predictability impose a practical limit to coherence. AMSE retains the units of MSE and has a similar magnitude, but it is no longer a proper metric because it does not satisfy the triangle inequality.

Unlike the mix of filtered and spectral loss functions used by NeuralGCM, (6) is parameter-free, requiring no selection of cutoff scales or scaled addition of qualitatively different terms. A parameter could be added to (6) to change the relative weights of its two terms, but that was not necessary in this work. Appendix A contemplates extending this framework to maximum likelihood estimation.

Equation 6 is defined for a single two-dimensional variable, but GraphCast produces several outputs per gridpoint. In the ¹/4°, 13-level version of the model considered here, there are six variables (geopotential, temperature, specific humidity, two components of horizontal wind, and vertical wind) produced at each of 13 atmospheric levels plus five variables (2-meter temperature, two components of 10-m horizontal wind, mean sea level pressure, and 6h-accumulated precipitation) at the surface. This work follows equation (A.19) of Lam et al. (2023) by aggregating each variable's error (MSE there, AMSE here) with a per-variable weight, level weighting proportional to the pressure level, and normalization Table 1. Fine-tuning curriculum for the ¹/₄°/13-level version of GraphCast trained for this study, including the peak and terminal learning rates (LR) of the cosine annealing schedule used at each stage. The batch size was 8 throughout, and each stage had a warm-up period of 64 batches.

Length	Batches	Peak/End LR	GPU Time
1 step (6h)	25,000	$2.5 \cdot 10^{-5} / 1.25 \cdot 10^{-7}$	7.7d
2 steps (12h)	2,500	$2.5 \cdot 10^{-6} / 7.5 \cdot 10^{-8}$	2.2d
4 steps (24h)	2,500	$2.5 \cdot 10^{-6} / 7.5 \cdot 10^{-8}$	4.3d
8 steps (48h)	1,250	$2.5 \cdot 10^{-6} / 7.5 \cdot 10^{-8}$	4.6d
12 steps (72h)	1,250	$2.5 \cdot 10^{-6} / 7.5 \cdot 10^{-8}$	7.4d

of the disparate units by a per-variable, per-level standard deviation.

2.4. Fine-Tuning Methodology

We demonstrate the efficacy of this loss function using a $\frac{1}{4}^{\circ}$, 13-level version of GraphCast. Based on the observation above that the model tends to rapidly adjust its per-scale smoothing to match its coherence, we treat this as a fine-tuning process and begin with the "operational" checkpoint provided by Lam et al. (2024), which is publicly available under a Creative Commons license.

Our fine-tuning methodology is summarized in table 1, and the overall approach is inspired by Subich (2024). While the baseline model checkpoint was trained over 72h (12 autoregressive steps of 6h each), in earlier testing at 1° we found it better to begin the fine-tuning with single-step forecasts and increase the forecast length in stages. Training over single steps is both faster per step and supports higher learning rates.

The other training hyperparameters, including AdamW (Loshchilov & Hutter, 2019) hyperparameter settings and per-variable, per-level loss weightings were identical to those described by Lam et al. (2023).

The 13-level GraphCast checkpoint that forms the base of our fine-tuned model was originally trained on the ERA5 reanalysis from 1979–2017, then itself fine-tuned on the initial conditions used for the contemporaneous HRES (IFS) model from ECMWF over 2016–2021. We used this latter dataset and training period in our work, and it is available from Rasp et al. (2024) as the "HRES-fc0" dataset. As described in Lam et al. (2023), we supplemented the HRES data with the accumulated precipitation field from the ERA5 reanalysis over the training period, since an initial conditions dataset has no accumulated precipitation by definition. The equivalent data for calendar year 2022 is also available, and we used this period for model evaluation.

We fine-tuned our model on 1-2 nodes of a cluster con-



Figure 3. Amplitude ratio (solid) and coherence (dashed) for all output variables and levels, weighted using the variable/level weights in the loss function, for the control model and this work after the 1-step training and after complete fine-tuning. Top: 6h lead time, middle: 120h (5d) lead time, bottom: 240h (10d) lead time. The dashed line is placed where a model would underrepresent the power spectral density by 25%.

taining 4 NVidia A100 40GiB GPUs per node, using data parallelism with the batch split across GPUs and the gradients accumulated via MPI. Overall, the fine-tuning process took about 26.2 GPU-days. Lam et al. (2023) does not disclose the total training time required to produce the model checkpoint from scratch, but other models of similar size (Bonev et al., 2023; Lang et al., 2024a; Bi et al., 2023) report training times of about 1 GPU-year using similar hardware.

3. Results

The fine-tuned model is evaluated against the control (unmodified) model over calendar year 2022 using the HRES dataset for initialization and as ground truth unless otherwise specified. As reported in Lam et al. (2023) and is typical in other deterministic data-driven models, forecast performance at longer lead times improves when the model is autoregressively trained over multiple steps, and the fullytuned model (trained over 12 forecast steps and labelled "AMSE AR12" in the figures and discussion below) is considered the primary model for evaluation.

Since multi-step training also tends to cause both fine-scale smoothing and a loss of variability in ensemble settings, these respective evaluations (sections 3.1 and 3.2) will also include the model checkpoint created after just single-step

fine-tuning, denoted "AMSE AR1."

3.1. Effective Resolution

Conventional, physics-based NWP models are widely understood to have an effective resolution that is coarser than the model's native grid resolution. Limits to effective resolution come from the limited fidelity of spatial or temporal discretization, from artificial diffusion or damping used to stabilize a model, and from sub-grid processes (such as turbulence) that must be imperfectly estimated rather than directly modelled. A model behaves unrealistically at scales finer than its effective resolution, typically providing insufficient variability and too-smooth solutions.

Deterministic data-driven NWP models do not have the same underlying numerical issues that result in reductions to effective resolution, but the smoothing produced by training with an MSE-based loss function acts in a very similar way. Figure 3 shows the amplitude ratio $(\sqrt{\text{PSD}_k(x) \text{PSD}_k^{-1}(y)})$ and coherence $(\text{Coh}_k(x, y))$ between each of the GraphCast models and the verifying analysis over calendar year 2022. To compute a combined curve despite the many per-gridpoint values predicted by the model, the statistics for each separate variable are combined using the same variable and level weighting used in the model's loss function².

The control model significantly smooths fine scales even after a single 6-hour forecast step, and that smoothing increases with the forecast lead time. If we somewhat arbitrarily draw the line of effective resolution at the point where the model has lost 25% of the per-wavenumber energy (corresponding to a ratio of power spectral densities of 0.75 or an amplitude ratio of $\sqrt{0.75}$), the 5-day predictions of the control model reach that cutoff at wavenumber 32, corresponding to oscillations with a wavelength of about 1250 km. Small changes in the target amplitude ratio will result in small changes to the derived effective resolution.

The models fine-tuned in this work do not show this type of fine-scale dissipation. The AMSE AR12 model has a small amount of smoothing at moderate scales, but the variability recovers again at finer scales, and a dissipation-based definition of effective resolution would be extremely sensitive to the cutoff value. Instead, we observe that for longer forecasts the model has more energy at small scales than in the ground-truth dataset, suggesting a "noise-based" definition of effective resolution. For long forecasts, the amplitude ratio rises above 1 around wavenumber 250, giving an effective resolution of about 160 km.

The AMSE AR1 model shows the same qualitative behaviour but generates this "noise" more strongly, leading to a reduced effective resolution of about 450 km (wavenumber 90). The forecasts produced by this version of the model are less coherent with the analysis, showing a reduced forecast skill at all scales for longer forecasts.

For illustration, appendix B.2 shows amplitude spectra for select variables at various lead times, without normalizing by the spectral magnitude of the ground truth. Appendix B.5 discusses the effective resolution of the model when trained with either mean squared error or mean absolute (L1) error.

3.2. Lagged Ensemble Verification

The observation that AMSE-based fine-tuning provides sharp forecasts is encouraging, but that alone is not enough to demonstrate utility. The model might have learned to match its expected variance by generating quasi-static noise that does not sufficiently depend on the surrounding flow, for example. The ideal way to measure this sort of forecast skill is in an ensemble setting, where the chaotic nature of the atmosphere is accounted for by evaluating the full distribution of plausible outputs given an initial condition.

Development of a full ensemble system is well beyond the scope of this work, but Brenowitz et al. (2025) provides a procedure to evaluate a deterministic model using an ensemble generated from time-separated initial conditions. The central idea of this method is that predictions initialized at different times should diverge, so several consecutively-initialized forecasts that are all valid at a shared time form an ad-hoc ensemble, without the need for an auxiliary method of defining an ensemble of initial conditions.

This approach is implemented here, using forecasts initialized at 12-hourly intervals in 2022 and evaluated from 10 January 2022 0:00 UTC to 31 December 2022 12:00 UTC. Each set of nine consecutively initialized forecasts (spanning four days from beginning to end) forms an ensemble, and the ensemble's notional lead time is that of its central member.

The primary evaluation metrics are the CRPS, ensemble root mean squared error (eRMSE), and spread/error ratio, with definitions given in appendix B.4. For an operational ensemble, a spread/error ratio close to 1 is considered ideal, but that is confounded here because the members of a lagged ensemble are not statistically interchangeable. Since deterministic data-driven NWP models are underdispersive, however, a larger spread/error ratio is generally better.

Figure 4 shows the evolution of these statistics versus lead time for a selection of variables and levels, and more detailed evaluation of CRPS and eRMSE are shown in figures 11 and 12. The AMSE AR12 model shows consistent improve-

²Normalization of the disparate variables by standard deviation was not required here, since the amplitude ratio and coherence are already dimensionless.



Figure 4. Lagged ensemble statistics for geopotential (z) at 500hPa, temperature (t) at 850hPa, specific humidity (q) at 700hPa, and 2-meter temperature (2t) from left to right. The statistics are the CRPS, root mean squared error of the ensemble mean, and spread-error ratio, from top to bottom.

ments to the CRPS while the eRMSE sees little change, indicating that the fine tuning process produces a bettercalibrated (more dispersive) ensemble without degrading overall predictive performance.

While the AMSE AR1 model shows greater ensemble spread, the less skillful forecast results in a significantly reduced CRPS. However, unlike the results of Brenowitz et al. (2025), the spread/error ratio of the AR1 and AR12 models converge for most variables at longer lead times, suggesting that multi-step training in this framework does not cause a collapse of variability in an ensemble setting.

3.3. Hurricane Prediction and Extreme Weather

The effect of improved effective resolution is most strongly apparent in the prediction of local extremes, and few weather events are more extreme than tropical cyclones.

Data-driven NWP models like GraphCast improve predictions of hurricane tracks relative to conventional NWP models (see for example figure 3A of Lam et al. (2023)). Since storms are guided by large-scale "steering flows" that have natural scales of thousands of kilometers, these predictions of storm position are relatively unaffected by the models' limited effective resolutions but benefit from improvements in large-scale forecast skill. However, cyclones themselves are comparatively small, and predictions of the storm intensity are significantly affected by MSE-induced smoothing.

Figure 5 depicts this situation for Hurricane Ian, the most intense Atlantic tropical cyclone of the 2022 season. Both the control version of GraphCast and the AMSE AR12

version produce a reasonable 5-day prediction of the storm's location (within about 125 km), but the control version of GraphCast predicts an unrealistically weak storm.

More quantitatively, figure 6 shows the mean intensity and mean absolute position errors for tropical cyclones over 20 June–19 September 2022 initializations, using the algorithm of Zadra et al. (2014) to compare against the International Best Track Archive for Climate Stewardship database (Knapp et al., 2010). Compared to these observations, even the HRES data is imperfect and shows a weak-intensity bias. The control model has a larger weak-intensity bias that increases with lead time, but the AMSE AR12 model retains the quality of the HRES dataset. The storm location predictions between the control and AMSE AR12 models are equivalent.

Extreme weather includes more than tropical cyclones, and appendix B.3 discusses quantile-quantile predictions of surface wind speed and temperature, validated against station observations. Both the control model and AMSE AR12 produce realistic temperature extremes, but the AMSE AR12 model provides more realistic predictions of wind-speed extremes.

4. Discussion & Limitations

Using the mean squared error as a model loss function asks the model to average away unpredictable scales. In weather forecasting, the unpredictable scales are generally the smaller scales that carry information about local variance, and this averaging process leads to data-driven weather forecasts that are far smoother than the grid resolution would



Figure 5. 10 m wind speed and mean sea level pressure for Hurricane Ian, 28 Sept 2022 at 12h UTC. Top: HRES data at ¼°, middle: 5d forecast produced by the control GraphCast model, bottom: the model after 12-step fine-tuning with AMSE.

suggest.

This is not a property inherent to data-driven NWP. The alternate loss function based on (6) uses a spectral transform to separate the loss attributable to amplitude error from that attributable to decorrelation, encouraging the model to reproduce a realistic spectrum even if it can't make an accurate prediction. When applied to the $1/4^\circ$, 13-level version of GraphCast with an abbreviated fine-tuning process, we recover a model that has a much finer effective resolution, has improved CRPS-based verification in a lagged ensemble setting, and fixes the weak intensity bias in the prediction of tropical cyclones.

When fine-tuned autoregressively over multiple forecast steps, the model suffers from a small amount of smoothing at mesoscales (intermediate scales). We speculate that this is because such autoregressive training has two objectives: forecasts are asked both to be accurate (and thus sharp, per (6)) and to be good initial conditions for the next forecast step. This latter goal is implicit, and it is not directly affected by the loss function used in training. Future work will consider the use of a replay buffer in training (like that of Chen et al. (2023)) to see if long-range forecast skill might be retained with even better prediction of amplitudes.



Figure 6. Predictions of tropical cyclone intensity ((a), mean maximum surface wind speed; (b) mean minimum central pressure) and mean absolute position error (c) for forecasts initialized 20 June–19 September 2022. Orange squares show statistically significant differences between the AMSE AR12 and control predictions.

Since the AMSE loss function (6) is zero if and only if the predicted field matches the ground truth, it may be useful throughout model training rather than just during a finetuning pass. However, a thorough test of this proposition would require a considerable computational budget, so it is left for future work. Use of the AMSE loss function throughout the training process might improve the coherence of fine-scale prediction by allowing the model to spend more of its training time "seeing" these modes, but on the other hand the coherence-dependent smoothing encouraged by the MSE loss function (figure 2) might act as an implicit regularization that smooths the model's gradients and speeds up training overall.

4.1. Effective Resolution

The ultimate conclusion of this work is that the AMSEbased error measure improves the effective resolution of NWP weather models, but the phrase "effective resolution" must always be accompanied by the question, "effective at what?"

We chose to define an effective resolution based on smoothing of fine scales, since a model that simply doesn't represent a scale cannot effectively model it. However, other definitions exist in the literature, and users of these models should keep their ultimate goals in mind. For example, Kent et al. (2014) studies various discretization schemes for numerical partial differential equations under both diffusion (smoothing) and dispersion (wave propagation) definitions of effective resolution.

4.2. Alternative Grids

Passing from equation (2) to (4) makes use of Parseval's theorem to give an exact relationship between the spatially-defined mean squared error and the equivalent in the spectral representation. Implementing this in a training cycle requires fast computation of spherical harmonic transforms. This is simple enough for global latitude/longitude grids, but it might be difficult for local-area models without a regular global grid structure.

In these cases, we think that the basic intuition behind (6) might still apply through other multiscale decompositions such as wavelet lifting (Sweldens & Schröder, 2000), provided suitable equivalents to scale-dependent variance and correlation could be found. The multiscale decomposition is critical in some form, however, since the method takes advantage of the approximate independence of scale-separated modes. Without such a decomposition (e.g. applying the adjustment of (6) globally, without the harmonic transform), the model might be able to "cover up" a lack of fine-scale variability by over-emphasizing coarser scales.

4.3. Applications to Other Domains

AMSE is a natural error function for weather prediction because the spectral decomposition is physically meaningful and relatively stable over time. A partially incorrect but realistic prediction of weather at 2000 km scales would not significantly change the amount of energy present at 100 km scales, just its relative location. The goal of a deterministic forecast is to be physically plausible, and a correct prediction of spectral amplitudes is a necessary condition for physical plausibility.

The method can be mechanically applied whenever a spectral decomposition is possible, but additional value is only likely when a sub-aggregation of that spectrum is meaningful. This is most obviously possible in other areas of fluid dynamics, particularly the modelling of turbulent flows. In that domain, Chakraborty et al. (2025) developed a binned spectral loss function (on a planar domain) that is reminiscent of the amplitude-only component of (6), but it discards the phase information. We are optimistic that integrating the spectral correlation along the lines of AMSE will make such models more robust.

4.4. Applications to Ensemble Modelling

We are particularly encouraged by the beneficial impact that AMSE-based training has on the spread of forecasts in an ensemble setting. Without any dedicated ensemble-based training we end up with a model that nonetheless produces a more realistic spread of forecasts. In future work, we hope to use this loss function as a basis for an ensemble forecast where each individual ensemble member produces a realistic trajectory, in addition to the whole-ensemble optimization encouraged by CRPS-like ensemble training.

Code and Data Availability

An implementation of the AMSE error function and the code used to train GraphCast for this work are available at https://github.com/csubich/graphcast/tree/amse under the Apache 2.0 license. The fine-tuned checkpoints produced for this study are available at https://huggingface.co/csubich/graphcast_amse under the CC-BY-ND-SA 4.0 license, as derivative works of the DeepMind "graphcast-operational" checkpoint.

Acknowledgements

The authors would like to thank Charlie Hébert-Pinard, Vikram Khade, and Hugo Vandenbroucke-Menu of the Canadian Centre for Meteorological and Environmental Prediction for access to the 1°-trained version of GraphCast used to produce the results of figure 2.

Impact Statement

Accurate weather forecasts are a vital public service, and its benefits are disproportionately concentrated in the extremes. Accurate forecasts of extreme weather such as tropical cyclones save lives. On one hand, this means that we should be eager to develop improvements to weather forecasting systems, but on the other hand it means that we should be very careful not to just "chase scores," confusing what's easy to calculate with what's truly important.

This work contributes to this field by introducing a way to make data-driven weather forecasting more realistic, with variability at moderate and fine scales that is much closer to reality. This improves various probabilistic scores and predictions of tropical cyclone intensity, but this is not a guarantee of complete physical plausibility. In particular, we have not yet shown that these forecasts are better-behaved "out of distribution," such as when simulating possible future climate paths.

Operational weather centres are very diligent about performing rigorous evaluation of models before making them operational, and we hope that this work can help ease the path towards the adoption of better-performing, data-driven forecasting systems in the near future.

References

Allen, A., Markou, S., Tebbutt, W., Requeima, J., Bruinsma, W. P., Andersson, T. R., Herzog, M., Lane, N. D., Chantry, M., Hosking, J. S., and Turner, R. E. End-to-end datadriven weather prediction. *Nature*, pp. 1–8, March 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-08897-0.

- Bauer, P. What if? Numerical weather prediction at the crossroads. *Journal of the European Meteorological Society*, 1:100002, December 2024. ISSN 2950-6301. doi: 10.1016/j.jemets.2024.100002.
- Berner, J., Fossell, K. R., Ha, S.-Y., Hacker, J. P., and Snyder, C. Increasing the Skill of Probabilistic Forecasts: Understanding Performance Improvements from Model-Error Representations. *Monthly Weather Review*, 143(4): 1295–1320, April 2015. ISSN 1520-0493, 0027-0644.
- Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., and Tian, Q. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970):533– 538, July 2023. ISSN 1476-4687. doi: 10.1038/ s41586-023-06185-3.
- Bonev, B., Kurth, T., Hundt, C., Pathak, J., Baust, M., Kashinath, K., and Anandkumar, A. Spherical Fourier Neural Operators: Learning Stable Dynamics on the Sphere. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 2806–2823. PMLR, July 2023. ISSN: 2640-3498.
- Brenowitz, N. D., Cohen, Y., Pathak, J., Mahesh, A., Bonev, B., Kurth, T., Durran, D. R., Harrington, P., and Pritchard, M. S. A Practical Probabilistic Benchmark for AI Weather Models. *Geophysical Research Letters*, 52(7):e2024GL113656, 2025. ISSN 1944-8007. doi: 10.1029/2024GL113656.
- Chakraborty, D., Mohan, A. T., and Maulik, R. Binned Spectral Power Loss for Improved Prediction of Chaotic Systems, February 2025. URL http://arxiv.org/ abs/2502.00472. arXiv:2502.00472 [cs].
- Chen, K., Han, T., Gong, J., Bai, L., Ling, F., Luo, J.-J., Chen, X., Ma, L., Zhang, T., Su, R., Ci, Y., Li, B., Yang, X., and Ouyang, W. FengWu: Pushing the Skillful Global Medium-range Weather Forecast beyond 10 Days Lead, April 2023. URL http://arxiv.org/abs/2304. 02948. arXiv:2304.02948 [physics].
- Couairon, G., Singh, R., Charantonis, A., Lessig, C., and Monteleoni, C. ArchesWeather & ArchesWeatherGen: a deterministic and generative model for efficient ML weather forecasting, December 2024. URL http:// arxiv.org/abs/2412.12971. arXiv:2412.12971 [cs].
- Ebert, E., Wilson, L., Weigel, A., Mittermaier, M., Nurmi, P., Gill, P., Göber, M., Joslyn, S., Brown, B., Fowler, T., and Watkins, A. Progress and challenges in forecast verification. *Meteorological Applications*, 20(2):130–139, 2013. ISSN 1469-8080. doi: 10.1002/met.1392.

- Gneiting, T. and Raftery, A. E. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, March 2007. ISSN 0162-1459. doi: 10.1198/016214506000001437.
- Han, T., Guo, S., Ling, F., Chen, K., Gong, J., Luo, J., Gu, J., Dai, K., Ouyang, W., and Bai, L. FengWu-GHR: Learning the Kilometer-scale Medium-range Global Weather Forecasting, January 2024. URL http://arxiv. org/abs/2402.00059. arXiv:2402.00059 [physics].
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. ISSN 1477-870X. doi: 10.1002/qj.3803.
- Hoffman, R. N., Liu, Z., Louis, J.-F., and Grassoti, C. Distortion Representation of Forecast Errors. *Monthly Weather Review*, 123(9):2758–2770, September 1995. ISSN 1520-0493, 0027-0644. doi: 10.1175/1520-0493(1995) 123(2758:DROFE)2.0.CO;2.
- 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 Zadra, A. Leveraging data-driven weather models for improving numerical weather prediction skill through large-scale spectral nudging, July 2024. URL http://arxiv. org/abs/2407.06100. arXiv:2407.06100 [physics].
- Keisler, R. Forecasting Global Weather with Graph Neural Networks, February 2022. URL http://arxiv.org/abs/2202.07575. arXiv:2202.07575 [physics].
- Kent, J., Whitehead, J. P., Jablonowski, C., and Rood, R. B. Determining the effective resolution of advection schemes. Part I: Dispersion analysis. *Journal of Computational Physics*, 278:485–496, December 2014. ISSN 0021-9991. doi: 10.1016/j.jcp.2014.01.043.
- Knapp, K. R., Kruk, M. C., Levinson, D. H., Diamond, H. J., and Neumann, C. J. The International Best Track Archive for Climate Stewardship (IBTrACS). *Bulletin* of the American Meteorological Society, 91(3):363–376, March 2010. doi: 10.1175/2009BAMS2755.1.
- Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., Klöwer, M., Lottes, J., Rasp, S., Düben,

P., Hatfield, S., Battaglia, P., Sanchez-Gonzalez, A., Willson, M., Brenner, M. P., and Hoyer, S. Neural general circulation models for weather and climate. *Nature*, 632 (8027):1060–1066, August 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07744-y.

- Kurth, T., Subramanian, S., Harrington, P., Pathak, J., Mardani, M., Hall, D., Miele, A., Kashinath, K., and Anandkumar, A. FourCastNet: Accelerating Global High-Resolution Weather Forecasting Using Adaptive Fourier Neural Operators. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, PASC '23, pp. 1–11, New York, NY, USA, June 2023. Association for Computing Machinery. ISBN 9798400701900. doi: 10.1145/3592979.3593412.
- Lagerquist, R. and Ebert-Uphoff, I. Can We Integrate Spatial Verification Methods into Neural Network Loss Functions for Atmospheric Science? *Artificial Intelligence for the Earth Systems*, 1(4), November 2022. ISSN 2769-7525. doi: 10.1175/AIES-D-22-0021.1.
- Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., and Battaglia, P. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, December 2023. doi: 10.1126/science.adi2336.
- Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., and Battaglia, P. GraphCast GitHub repository, July 2024. URL https://github. com/google-deepmind/graphcast. originaldate: 2023-07-14T11:07:57Z.
- Lang, S., Alexe, M., Chantry, M., Dramsch, J., Pinault, F., Raoult, B., Clare, M. C. A., Lessig, C., Maier-Gerber, M., Magnusson, L., Bouallègue, Z. B., Nemesio, A. P., Dueben, P. D., Brown, A., Pappenberger, F., and Rabier, F. AIFS - ECMWF's data-driven forecasting system, June 2024a. URL http://arxiv.org/abs/2406. 01465. arXiv:2406.01465 [physics].
- Lang, S., Alexe, M., Clare, M. C. A., Roberts, C., Adewoyin, R., Bouallègue, Z. B., Chantry, M., Dramsch, J., Dueben, P. D., Hahner, S., Maciel, P., Prieto-Nemesio, A., O'Brien, C., Pinault, F., Polster, J., Raoult, B., Tietsche, S., and Leutbecher, M. AIFS-CRPS: Ensemble forecasting using a model trained with a loss function based on the Continuous Ranked Probability Score, December 2024b. URL http://arxiv.org/abs/2412.15832. arXiv:2412.15832 [physics] version: 1.

- Leutbecher, M. and Palmer, T. N. Ensemble forecasting. Journal of Computational Physics, 227(7):3515–3539, March 2008. ISSN 0021-9991. doi: 10.1016/j.jcp.2007. 02.014.
- Li, L., Carver, R., Lopez-Gomez, I., Sha, F., and Anderson, J. Generative emulation of weather forecast ensembles with diffusion models. *Science Advances*, 10(13):eadk4489, March 2024. doi: 10.1126/sciadv.adk4489.
- Lippe, P., Veeling, B., Perdikaris, P., Turner, R., and Brandstetter, J. PDE-Refiner: Achieving Accurate Long Rollouts with Neural PDE Solvers. *Advances in Neural Information Processing Systems*, 36:67398–67433, December 2023.
- Loshchilov, I. and Hutter, F. Decoupled Weight Decay Regularization, January 2019. URL http://arxiv.org/ abs/1711.05101. arXiv:1711.05101 [cs, math].
- Mahesh, A., Collins, W., Bonev, B., Brenowitz, N., Cohen, Y., Elms, J., Harrington, P., Kashinath, K., Kurth, T., North, J., OBrien, T., Pritchard, M., Pruitt, D., Risser, M., Subramanian, S., and Willard, J. Huge Ensembles Part I: Design of Ensemble Weather Forecasts using Spherical Fourier Neural Operators, August 2024. arXiv:2408.03100 [physics] version: 1.
- NOAA. GraphCast with GFS input, 2024. URL https://registry.opendata.aws/ noaa-nws-graphcastgfs-pds/.
- Palmer, T. N. A nonlinear dynamical perspective on model error: A proposal for non-local stochastic-dynamic parametrization in weather and climate prediction models. *Quarterly Journal of the Royal Meteorological Society*, 127(572):279–304, 2001. ISSN 1477-870X. doi: 10.1002/qj.49712757202.
- Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., Ewalds, T., Stott, J., Mohamed, S., Battaglia, P., Lam, R., and Willson, M. Probabilistic weather forecasting with machine learning. *Nature*, 637 (8044):84–90, January 2025. ISSN 1476-4687. doi: 10.1038/s41586-024-08252-9.
- Rasp, S., Hoyer, S., Merose, A., Langmore, I., Battaglia, P., Russell, T., Sanchez-Gonzalez, A., Yang, V., Carver, R., Agrawal, S., Chantry, M., Ben Bouallegue, Z., Dueben, P., Bromberg, C., Sisk, J., Barrington, L., Bell, A., and Sha, F. WeatherBench 2: A Benchmark for the Next Generation of Data-Driven Global Weather Models. *Journal of Advances in Modeling Earth Systems*, 16(6):e2023MS004019, 2024. ISSN 1942-2466. doi: 10.1029/2023MS004019.

- Sadeghi Tabas, S., Wang, J., Lei, W., Row, M., Zhang, Z., Zhu, L., Peng, J., and Carley, J. R. GFS-Powered Machine Learning Weather Prediction: A Comparative Study on Training GraphCast with NOAA's GDAS Data for Global Weather Forecasts, 2025. URL https://repository.library. noaa.gov/view/noaa/67485.
- Slivinski, L. C., Whitaker, J. S., Frolov, S., Smith, T. A., and Agarwal, N. Assimilating Observed Surface Pressure Into ML Weather Prediction Models. *Geophysical Research Letters*, 52(6):e2024GL114396, 2025. ISSN 1944-8007. doi: 10.1029/2024GL114396.
- Subich, C. Efficient fine-tuning of 37-level GraphCast with the Canadian global deterministic analysis, August 2024. URL http://arxiv.org/abs/2408. 14587. arXiv:2408.14587 [cs].
- Sweldens, W. and Schröder, P. Building your own wavelets at home. In Klees, R. and Haagmans, R. (eds.), *Wavelets in the Geosciences*, Lecture Notes in Earth Sciences, pp. 72–107. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000. ISBN 978-3-540-46590-4. doi: 10.1007/ BFb0011093.
- Tödter, J. and Ahrens, B. Generalization of the Ignorance Score: Continuous Ranked Version and Its Decomposition. *Monthly Weather Review*, 140(6):2005– 2017, June 2012. ISSN 1520-0493, 0027-0644. doi: 10.1175/MWR-D-11-00266.1.
- Weyn, J. A., Durran, D. R., and Caruana, R. Improving Data-Driven Global Weather Prediction Using Deep Convolutional Neural Networks on a Cubed Sphere. *Journal of Advances in Modeling Earth Systems*, 12 (9):e2020MS002109, 2020. ISSN 1942-2466. doi: 10.1029/2020MS002109.
- Zadra, A., McTaggart-Cowan, R., Vaillancourt, P. A., Roch, M., Bélair, S., and Leduc, A.-M. Evaluation of Tropical Cyclones in the Canadian Global Modeling System: Sensitivity to Moist Process Parameterization. *Monthly Weather Review*, 142(3):1197–1220, March 2014.

A. Relationship to Maximum Likelihood Estimation

In developing the AMSE loss function, the transformation from ordinary, gridpoint-based MSE (2) to its spectral definition with power spectral densities and coherence (5) is algebraic in nature. The beneficial effect of the AMSE loss function's separation of spectral-amplitude and decoherence terms arises because the underlying spectral decomposition is physically meaningful. At fine enough scales, atmospheric dynamics are increasingly rotationally symmetric and position-invariant, with individual spectral amplitudes that look like draws from a Gaussian distribution.

If we elevate this property from a fortunate coincidence to a simplifying assumption, we can treat the set of modes corresponding to a particular total wavenumber as random variables and apply the machinery of ensemble verification to individual, deterministic forecasts. The goal of producing realistic forecasts despite limited predictability is conceptually similar to the goal of maximum-likelihood estimation, so we consider here the effect of Kullback-Leibler (KL) divergence minimization. In the meteorology literature, the KL divergence is named the continuous ignorance score (Tödter & Ahrens, 2012), and it is sometimes used for ensemble verification.

Treat the modes corresponding to a single total wavenumber k as a draw from a 2k - 1-dimensional normal random variable³ with mean zero and some finite standard deviation. In this interpretation, the ground-truth analysis is:

$$Y_k = \sigma_Y \, \mathcal{N}^{2k-1}(0,1). \tag{7}$$

The forecast is itself taken to be a normal random variable, but following the pattern of (5) it is partially correlated to Y and has its own standard deviation. Take the correlation to be ρ and the forecast standard deviation to be σ_X , and:

$$X_k = \sigma_X \left(\frac{\rho}{\sigma_Y} Y + \sqrt{1 - \rho^2} \mathcal{N}^{2k-1}(0, 1) \right), \tag{8}$$

noting for emphasis that this definition of X depends upon Y. With the assumption that each of the per-wavenumber modes are independently drawn from this distribution, we can also treat X and Y as a product of 2k - 1 independent, scalar random variables, which will simplify the following algebra.

The KL divergence of the data given the forecast is then given by:

$$D_{KL}(Y||X) = \int P_Y(y') \log\left(\frac{P_Y(y')}{P_X(y')}\right) dy',$$
(9)

for the respective probability density functions (PDFs) P_Y and P_X and integrating over the space of possible observations parameterized by y'. With these formulations, the PDF of Y is simple:

$$P_Y(y') = (2\pi\sigma_Y^2)^{-1/2} \exp\left(-\frac{y'^2}{2\sigma_Y^2}\right).$$
(10)

The PDF of X is more complicated because of its dependence on Y, but for any individual observation y (8) becomes a shifted Gaussian, giving:

$$P_X(x|y) = (2\pi\sigma_X^2(1-\rho^2))^{-1/2} \exp\left(-\frac{(x-\rho\frac{\sigma_X}{\sigma_Y}y)^2}{2\sigma_X^2(1-\rho^2)}\right), \text{ or}$$

$$P_X(y|y) = (2\pi\sigma_X^2(1-\rho^2))^{-1/2} \exp\left(-\frac{(1-\rho\frac{\sigma_X}{\sigma_Y})^2y^2}{2\sigma_X^2(1-\rho^2)}\right).$$
(11)

(9) then becomes:

$$D_{\rm KL}(Y||X) = const(Y) - \int P_Y(y) \log(P_X(y)) dy$$

= $const(Y) + \int (2\pi\sigma_Y^2)^{-1/2} \exp\left(-\frac{y^2}{2\sigma_Y^2}\right) \left(\log(2\pi\sigma_X^2(1-\rho^2)) + \frac{(1-\rho\frac{\sigma_X}{\sigma_Y})^2 y^2}{2\sigma_X^2(1-\rho^2)}\right) dy$

³That is, k independent complex-valued modes from $1 \dots k$ with independent real and imaginary parts and a single, real zerowavenumber mode.



Figure 7. Visualization of high-pass filtered forecast and analysis fields for the forecast shown in 1.

$$= const(Y) + \log(\sigma_X^2(1-\rho^2)) + \frac{(\sigma_Y - \rho\sigma_X)^2}{2\sigma_X^2(1-\rho^2)}.$$
(12)

Minimizing (12) for σ_X while holding σ_Y and ρ fixed is complicated, but solved numerically the optimal standard deviation ratio $\sigma_X \sigma_Y^{-1}$ is less than unity, reaching a minimum of about 0.66 near $\rho = 0.4$ and increasing for both lower and higher values of correlation. This is less intuitive than the $\sigma_X = \sigma_Y$ optimum of (6), but it still would smooth fine scales much less than the $\sigma_X = \sigma_Y \rho$ optimum of (2).

Implementing (12) as a loss function would be conceptually interesting, but this seems impractical because the expression has singular behaviour near $\rho = 1$, where the implied random part of the prediction collapses to zero variance.

B. Supplemental Verification

B.1. Visualization

Figure 7 visualizes the high-wavenumber components of a sample forecast matching the winter storm Eunice prediction shown in figure 1. The applied filter fourth-order in spherical harmonic space, with the functional form:

$$HPF(k) = 1 - \frac{k_0^4}{k_0^4 + k^4},$$
(13)

where k is the total wavenumber and $k_0 = 50$ is the cutoff number, chosen to emphasize modes with length scales of 800 km and shorter. Overall, the predictions of the control and AMSE-trained models show very similar structures, but training with (6) as the loss function enhances the high-mode variability of the forecasts.



Figure 8. Amplitude spectral density for the variables of figure 4 at 6h, 120h, and 240h lead times.

B.2. Spectra

Figure 8 shows the amplitude spectral density (square root of power spectral density, with units proportional to $1/\sqrt{cycle}$) at moderate to fine scales for several variables and lead times. Because of the energy cascade in the atmosphere, the spectra of most variables follow power-law distributions. Energy in the atmosphere is ultimately removed by turbulent, frictional dissipation, but no practical global atmospheric model can effectively resolve these scales.

Nonetheless, the available energy per total wavenumber varies over several orders of magnitude, and even large amplitude density differences between models can appear small on the typical log-log scales of these graphs. The 2m temperature field shows very little smoothing compared to the analysis regardless of model because it is strongly affected by the local elevation, which is always supplied as a constant field.

B.3. Quantile/Quantile Plots

Quantile-quantile plots show a joint cumulative density function, and we use them here to evaluate the overall realism of the forecasts produced by the control and AMSE AR12 models independently of the forecast skill. In figures 9 and 10, the x-location of each point is the labelled percentile of North American weather station observations for Northern Hemisphere winter and summer periods. The y-location of each point is the corresponding percentile for the HRES analysis or the 5-day forecasts produced by the control and AMSE AR12 models, interpolated to the station locations. For example, in the left panel of figure 9, the 98th percentile corresponds to an observed wind speed of about 11.5m/s, but the 98th percentile of the HRES analysis was about 10m/s.

The y = x line on the quantile-quantile plot, shown as a dashed line in each panel, suggests that the forecast and observations have the same unconditional distributions when aggregated, and departures from the diagonal line indicate systematic underprediction or overprediction of extreme values. In our case, figure 9 shows that the AMSE AR12 model has a more realistic representation of surface winds, matching the trends seen in the HRES data. The control model produces noticeably weaker winds at all percentiles, showing a systematic shift in the distribution towards weaker surface winds, particularly in summer.

In contrast, figure 10 shows that the models are essentially equivalent in the distribution of 2m temperatures. As discussed in section B.2, the 2m temperature field shows little smoothing in the control model, likely due to the strong influence of elevation on the surface temperature. Improvements to the forecast of 2m temperature in the AMSE AR12 model are found



Figure 9. Quantile-quantile plots of 10 m wind speed at surface station locations for the North American domain. At left, 1 Jan–30 March 2022 (boreal winter), and at right 20 June–19 September 2022 (boreal summer). The control and AMSE AR12 points show model evaluations for 5-day forecasts. The shaded region denotes confidence interval based on the Kolmogorov-Smirnov test.

more in the forecast skill (see figure 11) than in the unconditional distribution of temperatures.

B.4. Details of the Lagged Ensemble Verification

Brenowitz et al. (2025) uses several metrics to evaluate the quality of the lagged ensembles. In this work, we use the fair CRPS score, the ensemble root mean squared error (eRMSE), and the spread-error ratio (SER). The eRMSE statistic is derived from its squared version (ensemble mean squared error), evaluated pointwise and integrated over the grid. The SER statistic is the simple ratio of the integrated MSE and ensemble spread (unbiased estimate of variance), noting for emphasis that the ratio is taken after the grid-averaging. For an ensemble of N_e members $(x_{1...N_e})$ evaluated over N_{date} forecasts with verifying analysis y, the corresponding formulas are:

$$CRPS(x,y) = \frac{1}{N_{date}} \sum_{d=1}^{N_{date}} \sum_{i,j} dA(i,j) \left(\frac{1}{N_e} \sum_{k=1}^{N_e} |x_k(i,j) - y(i,j)| + \frac{1}{2N_e(N_e-1)} \sum_{k=1}^{N_e} \sum_{l=1}^{N_e} |x_k(i,j) - x_l(i,j)| \right),$$
(14)

$$eRMSE(x,y) = \left(\frac{1}{N_{date}} \sum_{d=1}^{N_{date}} \sum_{i,j} dA(i,j)(\bar{x}(i,j) - y(i,j))^2\right)^{1/2}, \text{ and}$$
(15)

$$\operatorname{SER}(x,y) = \left(\frac{1}{N_{date}} \sum_{d=1}^{N_{date}} \frac{1}{N_e - 1} \frac{\sum_{i,j} \mathrm{d}A(i,j) \sum_{k=1}^{N_e} (x_k(i,j) - \bar{x}(i,j))^2}{\sum_{i,j} \mathrm{d}A(i,j)(\bar{x}(i,j) - y(i,j))^2}\right)^{1/2},\tag{16}$$

where $\bar{x}(i,j) = N_e^{-1} \sum_k x_k(i,j)$ is the ensemble mean at the (i,j) gridpoint.

Figures 11 and 12 show the CRPS and eRMSE skill scores respectively of the lagged ensemble generated with AMSE AR12 compared to the lagged ensemble of the control model for the geopotential (z), temperature (t), specific humidity (q), and u-component of wind (u) at several elevations and for the mean sea level pressure (msl), 2-meter temperature (2t), u-component of 10m wind (10u), and 6h-accumulated precipitation (tp) at the surface.

For these figures, statistical significance was determined by bootstrapping, sampling 1/3 of the total dates in each sample to give an average gap between dates of 36h. The forecast skill of persistence (that is, the gain over a climatological forecast by



Figure 10. As in figure 9, for 2m temperature. Low percentiles (extreme cold) are shown for the Northern Hemisphere winter, and high percentiles (extreme heat) are shown for the Northern Hemisphere summer.

predicting that everything will remain constant) decays very quickly over 36h, so samples so-spaced apart are reasonably independent of each other.

Overall, AMSE AR12 shows CRPS skill improvements for most variables and most lead times, but total precipitation shows only small improvements at long lead times and degradation at short lead times. This is explained by the separation of modes in (5) not being a natural one for precipitation. Precipitation is often localized but always non-negative, and consequently its spectral decomposition does not really resemble the normally-distributed random values that give meaning to (5) and (6).

The eRMSE skill chart should be interpreted with caution. The scores of (14)–(16) were developed for the case of an ideal ensemble, where members are statistically indistinguishable from each other and equally accurate in expectation. This is not really the case for a lagged ensemble, where the shorter-duration members should be noticeably more accurate than longer-duration members and an ideal aggregation would separately weight each term. This is not done by Brenowitz et al. (2025) for simplicity and to avoid free parameters, but we believe that the early lead-time smoothing in the control model makes equal-weighting more optimal for its lagged ensemble than for the lagged ensemble of AMSE AR12.

For long lead times this advantage diminishes, where the relative degradation of forecast quality is much stronger between 0.5 days and 4.5 days than it is between 6 days and 10 days. In this regime, AMSE AR12 begins to show eRMSE skill over the control ensemble.

B.4.1. UNBIASED ENSEMBLE ROOT MEAN SQUARED ERROR

The eRMSE formula of (15) is a biased estimator of the true ensemble mean error, overestimating the error in proportion to the ensemble (sample) spread when the ensemble size is finite.

Consider N_e different realizations (x_i) of a single variable drawn from $\mathcal{N}(\mu, \sigma^2)$ when the ground-truth value is 0. Applying (15) to this gives:

$$\mathbb{E}\left(\mathrm{eMSE}(x,0)\right) = \mathbb{E}\left(\left(\frac{1}{N_e}\sum_i x_i\right)^2\right) = \mu^2 + \frac{\sigma^2}{N_e},\tag{17}$$

which overestimates the true ensemble mean squared error. This overestimate is more severe for small ensembles such as the lagged ensemble configuration of section 3.2, where the ensemble size cannot be easily increased.

Leutbecher & Palmer (2008) proposes correcting this overestimate by subtracting the standard error term to give an unbiased estimator of the ensemble mean squared error with a finite sample size. In the notation of (15), the corresponding root mean

	z -	8	9	10	10	10	10	10	10	11	11	11	11
hPa	t -		2	<u>3</u>	à	4	4	4	5	<u></u>	<u>6</u>	6	2
250	q -	Δ	Δ	Δ				2	<u>3</u>	<u>3</u>	4	4	4
	u -	Δ	Δ	Δ	Δ							2	2
	z -		2	<u>í</u> 3	4	4	4	<u></u>	5	<u></u>	5	5	6
hPa	t -	Δ				2	<u>/3</u>	<u>3</u>	Ä	Â	4	4	4
500	q -	<u>/3</u>	<u>3</u>	3	<u>3</u>	<u>3</u>	<u>/3</u>	<u>3</u>	<u>3</u>	Ä	<u>á</u>	3	4
	u -	Δ	Δ	Δ	Δ	Δ	Δ	Δ	Δ	Δ			
	z -	0	Δ			2	2	<u>3</u>	<u>3</u>	<u>3</u>	<u>/3</u>	<u>3</u>	<u>/3</u>
hPa	t-	2	<u>í</u>	4	5	5	6	6	6	6	7	2	2
850	q -	Δ	Δ	Δ	0	0	0	0	0	0	0	Δ	Δ
	u -	▼	▼	▼	▼	V	0	0	Δ	Δ	Δ		
n	nsl -			2	2	2	2	<u>3</u>	<u>3</u>	<u>3</u>	<u>/3</u>	<u>3</u>	<u>/3</u>
ace	2t -	V	V	0	Δ	Δ					2	2	2
Surf	0u -	V	V	▼	V	0	0	0	Δ	Δ	Δ		
	tp -	3	2	▼	▼	▼	V	V	0	0	Δ	Δ	
			3		4		5		6		7		8
							Lead t	ime (d)					

Figure 11. CRPS skill score (% improvement), measured as the relative difference between the CRPS (14) of the 12-step fine-tuned model and the CRPS of the control model, for a selection of variables and lead times. Orange up-arrows show where the fine-tuned model performs better, blue down-arrows show where the control model performs better. Hollow arrows represent a difference of less than 1%, and differences of 2% or larger are marked. Hollow circles mark values that are not statistically significant at the 90% level.

squared formula becomes:

ub_eRMSE
$$(x, y) = \left(\frac{1}{N_{date}} \sum_{d=1}^{N_{date}} \sum_{i,j} dA(i,j) \left(\left(\frac{1}{N_e} \sum_{k=1}^{N_e} x_k(i,j) - y(i,j)\right)^2 - \frac{1}{N_e(N_e-1)} \sum_{k=1}^{N_e} (x_k(i,j) - \bar{x}(i,j))^2 \right) \right)^{1/2},$$
(18)

which performs this correction pointwise on the grid before computing the spatial average and taking the square root.

Implementing this adjustment slightly improves the scores of the AMSE-tuned model compared to the control model, and the corresponding "scorecard" is shown in figure 13.

	z -											2	2
hPa	t -	▼	0	0	0	Δ	Δ	Δ	Δ				
2501	q -	6	6	6	6	6	5	5	4	3	3	2	▼
	u -	▼	V	V	V	V	V	0	0	0	0	Δ	Δ
	z -	▼	0	0	0	0	0	0	0	0	0		
500hPa	t -	V	V	V	0	0	0	0	0	0	0	Δ	Δ
	q -	3	3	3	3	3	3	2	▼	▼	▼	▼	V
	u -	▼	▼	▼	▼	▼	▼	V	V	V	0	0	0
	z -	▼	▼	V	0	0	0	0	0	0	0	Δ	
hPa	t -	▼	V	0	0	Δ	Δ	Δ	Δ	Δ	Δ		
850	q -	4	5	5	4	4	4	4	4	3	3	3	1 3.
	u -	2	2	▼	▼	▼	▼	▼	V	V	V	0	0
n	nsl -	0	0	0	0	0	0	0	0	0	0	Δ	Δ
ace	2t -	▼	▼	▼	▼	▼	▼	▼	▼	V	V	V	V
1 Surf	0u -	2	▼	▼	▼	▼	▼	▼	V	∇	V	0	0
	tp -	6	6	6	6	6	6	6	5	5	5	5	5
			3		4		5		6		7		8
							Lead t	ime (d)					

Figure 12. As figure 11, for ensemble root mean squared error (15).

	z -	2	2							2	2	3	<u>/3</u>
hPa	t -	Δ	Δ	Δ							2	2	Ä
250	q -	5	5	5	5	5	4	4	3	2	▼	▼	0
	u -	0	0	0	0	0	0	Δ	Δ	Δ	Δ	A	
	z -	0	0	0	0	0	0	0	Δ				
hPa	t -	0	Δ	Δ	Δ	Δ	Δ	Δ	Δ				
500	q -	▼	▼	▼	▼	▼	▼	▼	V	V	0	0	Δ
	u -	V	V	V	0	0	0	0	0	0	0	Δ	Δ
	z -	▼	0	0	0	0	0	0	0				
hPa	t -	0	Δ	Δ	Δ								
850	q -	4	4	4	4	3	3	3	3	3	2	▼	▼
	u -	▼	▼	▼	V	∇	V	0	0	0	0	Δ	Δ
m	ısl -	0	0	0	0	Δ	Δ	Δ					
ace	2t -	▼	▼	▼	▼	▼	▼	V	V	V	V	0	0
Ju ² l	0u -	▼	▼	▼	▼	∇	V	V	0	0	0	0	0
	tp -	5	5	5	5	5	5	4	4	4	4	4	3
			3		4		5		6		7		8
							Lead t	ime (d)					

Figure 13. As figures 11 and 12, for the unbiased ensemble root mean squared error (18).

Table 2. Cumulative ranked probability scores for the models fine-tuned in this study in the lagged ensemble configuration described in section 3.2, for the "headline" variables and levels in figure 1 of Rasp et al. (2024). Lower is better; the best score is bolded and the second-place score is italicized.

Model	$z 500hPa (m^2 s^{-2})$			t 850hPa (K)			q 700	hPa (g · l	kg^{-1})	u 850hPa ($m \cdot s^{-1}$)		
	2.5d	5.0d	7.5d	2.5d	5.0d	7.5d	2.5d	5.0d	7.5d	2.5d	5.0d	7.5d
Control	31.038	84.315	162.481	0.428	0.691	1.003	0.357	0.523	0.652	0.823	1.340	1.904
MSE AR12	31.285	82.100	155.703	0.419	0.664	0.949	0.356	0.526	0.652	0.819	1.335	1.886
MAE AR12	29.969	80.621	155.361	0.410	0.654	0.947	0.340	0.499	0.624	0.811	1.313	1.859
AMSE AR1	33.720	94.703	186.202	0.422	0.721	1.078	0.354	0.558	0.721	0.863	1.485	2.115
AMSE AR12	30.565	80.469	153.267	0.418	0.653	0.935	0.347	0.510	0.634	0.832	1.341	1.882

B.5. Ablation Studies



Figure 14. As figure 3, for the comparison models of section B.5. Only the model trained with the AMSE error function retains sharpness to fine scales.

To ensure that the results of this study are not simply an artifact of increasing the model's overall training time, we compare against two additional fine-tunings:

- 1. MSE AR12 implements the fine-tuning schedule of table 1 with the unmodified mean squared error loss function, as with GraphCast's principal training.
- 2. MAE AR12 implements the fine-tuning schedule with a mean absolute error loss function, preserving the per-variable and per-level weightings of error.

Figure 14 shows the aggregated per-wavenumber performance of these models, and table 2 evaluates their CRPS for a selection of variables, levels, and lead times in the lagged ensemble configuration.

Both models still show excessive smoothing of fine scales, but training with mean absolute error moderately improves sharpness at the medium scales (wavenumbers 32–200 for longer lead times, corresponding to length scales of 1250–250 kilometers).

The excessive smoothing of the MSE-trained model is expected from section 2.1, but that argument does not directly apply to the mean absolute error loss function. However, we can still understand this behaviour intuitively. A model that is optimal under the mean absolute error predicts the mean of a distribution, and at longer lead times fine scales are less predictable than coarser scales. Therefore, the prediction of the median future should be smoother than its realization.

Even the moderate improvement to sharpness for the MAE-trained model results in improvements to the CRPS of the lagged ensemble, as shown in table 2.