Geosci. Model Dev., 19, 10651–10667, 2026
https://doi.org/10.5194/gmd-19-10651-2026
© Author(s) 2026. This work is distributed under
the Creative Commons Attribution 4.0 License.

Geoscientific Model Development
Open Access
EGU

Evaluating extreme precipitation forecasts: a threshold-weighted, spatial verification approach for comparing an AI weather prediction model against a high-resolution NWP model

Nicholas Loveday[1] **and Tracy Hertneky**[2]

[1]Research, Science and Innovation Group, Bureau of Meteorology, Melbourne, Australia
[2]NSF National Center for Atmospheric Research (NCAR) and Developmental Testbed Center (DTC), Boulder, CO, USA

Correspondence: Nicholas Loveday (nicholas.loveday@bom.gov.au)

Received: 21 November 2025 – Discussion started: 4 February 2026
Revised: 9 September 2026 – Accepted: 23 October 2026 – Published: 10 November 2026

Abstract. Recent advances in AI-based weather prediction have led to the development of artificial intelligence weather prediction (AIWP) models with competitive forecast skill compared to traditional NWP models, but with substantially reduced computational cost. There is a strong need for appropriate methods to evaluate their ability to predict extreme weather events, particularly when spatial coherence is important, and grid resolutions differ between models.

We introduce a verification framework that combines spatial verification methods and weighted proper scoring rules. Specifically, the framework extends the High-Resolution Assessment (HiRA) approach with threshold-weighted scoring rules. It enables user-oriented evaluation consistent with how forecasts may be interpreted by operational meteorologists or used in simple post-processing systems. The method supports targeted evaluation of extreme events by allowing flexible weighting of the relative importance of different decision thresholds. We demonstrate this framework by evaluating 32 months of precipitation forecasts from an AIWP model and a high-resolution NWP model. Our results show that model rankings are sensitive to the choice of neighbourhood size. Increasing the neighbourhood size has a greater impact on scores evaluating extreme-event performance for the high-resolution NWP model than for the AIWP model. At near equivalent neighbourhood sizes, the empirical CDF of the high-resolution NWP model only outperformed the empirical CDF of the AIWP model in predicting extreme precipitation events at short lead times. We also demonstrate how this approach can be extended to evaluate discrimination ability in predicting heavy precipitation. We find that the high-resolution NWP model had superior discrimination ability at short lead times.

1 Introduction

In recent years, there has been rapid progress in the development of artificial intelligence weather prediction (AIWP) models, sometimes referred to as data-driven models (Ben Bouallègue et al., 2024). These models are typically based on neural networks trained on reanalysis datasets. This approach contrasts with traditional numerical weather prediction (NWP) models, which rely on equations that model physical processes. One major advantage of AIWP models is that they are significantly more computationally efficient than NWP models for making predictions once they have been trained.

Most early global AIWP models are trained on ERA5 (Hersbach et al., 2020) data, and therefore have a similar or coarser grid resolution of approximately 0.25°. These include deterministic models (e.g., Keisler, 2022; Lam et al., 2023; Bi et al., 2023; Lang et al., 2024) and ensemble models (e.g., Price et al., 2024; Lang et al., 2026). In general, these models operate at much coarser resolutions than the high-resolution physical NWP models available to many international weather forecasting centers. However, higher-resolution, limited-area AIWP models have started to emerge more recently (e.g., Nipen et al., 2026; Adamov et al., 2025; Abdi et al., 2025), and this trend is expected to continue.

As modelling centers have access to high-resolution NWP models and with the accelerating development of higher-resolution AIWP models, there is a strong need for appropriate methodologies to compare the performance of models with differing spatial resolutions. Such comparisons are important to help meteorological agencies make informed decisions regarding future modelling strategies and to assist operational meteorologists and forecast system developers in understanding each model's strengths and weaknesses.

This raises several key questions we aim to address.

First, most benchmarking efforts of global AIWP models to date have relied on point-to-point verification methods (e.g., Rasp et al., 2024), with perhaps the exception of Radford et al. (2025a), who applied an object-based verification approach. Point-to-point methods may be appropriate for evaluating forecasts at individual locations (e.g., for end-users accessing weather forecasts via mobile apps). However, meteorologists and forecast systems often use forecasts spatially, where spatial coherence and physical realism are important.

Point-to-point verification methods tend to favour smoother forecasts than are typically observed in gridded observations (Subich et al., 2025). When spatial structure is important, point-to-point metrics can suffer from the "double penalty" effect (Ebert, 2008). This can occur when a forecast feature matches the observed shape and intensity but is slightly displaced in location, leading to it being penalised twice. It is penalised once for the false prediction at the forecast location, and once for the missed event at the observed location. This issue is particularly problematic for high-resolution models, which aim to represent fine-scale features. Lower-resolution models that fail to capture the feature at all may score better under point-to-point metrics, because they incur only a single penalty. Thus, while high-resolution NWP models may produce forecasts that appear more realistic to meteorologists, they may not perform better under standard point-to-point verification metrics (Mass et al., 2002).

To address the limitations of point-based verification, a wide range of spatial verification methods have been developed over the past two and a half decades (Gilleland et al., 2009; Dorninger et al., 2018). These methods are typically grouped into five categories: neighbourhood methods, scale-separation methods, feature-based (or object-based) methods, field-deformation methods, and distance measures. In this paper, we focus on an approach within the neighbourhood methods category (Ebert, 2009; Schwartz and Sobash, 2017). Neighbourhood methods are verification approaches that evaluate forecasts using values within a local neighbourhood surrounding the observation location.

Second, there is a pressing need to assess how well both AIWP and high-resolution NWP models predict extreme weather events. Much of the existing literature has focused on a limited number of case studies (e.g., Charlton-Perez et al., 2024; Morisseau et al., 2025). While case studies are valuable, they may be prone to selection bias or may not generalise across more cases. Focusing solely on cases where extremes occur can lead to the "forecaster's dilemma" (Lerch et al., 2017). The forecaster's dilemma arises when forecasters or modellers must choose between issuing honest forecasts and adjusting their forecasts to optimise a particular scoring rule (we refer to the latter as "hedging"). Proper scores discourage hedging and avoid the "forecaster's dilemma". A scoring rule is proper if the expected score is optimised by producing a forecast that corresponds to the forecaster's true belief (Winkler and Murphy, 1968; Gneiting and Raftery, 2007).

There are some recent examples of broader evaluation of extremes across many events. For example, Lam et al. (2023) and Ben Bouallègue et al. (2024) evaluated the discrimination ability of AIWP models to differentiate between climatological extreme temperature events and non-events. Olivetti and Messori (2024) assessed performance in forecasting temperature and wind extremes primarily by conditioning on observed extremes, but also included an evaluation using a threshold-weighted squared error (Taggart, 2021). McGovern et al. (2026) recently introduced a benchmark that focuses on evaluating several models to predict extremes based on a comprehensive set of events and marginal events. Zhang et al. (2026) evaluated several AIWP models and ECMWF's HRES model against observations that are more extreme than the observations in the 1979–2017 training period of the AIWP models. Some of these studies have taken efforts to avoid, or practically reduce the impact of the "forecaster's dilemma", although some of these approaches still suffer from this problem.

These efforts have been limited to point-to-point verification methods, and there is a need for more comprehensive verification methods that assess the performance of models predicting extremes when they are used in a user-oriented spatial framework.

The appropriate verification approach depends on how forecasts are used in practice. Weather models are used in a wide variety of ways. For example, model output could be used at each grid point to produce forecasts for a user's location (e.g., via a mobile weather application). A high-resolution ensemble could be used by an operational meteorologist to infer the likelihood of different modes of convection during the afternoon. The dynamic tropopause output of a deterministic model (or single ensemble member) can also be used by a decision-support meteorologist to construct a conceptual model of the atmosphere, aiding rapid responses to bespoke queries in an emergency response centre. These use cases require different evaluation approaches. In the first example, point-to-point evaluation using proper or consistent scoring rules may be appropriate. In the second, point-to-point verification would not capture the spatial realism of convection within the ensemble, necessitating some form of spatial processing. Similarly, in the third case, evaluating the spatial structure of the dynamic tropopause may re-

quire a spatial verification method distinct from that used in the second example. These examples illustrate that no single metric can assess all the diverse ways that weather models are used.

It is desirable that forecasts are probabilistic and evaluated using proper scoring rules[1] (Gneiting and Katzfuss, 2014). This is because probabilistic forecasts support optimal decision-making, and proper scores reward honest forecasts while discouraging hedging. Common examples of probabilistic forecasts in meteorology include ensembles, predictive distributions, or binary classifiers (i.e., probability forecasts for binary events). These forecasts are traditionally evaluated point-to-point; however, spatial processing can also be applied to construct probabilistic forecasts. For example, the high-resolution ensemble output in the second example could be used to create a probability mass function of various modes of convection, which could then be evaluated using a proper scoring rule.

There have been at least two recent examples of using proper scoring rules to compare single-valued AIWP models against traditional NWP models. First, Brenowitz et al. (2025) created a "lagged" ensemble of single-valued AIWP forecasts that could be evaluated with the CRPS. Secondly, Gneiting et al. (2025) converted single-valued forecasts into a predictive distribution using Isotonic Distributional Regression (IDR) (Henzi et al., 2021) and evaluated the resulting distribution using the CRPS. Both approaches are point-to-point methods that do not account for any spatial information.

One practical use case for weather models is when meteorologists visually assess a deterministic precipitation forecast around a point to qualitatively infer the likelihood of different precipitation amounts. A quantitative forecast can be constructed by generating an empirical cumulative distribution function (CDF) from a neighbourhood pseudo-ensemble around the observation. Similarly, post-processing techniques often employ neighbourhood approaches on NWP output to generate probabilistic forecasts (Theis et al., 2005; Schwartz and Sobash, 2017). In this paper, we demonstrate a verification approach that merges two existing methods and assesses the model in one particular framework that aligns with a specific way that operational meteorologists may use a model. That is, we evaluate how well extremes are predicted when the model's spatial neighbourhood around a point is treated as an empirical CDF. The approach has several strengths including:

1. It is a spatial verification method that aligns with a particular way that a model may be used.
2. It uses proper scoring rules within a user-oriented framework.
3. Threshold-weighting of proper scoring rules allows for the evaluation of extremes and other important decision thresholds.
4. It permits each forecast system to be evaluated on its native grid when the forecast is defined as the finite neighbourhood empirical distribution.
5. Models can be evaluated against station-based observations which may be desirable in some cases.
6. It can be extended to measure the discrimination ability.

Pagano et al. (2024) raised the research question, "Can spatial verification methods be extended to emphasize predictive performance for extremes without creating forecaster's dilemmas...?" The approach presented here provides one response to this question by evaluating extremes within a user-oriented spatial framework that uses proper scoring rules. Specifically, we compare the native-grid neighbourhood forecast products derived from GraphCast-GFS and HRRR over 32 months. This comparison includes, rather than removes, differences associated with the grid spacing and effective resolution of the two forecast systems.

2 Data

2.1 Observations

Station data were used to assess the performance of the forecast models using a spatial verification approach. Most evaluation of AIWP models has been done against ERA5, which has some biases in its precipitation fields, with the United States having a dry bias (Lavers et al., 2022). Additionally, most global AIWP models are trained against ERA5, and biases in ERA5 may propagate into these models and may go undetected if evaluated against ERA5. Some countries have alternative gridded rainfall datasets, which may serve as suitable substitutes, though these are not available everywhere. For this reason, some countries may place greater trust in verification using weather station data.

While most past efforts have evaluated these models against gridded analysis, there have been limited examples of verification of AIWP models against station data (e.g., Ben-Bouallègue et al., 2026). WeatherReal (Jin et al., 2024) was recently developed to evaluate AIWP models against in situ observations. One limitation of this dataset is that it is only available for 2023, and our aim was to evaluate model performance over a longer period to assess extreme event prediction. Instead, one-minute Automated Surface Observing System (ASOS) (NWS, 1998) precipitation data were retrieved for the contiguous United States (CONUS). A map of ASOS weather station data is shown in Fig. 1. From this

[1] In some cases, only a single-valued forecast is required for the forecast service rather than a full predictive distribution. Consistent scoring functions can then be used to evaluate forecasts expressed through a directive in the form of a statistical functional (Gneiting, 2011); for example, a 90th percentile forecast can be evaluated using a quantile loss.

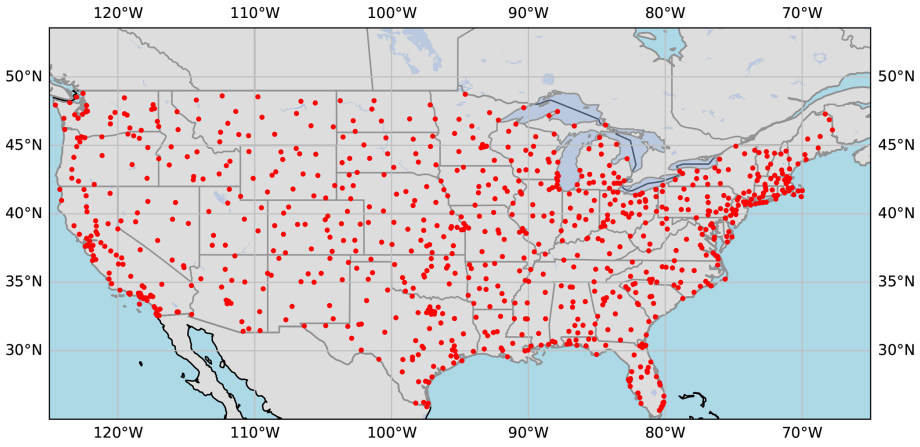


Figure 1. A map of ASOS station locations used in this study.

data, 6 h precipitation accumulations were derived. As a basic quality control measure, data were removed when one-minute accumulations exceeded 38 mm or 6 h accumulations exceeded 840 mm, which correspond to world record values (WMO, 1994). Additionally, 6 h accumulations were treated as missing when fewer than 5 h of valid data existed within the 6 h period. ASOS data have previously been used to evaluate the HRRR model (e.g., Ikeda et al., 2013; Fovell and Gallagher, 2022), but, to our knowledge, have not been used to evaluate an AIWP model.

To define thresholds for extreme precipitation events, several options can be considered, each addressing slightly different questions. These include using fixed thresholds across all stations, annual climatological thresholds for each station, or seasonally varying climatological thresholds at each station. In this study, we adopt annual climatological thresholds for each station, which aligns with approaches used by some meteorological agencies in operational warning services. Specifically, we define extreme precipitation events using the 99th and 99.9th percentile thresholds of 6 h precipitation accumulations at each station. As long observational time series are not available for all stations in the ASOS dataset, these thresholds are derived from ERA5 reanalysis data (1990–2020) at the grid point corresponding to each station location. These thresholds may differ from those derived directly from station observations; however, in the absence of a universally defined threshold for extreme precipitation, the exact choice of threshold is inherently somewhat subjective.

2.2 Forecasts

We used the 00:00 UTC run of two different models, with base-run times spanning from 1 January 2022 to 30 August 2024 (973 model runs).

The first model we evaluated was GraphCast, with initial conditions from NOAA's Global Forecast System (GFS), using the reforecast archive generated by Radford et al. (2025b), hereafter referred to as GraphCast-GFS. The grid resolution of GraphCast-GFS is 0.25°. Radford et al. (2025a) found that, while GraphCast-GFS tended to overestimate lower rainfall amounts compared to GraphCast initialised with ECMWF's Integrated Forecasting System (GraphCast-IFS), its performance for higher rainfall amounts was similar. GraphCast-GFS was selected because it was the only AIWP model output available at the time of analysis that produced 6 h precipitation accumulations over a sufficiently long period to assess performance in predicting extremes.

Our evaluation begins in 2022 because the operational version of GraphCast was fine-tuned on data through 2021, and we sought to avoid testing the model on data used in its development. Since AIWP models can produce negative precipitation values, all negative values were set to zero, as this is a trivial post-processing step.

To compare an AIWP model with a high-resolution physical numerical weather prediction (NWP) model, we also evaluated version 4 of the High-Resolution Rapid Refresh (HRRR) model (Dowell et al., 2022). HRRR is a convection-allowing, cloud-resolving physical NWP model with a horizontal grid spacing of 3 km, which runs hourly. HRRR was selected due to its widespread use in the United States as a high-resolution operational model.

Prior to score calculation, valid cases were matched between GraphCast-GFS and HRRR so that every paired comparison used the same observation times and stations. Since HRRR produces forecasts only up to 48 h from the 00:00 UTC run, we limited the GraphCast-GFS evaluation to its first 48 forecast hours as well.

3 Using HiRA and twCRPS

This work unifies two different verification methods. We give an overview of each method.

3.1 HiRA

One neighbourhood verification method that can be used with point observations is the High-Resolution Assessment (HiRA) framework (Mittermaier, 2014). One of the motivations for the development of the HiRA framework was to evaluate models of varying resolutions against site-based observations while avoiding the double penalty effect within a specified distance.

The HiRA framework can be summarised as follows: for a given point observation, several squares (or sometimes circles) of varying sizes are used to define neighbourhoods of forecast grid cells surrounding the observation site. Each neighbourhood can then be used to generate a pseudo-ensemble, where each grid point within the neighbourhood is considered to have an equal probability of occurring. The pseudo-ensemble can subsequently be evaluated using a probabilistic score, such as the Brier score (Brier, 1950), the ranked probability score (RPS) (Epstein, 1969), or the continuous ranked probability score (CRPS) (Matheson and Winkler, 1976).

3.2 Threshold-Weighted Continuous Ranked Probability Score

In the context of evaluating the performance of forecasts for extreme events, threshold-weighted proper scoring rules are particularly important for forecast system developers and forecasters, as they help to avoid the "forecaster's dilemma". A scoring rule is considered proper if the expected score is optimised when the forecaster issues a probability forecast that corresponds to their true belief. The continuous ranked probability score (CRPS) is a strictly proper scoring rule that can be expressed as the integral of the Brier score over all possible thresholds (Matheson and Winkler, 1976; Gneiting and Raftery, 2007). The CRPS for evaluating a cumulative distribution function F is defined as

$$\mathrm{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(z) - \mathbb{1}\{y \le z\})^2 \mathrm{d}z \tag{1}$$

$$= \mathbb{E}_F|X - y| - \frac{1}{2}\mathbb{E}_F|X - X'|. \tag{2}$$

In the first expression, y is the observation, z is the decision threshold, and $\mathbb{1}\{y \le z\}$ is the indicator function, which equals 1 if $y \le z$ and 0 otherwise. In the second expression, X and X' are independent random variables with distribution F, and $\mathbb{E}_F$ is expectation with respect to the probability distribution F. Lemma 2.2 in Baringhaus and Franz (2004) and Eq. (17) in Székely and Rizzo (2005) show that the two expressions above are equal. When used to evaluate an ensemble forecast F_{ens} as an empirical CDF, the CRPS can also be expressed as

$$\mathrm{CRPS}(F_{\mathrm{ens}}, y) = \frac{1}{M}\sum_{m=1}^{M}|x_m - y| - \frac{1}{2M^2}\sum_{m=1}^{M}\sum_{j=1}^{M}|x_m - x_j|, \tag{3}$$

where M is the number of ensemble members, x_m and x_j are the forecasts of a single ensemble member. Ferro (2013) proposed the "fair" CRPS for evaluating an ensemble which is defined as

$$\mathrm{CRPS}_{\mathrm{fair}}(F_{\mathrm{ens}}, y) = \frac{1}{M}\sum_{m=1}^{M}|x_m - y| - \frac{1}{2M(M-1)}\sum_{m=1}^{M}\sum_{j=1}^{M}|x_m - x_j|. \tag{4}$$

The fair correction addresses the effect of estimating an underlying predictive distribution with a finite number of ensemble members.

Weighted scoring rules have been developed to evaluate forecasts with an emphasis on specific ranges of decision thresholds, such as those that lie in the extremes of a climatological distribution. These scoring rules have proven useful in several contexts for assessing the ability to predict extreme events. For example, Loveday et al. (2024) demonstrated that the standard mean squared error (MSE) showed no statistically significant difference in temperature forecast skill between meteorologists and automated guidance. However, using a threshold-weighted MSE revealed that meteorologists performed better at predicting temperature extremes than automated guidance. Similarly, Wessel et al. (2025) showed that threshold-weighted scores can be used as a loss function to improve the predictive performance for extremes. Allen et al. (2023) illustrated how multivariate threshold-weighted scores can be applied to evaluate forecasts of rainfall accumulation over consecutive days.

Among the wide variety of threshold-weighted scores, the threshold-weighted continuous ranked probability score (twCRPS) has been the most extensively used and is defined as

$$\mathrm{twCRPS}(F, y) = \int_{-\infty}^{\infty} (F(z) - \mathbb{1}\{y \le z\})^2 w(z)\mathrm{d}z, \tag{5}$$

where $w(z)$ is a non-negative weight function (Gneiting and Ranjan, 2011). The choice of $w(z)$ depends on the user's decision thresholds of interest. For example, if the aim is to evaluate a model's ability to produce accurate forecasts for users whose decision thresholds exceed 40 °C, a suitable threshold weight function could be $w(z) = \mathbb{1}(z > 40)$. Another example might involve choosing $w(z)$ to focus on temperature ranges associated with aircraft icing, in which case

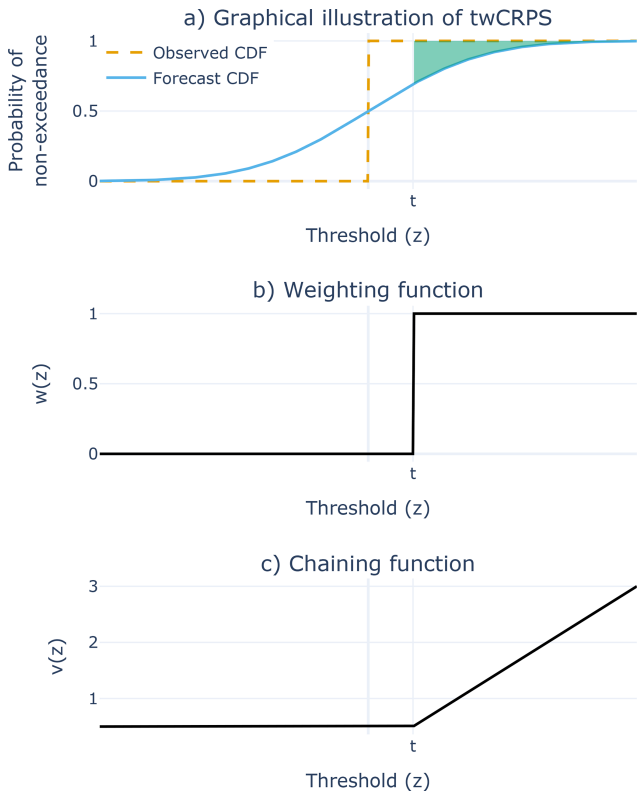


Figure 2. Graphical illustration of the threshold-weighted continuous ranked probability score (twCRPS) with a uniform weight of 1 applied to all thresholds $z > t$ and a weight of 0 applied elsewhere. **(a)** The solid blue curve shows the forecast CDF and the dashed orange line represents the Heaviside step function of the observation. The twCRPS is the integrated squared difference between the solid blue curve and the dashed orange line, with interval of integration $(z > t)$ indicated by the green shaded region. **(b)** The threshold weight function $w(z) = \mathbb{1}(z > t)$. **(c)** The corresponding chaining function $v(z) = \max(z, t)$.

the weights could vary smoothly to reflect the relative importance of different thresholds z.

Figure 2 provides a graphical illustration of how the twCRPS is computed when a constant weight is applied to the upper tail of decision thresholds, aiding interpretation.

Recently, Allen et al. (2023) showed that the twCRPS can be adapted for use with ensemble forecasts. The twCRPS for evaluating an ensemble forecast F_{ens} is defined as

$$\text{twCRPS}(F_{\text{ens}}, y; v) = \frac{1}{M}\sum_{m=1}^{M}|v(x_m) - v(y)| - \frac{1}{2M^2}\sum_{m=1}^{M}\sum_{j=1}^{M}|v(x_m) - v(x_j)|, \quad (6)$$

where v is the *chaining function*. The chaining function v is an antiderivative of the threshold weight function $w(z)$, such that

$$v(z) - v(z') = \int_{z'}^{z} w(z)\mathrm{d}z. \quad (7)$$

For example, if we wish to assign a threshold weight of 1 to thresholds above a specified threshold t, and a weight of 0 below t, the threshold weight function would be $w(z) = \mathbb{1}(z > t)$, where $\mathbb{1}$ is the indicator function that returns 1 if the condition is true and 0 otherwise. A corresponding chaining function is then $v(z) = \max(z, t)$. Allen et al. (2023) and Allen (2024) provide further examples illustrating the relationship between chaining functions and threshold weight functions.

Similarly to the fair CRPS, the fair twCRPS is defined as

$$\text{twCRPS}_{\text{fair}}(F_{\text{ens}}, y; v) = \frac{1}{M}\sum_{m=1}^{M}|v(x_m) - v(y)| - \frac{1}{2M(M-1)}\sum_{m=1}^{M}\sum_{j=1}^{M}|v(x_m) - v(x_j)|. \quad (8)$$

3.3 Constructing and interpreting neighbourhood pseudo-ensembles

When HiRA is used to compare forecast systems with different grid spacings, neighbourhoods can be selected to cover approximately the same physical area. For example, if Model A has a grid spacing of 3 km and Model B has a grid spacing of 9 km, a 27 km × 27 km neighbourhood can be represented by a 9×9 grid of Model A and a 3×3 grid of Model B, both centred on the observation point.

The score used to evaluate a neighbourhood pseudo-ensemble depends on how the neighbourhood forecast is interpreted. In this study, the M native-grid values within a specified neighbourhood $\mathcal{N}$ define the equally weighted empirical CDF

$$\widehat{F}_{\mathcal{N}}(z) = \frac{1}{M}\sum_{m=1}^{M}\mathbb{1}\{x_m \le z\}. \quad (9)$$

Under this interpretation, the empirical CDF is the forecast being evaluated. Its CRPS and twCRPS are therefore calculated using Eqs. (3) and (6). This corresponds to using the values surrounding an observation site to construct a predictive distribution for that site.

Alternatively, the neighbourhood values can be treated as a finite sample from an underlying predictive distribution. The fair CRPS and twCRPS in Eqs. (4) and (8) account for the finite number of members under this interpretation, provided that the values can be treated as suitable draws from the underlying distribution. Because neighbouring grid values are spatially dependent, the suitability of this interpretation should be considered for each application. The fair correction does not remove differences in grid spacing or spatial scale.

Table 1. Neighbourhood sizes used for HRRR and GraphCast-GFS, paired to cover similar physical areas.

Approximate physical dimensions (km)	HRRR grid points	GraphCast-GFS grid points
3 km × 3 km	1 × 1	–
21 km × 27 km	7 × 9	1 × 1
63 km × 81 km	21 × 27	3 × 3

Evaluating native-grid neighbourhoods compares the resulting forecast products, including effects associated with grid spacing, effective resolution, and spatial sampling. If the aim is instead to compare forecast systems at a common spatial scale, the forecasts should first be re-gridded to a common grid and evaluated against observations at a comparable spatial scale. In this study, we use the native-grid empirical-CDF interpretation and use the corresponding empirical scores. The Brier scores are also calculated from the same empirical forecast probabilities.

4 CRPS and twCRPS HiRA results

We select HRRR and GraphCast-GFS neighbourhoods that cover similar physical areas. Because GraphCast-GFS uses a latitude–longitude grid, its physical grid spacing varies with latitude, and no single HRRR neighbourhood covers exactly the same physical area at every station. We therefore use rectangular HRRR neighbourhoods that approximate the physical areas covered by the GraphCast-GFS neighbourhoods over CONUS. This differs from most HiRA applications, which typically use square or circular neighbourhoods. Table 1 lists the paired neighbourhood sizes and their approximate physical dimensions.

Note that, although there is no neighbourhood size for GraphCast-GFS that is equivalent to the 3 km × 3 km point forecast for HRRR, we still evaluate the point-based HRRR to understand the impact of neighbourhood size on the scores because of the double penalty effect. The CRPS of a single ensemble member (i.e., the point-based forecast using a neighbourhood size of 1) is equivalent to the absolute loss. Likewise, threshold-weighted MAE (or threshold-weighted absolute loss) (Taggart, 2021) is equivalent to twCRPS for a single ensemble member.

When aggregating results spatially, we weight the results based on station density by following the approach in Rodwell et al. (2010). This ensures that the verification results are not overly influenced by geographical regions that have a higher concentration of stations compared to others. The station density weights are allowed to vary for each timestep to account for non-continuous reporting from weather stations.

4.1 CRPS results

We first calculate the mean CRPS of the specified neighbourhood-derived empirical forecasts to assess performance when all precipitation thresholds are weighted equally. Confidence intervals for the difference in mean scores between the two models are calculated as follows: spatial means are first taken across stations to account for spatial correlation, and the Hering–Genton modification of the Diebold–Mariano test (Diebold and Mariano, 1995; Hering and Genton, 2011) is then applied to account for temporal correlation.

Figure 3a shows the mean CRPS for GraphCast-GFS and HRRR across various neighbourhood sizes. Increasing the neighbourhood size leads to a better score for both models. While one might expect a large improvement in performance when increasing the neighbourhood size in the HRRR as it is a high-resolution physical NWP model, an improvement also occurs with GraphCast-GFS despite its smoother appearance. A diurnal trend in model performance can also be seen. When the two models are compared gridpoint-to-gridpoint (Fig. 3b), HRRR outperforms GraphCast-GFS during the first 24 h, with mixed results at longer lead times. However, when the models are assessed over comparable neighbourhood sizes (Figs. 3c and d), the HRRR demonstrated better performance across all lead times.

4.2 twCRPS results

To assess performance in predicting extremes, we construct a twCRPS weight function that takes the form $w(z) = \mathbb{1}(z > q_\alpha)$, where q_α is the α quantile of the climatological values in the ERA5 dataset for the point at which the ASOS station is located. We set $\alpha = 0.99$ to focus on performance above the climatological 99th percentile. The corresponding chaining function that we use is $v(z) = \max(z, q_\alpha)$. When the mean twCRPS is calculated across time and space, q_α varies by station but is fixed in time. To test the impact of a more extreme threshold, we repeat the computation in Appendix A with $\alpha = 0.999$ to examine performance above the climatological 99.9th percentile.

In contrast to Fig. 3a, Fig. 4a shows that increasing the neighbourhood size had a larger impact on the mean twCRPS for the HRRR compared to GraphCast-GFS. When evaluated on a point-to-point basis (Fig. 4b), GraphCast-GFS consistently achieved lower (better) twCRPS scores across all lead times. When evaluated across approximately equivalent neighbourhood sizes (Fig. 4c and d), the HRRR performs better at shorter lead times, but not at longer lead times. Potentially, this may be attributed to the HRRR's assimilation of

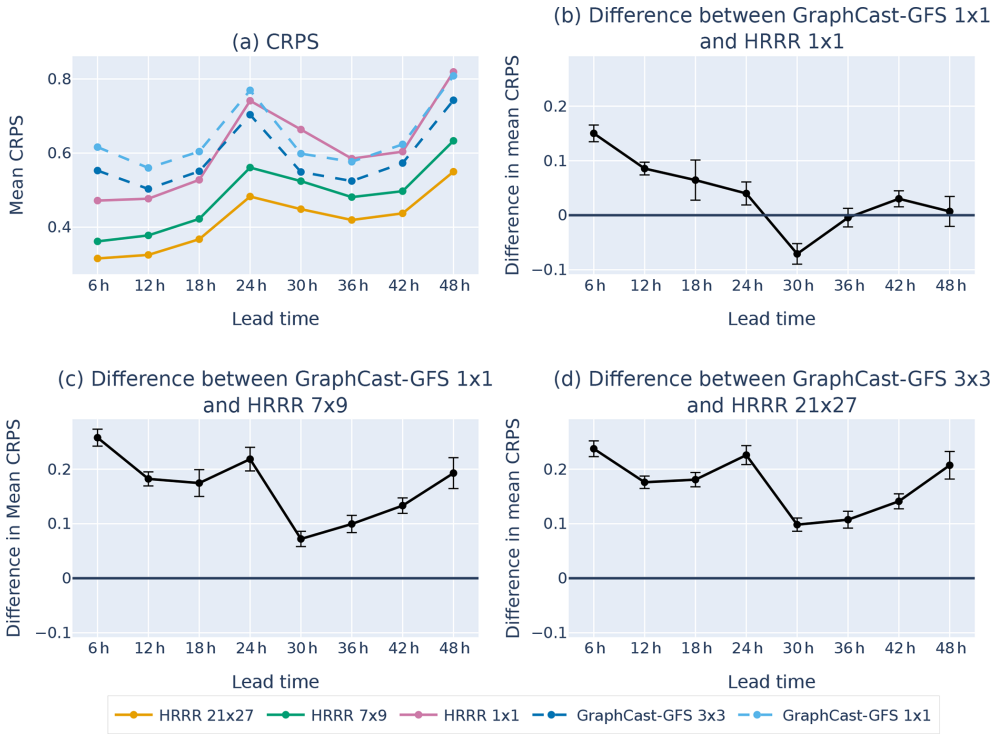


Figure 3. **(a)** Mean CRPS results aggregated across all stations and timesteps. Lower scores are better. **(b)** Difference between GraphCast-GFS 1×1 and HRRR 1×1 with 99 % confidence intervals. **(c)** Difference between GraphCast-GFS 1×1 and HRRR 7×9 (21×27 km equivalent) with 99 % confidence intervals. **(d)** Difference between GraphCast-GFS 3×3 and HRRR 21×27 (63×81 km equivalent) with 99 % confidence intervals. In subfigures **(b–d)**, positive values indicate that HRRR performed better than GraphCast-GFS for the specified neighbourhoods.

radar data, enhancing its short-term prediction of heavy precipitation. We leave a detailed investigation of this behaviour for future research.

5 Decomposing CRPS across decision thresholds

Since the CRPS is the integral of the Brier score across all thresholds (Eq. 1), we decompose it to visualise the Brier scores across a range of thresholds. This provides us with greater insight within the HiRA framework as to how different models provide value at different decision thresholds. We display results for two lead times in Fig. 5.

Figure 5 shows that the mean CRPS is dominated by non-extreme precipitation amounts (i.e., mean Brier score values from lower thresholds), which is unsurprising given their higher climatological frequency. To account for this, results are split at the 30 mm threshold, and the right-hand panels use a logarithmic vertical axis to better visualise the forecast rankings.

At a 6 h lead time (top panels), HRRR had lower mean Brier scores than GraphCast-GFS for precipitation thresholds below approximately 5 mm across all neighbourhood sizes. However, at the 30 h lead time, this is only true for some neighbourhood sizes. In cases where the mean Brier score curves for HRRR and GraphCast-GFS intersect, the crossing point typically occurs at higher precipitation thresholds for larger HRRR neighbourhood sizes. This may reflect the increased likelihood of double-penalty effects at higher precipitation amounts when using point-to-point verification.

For the upper precipitation thresholds, HRRR 21×27 had a slightly lower mean Brier score than GraphCast-GFS 3×3 at the 6 h lead time and a slightly higher mean Brier score at the 30 h lead time. This is consistent with the corresponding twCRPS results in Fig. 4.

6 Model climatology

We now assess the agreement between the model climatology and observations to understand any model biases. Quantile–Quantile (Q–Q) plots for the 6 h and 30 h lead-time point-based forecasts and the observations are shown in Fig. 6. They show that the climatology of HRRR 1×1 closely matches the observed climatology from the ASOS observations, whereas GraphCast-GFS 1×1 is substantially less likely to predict heavier precipitation. This may partly reflect differences in grid resolution, with the HRRR model producing forecasts that correspond more closely to station-based observations, while GraphCast-GFS produces fore-

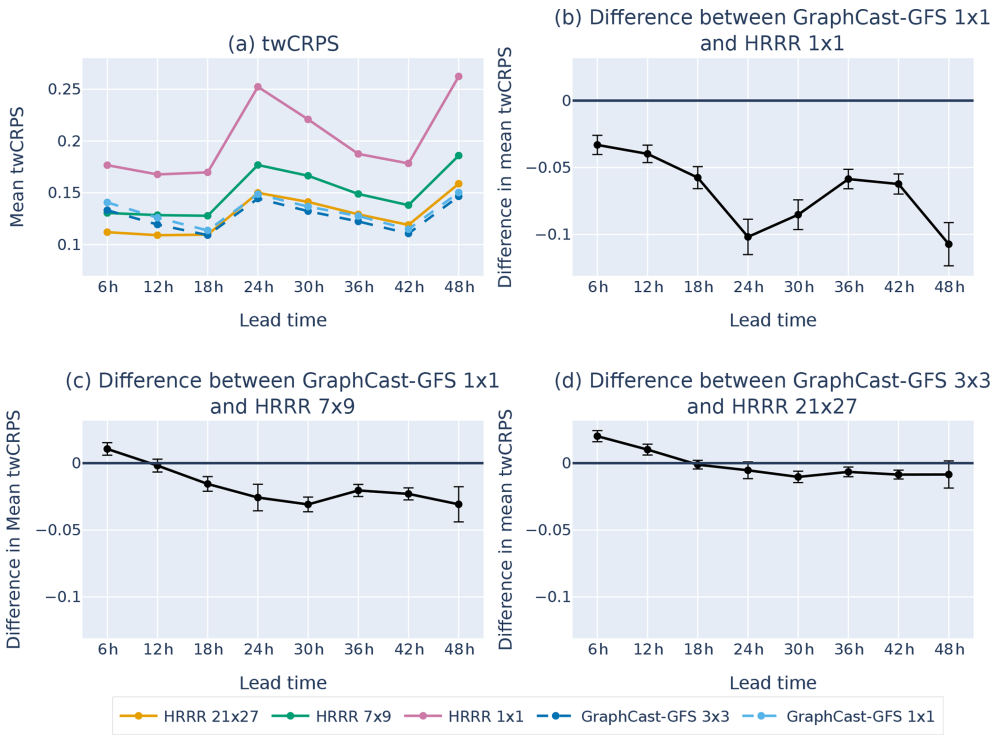


Figure 4. As for Fig. 3 but for the twCRPS with a threshold weight function of $w(z) = \mathbb{1}(z > q_{0.99})$.

casts that may be more representative of a 0.25° grid. This is consistent with the issues involved in comparing gridded model precipitation forecasts against rain gauge observations highlighted by Tustison et al. (2001). To explore the relationship between grid resolution and representativeness, we also take the neighbourhood mean of HRRR 7×9. The mean of the HRRR 7×9 neighbourhood shifts the distribution away from the diagonal line, but not nearly to the same extent as GraphCast-GFS 1×1, suggesting that differences in the climatology of model extremes are partly driven by the smoothness of GraphCast-GFS and not just because of differences in grid resolution.

7 Discrimination ability

As demonstrated in the preceding section, the behaviour of GraphCast-GFS may not be consistent with in-situ observations. For this reason, it is important to understand the discrimination ability (i.e., potential predictive ability) of the models as a way of accounting for the models representing the precipitation across a grid cell rather than at a specific point. This is because meteorologists may be able to learn how the model behaves and use its output accordingly. Alternatively, a model with good discrimination ability could easily be post-processed to correct for any conditional biases. The concept of "potential" scores as a measure of potential predictive or discrimination ability has been proposed and extended multiple times, including for the CRPS (e.g., Murphy and Epstein, 1989; Hersbach, 2000; Bröcker, 2011; Siegert, 2017). Recently Gneiting et al. (2025) introduced a version of a potential CRPS to measure and compare the discrimination ability of a deterministic NWP model to that of an AIWP model. This was done by calibrating the model (in-sample) using IDR via the EasyUQ method (Walz et al., 2024) to convert the single-valued forecasts to probabilistic forecasts. We adopt a related, but distinct approach for calculating the "potential" CRPS and twCRPS. Instead of using IDR on each neighbourhood member, we apply isotonic regression (in-sample) to each neighbourhood member. Isotonic regression is a method for fitting a nondecreasing function to a set of forecast-observation pairs (Ayer et al., 1955). Specifically, we use isotonic median regression (isotonic regression under L1 loss), which targets the median functional. We adopt this approach to approximate the behaviour of an operational meteorologist who is familiar with typical model biases and applies simple calibrations to the model output. Additionally, it partly circumvents the issue in neighbourhood verification approaches where adjacent grid cells in the neighbourhood are not representative of the station site (e.g., in areas of varying topography). This approach, however, is unlikely to be as effective as IDR in producing a well-calibrated predictive distribution if neighbourhood members were used as covariates. We then calculate a potential CRPS and potential twCRPS using the forecasts calibrated via isotonic regression as shown in Figs. 7 and 8.

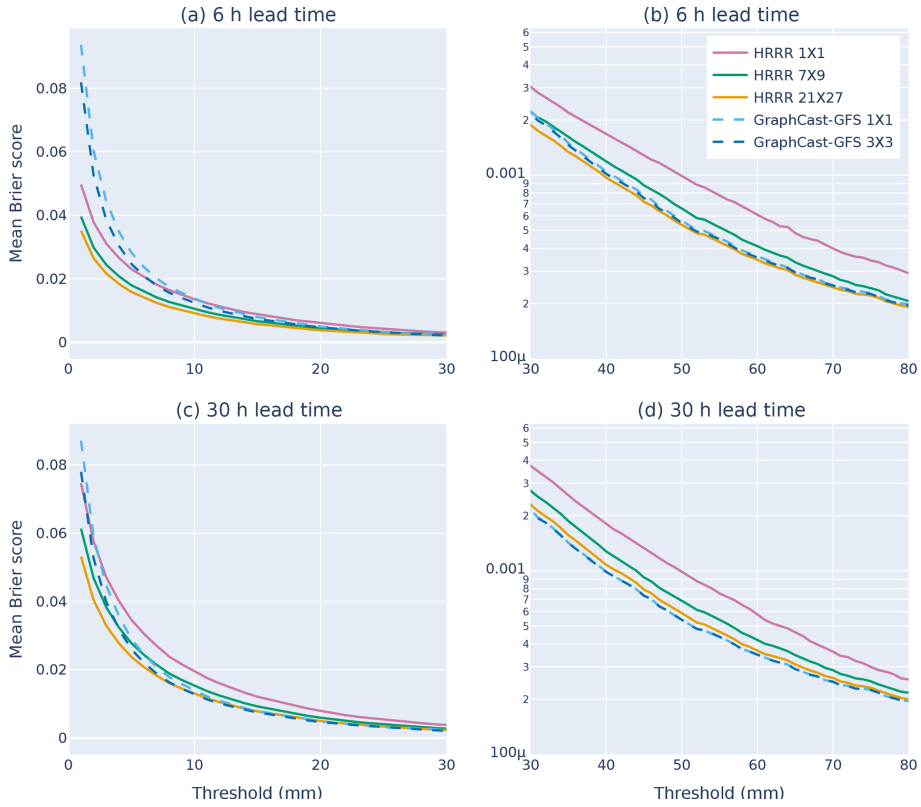


Figure 5. Brier score decomposition of the CRPS within the HiRA framework. Lower scores are better. The left panels **(a)** and **(c)** show the mean Brier score for thresholds below 30 mm, while the right panels **(b)** and **(d)** show the mean Brier score for threshold between 30 and 80 mm with a logarithmic vertical axis. Results are shown for 6 h lead time forecasts in panels **(a)** and **(b)**, and lead time 30 h forecasts in **(c)** and **(d)**.

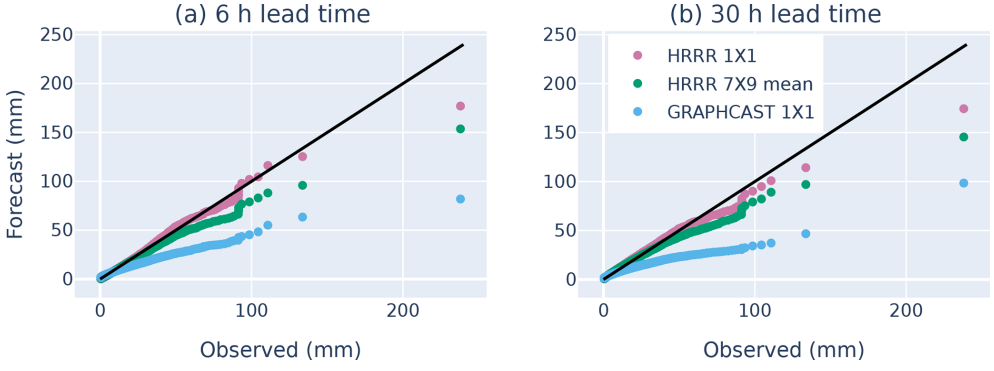


Figure 6. Q–Q plots of observations against forecasts. **(a)** shows results for 6 h lead time forecasts and **(b)** shows results for 30 h lead times.

The potential CRPS and potential twCRPS results, which measure the discrimination ability, show that while increasing neighbourhood size leads to better discrimination ability, the HRRR and GraphCast-GFS are sometimes ranked differently compared to Figs. 3 and 4, which measure overall predictive performance rather than discrimination ability.

The potential CRPS results indicate that, when all decision thresholds are weighted equally, the HRRR 7×9 shows stronger discrimination ability than GraphCast-GFS 1×1 across all lead times, and the HRRR 21×27 exhibits stronger discrimination ability than GraphCast-GFS 3×3 at short lead times, with the discrimination ability of the two models converging thereafter.

Examining discrimination ability for extreme decision thresholds based on the potential twCRPS results, the HRRR model outperforms GraphCast-GFS at a lead time of 6 h for all equivalent neighbourhood sizes; however, this advantage diminishes with increasing lead time, and GraphCast-GFS

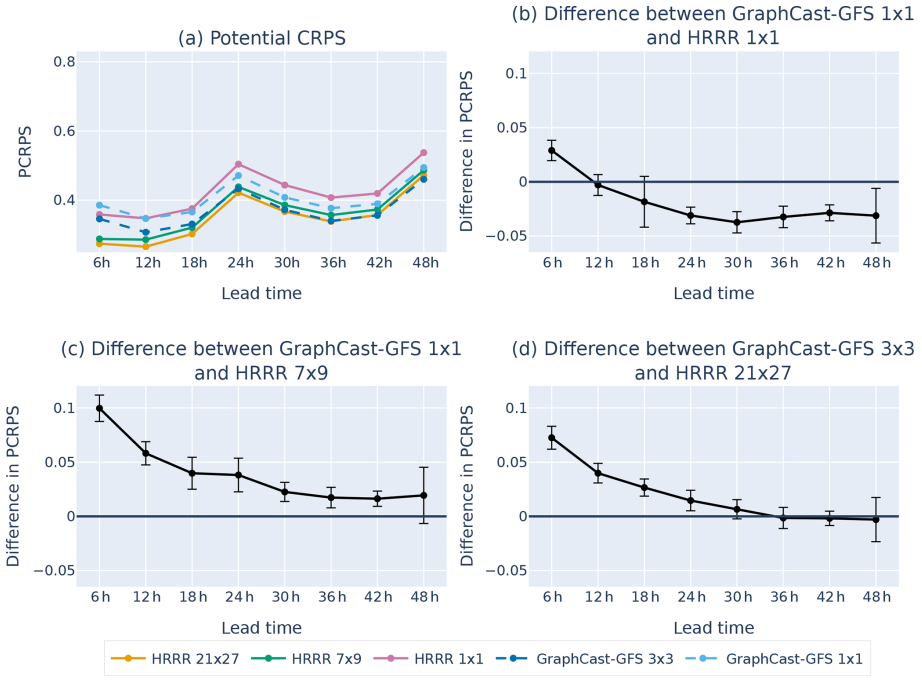


Figure 7. **(a)** Potential CRPS results. Lower values indicate more discrimination ability when all decision thresholds are weighted equally. **(b)** Difference between GraphCast-GFS 1×1 and HRRR 1×1 with 99 % confidence intervals. **(c)** Difference between GraphCast-GFS 1×1 and HRRR 7×9 with 99 % confidence intervals. **(d)** Difference between GraphCast-GFS 3×3 and HRRR 21×27 with 99 % confidence intervals. In subfigures **(b–d)**, positive values indicate that HRRR had more discrimination ability than GraphCast-GFS for the specified neighbourhoods.

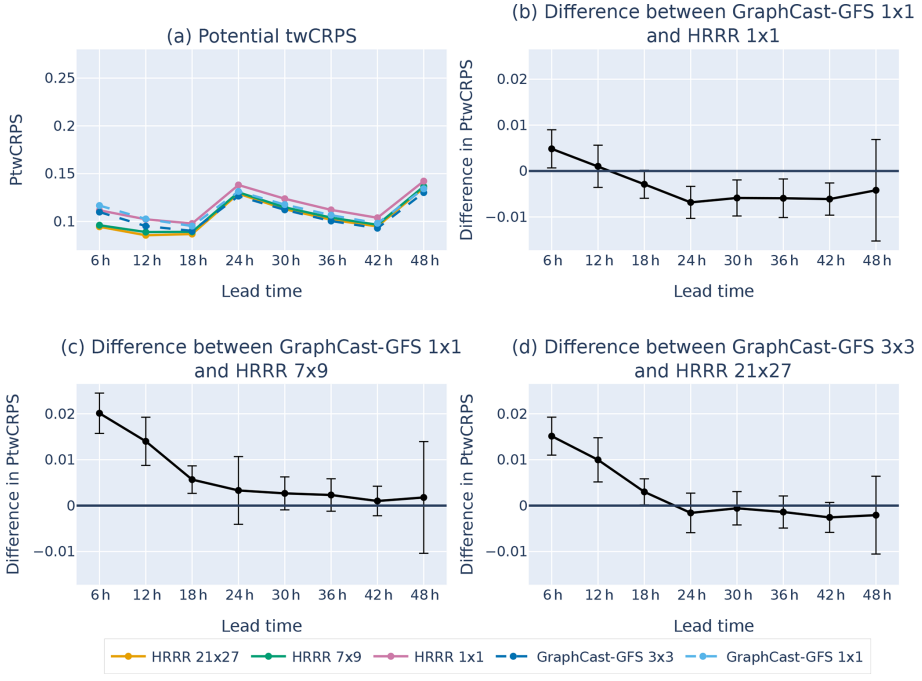


Figure 8. **(a)** As for Fig. 7, but for potential twCRPS. Lower values indicate more discrimination ability for predicting extremes (all thresholds above the climatological 99th percentile).

exhibits greater discrimination ability at lead times of 24 h and beyond for the GraphCast-GFS 3×3 versus HRRR 21×27 comparison. Differences in discrimination ability for the latter comparison were not statistically significant.

8 Conclusions

This paper demonstrates an approach to evaluating how an AIWP model compares to a high-resolution physical NWP model in predicting climatologically extreme 6-hourly precipitation. It combines two existing techniques: the HiRA framework and twCRPS. Model performance was assessed within a framework representative of potential operational use by meteorologists or simple post-processing systems. As with all NWP verification methods, this approach does not definitively determine which model is superior across all use cases. The comparison includes, rather than removes, differences associated with grid spacing, effective resolution, and spatial sampling. It is not a comparison after re-gridding the forecast systems to a common spatial scale.

When models were compared using neighbourhood sizes that covered a similar physical area to generate empirical CDFs, HRRR consistently outperformed GraphCast-GFS across all lead times as measured by CRPS. However, when focusing on predictive performance of extreme precipitation, HRRR only outperformed GraphCast-GFS at short lead times.

Decomposing CRPS by decision threshold provided further insight into which thresholds each model handled more effectively. Where Brier score curves for HRRR and GraphCast-GFS intersected, this generally occurred at higher precipitation thresholds for larger HRRR neighbourhood sizes.

The approach was extended to measure the discrimination ability of the models to predict extreme precipitation. This was important to measure since GraphCast-GFS forecasts of heavy precipitation showed an under-forecast bias when evaluated against rain gauge data. Increasing the neighbourhood size led to superior discrimination ability. When GraphCast-GFS 3×3 versus HRRR 21×27 were compared, the high-resolution model had superior discrimination ability at short lead times, but not beyond 24 h.

This verification approach has several strengths:

1. It is a spatial method that can address the double penalty effect within a specified distance.

2 It allows forecast systems with different grid spacings to be evaluated on their native grids when the forecast being assessed is defined as the finite neighbourhood empirical distribution.

3. It can be used to evaluate models against in-situ observations.

4. It can be threshold-weighted to focus on extremes or other important decision thresholds.

5. It uses proper scoring rules within a user-focused framework, avoiding selection bias associated with conditioning on extreme observed or forecast events.

6. It supports aggregation across domains by accounting for climatological differences via threshold weighting if required.

As noted by Pic et al. (2025), there is an opportunity to bridge the gap between the spatial verification methods and proper scoring rules communities. This work offers one such integration, and raises several opportunities for future research:

- The approach could be extended to evaluate multivariate forecasts by using threshold-weighted multivariate scores, such as the threshold-weighted variogram score (Allen et al., 2023) which is a proper scoring rule. This approach could be used to evaluate compound events as well as different variables simultaneously (e.g., surface wind speed, temperature and relative humidity for fire weather).

- Comparing this approach to other spatial verification methods.

- Applying threshold-weighted scoring rules within HiRA to ensemble forecasts, as in Mittermaier and Csima (2017).

- Extending the approach to also account for timing errors.

- Applying neighbourhood twCRPS to forecasts evaluated with gridded rainfall observations, which is similar to Stein and Stoop (2022) who used a neighbourhood CRPS with gridded data.

- Comparing a potential CRPS measure (Gneiting et al., 2025) that applies IDR using neighbourhood members as covariates. This may be useful when models are spatially sharper.

Finally, while most prior comparisons between AIWP and traditional NWP models have relied on point-to-point verification, this study shows that model rankings can differ when assessed from a spatial perspective, consistent with findings from Radford et al. (2025a). If AIWP models are to be operationalised, adopting spatial verification methods aligned with their practical use will be essential for accurately understanding their performance.

Appendix A: twCRPS results for the 0.999 climatological threshold

Results for the twCRPS using a 0.999 climatological threshold are shown in Fig. A1. The only instance in which the HRRR was significantly better than GraphCast-GFS was the 6 h lead time when GraphCast-GFS 3×3 and HRRR 21×27 were compared. Results were also calculated using a fixed 50 mm threshold at all locations and were similar (not shown).

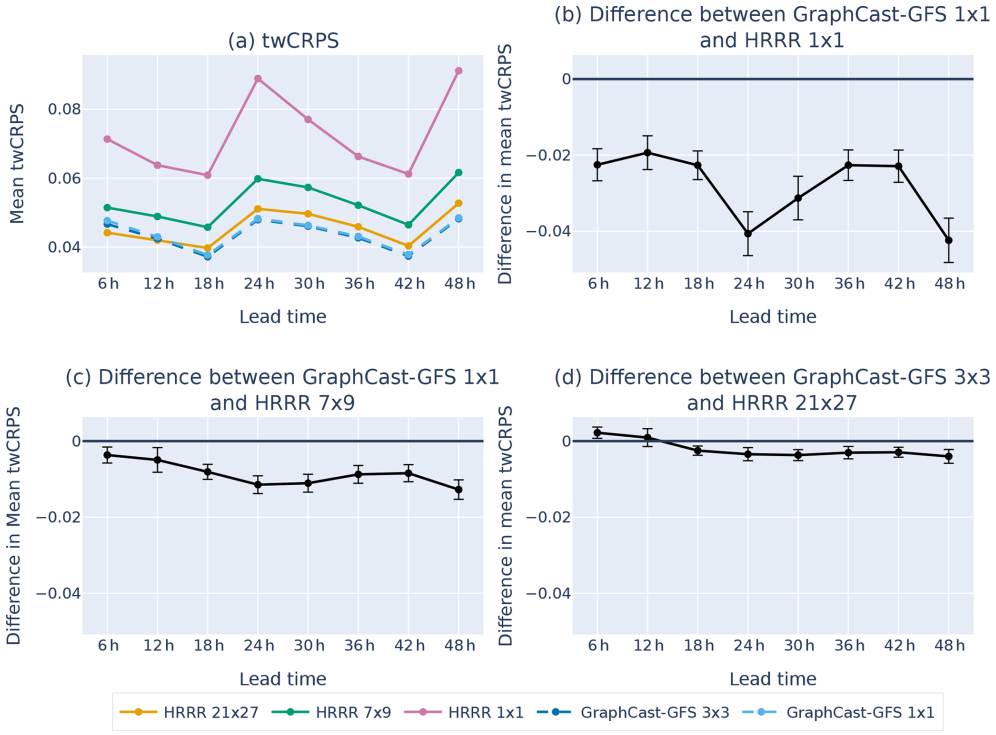


Figure A1. As for Fig. 3 but for the twCRPS with a threshold weight function of $w(z) = \mathbb{1}(z > q_{0.999})$.

Code and data availability. All code to reproduce the analysis and figures in this paper is available at https://doi.org/10.5281/zenodo.20438050 (Loveday, 2026).

All data used to reproduce this work is available at https://doi.org/10.5281/zenodo.19672887 (Loveday and Hertneky, 2026).

The original data sources for the data stored in the above record were retrieved from the following sources:

- One minute ASOS data were retrieved from https://mesonet.agron.iastate.edu/request/asos/1min.phtml (last access: 9 September 2026) which contains an archive of data provided by the National Climatic Data Center.
- HRRR data was retrieved from https://hrrrzarr.s3.amazonaws.com/index.html (last access: 9 September 2026).
- GraphCast-GFS data was retrieved from https://noaa-oar-mlwp-data.s3.amazonaws.com/index.html (last access: 9 September 2026).
- ERA5 data was retrieved from https://console.cloud.google.com/storage/browser/weatherbench2/data/era5 (last access: 9 September 2026).

The twCRPS, CRPS for ensembles, Brier score for ensembles, and Q–Q plots were added to the *scores* Python package (Leeuwenburg et al., 2024) as part of this work and can also be found at https://github.com/nci/scores (last access: 9 September 2026). The specific version of *scores* used in this work was 2.5.0 (https://doi.org/10.5281/zenodo.18638494; Leeuwenburg et al., 2026). Other methods including the Diebold Mariano (with the Hering–Genton modification) test and isotonic regression can also be found in the *scores* Python package.

Author contributions. NL Conceptualisation, Data Curation, Formal Analysis, Investigation, Methodology, Software, Visualisation, Writing – Original Draft Preparation, Writing – Review and Edit-

ing. TH Writing – Review and Editing, Project Administration, Resources.

Competing interests. The contact author has declared that neither of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Acknowledgements. The authors thank Louisa Nance (NSF NCAR and DTC), Belinda Trotta (BoM), Robert Taggart (BoM), and two anonymous reviewers for feedback on earlier versions of this manuscript. Nicholas Loveday would like to thank everyone at NCAR and NOAA GSL who supported his visit as a DTC Visiting Scientist.

Financial support. This research was funded by NSF NCAR's Developmental Testbed Center (DTC) Visiting Scientist Program.

Review statement. This paper was edited by Tao Zhang and reviewed by three anonymous referees.

References

Abdi, D., Jankov, I., Madden, P., Vargas, V., Smith, T. A., Frolov, S., Flora, M., and Potvin, C.: HRRRCast: a data-driven emulator for regional weather forecasting at convection allowing scales, arXiv [preprint], arXiv:2507.05658, 2025.

Adamov, S., Oskarsson, J., Denby, L., Landelius, T., Hintz, K., Christiansen, S., Schicker, I., Osuna, C., Lindsten, F., Fuhrer, O., and Schemm, S.: Building Machine Learning Limited Area Models: Kilometer-Scale Weather Forecasting in Realistic Settings, arXiv [preprint], arXiv:2504.09340, 2025.

Allen, S.: Weighted scoringRules: Emphasizing Particular Outcomes When Evaluating Probabilistic Forecasts, J. Stat. Softw., 110, https://doi.org/10.18637/jss.v110.i08, 2024.

Allen, S., Ginsbourger, D., and Ziegel, J.: Evaluating Forecasts for High-Impact Events Using Transformed Kernel Scores, SIAM/ASA Journal on Uncertainty Quantification, 11, 906–940, https://doi.org/10.1137/22m1532184, 2023.

Ayer, M., Brunk, H. D., Ewing, G. M., Reid, W. T., and Silverman, E.: An Empirical Distribution Function for Sampling with Incomplete Information, Ann. Math. Stat., 26, 641–647, https://doi.org/10.1214/aoms/1177728423, 1955.

Baringhaus, L. and Franz, C.: On a new multivariate two-sample test, J. Multivariate Anal., 88, 190–206, https://doi.org/10.1016/s0047-259x(03)00079-4, 2004.

Ben Bouallègue, Z., Clare, M. C. A., Magnusson, L., Gascón, E., Maier-Gerber, M., Janoušek, M., Rodwell, M., Pinault, F., Dramsch, J. S., Lang, S. T. K., Raoult, B., Rabier, F., Chevallier, M., Sandu, I., Dueben, P., Chantry, M., and Pappenberger, F.: The Rise of Data-Driven Weather Forecasting: A First Statistical Assessment of Machine Learning–Based Weather Forecasts in an Operational-Like Context, B. Am. Meteorol. Soc., 105, E864–E883, https://doi.org/10.1175/bams-d-23-0162.1, 2024.

Ben-Bouallègue, Z., Prieto-Nemesio, A., Iza Wong, A., Pinault, F., van der Schee, M., and Modigliani, U.: SEEPS4ALL: an open dataset for the verification of daily precipitation forecasts using station climate statistics, Earth Syst. Sci. Data, 18, 713–720, https://doi.org/10.5194/essd-18-713-2026, 2026.

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, 533–538, https://doi.org/10.1038/s41586-023-06185-3, 2023.

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, Geophys. Res. Lett., 52, https://doi.org/10.1029/2024gl113656, 2025.

Brier, G. W.: VERIFICATION OF FORECASTS EXPRESSED IN TERMS OF PROBABILITY, Mon. Weather Rev., 78, 1–3, https://doi.org/10.1175/1520-0493(1950)078<0001:vofeit>2.0.co;2, 1950.

Bröcker, J.: Estimating reliability and resolution of probability forecasts through decomposition of the empirical score, Clim. Dynam., 39, 655–667, https://doi.org/10.1007/s00382-011-1191-1, 2011.

Charlton-Perez, A. J., Dacre, H. F., Driscoll, S., Gray, S. L., Harvey, B., Harvey, N. J., Hunt, K. M. R., Lee, R. W., Swaminathan, R., Vandaele, R., and Volonté, A.: Do AI models produce better weather forecasts than physics-based models? A quantitative evaluation case study of Storm Ciarán, npj Clim. Atmos. Sci., 7, https://doi.org/10.1038/s41612-024-00638-w, 2024.

Diebold, F. X. and Mariano, R. S.: Comparing Predictive Accuracy, J. Bus. Econ. Stat., 13, 253–263, https://doi.org/10.1080/07350015.1995.10524599, 1995.

Dorninger, M., Gilleland, E., Casati, B., Mittermaier, M. P., Ebert, E. E., Brown, B. G., and Wilson, L. J.: The Setup of the MesoVICT Project, B. Am. Meteorol. Soc., 99, 1887–1906, https://doi.org/10.1175/bams-d-17-0164.1, 2018.

Dowell, D. C., Alexander, C. R., James, E. P., Weygandt, S. S., Benjamin, S. G., Manikin, G. S., Blake, B. T., Brown, J. M., Olson, J. B., Hu, M., Smirnova, T. G., Ladwig, T., Kenyon, J. S., Ahmadov, R., Turner, D. D., Duda, J. D., and Alcott, T. I.: The High-Resolution Rapid Refresh (HRRR): An Hourly Updating Convection-Allowing Forecast Model. Part I: Motivation and System Description, Weather Forecast., 37, 1371–1395, https://doi.org/10.1175/waf-d-21-0151.1, 2022.

Ebert, E. E.: Fuzzy verification of high-resolution gridded forecasts: a review and proposed framework, Meteorol. Appl., 15, 51–64, https://doi.org/10.1002/met.25, 2008.

Ebert, E. E.: Neighborhood Verification: A Strategy for Rewarding Close Forecasts, Weather Forecast., 24, 1498–1510, https://doi.org/10.1175/2009waf2222251.1, 2009.

Epstein, E. S.: A Scoring System for Probability Forecasts of Ranked Categories, J. Appl. Me-

teorol., 8, 985–987, https://doi.org/10.1175/1520-0450(1969)008<0985:assfpf>2.0.co;2, 1969.

Ferro, C. A. T.: Fair scores for ensemble forecasts: Fair Scores for Ensemble Forecasts, Q. J. Roy. Meteor. Soc., 140, 1917–1923, https://doi.org/10.1002/qj.2270, 2013.

Fovell, R. G. and Gallagher, A.: An Evaluation of Surface Wind and Gust Forecasts from the High-Resolution Rapid Refresh Model, Weather Forecast., 37, 1049–1068, https://doi.org/10.1175/waf-d-21-0176.1, 2022.

Gilleland, E., Ahijevych, D., Brown, B. G., Casati, B., and Ebert, E. E.: Intercomparison of Spatial Forecast Verification Methods, Weather Forecast., 24, 1416–1430, https://doi.org/10.1175/2009waf2222269.1, 2009.

Gneiting, T.: Making and Evaluating Point Forecasts, J. Am. Stat. Assoc., 106, 746–762, https://doi.org/10.1198/jasa.2011.r10138, 2011.

Gneiting, T. and Katzfuss, M.: Probabilistic Forecasting, Annu. Rev. Stat. Appl., 1, 125–151, https://doi.org/10.1146/annurev-statistics-062713-085831, 2014.

Gneiting, T. and Raftery, A. E.: Strictly Proper Scoring Rules, Prediction, and Estimation, J. Am. Stat. Assoc., 102, 359–378, https://doi.org/10.1198/016214506000001437, 2007.

Gneiting, T. and Ranjan, R.: Comparing Density Forecasts Using Threshold- and Quantile-Weighted Scoring Rules, J. Bus. Econ. Stat., 29, 411–422, https://doi.org/10.1198/jbes.2010.08110, 2011.

Gneiting, T., Biegert, T., Kraus, K., Walz, E.-M., Jordan, A. I., and Lerch, S.: Probabilistic measures afford fair comparisons of AIWP and NWP model output, arXiv [preprint], arXiv:2506.03744, 2025.

Henzi, A., Ziegel, J. F., and Gneiting, T.: Isotonic Distributional Regression, J. Roy. Stat. Soc. B, 83, 963–993, https://doi.org/10.1111/rssb.12450, 2021.

Hering, A. S. and Genton, M. G.: Comparing Spatial Predictions, Technometrics, 53, 414–425, https://doi.org/10.1198/tech.2011.10136, 2011.

Hersbach, H.: Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems, Weather Forecast., 15, 559–570, https://doi.org/10.1175/1520-0434(2000)015<0559:dotcrp>2.0.co;2, 2000.

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, Q. J. Roy. Meteor. Soc., 146, 1999–2049, https://doi.org/10.1002/qj.3803, 2020.

Ikeda, K., Steiner, M., Pinto, J., and Alexander, C.: Evaluation of Cold-Season Precipitation Forecasts Generated by the Hourly Updating High-Resolution Rapid Refresh Model, Weather Forecast., 28, 921–939, https://doi.org/10.1175/waf-d-12-00085.1, 2013.

Jin, W., Weyn, J., Zhao, P., Xiang, S., Bian, J., Fang, Z., Dong, H., Sun, H., Thambiratnam, K., and Zhang, Q.: WeatherReal: A Benchmark Based on In-Situ Observations for Evaluating Weather Models, arXiv [preprint], arXiv:2409.09371, 2024.

Keisler, R.: Forecasting Global Weather with Graph Neural Networks, arXiv [preprint], arXiv:2202.07575, 2022.

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, 1416–1421, https://doi.org/10.1126/science.adi2336, 2023.

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, arXiv [preprint], arXiv:2406.01465, 2024.

Lang, S., Alexe, M., Clare, M. C. A., Roberts, C., Adewoyin, R., Ben Bouallègue, Z., 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, npj Artificial Intelligence, 2, https://doi.org/10.1038/s44387-026-00073-7, 2026.

Lavers, D. A., Simmons, A., Vamborg, F., and Rodwell, M. J.: An evaluation of ERA5 precipitation for climate monitoring, Q. J. Roy. Meteor. Soc., 148, 3152–3165, https://doi.org/10.1002/qj.4351, 2022.

Leeuwenburg, T., Loveday, N., Ebert, E. E., Cook, H., Khanarmuei, M., Taggart, R. J., Ramanathan, N., Carroll, M., Chong, S., Griffiths, A., and Sharples, J.: scores: A Python package for verifying and evaluating models and predictions with xarray, Journal of Open Source Software, 9, 6889, https://doi.org/10.21105/joss.06889, 2024.

Leeuwenburg, T., Loveday, N., Ramanathan, N., Chong, S., Taggart, R. J., Shrestha, D., Khanarmuei, M., Cook, H., Bluett, L., Ebert, E. E., Carroll, M., Trotta, B., Sharples, J., Bishop, S., Squire, D. T., Griffiths, A., Pagano, T. C., Fisher, A., Mandelbaum, T., Jinghan, F., Smith, P. R., Abellan, E., Beunk, J., Esperson, F., Smallwood, J., and Wu, X.: scores: Metrics for the verification, evaluation and optimisation of forecasts, predictions or models, Zenodo [code], https://doi.org/10.5281/zenodo.18638494, 2026.

Lerch, S., Thorarinsdottir, T. L., Ravazzolo, F., and Gneiting, T.: Forecaster's Dilemma: Extreme Events and Forecast Evaluation, Stat. Sci., 32, https://doi.org/10.1214/16-sts588, 2017.

Loveday, N.: Evaluating Extreme Precipitation Forecasts: A Threshold-Weighted, Spatial Verification Approach for Comparing an AI Weather Prediction Model Against a High-Resolution NWP Model, Zenodo [code], https://doi.org/10.5281/zenodo.20438050, 2026.

Loveday, N. and Hertneky, T.: Data for: Evaluating Extreme Precipitation Forecasts: A Threshold-Weighted, Spatial Verification Approach for Comparing an AI Weather Prediction Model Against a High-Resolution NWP Model, Zenodo [data set], https://doi.org/10.5281/zenodo.19672887, 2026.

Loveday, N., Griffiths, D., Leeuwenburg, T., Taggart, R. J., Pagano, T. C., Cheng, G., Plastow, K., Ebert, E. E., Templeton, C., Carroll, M., Khanarmuei, M., and Nagpal, I.: The Jive Verification System and Its Transformative Impact on Weather Fore-

casting Operations, B. Am. Meteorol. Soc., 105, E2047–E2063, https://doi.org/10.1175/bams-d-23-0267.1, 2024.

Mass, C. F., Ovens, D., Westrick, K., and Colle, B. A.: Does Increasing Horizontal Resolution Produce More Skillful Forecasts?, B. Am. Meteorol. Soc., 83, 407–430, https://doi.org/10.1175/1520-0477(2002)083<0407:dihrpm>2.3.co;2, 2002.

Matheson, J. E. and Winkler, R. L.: Scoring Rules for Continuous Probability Distributions, Management Science, 22, 1087–1096, https://doi.org/10.1287/mnsc.22.10.1087, 1976.

McGovern, A., Mandelbaum, T., Rothenberg, D., Loveday, N., Potvin, C., Flora, M., Magnusson, L., Gilleland, E., and Allen, J.: Extreme Weather Bench: A framework and benchmark for evaluation of high-impact weather, arXiv [preprint], arXiv:2605.01126, 2026.

Mittermaier, M. P.: A Strategy for Verifying Near-Convection-Resolving Model Forecasts at Observing Sites, Weather Forecast., 29, 185–204, https://doi.org/10.1175/waf-d-12-00075.1, 2014.

Mittermaier, M. P. and Csima, G.: Ensemble versus Deterministic Performance at the Kilometer Scale, Weather Forecast., 32, 1697–1709, https://doi.org/10.1175/waf-d-16-0164.1, 2017.

Morisseau, H., Zhu, H., Hudson, D., and de Burgh-Day, C.: Object-oriented verification of TC-Jasper rainfall forecasts: Machine learning, Tech. Rep. Bureau Research Report No. 106, Bureau of Meteorology, http://www.bom.gov.au/research/publications/researchreports/BRR-106.pdf (last access: 21 November 2025), 2025.

Murphy, A. H. and Epstein, E. S.: Skill Scores and Correlation Coefficients in Model Verification, Mon. Weather Rev., 117, 572–582, https://doi.org/10.1175/1520-0493(1989)117<0572:ssacci>2.0.co;2, 1989.

Nipen, T. N., Haugen, H. H., Ingstad, M. S., Nordhagen, E. M., Salihi, A. F. S., Tedesco, P., Seierstad, I. A., Kristiansen, J., Lang, S., Alexe, M., Dramsch, J., Raoult, B., Mertes, G., and Chantry, M.: Regional Data-Driven Weather Modeling with a Global Stretched Grid, Artificial Intelligence for the Earth Systems, 5, https://doi.org/10.1175/aies-d-25-0001.1, 2026.

NWS: Automated Surface Observing System (ASOS) user's guide, Tech. rep., NOAA, 1998.

Olivetti, L. and Messori, G.: Do data-driven models beat numerical models in forecasting weather extremes? A comparison of IFS HRES, Pangu-Weather, and GraphCast, Geosci. Model Dev., 17, 7915–7962, https://doi.org/10.5194/gmd-17-7915-2024, 2024.

Pagano, T. C., Casati, B., Landman, S., Loveday, N., Taggart, R., Ebert, E. E., Khanarmuei, M., Jensen, T. L., Mittermaier, M., Roberts, H., Willington, S., Roberts, N., Sowko, M., Strassberg, G., Kluepfel, C., Bullock, T. A., Turner, D. D., Pappenberger, F., Osborne, N., and Noble, C.: Challenges of Operational Weather Forecast Verification and Evaluation, B. Am. Meteorol. Soc., 105, E789–E802, https://doi.org/10.1175/bams-d-22-0257.1, 2024.

Pic, R., Dombry, C., Naveau, P., and Taillardat, M.: Proper scoring rules for multivariate probabilistic forecasts based on aggregation and transformation, Adv. Stat. Clim. Meteorol. Oceanogr., 11, 23–58, https://doi.org/10.5194/ascmo-11-23-2025, 2025.

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, 84–90, https://doi.org/10.1038/s41586-024-08252-9, 2024.

Radford, J. T., Ebert-Uphoff, I., and Stewart, J. Q.: A Comparison of AI Weather Prediction and Numerical Weather Prediction Models for 1–7-Day Precipitation Forecasts, Weather Forecast., 40, https://doi.org/10.1175/waf-d-24-0081.1, 2025a.

Radford, J. T., Ebert-Uphoff, I., Stewart, J. Q., Musgrave, K. D., DeMaria, R., Tourville, N., and Hilburn, K.: Accelerating Community-Wide Evaluation of AI Models for Global Weather Prediction by Facilitating Access to Model Output, B. Am. Meteorol. Soc., 106, E68–E76, https://doi.org/10.1175/bams-d-24-0057.1, 2025b.

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, J. Adv. Model. Earth Sy., 16, https://doi.org/10.1029/2023ms004019, 2024.

Rodwell, M. J., Richardson, D. S., Hewson, T. D., and Haiden, T.: A new equitable score suitable for verifying precipitation in numerical weather prediction, Q. J. Roy. Meteor. Soc., 136, 1344–1363, https://doi.org/10.1002/qj.656, 2010.

Schwartz, C. S. and Sobash, R. A.: Generating Probabilistic Forecasts from Convection-Allowing Ensembles Using Neighborhood Approaches: A Review and Recommendations, Mon. Weather Rev., 145, 3397–3418, https://doi.org/10.1175/mwr-d-16-0400.1, 2017.

Siegert, S.: Simplifying and generalising Murphy's Brier score decomposition, Q. J. Roy. Meteor. Soc., 143, 1178–1183, https://doi.org/10.1002/qj.2985, 2017.

Stein, J. and Stoop, F.: Neighborhood-Based Ensemble Evaluation Using the CRPS, Mon. Weather Rev., 150, 1901–1914, https://doi.org/10.1175/mwr-d-21-0224.1, 2022.

Subich, C., Husain, S. Z., Separovic, L., and Yang, J.: Fixing the Double Penalty in Data-Driven Weather Forecasting Through a Modified Spherical Harmonic Loss Function, arXiv [preprint], arXiv:2501.19374, 2025.

Székely, G. J. and Rizzo, M. L.: A new test for multivariate normality, J. Multivariate Anal., 93, 58–80, https://doi.org/10.1016/j.jmva.2003.12.002, 2005.

Taggart, R.: Evaluation of point forecasts for extreme events using consistent scoring functions, Q. J. Roy. Meteor. Soc., 148, 306–320, https://doi.org/10.1002/qj.4206, 2021.

Theis, S. E., Hense, A., and Damrath, U.: Probabilistic precipitation forecasts from a deterministic model: a pragmatic approach, Meteorol. Appl., 12, 257–268, https://doi.org/10.1017/s1350482705001763, 2005.

Tustison, B., Harris, D., and Foufoula-Georgiou, E.: Scale issues in verification of precipitation forecasts, J. Geophys. Res.-Atmos., 106, 11 775–11 784, https://doi.org/10.1029/2001jd900066, 2001.

Walz, E.-M., Henzi, A., Ziegel, J., and Gneiting, T.: Easy Uncertainty Quantification (EasyUQ): Generating Predictive Distributions from Single-Valued Model Output, SIAM Rev., 66, 91–122, https://doi.org/10.1137/22m1541915, 2024.

Wessel, J. B., Ferro, C. A. T., Evans, G. R., and Kwasniok, F.: Improving probabilistic forecasts of extreme wind speeds by training statistical post-processing models with weighted scor-

ing rules, Mon. Weather Rev., https://doi.org/10.1175/mwr-d-24-0151.1, 2025.

Winkler, R. L. and Murphy, A. H.: "Good" Probability Assessors, J. Appl. Meteorol., 7, 751–758, https://doi.org/10.1175/1520-0450(1968)007<0751:pa>2.0.co;2, 1968.

WMO: Guide to Hydrological Practices, Tech. Rep. 168, World Meteorological Organization, 1994.

Zhang, Z., Fischer, E., Zscheischler, J., and Engelke, S.: Physics-based models outperform AI weather forecasts of record-breaking extremes, Science Advances, 12, https://doi.org/10.1126/sciadv.aec1433, 2026.