

ATMOSPHERIC SCIENCE

Physics-based models outperform AI weather forecasts of record-breaking extremes

Zhongwei Zhang^{1,2*}, Erich Fischer³, Jakob Zscheischler^{4,5,6}, Sebastian Engelke^{2*}

Artificial intelligence (AI)-based models are revolutionizing weather forecasting and have surpassed leading numerical weather prediction systems on various benchmark tasks. However, their ability to extrapolate and reliably forecast unprecedented extreme events remains unclear. Here, we show that for record-breaking weather extremes, the physics-based numerical model High RESolution forecast (HRES) from the European Centre for Medium-Range Weather Forecasts still consistently outperforms state-of-the-art AI models GraphCast, GraphCast operational, Pangu-Weather, Pangu-Weather operational, and Fuxi. We demonstrate that forecast errors in AI models are consistently larger for record-breaking heat, cold, and wind than in HRES across nearly all lead times. We further find that the examined AI models tend to underestimate both the frequency and intensity of record-breaking events, and they underpredict hot records and overestimate cold records with growing errors for larger record exceedance. Our findings underscore the current limitations of AI weather models in extrapolating beyond their training domain and in forecasting the potentially most impactful record-breaking weather events that are particularly frequent in a rapidly warming climate. Further rigorous verification and model development is needed before these models can be solely relied upon for high-stakes applications such as early warning systems and disaster management.

INTRODUCTION

Record-breaking weather extremes, such as the 2021 Pacific Northwest, 2010 Russian and 2003 European heatwaves, and winter storms Lothar in 1999 and Kyrill in 2007, have caused numerous fatalities and severe impacts on society, the economy, and ecosystems (1–5). The level of disaster preparedness and adaptation to extreme events is strongly influenced by events observed in recent decades. Consequently, after extended periods without major events, or when events substantially exceed previous record levels, socioeconomic impacts tend to be particularly large.

In addition to long-term disaster preparedness (6), accurate physics-based numerical weather prediction (NWP) is critical for early-warning systems to save lives and reduce the impacts of climate extremes (7). Recently, a new generation of artificial intelligence (AI) weather models has reached and sometimes exceeded forecast skills of state-of-the-art physics-based NWP systems (8–11). These models offer considerable advantages in speed and energy efficiency, raising important questions about their potential to supplement or eventually replace traditional physics-based NWP systems (12).

Before warnings for population and critical infrastructure are routinely based on AI models, their performance needs to be further evaluated. In particular, their reliability in forecasting extreme events remains less well understood. Such events are, by definition, rare and contribute little to aggregated overall skill metrics (13).

Nevertheless, recent studies suggest that AI models perform well—and in some cases even better than numerical models—in forecasting extreme weather events (14, 15), particularly for longer lead times (8, 9).

Current forecast evaluation approaches for extreme events typically focus on extreme events exceeding a certain threshold for one or several given variables, such as extreme wind speeds (15), tropical cyclones (8, 9, 14), and high and low temperatures (9, 14–16). However, due to small sample sizes, the thresholds are often set to, say, the 95th percentile of the test data, thus capturing mostly moderate extremes. Much less is known about record-breaking events, a subset of extreme events that are unprecedented in the observational record. Given the current high rate of global warming, record-breaking events sometimes exceed previous record levels by large margins and have been referred to as black or gray swans (17), or record-shattering extremes (18, 19).

A number of case studies have shown mixed results on the ability of AI weather models to extrapolate beyond the range of their training data. For instance, a seasonal AI forecasting model (20) struggled to predict North Atlantic Oscillation values that extended outside its training distribution (21). While AI models appear to outperform traditional physics-based NWP models on tracking tropical cyclones (8, 9, 22), they tend to underpredict the intensity of the most extreme storms, as measured by mean sea-level pressure (17, 23, 24). Similar limitations in reaching unprecedented amplitudes have also been observed in other high-impact events such as heatwaves, winter storms, or compound extremes (25). On the other hand, the unprecedented 2024 rainfall in Dubai was well predicted by GraphCast, suggesting that generalization to new events may be possible if they share dynamical similarity with past extremes from other regions (26).

However, these insights primarily rely on isolated case studies of specific events, whose conclusions are inherently difficult to generalize due to the unique features of the analyzed events and models. To systematically evaluate extrapolation in state-of-the-art AI weather

¹Institute of Statistics, Karlsruhe Institute of Technology, Karlsruhe, Germany. ²Research Institute for Statistics and Information Science, Geneva School of Economics and Management, University of Geneva, Geneva, Switzerland. ³Institute for Atmospheric and Climate Science, Department of Environmental Systems Science, ETH Zurich, Zurich, Switzerland. ⁴Department of Compound Environmental Risks, Helmholtz Centre for Environmental Research—UFZ, Leipzig, Germany. ⁵Department of Hydro Sciences, TUD Dresden University of Technology, Dresden, Germany. ⁶Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, Germany.

*Corresponding author. Email: zhongwei.zhang@kit.edu (Z.Z.); sebastian.engelke@unige.ch (S.E.)

Copyright © 2026 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution License 4.0 (CC BY).

Downloaded from https://www.science.org at KIT Library on May 04, 2026

models, we construct a benchmark dataset consisting of record-breaking events for heat, cold, and wind extremes. This dataset includes all observations during the test years 2018 and 2020 that exceed the respective historical records from the training data of all considered AI models. The records are defined per variable, per grid cell, and per calendar month by using the ERA5 reanalysis data (27) from 1979–2017 with daily observations at 00, 06, 12, and 18 UTC

time, yielding a large sample size of record-breaking events even in individual test years (see Materials and Methods). For the year 2020, this yields 162,751 heat, 32,991 cold, and 53,345 wind records, which are spread across different seasons and climatic zones from tropics to high latitudes (Fig. 1, A and B, and fig. S1, A to D). The dataset includes many prominent record-breaking events, such as the Siberian heatwave in early 2020 (28) and the U.S. heatwave of August 2020

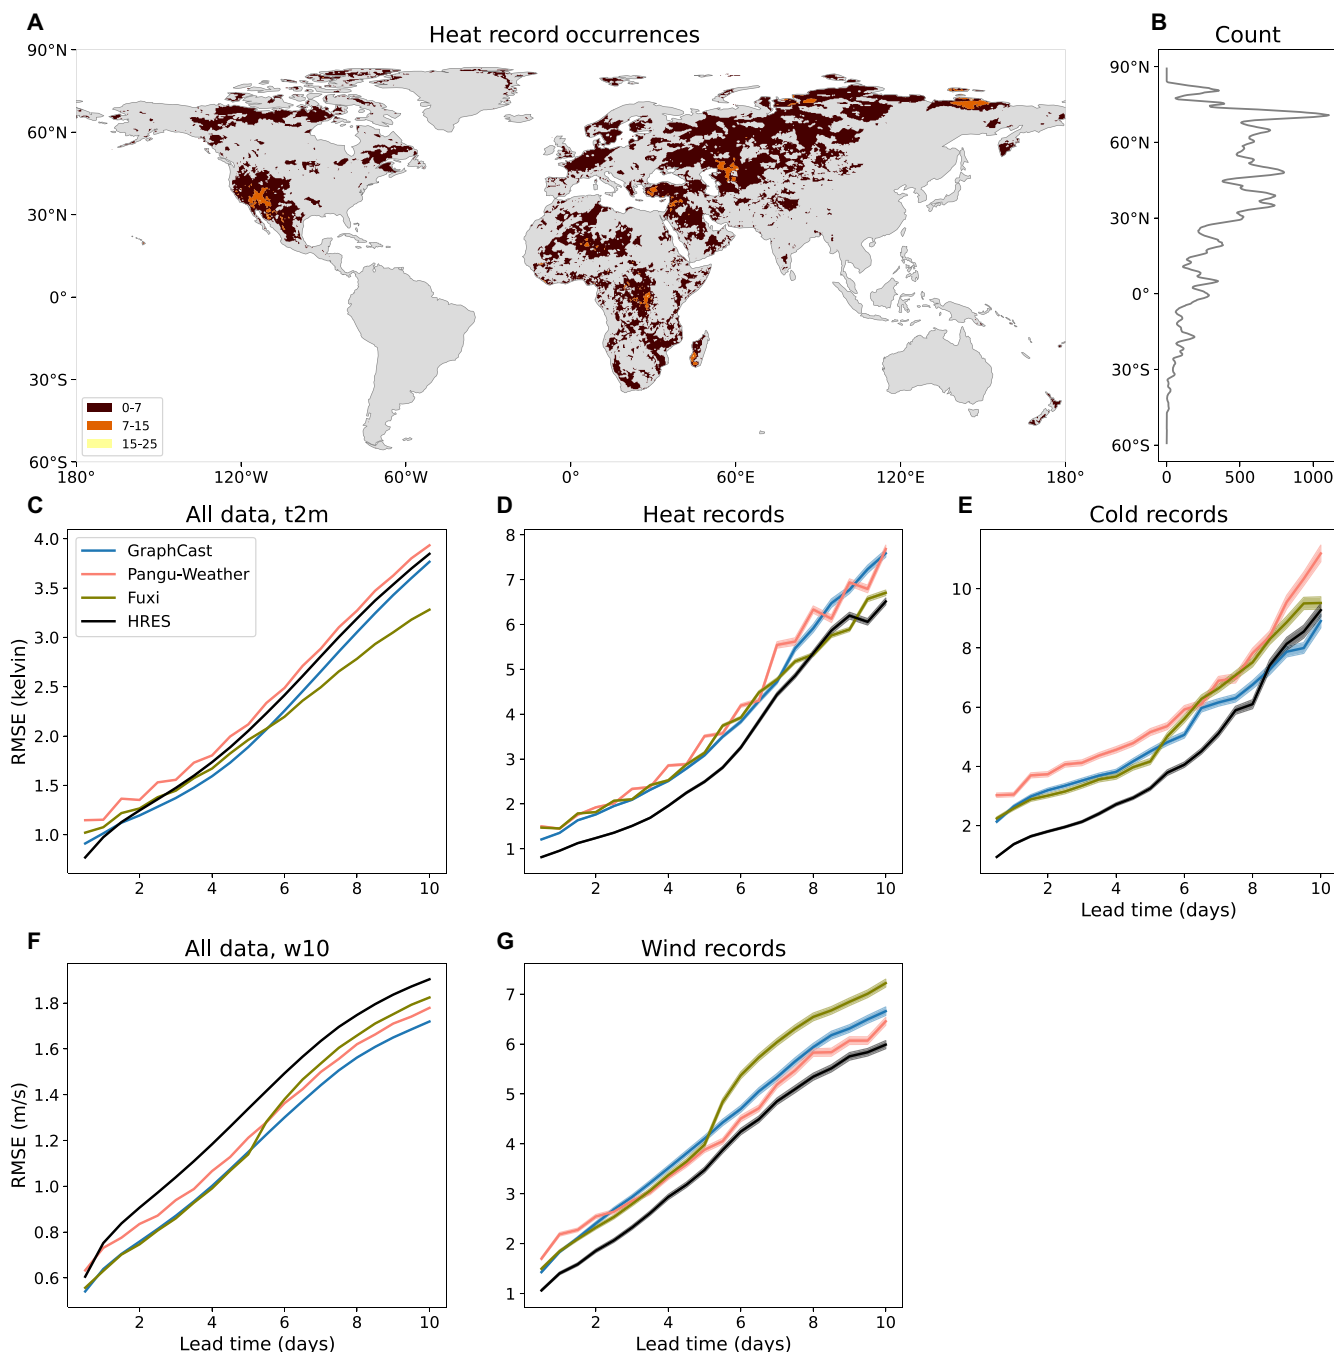


Fig. 1. Model performance on all events and record-breaking events. (A) Number of heat records in 2020 in ERA5. (B) Number of heat records per latitude. (C to G) Root mean square error (RMSE) of forecasted 2-m temperature and 10-m wind speed over land (excluding the Antarctic region) of HRES, Pangu-Weather, GraphCast, and Fuxi for all data [(C) and (F)] and only record-breaking events [(D), (E), and (G)] in 2020 for different lead times. The transparent shaded areas indicate 95% confidence bands.

(29). Evaluating AI models on this record dataset challenges them to forecast on out-of-distribution data, which is known to be difficult for neural networks in the machine learning literature.

We assess the extrapolation performance on our benchmark dataset of record-breaking events of three leading deterministic AI weather models: GraphCast (9), Pangu-Weather (8), and Fuxi (30), as well as the operational variants of GraphCast and Pangu-Weather. Their performance is compared to the physics-based model High Resolution forecast (HRES), which is the deterministic high-resolution configuration of the operational Integrated Forecasting System of the European Centre for Medium-Range Weather Forecasts (ECMWF) and is widely considered as the leading physics-based NWP model.

RESULTS

Model comparison on records' intensity

Consistent with previous studies (8, 9, 30, 31), we find that, on overall performance, all AI models—except Pangu-Weather—outperform the physics-based ECMWF model HRES in forecasting 2-m temperature across most lead times (Fig. 1C). Forecast accuracy is quantified using root mean square errors (RMSEs), computed over all 00 and 12 UTC time steps in test year 2020 and over all land grid points (excluding the Antarctic region; see Materials and Methods). For 10-m wind speed, all AI models consistently outperform HRES across nearly all lead times (Fig. 1F).

However, the predictive skill is drastically different for record-breaking temperature and wind events in 2020. Restricting the RMSE to record-breaking events, the physics-based HRES model consistently outperforms all AI models for hot and cold temperature records as well as wind speed records across almost all lead times (Fig. 1, D, E, and G). The performance gap is most pronounced for short lead times. For lead times beyond 5 days, HRES still generally performs better but to a lesser extent. This aligns with previous findings that AI models tend to perform relatively better at longer lead times (9).

Because of limited data availability [GraphCast forecasts are only available for 2018 and 2020 on WeatherBench 2 (31), and Fuxi forecasts are only available for 2020], the evaluation is shown for a single year only as in most previous studies (8, 15). We observe the same pattern in 2018 (fig. S2). The years 2018 and 2020 are distinctly different in terms of El Niño–Southern Oscillation (ENSO) conditions, with 2018 transitioning from La Niña to El Niño and 2020 undergoing a strong El Niño to La Niña shift. Since ENSO strongly influences the occurrence of temperature records (32), particularly in the tropics, the consistent outperformance of HRES across both years shows the robustness of the results. The better skill of HRES in predicting record-breaking events is further consistent across different seasons and a wide range of different climate zones, including tropics, subtropics, mid-latitudes, and northern high latitudes (Fig. 1A and figs. S3 and S4), although there are few or no record-breaking events in South America, Southeast Asia, maritime continent, or Australia. To remove the temporal dependence in our test records data, we further evaluated the forecasts of record-breaking events that have the largest exceedances per month at each grid point. Again, HRES consistently outperforms the three AI models for almost all lead times (fig. S5).

While it is common to evaluate ERA5-trained AI models against ERA5 reanalysis, and HRES against its own analysis at lead time 0 (HRES-fc0) (8, 9, 30) (see Materials and Methods), this approach

can complicate comparisons due to different horizontal resolution: ERA5 has a resolution of 0.25°, whereas HRES operates at 0.1°. To assess the sensitivity of our findings to the choice of different reference datasets, we also evaluate operational versions of GraphCast and Pangu-Weather against HRES on a common test dataset of record-breaking events identified using HRES-fc0 as observational ground truth. Also in this setting, HRES consistently outperforms the AI models on the records (fig. S6).

Following (9), we have focused on lead times that are multiples of 12 hours, which lead to record-breaking events at 00/12 UTC [as all AI forecasts are only available for initializations at 00/12 UTC on WeatherBench 2 (31)]. When considering lead times at 6 hours, 18 hours, etc., more record-breaking events appear in regions such as South America, Southeast Asia, and Australia (figs. S7A and S8A). As shown in previous studies (8, 30), in this case, the ranking of competing model forecasts remains the same and HRES consistently outperforms the AI models (figs. S7, C to G, and S8, C to G).

Selecting a subset of extreme events based on observations can favor models that produce too many extreme forecasts—a problem known as the forecaster's dilemma (33) (see Materials and Methods for discussion). Thus, we construct an alternative benchmark avoiding the forecaster's dilemma, based on events where the forecast itself, rather than the observation, exceeds the training record (34). Results from this forecast-conditioned evaluation (fig. S9) are consistent with the previous conclusion that HRES outperforms current AI models in forecasting records.

AI models underestimate intensities of records

While we demonstrate that AI models underperform compared to HRES in forecasting record-breaking events, their errors may arise from over- or underprediction of event intensity. When considering all data of the test year 2020, all models have relatively small biases (GraphCast slightly underestimates 2-m temperature, while Fuxi overestimates 2-m temperature for lead times longer than 7 days; HRES and Pangu-Weather also underestimate 2-m temperature for long lead times, but with a smaller bias than GraphCast; all models slightly underestimate 10-m wind speeds) (fig. S10). To better understand model behavior beyond their training domain, we compare forecast accuracy and bias against the record exceedance, that is, the margin by which a record is exceeded. We find that AI models generally underpredict temperature during high records and overpredict during low records. This pattern is shown for GraphCast and heat records (Fig. 2A). The systematic underprediction is remarkably consistent across regions, seasons, and location in tropics, subtropics, and mid- to high-latitudes, despite the fact that the physical drivers of heat records vary substantially across regions. This behavior is not limited to a single model: Other AI models show similar patterns of intensity underestimation, while HRES demonstrates a more balanced distribution of over- and underpredictions (fig. S11). These results strongly suggest that AI model forecast errors are at least partly due to systematic extrapolation limitations.

For all record types, the errors of the three AI models seem to grow almost linearly with respect to the degree of record exceedance (Fig. 2, B to D, for a lead time of 2 days; additional lead times in fig. S12). This trend indicates that forecast bias is the primary driver of error (Fig. 2, E to G, and fig. S10): The greater the record exceedance, the larger the underestimation of event intensity. The models behave as if their predictions have an implicit (soft) cap at a certain

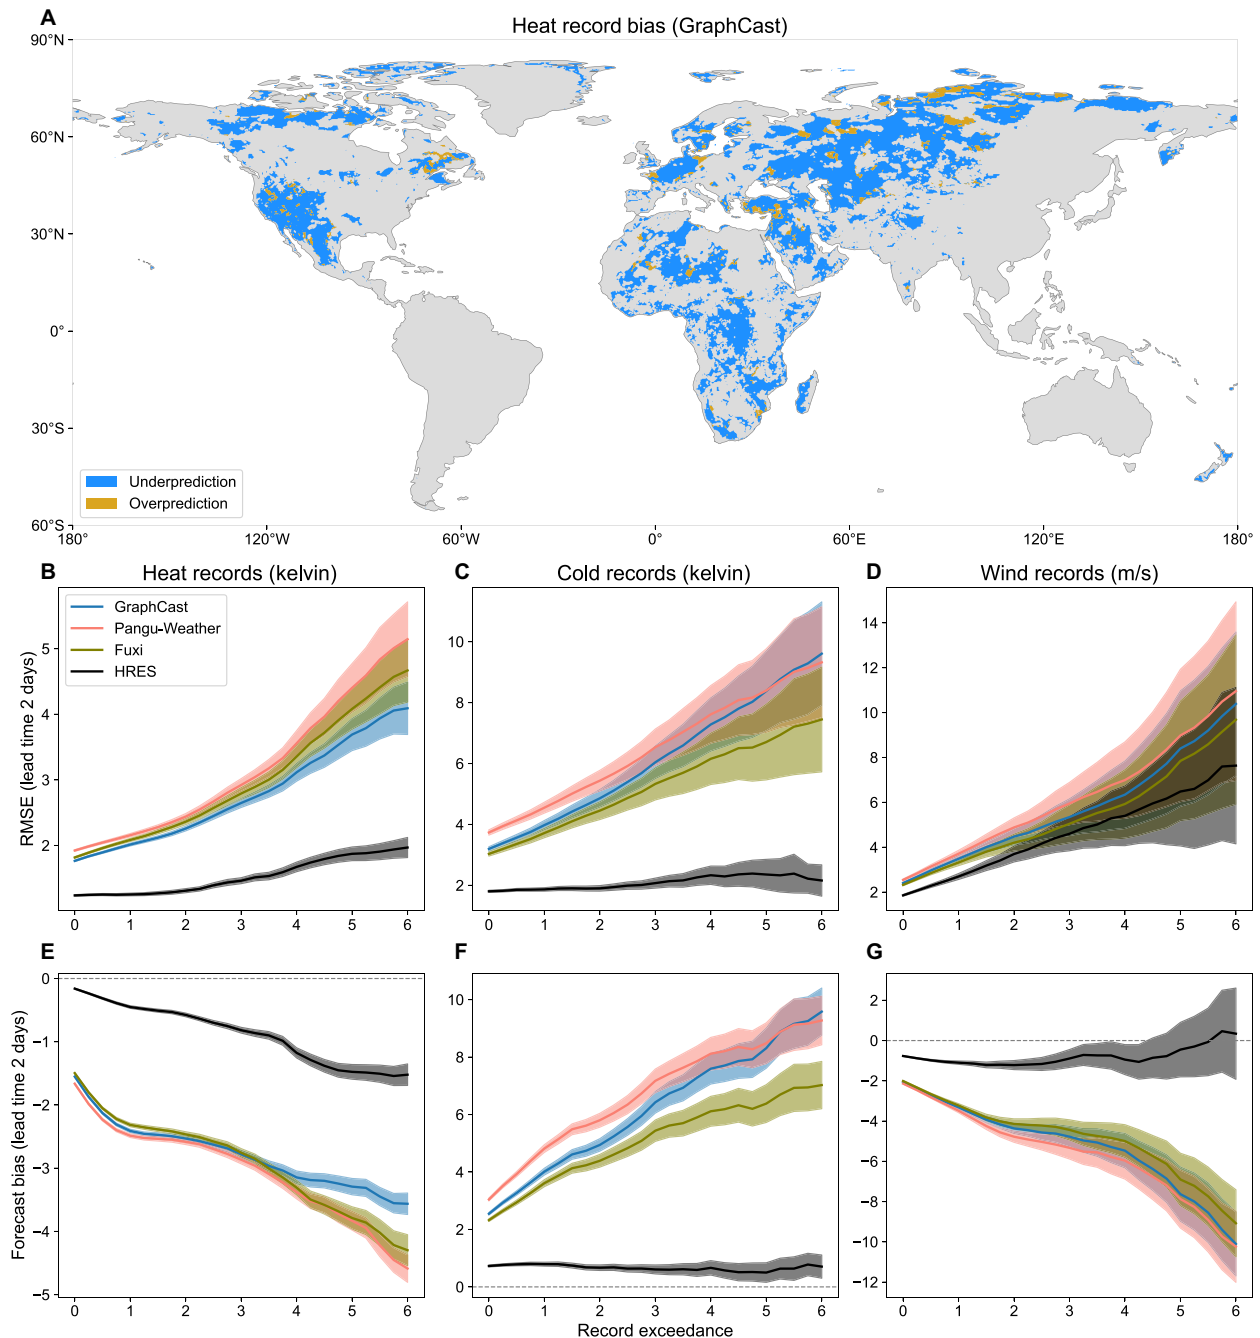


Fig. 2. Forecast bias against record exceedance. (A) Forecast bias of lead time 2 days of the maximum heat records (GraphCast). (B to D) RMSE of 2-m temperature for heat and cold records, and 10-m wind speed for wind records for events in 2020 that exceed the record by at least a certain margin (x axis). Only land pixels (excluding the Antarctic region) are considered. (E to G) Forecast bias of heat, cold, and wind records, for events that exceed the record by at least a certain margin. The transparent shaded areas indicate 95% confidence bands.

local value. In contrast, the physics-based HRES model is more robust to extreme record exceedances. For cold records, HRES exhibits a nearly constant error across increasing exceedances. For heat and wind records, it shows a mild tendency of underestimation, though far less so than AI models. Overall, HRES exhibits lower forecast bias for all record types, and bias is not the dominant source of error, particularly for cold and wind records.

This behavior, shown here for the evaluation year 2020, is fully consistent with results from both the operational forecasts in 2020 and non-operational forecasts in 2018 (figs. S13 and S14). The systematic, one-sided bias observed across event types, lead times, regions, and independent years provides strong evidence that current AI models have a structural extrapolation problem when forecasting record-breaking events.

Model comparison on records' occurrence

We further test the ability of AI models to predict not only the intensity but also the frequency of record-breaking events. We find that, in addition to underestimating event intensity, AI models systematically underpredict the number of records relative to their ERA5 ground truth (Fig. 3, A to C). This underestimation results in a low number of true positives and a high number of false negatives, and consequently low recall (defined as the ratio of true positives to the observed positives) (fig. S15, A to C). In contrast, HRES forecasts a number of records comparable to its HRES-fc0 ground truth, with a slight overestimation for heat records at smaller lead times.

Correctly predicting the number of record-breaking events does not imply accurate timing. In risk management, the trade-off between false positives and false negatives is typically evaluated using precision-recall curves, where high precision corresponds to low numbers of false positives and high recall corresponds to low numbers of false negatives, and the curve is obtained by comparing the forecast to different levels of thresholds (see Materials and Methods). Across all record types and lead times, HRES's precision-recall curves are consistently better than GraphCast's, in the sense that they have higher precisions for the same recall (or higher recalls for the same precision), indicating superior classification performance

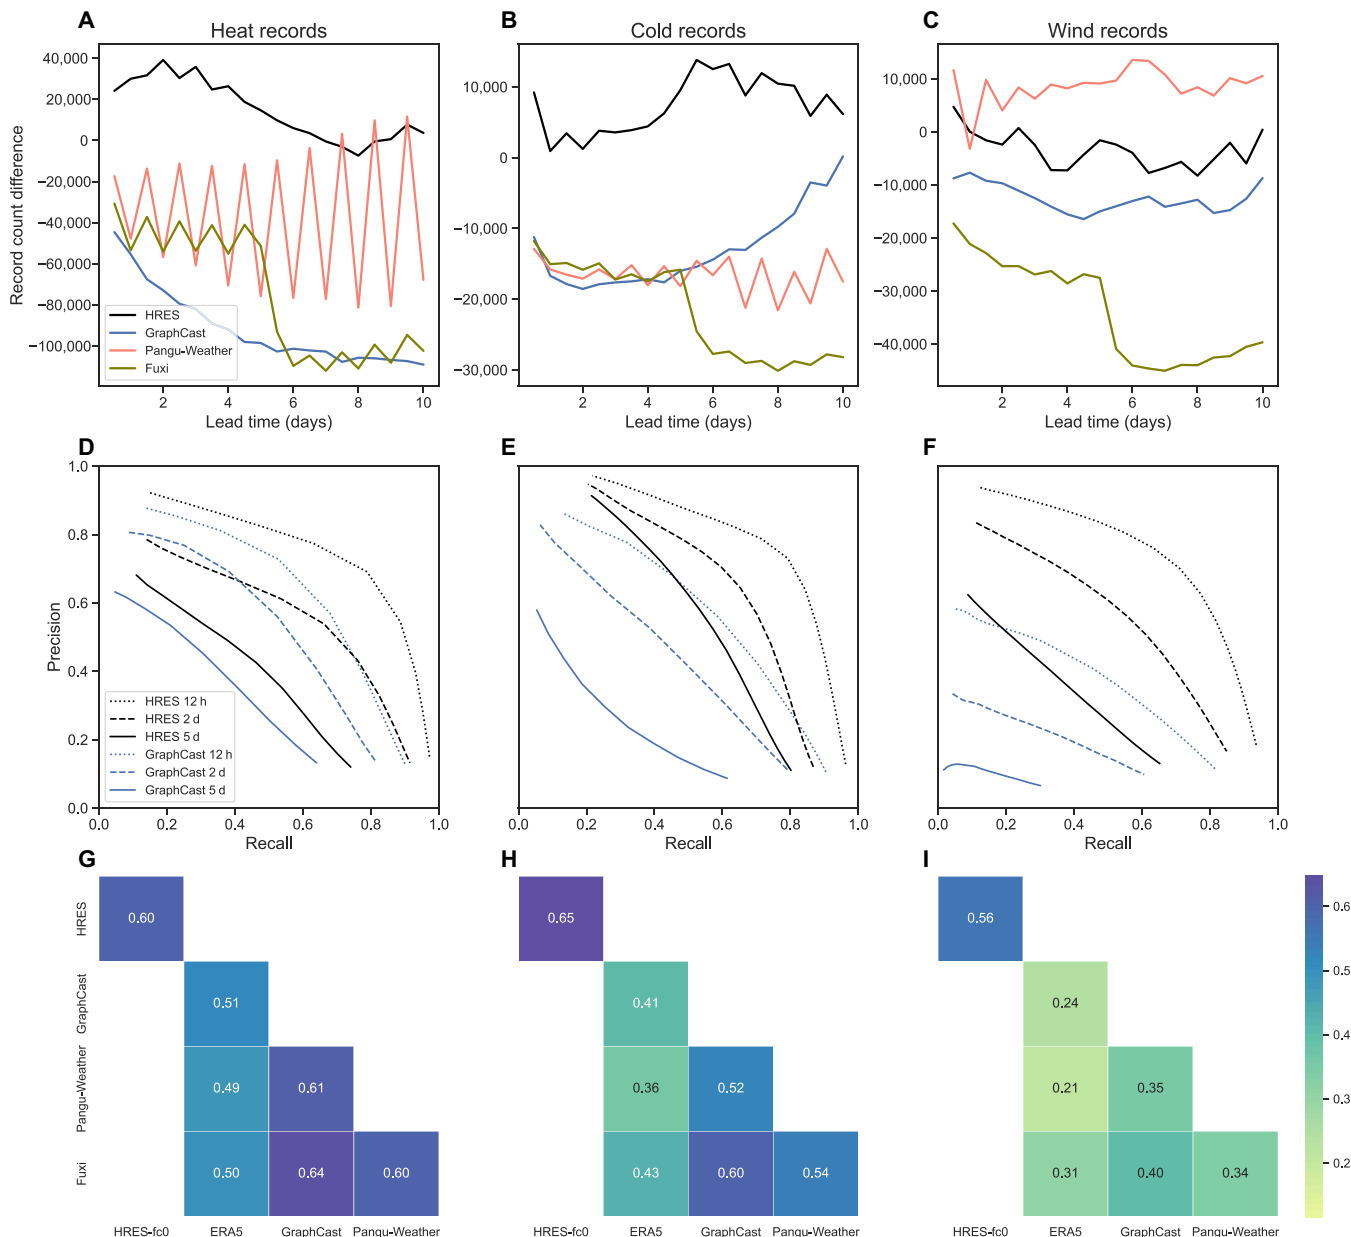


Fig. 3. Prediction of occurrence of record-breaking events. (A to C) Difference between the number of heat, cold, and wind records in the forecasts and in their ground truth data. Only land pixels (excluding the Antarctic region) are considered. (D to F) Precision and recall curves of GraphCast and HRES forecasts when the records are used as the threshold for different lead times. (G to I) Correlations between the indicator functions of whether the ground truth or 2-day forecasts exceeded the record. Pangu-Weather uses two different models (6 and 24 hours) for different lead times, resulting in the zigzag pattern of its record counts.

Downloaded from https://www.science.org at KIT Library on May 04, 2026

for heat, cold, and wind records (Fig. 3, D to F). This is in contrast with earlier results that demonstrate that GraphCast outperforms physics-based models for more moderate extreme events (9). Similar results are observed for Pangu-Weather and Fuxi, where HRES again shows a better classification skill across all lead times (figs. S16 and S17).

As an additional evaluation, we convert both forecast and ground truth into binary variables (1 if a record is exceeded and 0 otherwise) and compute the correlation between them (see Materials and Methods). This metric complements the precision-recall analysis by incorporating true negatives and measuring the degree of dependence between different models' forecasts. HRES has a higher correlation with its ground truth HRES-fc0 than the AI models with their ground truth ERA5, reaffirming its superior performance in forecasting record-breaking events (Fig. 3, G to I). All AI models are positively correlated with each other, showing that they tend to make errors on the same events. This may be due to shared biases learned from their common training data.

DISCUSSION

Our findings consistently show that current AI models underperform HRES in forecasting record-breaking events. They tend to underpredict the intensity and frequency of heat, cold, and wind speed records, with greater forecast biases the larger the record margin. This strongly suggests a systematic extrapolation problem in these models.

Although our evaluation study is restricted to 2018 and 2020, our results are likely to hold in more recent years, as AI models tend to perform worse when the test year is further from the training years (9). Because of substantial regional biases and high resolution dependence in the ERA5 precipitation data (35), here we followed (9) and excluded precipitation in the evaluation. While in this work we evaluated AI weather forecasts against their respective ground truth data ERA5 or HRES-fc0, it would be interesting to evaluate against other data such as in situ observations, to assess their robustness. Since our records are defined locally per grid cell, one extreme event might be counted multiple times across neighboring locations during evaluation. Therefore, it would be worthwhile to investigate methods that ensure spatial independence among the evaluated events. We leave these open questions for future research.

All current state-of-the-art AI weather models are built on neural network architectures such as transformers (8, 30) or graph neural networks (9, 10, 36). In machine learning, extrapolation, also referred to as out-of-distribution generalization, is a well-known fundamental challenge in these models. It has been observed in a range of applications, including image classification (37), protein fitness prediction (38), and large language models (39). Our record benchmark dataset is explicitly designed to test this out-of-distribution problem within AI weather models (see Materials and Methods for discussion).

The AI models studied here do not use any knowledge of physical principles and do not explicitly enforce energy balances or other physical constraints (40, 41). They are purely data-driven and essentially interpolate between observed historical weather patterns in the training period 1979–2017 to produce forecasts for new initial conditions in the test period. This is in stark contrast to physics-based numerical models like HRES that strongly rely on partial differential equations describing the evolution of the atmosphere based on our understanding of physics. This fundamental difference in modeling philosophy likely explains the discrepancy in performance between AI and physics-based NWP

models for record-breaking events (Fig. 1, C to G). While AI models excel when the test set closely resembles the training distribution, capturing complex atmospheric patterns and improving skill on average conditions, they struggle when forecasting unprecedented events outside the training domain, even at short lead times. The nearly linear increase of the biases with record exceedance (Fig. 2, E to G) suggests an implicit cap in AI forecasts around the most extreme training observation. Physics-based models do not have such a bound since physical principles allow them to extrapolate, and, consequently, they exhibit less bias across record magnitudes.

We have focused on deterministic AI weather forecasting models, which issue a point forecast for the mean of future weather states. To account for the uncertainty associated with the point forecast arising from the initialization and the model, a number of probabilistic AI weather models have been developed recently (36, 42, 43). Deterministic AI models are often trained by minimizing the RMSE loss function and are designed to predict the mean of the distribution. Thus, they tend to smooth out fine-scale spatial features such as sharp wind peaks. By contrast, probabilistic AI weather models are trained by minimizing proper scoring rules (42, 43), aiming to forecast the whole distribution and avoid such smoothing. However, both deterministic and probabilistic AI models are trained on the same historical ERA5 reanalysis data, meaning that even these probabilistic models likely face similar extrapolation challenges when forecasting out-of-distribution, record-breaking events.

Several promising avenues exist to address this shortcoming in future generations of AI weather models. One strategy is data augmentation, a widely used technique in machine learning to improve robustness to unseen scenarios by enriching the training data (44). In weather and climate modeling, a key advantage is that numerical climate models can produce very large amounts of physically plausible extreme events outside the training domain. Augmenting training with simulations from different climate regimes (11) or record-breaking events from ensemble boosting (19) could allow AI models to learn from more extreme events than in the original training data. This approach has already shown promise: FourCastNet's (22) performance on tropical cyclones improves substantially when trained on datasets that include such events (17). Another promising direction involves hybrid models and physics-informed neural networks, where specific parameterizations in physical climate models are replaced with AI components (45) or neural networks are trained while respecting specific physical laws described by nonlinear partial differential equations (46). These models combine the efficiency and learning capacity of AI models with the physical consistency and extrapolation ability of physical models. Finally, to improve extrapolation performance on extremes, it may be possible to adapt principles from statistical learning and extreme value theory (47–49).

Given the remarkably fast evolution of AI models in recent years, there are promising ways to further improve these models even for forecasting record-breaking extremes that will continue to frequently occur in a rapidly warming climate. Nevertheless, the current generation still underperforms HRES exactly during the potentially most impactful weather events, including record-breaking heat and cold events as well as wind storms. Thus, it remains vital to fund and run physics-based NWP and AI weather models in parallel and to rigorously evaluate their performance for the most impactful type of weather events.

MATERIALS AND METHODS

Models and data

For the definition of records, we use the ECMWF's ERA5 reanalysis data (27) from 1979–2017 with daily observations at 00, 06, 12, and 18 UTC times. This dataset coincides with the training data of almost all AI models considered in this paper. The time points in this training data are denoted by T_{train} . The ERA5 data are available on a 0.25° by 0.25° latitude-longitude grid. Throughout the paper, we only consider data over land. We use the land-sea mask from the ERA5 and follow ECMWF (50) by defining a grid cell as land if more than 50% of the cell is covered by land; otherwise, it is considered as sea. We exclude the Antarctic region (grid cells with latitude in the range $(-60^\circ, -90^\circ]$) due to aberrant behavior exhibited by some AI models in this region, and denote the remaining set of land grid cells (244,450 grid cells in total) from the ERA5 dataset by $G_{0.25^\circ}$.

We use forecasts from the state-of-the-art AI models GraphCast (9), Pangu-Weather (8), and Fuxi (30) from a test period T_{test} , which is either of the year 2018 or 2020 in our analyses. For the same period, we use forecasts from the physics-based HRES model of ECMWF for comparison. All the forecast data are publicly available from WeatherBench 2 (31). Pangu-Weather and Fuxi are trained and validated on ERA5 data from 1979–2017; the GraphCast forecast data for years 2018 and 2020 are produced by two slightly different versions of GraphCast, i.e., the 2018 data are generated by the GraphCast model trained on ERA5 data from 1979–2017, while the 2020 data are generated by the GraphCast model trained with ERA5 data from a slightly extended period 1979–2019. In addition, we also use the operational versions of GraphCast and Pangu-Weather. The former has been fine-tuned on the HRES-fc0 data from 2016–2021, while the latter was used in an operational setting without fine-tuning.

As ground truth for the AI models, we use ERA5 data with locations in $G_{0.25^\circ}$ in the test period. For HRES and the operational AI models, we use HRES-fc0 as ground truth. Using these two different datasets to evaluate the forecasts against is the standard approach in the literature of AI weather models to avoid unfair comparisons (8, 9, 30).

A benchmark dataset of record-breaking events

To define a dataset of record-breaking events in a given year (e.g., 2020) for a variable x of interest (e.g., 2-m temperature), we first compute the corresponding record in the ERA5 data T_{train} in the training period of the AI models from 1979–2017. We specify whether we consider records in the positive direction (e.g., heat records) or the negative direction (e.g., cold records) by superscripts max or min, respectively. A record $r_{s,m}^{x,\text{max}}$ of variable x is defined locally per grid cell $s \in G_{0.25^\circ}$ and per month $m \in \{\text{January}, \dots, \text{December}\}$. More precisely, we define

$$r_{s,m}^{x,\text{max}} = \max_{t \in T_{\text{train}}: t \in m} x_{s,t} \quad (1)$$

where $x_{s,t}$ is the value of variable x at location s and time t , and $t \in m$ indicates that only time points in month m are considered.

We define the set $R^{x,\text{max}} \subseteq G_{0.25^\circ} \times T_{\text{test}}$ of record-breaking events of variable x consisting of location-time pairs encoding where and when the event occurred. The test period T_{test} contains all time points at 00 and 12 UTC in the test year, i.e., the year 2018 or 2020 in our analyses. We denote by $m(t)$ the month corresponding to a time $t \in T_{\text{test}}$ so that $x_{s,t} > r_{s,m(t)}^{x,\text{max}}$ means that observation $x_{s,t}$ exceeds its respective monthly historical record. With this, we have

$$R^{x,\text{max}} = \left\{ (s, t) \in G_{0.25^\circ} \times T_{\text{test}} : x_{s,t} > r_{s,m(t)}^{x,\text{max}} \right\} \quad (2)$$

We do not evaluate forecasts initiated at 06 and 18 UTC since the HRES forecasts with these initializations are only available for 3.75 days at ECMWF, and all AI-based forecasts are only available for initializations at 00 and 12 UTC on WeatherBench 2 (31). In addition, the data assimilation windows for ERA5 and HRES-fc0 are different. ERA5 assimilates observations using +9-hour/−3-hour windows centered at 00 and 12 UTC, while using +3-hour/−9-hour windows centered at 06 and 18 UTC. By contrast, HRES-fc0 has a consistent assimilation window of +3 hours/−3 hours. Therefore, to ensure a fair comparison between AI models and HRES (9), treated forecasts initialized at 00/12 UTC and 06/18 UTC separately during evaluation and only considered lead times that are multiples of 12 hours. Following [9], here, we use forecast data initialized at 00/12 UTC, and restrict lead times to multiples of 12 hours. In this case, both the forecast initialization time and target time are 00/12 UTC. Consequently, this comparison setup disadvantages HRES due to the mismatch between the +9-hour lookahead of ERA5 (as its assimilation window is +9 hours/−3 hours) that is used to initialize AI models and +3-hour lookahead of HRES-fc0 used as input for HRES, thereby even strengthening our main result that HRES outperforms AI models on record-breaking events.

To test the sensitivity of our results on record-breaking events at 06/18 UTC, we further evaluated HRES and AI weather forecasts with lead times 6 hours, 18 hours, etc. As expected and consistent with previous studies (8, 30), the same ranking is observed: HRES outperforms the AI models and their operational variants in predicting the records' intensity for almost all lead times (figs. S7 and S8).

Note that the notion of a record-breaking event is to be understood relative to the training period. We do not update the record if a larger event has occurred after 2017 (the end of the training period). The reason is that AI models are not retrained and a record-breaking event in the test period will not inform or improve the model for later time steps.

Using as test data the ERA5 ground truth in 2020, we obtain 162,751 records for heat, 32,991 for cold, and 53,345 for wind (see the geographical distribution of these records in the map in Fig. 1A and fig. S1, respectively). For the analysis of operational models, we define the set of record-breaking events as those where HRES-fc0 exceeds the training record, yielding 170,136 records for heat, 109,155 for cold, and 338,235 for wind (fig. S18). HRES-fc0 has more record-breaking events particularly for cold and wind records, possibly due to higher horizontal resolutions. The intensity of those events also seems to differ slightly in the two ground truths. For the record-breaking events identified by ERA5, HRES-fc0 seems to have slightly lower intensity for heat and cold records (fig. S15, D to F). However, this could be a result of selection bias, as HRES-fc0 exhibits its higher intensity than ERA5 for the records data identified by HRES-fc0 (fig. S15, G to I). This implies the importance of evaluating both the AI models and their operational variants, where the record-breaking events in ERA5 and HRES-fc0 are used for evaluation, respectively.

We further tested whether defining the records with a running-window approach would alter the results. Specifically, we have extracted the records for each day of the year using a 31-day running window, i.e., 15 days before the day of interest and 15 days after. For

ERA5 in 2020, this resulted in fewer record-breaking events (90,471 heat records and 18,054 cold records). Among these events, approximately 70% also break the monthly records. When using the running-window approach, we observe similar results that HRES outperforms the AI weather models on predicting the record-breaking events (fig. S19).

Extrapolation in AI models

Extrapolation or out-of-distribution generalization in AI models refers to the situation where a test predictor is far away from the distribution of the training predictors. In high-dimensional predictor spaces, it is not trivial to mathematically describe such points. One way of framing extrapolation is to require that the predictor is outside of the convex hull (blue line in Fig. 4) formed by the training data. Balestrierio *et al.* (51) argue that with this definition of training domain it is in fact very likely that test points need extrapolation. However, convex hulls are computationally prohibitive in high dimensions since the number of facets grows rapidly with the dimension. Our record dataset therefore considers a stronger yet simpler definition, namely, all points where at least one test variable is beyond its univariate training range. In Fig. 4, this corresponds to all test points outside of the green rectangle. All events in the record set $R^{x,max}$ in Eq. 2 satisfy this strong definition of out-of-distribution samples.

Root mean square error

We quantify the forecast error with the RMSE. For a target variable x of interest (e.g., 2-m temperature T_{2m}), the RMSE on a subset of location-initialization pairs for lead time τ is defined as

$$RMSE_I(\tau) = \sqrt{\frac{1}{\sum_{(s,t_0) \in I} \omega_s} \sum_{(s,t_0) \in I} \omega_s (\hat{x}_{s,t_0}^\tau - x_{s,t_0+\tau})^2} \quad (3)$$

where (i) $G_{0.25^\circ}$ is the set of locations/grid cells, (ii) $I \subseteq G_{0.25^\circ} \times T_{test}$ is the set of location-initialization pairs of interest, (iii) \hat{x}_{s,t_0}^τ is a forecast of variable x with lead time τ at location $s \in G_{0.25^\circ}$ and initialization

time $t_0 \in T$, and $x_{s,t_0+\tau}$ is the corresponding ground truth, (iv) ω_s is the latitude-based weight chosen as the one used in previous studies (9)

$$\omega_s = \begin{cases} \cos(\theta_{lat(s)}) \sin(\theta_{0.25^\circ/2}), & \text{if } |\theta_{lat(s)}| < \pi/2 \\ \sin^2(\theta_{0.25^\circ/4}), & \text{if } |\theta_{lat(s)}| = \pi/2 \end{cases}$$

with θ_a as the radian associated with degree a .

Our definition of RMSE is more general than the conventional one (31) in the sense that we allow to focus on a subset I of location-initialization pairs (s, t_0) . If we set I as the product of the set of all grid cells over the globe and all time points in T_{test} , we recover the traditional (latitude-weighted) RMSE on all test locations and initialization times.

In the computation of RMSE on record-breaking events such as shown in Fig. 1, we choose the set I in Eq. 3 in the following way. Recall the set $R^{x,max} \subseteq G_{0.25^\circ} \times T_{test}$ of location-time pairs of all records in a given time period (e.g., the year 2020). We choose all location-initialization pairs such that the target of the forecasting with lead time τ corresponds to a record, i.e.

$$I_\tau = \{ (s, t_0) \in G_{0.25^\circ} \times T_{test} : (s, t_0 + \tau) \in R^{x,max} \} \quad (4)$$

The corresponding $RMSE_I(\tau)$ is the error of a model made in forecasting records with lead time τ .

To construct a confidence interval for the $RMSE_I(\tau)$ in Eq. 3, we assume that the central limit theorem holds for the weighted square errors, i.e.

$$\sqrt{|I|} \left[\frac{1}{\sum_{(s,t_0) \in I} \omega_s} \sum_{(s,t_0) \in I} \omega_s (\hat{x}_{s,t_0}^\tau - x_{s,t_0+\tau})^2 - \mu_\tau \right] \xrightarrow{d} \mathcal{N}(0, \sigma_\tau^2),$$

$|I| \rightarrow \infty$

where μ_τ denotes the true mean of weighted square error and σ_τ^2 its asymptotic variance, and \xrightarrow{d} means convergence in distribution.

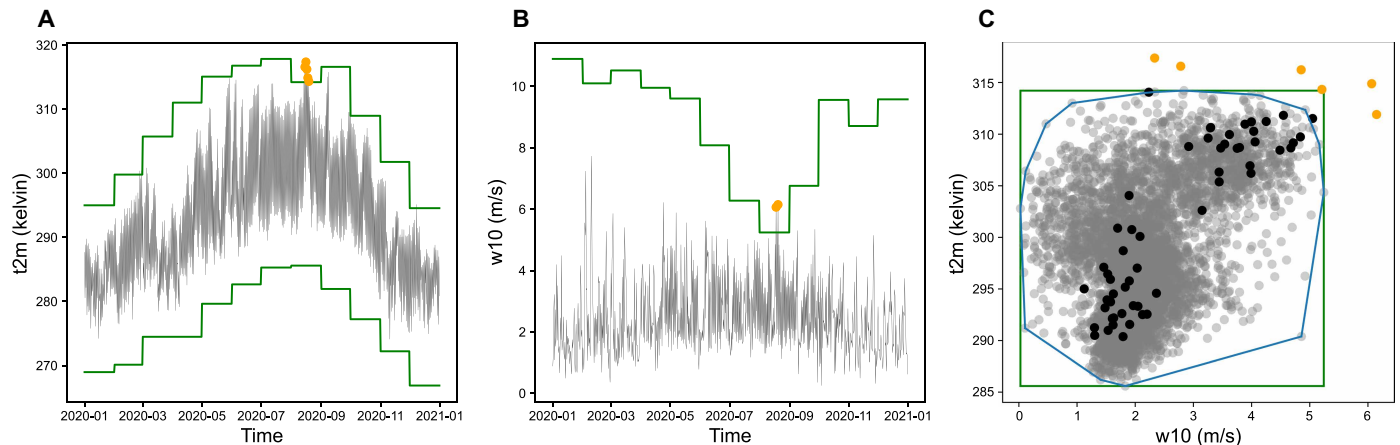


Fig. 4. Illustration of our definitions of record and extrapolation. (A) Daily time series of 2-m temperature at the location with latitude 37.5 and longitude -121 in 2020 (black), and monthly max/min records (in green) at this location, where orange points indicate the record-breaking events in August. (B) Daily time series of 10-m wind speed and monthly max records at the same location. (C) Scatter plots of 2-m temperature and 10-m wind speed in August in the training period from 1979–2017 (in gray) and in the evaluation year 2020 (in black) at this location. The blue line represents the convex hull formed by the training data, while the green rectangle shows the max/min records in the training period. Orange points indicate the record-breaking events in the evaluation period.

Then, by the Delta method, we obtain the asymptotic distribution of $RMSE_I(\tau)$

$$\sqrt{|I|}(RMSE_I(\tau) - \sqrt{\mu_\tau}) \xrightarrow{d} \mathcal{N}[0, \sigma_\tau^2 / (4\mu_\tau)]$$

Hence, an approximate α -level confidence interval for $RMSE_I(\tau)$ can be constructed as $[\mu_\tau - q_{(1+\alpha)/2}\sigma_\tau/\sqrt{4\mu_\tau|I|}, \mu_\tau + q_{(1+\alpha)/2}\sigma_\tau/\sqrt{4\mu_\tau|I|}]$, where $q_{(1+\alpha)/2}$ is the $(1 + \alpha)/2$ quantile of a standard normal distribution. Alternatively, bootstrap can be used to construct the confidence bands (II). We tried the nonparametric bootstrap with 1000 resampling, which yielded similar confidence bands to the normal ones. For the sake of computational feasibility, we use the above normal confidence levels throughout the paper.

Forecast bias

To complement RMSE and investigate whether a forecasting model under- or overpredicts the ground truth, we consider the latitude-weighted forecast bias

$$FB_I(\tau) = \frac{1}{\sum_{(s,t_0) \in I} \omega_s} \sum_{(s,t_0) \in I} \omega_s (\hat{x}_{s,t_0}^\tau - x_{s,t_0+\tau})$$

where the notation is the same as in Eq. 3. Confidence intervals for the forecast bias are computed in the same way as for RMSE based on asymptotic normality.

Precision and recall curves

For early warning systems, it is crucial that a weather forecasting model is able to predict the occurrence of an extreme event accurately. We therefore consider record-breaking event forecasting as a binary classification problem by assessing whether a forecasting model can predict the exceedance of a variable over its previous record in the sense of Eq. 1. Since this classification problem is strongly imbalanced, similar to previous studies (9), we use precision-recall curves that are well suited for such cases since they account for both false positives and false negatives.

For a set $I \subseteq G_{0.25^\circ} \times T_{\text{test}}$ of location-initialization pairs of interest, we compute the precision and recall for variable x at lead time τ as (we set $r_{s,m} = r_{s,m}^{x,\max}$ to simplify notation)

$$\text{Precision}_I(\tau) = \frac{\sum_{(s,t_0) \in I} 1 \{ \hat{x}_{s,t_0}^\tau > r_{s,m(t_0+\tau)} \} 1 \{ x_{s,t_0+\tau} > r_{s,m(t_0+\tau)} \}}{\sum_{(s,t_0) \in I} 1 \{ \hat{x}_{s,t_0}^\tau > r_{s,m(t_0+\tau)} \}}$$

and

$$\text{Recall}_I(\tau) = \frac{\sum_{(s,t_0) \in I} 1 \{ \hat{x}_{s,t_0}^\tau > r_{s,m(t_0+\tau)} \} 1 \{ x_{s,t_0+\tau} > r_{s,m(t_0+\tau)} \}}{\sum_{(s,t_0) \in I} 1 \{ x_{s,t_0+\tau} > r_{s,m(t_0+\tau)} \}}$$

where \hat{x}_{s,t_0}^τ denotes a forecast of variable x at location s initialized at time t_0 with lead time τ , and $x_{s,t_0+\tau}$ is the corresponding ground truth. As above, $m(t_0 + \tau)$ is the month corresponding to time $t_0 + \tau$.

To produce a precision-recall curve from a deterministic forecast, we introduce a common ‘‘gain’’ parameter to define scaled forecasts by

$$\text{scaled forecast} = \text{forecast} + \text{gain} \times \text{forecast std. deviation} \quad (5)$$

Using these scaled forecasts in the precision and recall formulae instead of only \hat{x}_{s,t_0}^τ and varying the gain parameter in a suitable range

$([-1.5,1.5]$ in our case) yields a precision-recall curve. The scaling allows the study of different trade-offs between false positives and false negatives, and using a common gain parameter enables averaging over all spatial locations $s \in G_{0.25^\circ}$. Our parameterization of the scaled forecasts is slightly different from the one in previous studies (9), but it is theoretically more justified. For a probabilistic forecast from a location-scale family, Eq. 5 corresponds to choosing the same quantile of the forecast distribution at all locations.

For each variable x , each location $s \in G_{0.25^\circ}$, each month m , and each lead time τ , we estimate the forecast standard deviations in Eq. 5 from forecasts in the year 2020 for the different models. We assume that this standard deviation is constant for time points in the same month so that we have enough data for the estimation.

Correlation between record forecasts

To compute the correlation between the different model forecasts and ground truths, we define suitable functions indicating whether the corresponding record is exceeded. Given a lead time τ and, for instance, a variable x with maximum records abbreviated by $r_{s,m} = r_{s,m}^{x,\max}$ for a forecast $\hat{x}_{s,t}^\tau$ from some model define for each time point $t_0 \in T_{\text{test}}$ the indicator $1 \{ \hat{x}_{s,t_0}^\tau > r_{s,m(t_0+\tau)} \}$ that takes value 1 if \hat{x}_{s,t_0}^τ exceeds the record $r_{s,m(t_0+\tau)}$ and 0 otherwise. We can now compute the correlation between these indicators, indexed by all $t_0 + \tau \in T_{\text{test}}$ and $s \in G_{0.25^\circ}$, for two different forecast models. Similarly, for a ground truth (either ERA5 or HRES-fc0), we define for each time point $t_0 + \tau \in T_{\text{test}}$ the indicator $1 \{ x_{s,t_0+\tau} > r_{s,m(t_0+\tau)} \}$. We then compute correlations of these ground truths with the forecast indicators, and between forecast indicators from different models. The resulting correlation matrix is shown for the ERA5-trained AI models in Fig. 3 (G to I), for the operational AI models in fig. S20 (G to I), and for the 2018 forecasts in fig. S21 (G to I). We further use the Pearson chi-square test (52) to test the independence between these indicators (note that independence is equivalent to zero correlation for binary random variables). Notably, the P values of all tests are less than 10^{-10} and the null hypotheses of independence are thus rejected, meaning that all pairs shown in Fig. 3 (G to I), fig. S20 (G to I), and fig. S21 (G to I) are significantly dependent.

Forecaster’s dilemma

In the theory of forecast evaluation, the forecaster’s dilemma (33) shows that computing an evaluation score only on a subset of observations can incentivize suboptimal forecasts. Such conditioning appears in the RMSE defined in (3) if the set I depends on the observations, as, for instance, in the case of record-breaking events defined in (4). This metric should therefore not be used as the sole evaluation criterion, but rather in combination with others. We therefore also report the overall RMSE on all events in Fig. 1 and figs. S2 and S6, which show that all methods yield errors on a comparable scale and do not appear to artificially hedge forecasts of extreme events. In addition, we consider different evaluation criteria such as precision-recall curves that take into account both false positives and false negatives (Fig. 3, D to F, and figs. S20, D to F, and S21, D to F).

Computing the RMSE on a subset of extreme observations is common in the literature of AI weather forecasts (9, 11, 15). Another approach that avoids the forecaster’s dilemma completely is to condition on the forecasts instead of the observations (34). We follow this approach to compare the operational version of GraphCast with HRES. We choose as the set of record-breaking events in (2) all

Downloaded from https://www.science.org at KIT Library on May 04, 2026

location-initialization pairs such that forecasts with lead time τ from both GraphCast operational and HRES exceed the training record. For maximum records of variable x , for instance, this yields an index set

$$I_{\tau} = \left\{ (s, t_0) \in G_{0.25^{\circ}} \times T_{\text{test}} : \hat{x}_{s,t_0}^{\text{HRES},\tau} > r_{s,m}(t_0+\tau), \hat{x}_{s,t_0}^{\text{GraphCast},\tau} > r_{s,m}(t_0+\tau) \right\}$$

to be used in the RMSE in (3). The results (fig. S9) look qualitatively similar to those from conditioning on the observations, except that we have a much smaller set of events that are jointly forecasted to be record-breaking by both models compared to the original record dataset.

Supplementary Materials

This PDF file includes:

Figs. S1 to S21

REFERENCES

- R. H. White, S. Anderson, J. F. Booth, G. Braich, C. Draeger, C. Fei, C. D. G. Harley, S. B. Henderson, M. Jakob, C. A. Lau, L. Mareshet Admasu, V. Narinesingh, C. Rodell, E. Roodcroft, K. R. Weinberger, G. West, The unprecedented Pacific Northwest heatwave of June 2021. *Nat. Commun.* **14**, 727 (2023).
- D. Barriopedro, E. M. Fischer, J. Luterbacher, R. M. Trigo, R. García-Herrera, The hot summer of 2010: Redrawing the temperature record map of Europe. *Science* **332**, 220–224 (2011).
- R. García-Herrera, J. Díaz, R. M. Trigo, J. Luterbacher, E. M. Fischer, A review of the European summer heat wave of 2003. *Crit. Rev. Environ. Sci. Technol.* **40**, 267–306 (2010).
- H. Wernli, S. Dirren, M. A. Liniger, M. Zillig, Dynamical aspects of the life cycle of the winter storm ‘Lothar’ (24–26 December 1999). *Q. J. R. Meteorol. Soc.* **128**, 405–429 (2002).
- A. H. Fink, T. Brücher, V. Emert, A. Krüger, J. G. Pinto, The European storm Kyrill in January 2007: Synoptic evolution, meteorological impacts and some considerations with respect to climate change. *Nat. Hazards Earth Syst. Sci.* **9**, 405–423 (2009).
- T. Kelder, D. Heinrich, L. Klok, V. Thompson, H. M. D. Goulart, E. Hawkins, L. J. Slater, L. Suarez-Gutierrez, R. L. Wilby, E. Coughlan de Perez, E. M. Stephens, S. Burt, B. van den Hurk, H. de Vries, K. van der Wiel, E. L. F. Schipper, A. Carmona Baéz, E. van Bueren, E. M. Fischer, How to stop being surprised by unprecedented weather. *Nat. Commun.* **16**, 2382 (2025).
- P. Bauer, A. Thorpe, G. Brunet, The quiet revolution of numerical weather prediction. *Nature* **525**, 47–55 (2015).
- K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, Q. Tian, Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **619**, 533–538 (2023).
- R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, P. Battaglia, Learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023).
- S. Lang, M. Alexe, M. Chantry, J. Dramsch, F. Pinault, B. Raoult, M. C. A. Clare, C. Lessig, M. Maier-Gerber, L. Magnusson, Z. B. Bouallègue, A. P. Nemesio, P. D. Dueben, A. Brown, F. Pappenberger, F. Rabier, AIFS—ECMWF’s data-driven forecasting system. arXiv:2406.01465 [physics.ao-ph] (2024).
- C. Bodnar, W. P. Bruinsma, A. Lucic, M. Stanley, A. Allen, J. Brandstetter, P. Garvan, M. Riechert, J. A. Weyn, H. Dong, J. K. Gupta, K. Thambiratnam, A. T. Archibald, C.-C. Wu, E. Heider, M. Welling, R. E. Turner, P. Perdikaris, A foundation model for the Earth system. *Nature* **641**, 1180–1187 (2025).
- M. G. Schultz, C. Betancourt, B. Gong, F. Kleinert, M. Langguth, L. H. Leufen, A. Mozaffari, S. Stadler, Can deep learning beat numerical weather prediction? *Philos. Trans. R. Soc. A* **379**, 20200097 (2021).
- P. A. G. Watson, Machine learning applications for weather and climate need greater focus on extremes. *Environ. Res. Lett.* **17**, 111004 (2022).
- Z. Ben Bouallègue, M. C. A. Clare, L. Magnusson, E. Gascón, M. Maier-Gerber, M. Janoušek, M. Rodwell, F. Pinault, J. S. Dramsch, S. T. K. Lang, B. Raoult, F. Rabier, M. Chevallier, I. Sandu, P. Dueben, M. Chantry, F. Pappenberger, The rise of data-driven weather forecasting: A first statistical assessment of machine learning-based weather forecasts in an operational-like context. *Bull. Am. Meteorol. Soc.* **105**, E864–E883 (2024).
- L. Olivetti, G. Messori, 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 (2024).
- Z. Meng, G. J. Hakim, W. Yang, G. A. Vecchi, Deep learning atmospheric models reliably simulate out-of-sample land heat and cold wave frequencies. *Geophys. Res. Lett.* **53**, e2025GL117990 (2026).
- Y. Q. Sun, P. Hassanzadeh, M. Zand, A. Chattopadhyay, J. Weare, D. S. Abbot, Can AI weather models predict out-of-distribution gray swan tropical cyclones? *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2420914122 (2025).
- E. M. Fischer, S. Sippel, R. Knutti, Increasing probability of record-shattering climate extremes. *Nat. Clim. Change* **11**, 689–695 (2021).
- E. M. Fischer, U. Beyerle, L. Bloin-Wibe, C. Gessner, V. Humphrey, F. Lehner, A. G. Pendergrass, S. Sippel, J. Zeder, R. Knutti, Storylines for unprecedented heatwaves based on ensemble boosting. *Nat. Commun.* **14**, 4643 (2023).
- O. Watt-Meyer, B. Henn, J. McGibbon, S. K. Clark, A. Kwa, W. A. Perkins, E. Wu, L. Harris, C. S. Bretherton, ACE2: Accurately learning subseasonal to decadal atmospheric variability and forced responses. *npj Clim. Atmos. Sci.* **8**, 205 (2025).
- C. Kent, A. A. Scaife, N. J. Dunstone, D. Smith, S. C. Hardiman, T. Dunstan, O. Watt-Meyer, Skillful global seasonal predictions from a machine learning weather model trained on reanalysis data. *npj Clim. Atmos. Sci.* **8**, 314 (2025).
- J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, A. Anandkumar, FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. arXiv:2202.11214 [physics.ao-ph] (2022).
- M. DeMaria, J. L. Franklin, G. Chirokova, J. Radford, R. DeMaria, K. D. Musgrave, I. Ebert-Uphoff, An operations-based evaluation of tropical cyclone track and intensity forecasts from artificial intelligence weather prediction models. *Artif. Intell. Earth Syst.* **4**, 240085 (2025).
- A. J. Charlton-Perez, H. F. Dacre, S. Driscoll, S. L. Gray, B. Harvey, N. J. Harvey, K. M. R. Hunt, R. W. Lee, R. Swaminathan, R. Vandaele, A. Volonté, Do AI models produce better weather forecasts than physics-based models? A quantitative evaluation case study of Storm Ciarán. *npj Clim. Atmos. Sci.* **7**, 93 (2024).
- O. C. Pasche, J. Wider, Z. Zhang, J. Zscheischler, S. Engelke, Validating deep learning weather forecast models on recent high-impact extreme events. *Artif. Intell. Earth Syst.* **4**, e240033 (2025).
- Y. Q. Sun, P. Hassanzadeh, T. Shaw, H. A. Pahlavan, Predicting beyond training data via extrapolation versus translocation: AI weather models and Dubai’s unprecedented 2024 rainfall. arXiv:2505.10241 [physics.ao-ph] (2025).
- H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. de Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, J. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, J. N. Thépaut, The ERA5 global reanalysis. *Q. J. Roy. Meteorol. Soc.* **146**, 1999–2049 (2020).
- J. E. Overland, M. Wang, The 2020 Siberian heat wave. *Int. J. Climatol.* **41**, E2341–E2346 (2021).
- X. Li, Y. Ryu, J. Xiao, B. Dechant, J. Liu, B. Li, S. Jeong, P. Gentine, New-generation geostationary satellite reveals widespread midday depression in dryland photosynthesis during 2020 western U.S. heatwave. *Sci. Adv.* **9**, eadi0775 (2023).
- L. Chen, X. Zhong, F. Zhang, Y. Cheng, Y. Xu, Y. Qi, H. Li, FuXi: A cascade machine learning forecasting system for 15-day global weather forecast. *npj Clim. Atmos. Sci.* **6**, 190 (2023).
- S. Rasp, S. Hoyer, A. Merose, I. Langmore, P. Battaglia, T. Russell, A. Sanchez-Gonzalez, V. Yang, R. Carver, S. Agrawal, M. Chantry, Z. Ben Bouallegue, P. Dueben, C. Bromberg, J. Sisk, L. Barrington, A. Bell, F. Sha, WeatherBench 2: A benchmark for the next generation of data-driven global weather models. *J. Adv. Model. Earth Syst.* **16**, e2023MS004019 (2024).
- E. M. Fischer, M. Bador, R. Usher, E. J. Kendon, A. Robinson, S. Sippel, Record-breaking extremes in a warming climate. *Nat. Rev. Earth Environ.* **6**, 456–470 (2025).
- S. Lerch, T. L. Thorarindottir, F. Ravazzolo, T. Gneiting, Forecaster’s dilemma: Extreme events and forecast evaluation. *Stat. Sci.* **32**, 106–127 (2017).
- H. Holzmann, M. Eulert, The role of the information set for forecasting—With applications to risk management. *Ann. Appl. Stat.* **8**, 595–621 (2014).
- D. A. Lavers, A. Simmons, F. Vamborg, M. J. Rodwell, An evaluation of ERA5 precipitation for climate monitoring. *Q. J. R. Meteorol. Soc.* **148**, 3152–3165 (2022).
- I. Price, A. Sanchez-Gonzalez, F. Alet, T. R. Andersson, A. el-Kadi, D. Masters, T. Ewalds, J. Stott, S. Mohamed, P. Battaglia, R. Lam, M. Willson, Probabilistic weather forecasting with machine learning. *Nature* **637**, 84–90 (2025).
- Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, V. Lempitsky, Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* **17**, 1–35 (2016).
- C. R. Freschlin, S. A. Fahlberg, P. Heinzelman, P. A. Romero, Neural network extrapolation to distant regions of the protein fitness landscape. *Nat. Commun.* **15**, 6405 (2024).
- D. Hupkes, M. Giulianelli, V. Dankers, M. Artetxe, Y. Elazar, T. Pimentel, C. Christodoulopoulos, K. Lasri, N. Saphra, A. Sinclair, D. Ulmer, F. Schottmann,

- K. Batsuren, K. Sun, K. Sinha, L. Khalatbari, M. Ryskina, R. Frieske, R. Cotterell, Z. Jin, A taxonomy and review of generalization research in NLP. *Nat. Mach. Intell.* **5**, 1161–1174 (2023).
40. T. Selz, G. C. Craig, Can artificial intelligence-based weather prediction models simulate the butterfly effect? *Geophys. Res. Lett.* **50**, e2023GL105747 (2023).
41. M. Bonavita, On some limitations of current machine learning weather prediction models. *Geophys. Res. Lett.* **51**, e2023GL107377 (2024).
42. S. Lang, M. Alexe, M. C. A. Clare, C. Roberts, R. Adewoyin, Z. B. Bouallègue, M. Chantry, J. Dramsch, P. D. Dueben, S. Hahner, P. Maciel, A. Prieto-Nemesio, C. O'Brien, F. Pinault, J. Polster, B. Raoult, S. Tietsche, M. Leutbecher, AIFS-CRPS: Ensemble forecasting using a model trained with a loss function based on the continuous ranked probability score. *npj Artif. Intell.* **2**, 18 (2026).
43. F. Alet, I. Price, A. El-Kadi, D. Masters, S. Markou, T. R. Andersson, J. Stott, R. Lam, M. Willson, A. Sanchez-Gonzalez, P. Battaglia, Skillful joint probabilistic weather forecasting from marginals. *arXiv:2506.10772 [cs.LG]* (2025).
44. C. Shorten, T. M. Khoshgoftaar, A survey on image data augmentation for deep learning. *J. Big Data* **6**, 60 (2019).
45. T. A. Shaw, B. Stevens, The other climate crisis. *Nature* **639**, 877–887 (2025).
46. M. Raissi, P. Perdikaris, G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
47. X. Shen, N. Meinshausen, Engression: Extrapolation through the lens of distributional regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **87**, 653–677 (2024).
48. G. Buriticá, S. Engelke, Progression: An extrapolation principle for regression. *arXiv:2410.23246 [stat.ME]* (2024).
49. Y. Boulaguiem, J. Zscheischler, E. Vignotto, K. van der Wiel, S. Engelke, Modeling and simulating spatial extremes by combining extreme value theory with generative adversarial networks. *Environ. Data Sci.* **1**, e5 (2022).
50. R. Owens, T. Hewson, *ECMWF Forecast User Guide* (ECMWF, 2018).
51. R. Balestriero, J. Pesenti, Y. LeCun, Learning in high dimension always amounts to extrapolation. *arXiv:2110.09485 [cs.LG]* (2021).
52. A. Agresti, *An Introduction to Categorical Data Analysis*, Wiley Series in Probability and Mathematical Statistics (Wiley-Interscience, ed. 2, 2007).

Acknowledgments: We are grateful for the computing time provided by the Baobab and Bamboo HPC service at the University of Geneva, and by the high-performance computer HoreKa at the National High-Performance Computing Center at KIT. We also thank the participants in the EXCLAIM Symposium at ETH Zurich in June 2025 for the helpful discussions.

Funding: Z.Z. and S.E. are grateful for the funding from the Swiss National Science Foundation Eccellenza grant “Graph structures, sparsity and high-dimensional inference for extremes” (grant no. 186858), and E.F. and J.Z. have received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 101003469 (XAIDA).

Author contributions: Z.Z. and S.E. proposed the idea and initiated the project. Z.Z., E.F., J.Z., and S.E. developed the methodology. Z.Z. implemented the methodology and conducted the analysis. Z.Z., E.F., J.Z., and S.E. interpreted the results, which led to refinements of the method. Z.Z., E.F., J.Z., and S.E. contributed to the writing of the paper. **Competing interests:** The authors declare that they have no competing interests. **Data, code, and materials availability:** HRES and all AI weather forecast data, as well as the ERA5 data analyzed in this study are publicly available on WeatherBench 2 at <https://weatherbench2.readthedocs.io/en/latest/data-guide.html>. This study did not generate new materials. The code and data for reproducing the analysis are provided at <https://doi.org/10.5281/zenodo.18929001>.

Submitted 8 September 2025

Accepted 25 March 2026

Published 29 April 2026

10.1126/sciadv.aec1433

Physics-based models outperform AI weather forecasts of record-breaking extremes

Zhongwei Zhang, Erich Fischer, Jakob Zscheischler, and Sebastian Engelke

Sci. Adv. **12** (18), eaec1433. DOI: 10.1126/sciadv.aec1433

View the article online

<https://www.science.org/doi/10.1126/sciadv.aec1433>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2026 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution License 4.0 (CC BY).