PNAS RESEARCH ARTICLE | EARTH, ATMOSPHERIC, AND PLANETARY SCIENCES

Dynamical-generative downscaling of climate model ensembles

Ignacio Lopez-Gomez[a,1], Zhong Yi Wan[a], Leonardo Zepeda-Núñez[a], Tapio Schneider[a,b], John Anderson[a], and Fei Sha[a,1]

Affiliations are included on p. 8.

Edited by Noah S. Diffenbaugh, Stanford University, Stanford, CA; received October 2, 2024; accepted March 19, 2025 by Editorial Board Member Robert E. Dickinson

Regional high-resolution climate projections are crucial for many applications, such as agriculture, hydrology, and natural hazard risk assessment. Dynamical downscaling, the state-of-the-art method to produce localized future climate information, involves running a regional climate model (RCM) driven by an Earth System Model (ESM), but it is too computationally expensive to apply to large climate projection ensembles. We propose an approach combining dynamical downscaling with generative AI to reduce the cost and improve the uncertainty estimates of downscaled climate projections. In our framework, an RCM dynamically downscales ESM output to an intermediate resolution, followed by a generative diffusion model that further refines the resolution to the target scale. This approach leverages the generalizability of physics-based models and the sampling efficiency of diffusion models, enabling the downscaling of large multimodel ensembles. We evaluate our method against dynamically downscaled climate projections from the Coupled Model Intercomparison Project 6 (CMIP6) ensemble. Our results demonstrate its ability to provide more accurate uncertainty bounds on future regional climate than alternatives such as dynamical downscaling of smaller ensembles, or traditional empirical statistical downscaling methods. We also show that dynamical-generative downscaling results in significantly lower errors than popular statistical downscaling techniques, and captures more accurately the spectra, tail dependence, and multivariate correlations of meteorological fields. These characteristics make the dynamical-generative framework a flexible, accurate, and efficient way to downscale large ensembles of climate projections, currently out of reach for pure dynamical downscaling.

climate change | regional climate | climate risk

Regional climate projections below the 10 km scale represent a valuable source of information to stakeholders in need of climate risk assessments. Higher-resolution projections enable a more faithful representation of orography (1), land–atmosphere interactions (2), and mesoscale convective systems (3), all of which greatly influence the magnitude and frequency of local extreme weather events (4, 5). Localized climate data are particularly necessary in coastal and mountainous regions, where landscape changes at the kilometer scale largely shape the local climatology (6–8). Sectors in need of this kind of granular information include agriculture (9), hydrology (10), energy (11), and natural hazard risk assessment (12, 13).

The need for high-resolution data has prompted the creation of downscaling frameworks, where a statistical or dynamical model refines the projections provided by a coarser-resolution Earth System Model (ESM) over an area of interest (14). Projecting regional climate change through downscaling is not without caveats: biases in the driving ESM can be exacerbated by the regional model—the so-called "garbage in, garbage out" problem (14, 15). Nevertheless, techniques to alleviate such biases are maturing (16, 17), and downscaling remains the best source of climate data at impact-relevant kilometer scales (18). Dynamical downscaling, in which a high-resolution regional climate model (RCM) is driven with the large-scale and boundary conditions from a global model, is widely recognized as the state-of-the-art method to obtain regional information about future climates (19). Dynamical downscaling is often performed in nested stages, each one increasing the resolution and reducing the spatial extent of the simulated domain (20–23). While dynamical downscaling enables trading spatial extent for resolution, it remains a computationally intensive task: a computational budget sufficient to simulate the global climate at 100 km resolution would fall short of that required to downscale a region the size of Spain to 10 km resolution. Computational constraints also limit the number of global projections that can be dynamically downscaled, resulting in data coverage gaps and an incomplete assessment of model uncertainty and internal variability in regional climate projections (24, 25).

Significance

Regional climate risk assessments serve as a crucial source of information for climate resilience and adaptation policies. The current regional climate modeling paradigm, which leverages physics-based models to downscale climate projections over limited areas, is too costly to apply to large climate projection ensembles. This hinders our ability to capture the uncertainty in regional climate projections. Alternative statistical downscaling methods, while efficient, often fail to capture compound extremes or generalize to unseen climate conditions. We propose a paradigm that jointly exploits physics-based models and generative AI to drastically reduce the cost of downscaling climate projections, while retaining the skill of physics-based approaches. This framework enables translating large climate projection ensembles into impact-relevant climate risk assessments.

Author contributions: I.L.-G., T.S., J.A., and F.S. designed research; I.L.-G. and Z.Y.W. performed research; I.L.-G. analyzed data; I.L.-G., Z.Y.W., and L.Z.-N. developed software and designed models; and I.L.-G., L.Z.-N., T.S., and F.S. wrote the paper.

Competing interest statement: The authors are employees of Google Limited Liability Company (LLC) and own Alphabet stock as part of the standard compensation package.

This article is a PNAS Direct Submission. N.S.D. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: ilopezgp@google.com or fsha@google.com.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2420288122/-/DCSupplemental.

Published April 25, 2025.

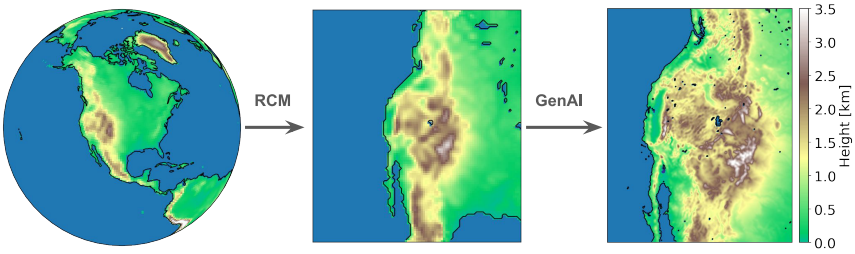

Fig. 1. Schematic of the dynamical-generative downscaling framework. A regional climate model (RCM) is used to downscale global simulations from different ESMs to an intermediate grid. A generative AI system (GenAI), such as a diffusion model, is then used to further downscale the RCM output to the desired resolution. The topographic height is shown at 100 km (*Left*), 45 km (*Middle*), and 9 km (*Right*) resolution, to showcase landscape changes at the different scales of the process. Water bodies are highlighted in blue.

To address these issues, we propose to leverage generative AI jointly with dynamical downscaling to make downscaling multimodel climate projection ensembles feasible. In our proposed framework, sketched in Fig. 1, an RCM is first used to downscale ESM output to an intermediate, but still coarse and inexpensive, resolution. In a second stage, a probabilistic diffusion model is used to efficiently downscale the intermediate RCM fields to the target resolution. Notably, the first dynamical downscaling stage maps input ESM data to atmospheric states that are consistent with the dynamics, resolution, and subgrid-scale parameterizations of a single RCM. This alleviates the need of the generative stage to generalize across the wide range of resolutions and physics of independent ESMs, and enables us to learn a generative model capable of downscaling multimodel ensembles using as training data dynamically downscaled output from a single ESM.

The output from the first stage is then fed into a generative diffusion model, which completes the downscaling process inexpensively, making efficient use of modern accelerators. Moreover, the generative stage enables sampling the uncertainty of the high-resolution fields given their large-scale context, retaining variable and spatial correlations that are crucial to assess the likelihood of compound extreme events (26, 27). This property provides additional downstream value beyond the computational benefits of deterministic statistical downscaling (28–30) and RCM emulators (31–34), neither of which provide a measure of downscaling error. Altogether, dynamical-generative downscaling enables obtaining regional climate projections with more accurate uncertainty bounds than those afforded by statistical downscaling methods or computationally limited dynamical downscaling approaches.

Our methodology is inspired by the remarkable ability of probabilistic diffusion models to perform conditional sampling of high-dimensional meteorological fields (35–38). Prior work has established the utility of generative models for downscaling short-time precipitation forecasts using radar data (39), and for downscaling historical weather reanalyses (36, 40). Here, we demonstrate how probabilistic diffusion models trained on data from a single ESM can be used to downscale multimodel climate projections, relying on the proven generalization abilities of physical RCMs. This task requires generalization to unseen climate forcings and is therefore more challenging than downscaling present weather. We evaluate the merits of our approach using as ground truth dynamically downscaled climate projections from the CMIP6 ensemble.

Results

We demonstrate our dynamical-generative downscaling framework using hourly data from the recently developed Western United States Dynamically Downscaled Dataset (WUS-D3) (41). WUS-D3 contains nested downscaled climate projections of a CMIP6 multimodel ensemble over the western United States. The first downscaling stage covers the central domain shown in Fig. 1 along with a 675 km east Pacific extension, at an intermediate resolution of 45 km. The second and final stage yields climate projections covering the rightmost domain depicted in Fig. 1 at 9 km resolution. The geographical diversity of this region cannot be fully captured by coarse-resolution simulations and highlights the need for downscaled climate projections.

The CMIP6 climate projections in WUS-D3 are taken from the ScenarioMIP intercomparison project (42) and follow the anthropogenic forcing conditions specified by the Shared Socioeconomic Pathway 3 (SSP3-7.0) (43). Our goal is to capture the internal variability and model uncertainty of climate projections under this forcing. From the WUS-D3 multimodel ensemble we select 8 ESMs for which historical biases, as well as interpolation errors in the sea surface temperature of the Gulf of California, were removed prior to downscaling (see *Materials and Methods* for more details). The RCM used for dynamical downscaling is the Weather Research and Forecasting Model (WRF), in its version 4.1.3 (44).

We seek to substitute the dynamical downscaling stage from the intermediate 45 km resolution to the final 9 km resolution with a generative model. This dynamical downscaling stage is roughly 40 times more computationally intensive than the first downscaling stage to 45 km resolution, covering about a third of the area with 5 times higher resolution. Therefore, our dynamical-generative downscaling framework is designed to overhaul a component that represents more than 97.5% of the cost of the original system. To this end, we train a Regional Residual Diffusion-based Downscaling (R2-D2) model to sample the probability density function of the residual between the fine (9 km) and coarse (45 km) resolution values of the meteorological fields specified in Table 1. Sampling with R2-D2 is highly efficient: with a batch size of 32 samples on 16 NVIDIA A100 GPUs, the model can generate over 800 downscaled samples per hour.

The R2-D2 model is a conditional probabilistic diffusion model that samples high-resolution residuals conditioned on the coarse resolution input data, as well as additional static input fields such as the topographic height at the target 9 km resolution. As in previous studies, we target the appropriate residuals in the generative learning task to facilitate learning and greatly improve generalization (35, 36). The residual modeling approach, combined with the mapping of the parent ESM fields to a common effective resolution provided by WRF in the first dynamical downscaling stage, enables us to train a general generative downscaling model using data from a single CMIP6 model.

Table 1. List of fields used as inputs and outputs to the generative model

Field	Type (resolution)
Temperature at 2 m	Input, Output
Specific humidity at 2 m	Input, Output
Zonal wind at 10 m	Input, Output
Meridional wind at 10 m	Input, Output
Surface pressure	Input, Output
Precipitation over last 24 h	Input, Output
Precipitation over last 12 h	Input
Precipitation over last 6 h	Input
Surface downwelling longwave flux	Input
Surface upwelling longwave flux	Input
Surface downwelling shortwave flux	Input
Surface upwelling shortwave flux	Input
Surface runoff	Input
Snow water equivalent	Input
Land mask, terrain height	Input (9 km)
Latitude, longitude	Input (9 km)
Orographic variance	Input (9 km)

All inputs are taken from the 45 km simulation unless otherwise noted.

We present dynamical-generative downscaling results from an R2-D2 model trained on 80 y (2014–2094) of CanESM5 SSP3-7.0 climate projection data (45). As shown in *SI Appendix*, our results are robust to the choice of training ESM. Model skill and generalization to unseen ESMs are assessed by downscaling the full 8-model ensemble, using the dynamically downscaled ensemble as the reference. The evaluation period spans 2095–2097, equivalent to 24 simulation years from a single climate projection. Two statistical downscaling baselines are included for comparison: Bias Correction and Spatial Disaggregation (BCSD) (46–48), and the STAR-ESDM method used in the Fifth US National Climate Assessment (49, 50). Both baselines utilize input data from the 45 km simulations and share training data with the R2-D2 model. Finally, bicubic interpolation of coarse-resolution fields to the target resolution is included as a non-value-added baseline.

Generative Downscaling Skill. The generative downscaling operator in our framework is inherently probabilistic, since the high-resolution fields in WUS-D3 are only constrained to follow their coarse-resolution counterpart at the domain boundaries. Boundary coupling prevents chaotic divergence of the large-scale fields but allows the finer scales to evolve freely (51). We assess the probabilistic skill of the R2-D2 model by generating 32-member ensembles for each evaluation date and employing the continuous ranked probability score (CRPS) as the error metric. The CRPS reduces to the mean absolute error (MAE) in the deterministic limit, facilitating comparisons with deterministic systems (52). The preservation of climate change signals by R2-D2 is demonstrated by considering the end-of-century and climate change bias. *SI Appendix* provides evaluations of additional fields and metrics, as well as individual generative samples.

Figs. 2 and 3 demonstrate that R2-D2 ensembles provide significant added value, reducing downscaling CRPS by over 40% compared to the baselines across all fields considered. R2-D2 also preserves realistic multivariate correlations, as evidenced by the kernel density estimates of near-surface winds at the Shepherds Flat wind farm in Oregon (Fig. 2*I*). Capturing multivariate correlations ensures that downscaling skill is maintained for derived fields, such as relative humidity (Fig. 2*B*). Importantly, correlations between extremes are also realistic. Fig. 2 *J*–*L* illustrates this in terms of the tail dependence (53) of hot and dry summer extremes in the southwestern United States, which have increased in frequency in recent decades (54). Many downstream climate downscaling applications, like hydrologic forecasting (55, 56), rely on realistic spatial structure in meteorological inputs. The radially averaged energy spectrum of the generative samples, which characterizes the realism of the downscaled output, follows closely that of the dynamically downscaled fields (Fig. 2 *E*–*H*). Minor differences are observed for precipitation, but these are still smaller than those observed for the baselines considered.

A fraction of the error reduction afforded by downscaling methods can be attributed to the correction of systematic climatological differences between the coarse and high-resolution fields. The reduction of these differences can be assessed in terms of the long-term bias, illustrated in Fig. 3 for near-surface temperature, precipitation, and wind speed. Quantile-mapping methods like BCSD and STAR-ESDM effectively correct biases in directly modeled variables by design, achieving average bias reductions over 57% for precipitation and 77% for temperature relative to interpolation (*SI Appendix*). R2–D2 matches and can even exceed the debiasing skill of quantile-based methods, as demonstrated by the noticeable reduction in near-surface temperature bias near the Great Salt Lake (Fig. 3 *A* and *M*). It also provides unbiased estimates for derived variables like wind speed, which simpler statistical methods may struggle with (Fig. 3*E*). R2-D2 ensembles preserve climate change signals in multimodel projections, as shown in Fig. 2 *C* and *D* for near-surface temperature and precipitation.

In addition to debiasing, the R2-D2 model further reduces the downscaling error by capturing the distribution of high-resolution anomalies conditioned on the coarse-resolution fields. This amounts to a general reduction of the downscaling error for all fields and over all regions compared to quantile-based methods, as presented in Fig. 3, and in *SI Appendix* in terms of the root mean square error. The enhanced conditional sampling skill provided by the R2-D2 model is important for applications that focus on particular events, such as extreme event attribution (57, 58) and climate storyline analysis (59, 60).

Quantifying Uncertainty in Multimodel Climate Projections. Quantifying model uncertainty and internal variability is crucial for accurate regional climate risk assessments but generally unaffordable through pure dynamical downscaling (24). We analyze the ability of dynamical-generative downscaling to capture the full distribution of regional climate projections by downscaling end-of-century multimodel projections under the SSP3-7.0 forcing scenario. Reliable estimates of the regional projection quantiles are particularly important, as they directly impact the accuracy of climate risk forecasts.

Alongside BCSD and STAR-ESDM, we compare our framework to a cost-saving strategy common in dynamical downscaling: a priori ESM selection (25). This method involves dynamically downscaling climate projections from a select subset of ESMs to reduce computational costs, albeit at the expense of diminished model uncertainty quantification. To emulate this approach, we create and evaluate 4-member subensembles from the original model ensemble, effectively halving the downscaled ESM ensemble size. This alternative remains significantly more computationally intensive than our hybrid approach, which only requires dynamical downscaling of a single ESM to 9

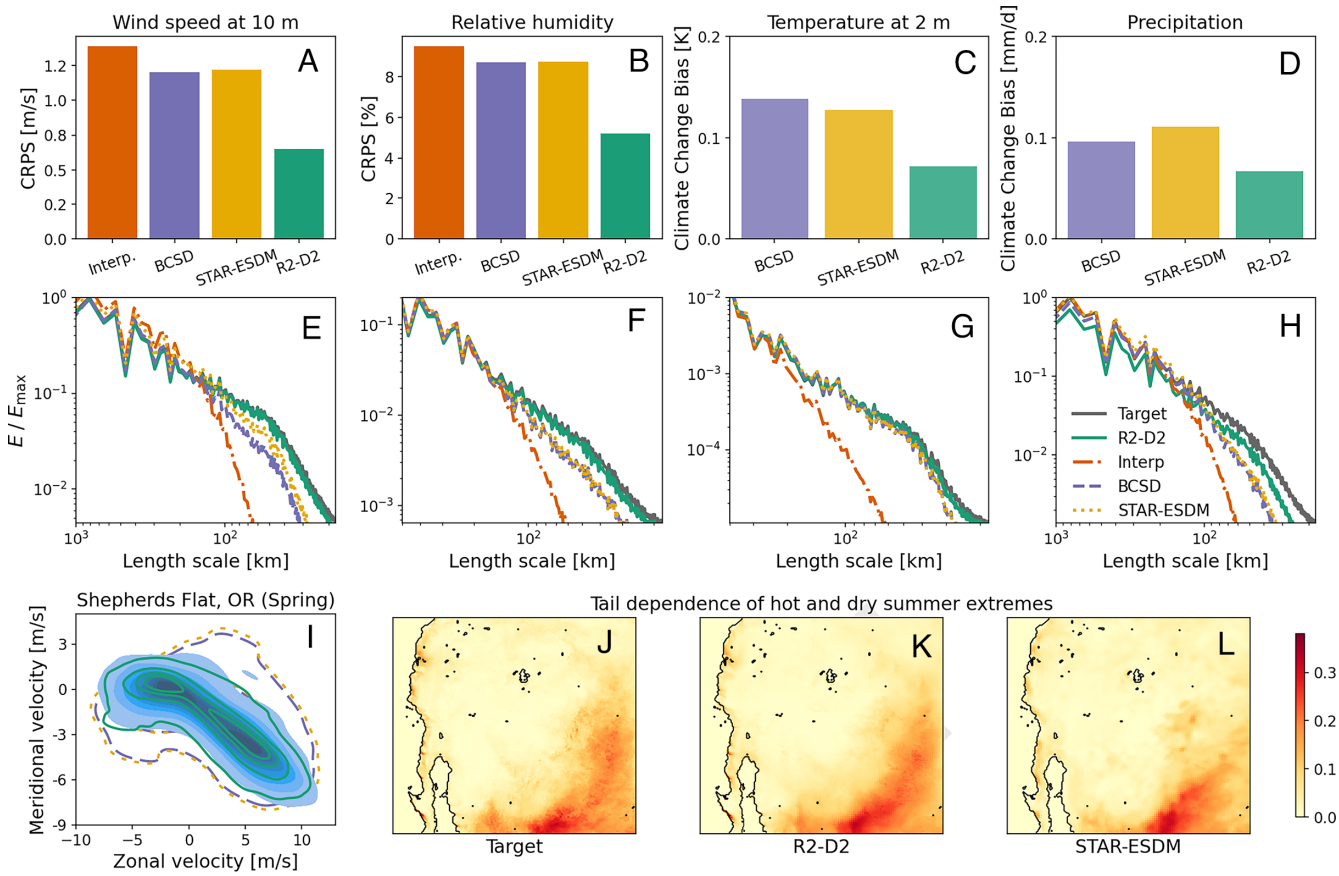

Fig. 2. Downscaling skill measured by land-averaged CRPS (*A* and *B*), absolute climate change bias (*C* and *D*), and energy spectra (*E*–*H*) for selected fields. The last row illustrates (*I*) the kernel density estimates (KDEs) of spring (March–May) near-surface winds near Shepherds Flat, and (*J*–*L*) the tail dependence of summer (June–August) hot and dry extremes. KDEs are shown for dynamical downscaling (blue shade), R2-D2, BCSD, and STAR-ESDM (legend in *H*). Results are computed from 4-hourly (*A*–*H*) or daily (*I*–*L*) data spanning years 2095–2097 of the multimodel SSP3-7.0 projection ensemble. The reference climate in (panels *C* and *D*) is evaluated over 2025–2027. R2-D2 employs 32-member ensembles.

km resolution for generative model training. We construct four such subensembles to ensure coverage of all available ESMs and mitigate the influence of model selection criteria.

Fig. 4 *A*–*C* evaluates the land-averaged MAE of the multimodel projection quantiles predicted by the dynamical-generative framework for selected variables and seasons, with respect to the full dynamically downscaled multimodel ensemble. Multimodel ensemble quantiles are computed using daily snapshots from all ESMs over the periods of June–August (Summer), September–November (Fall), and December–February (Winter). Daily frequency is chosen to prevent the diurnal cycle from dominating the projection distributions in variables like temperature. Therefore, the distributions shown capture the spread from internal variability and model uncertainty, and the lowest and highest quantiles represent extreme conditions with respect to the seasonal climatology. *SI Appendix* illustrates the spatial distribution of end-of-century extremes, including their projected climate change shifts, as determined by the different downscaling methods.

Fig. 4*A* presents results for the summer near-surface wet-bulb globe temperature (WBGT), computed following the simplified definition used by the Australian Bureau of Meteorology (61); all derived variables are defined in *SI Appendix*. The WBGT measures heat stress, accounting for both temperature and humidity. Accurately capturing its summer daytime distribution, especially its upper quantiles, is important for extreme heat risk assessment. The dynamical-generative approach surpasses all baselines in this task. Its downscaled climate projections match the target distribution better at all quantiles, reducing the 99%-quantile error by over 20% compared to 4-member subensembles and over 40% compared to the statistical downscaling baselines.

Quantile errors are also shown for fall near-surface wind speed and winter precipitation in Fig. 4 *B* and *C*. The proposed framework yields 99%-quantile error reductions of over 10% for wind speed and precipitation with respect to the best performing baseline, demonstrating the consistently superior performance of R2-D2 across seasons and downscaled fields. Results for additional variables and spring can be found in *SI Appendix*. The difference between dynamical-generative downscaling and subensemble dynamical downscaling is more substantial in variables for which model uncertainty is an important contributor to future climate uncertainty. In the case of temperature, statistical downscaling can also outperform pure dynamical downscaling of smaller ensembles. However, for variables that show a less pronounced intermodel spread contribution to the total uncertainty, such as wind speed, pure dynamical downscaling of subensembles is superior to BCSD or STAR-ESDM. The dynamical-generative approach improves upon the baseline approaches in most cases, particularly for extreme quantiles.

Local future climatological distributions projected with our framework are compared to the target distribution and the baselines in Fig. 4 *D*–*F* for selected regions. Summer temperatures in the coastal city of Ensenada, in Baja California, are strongly

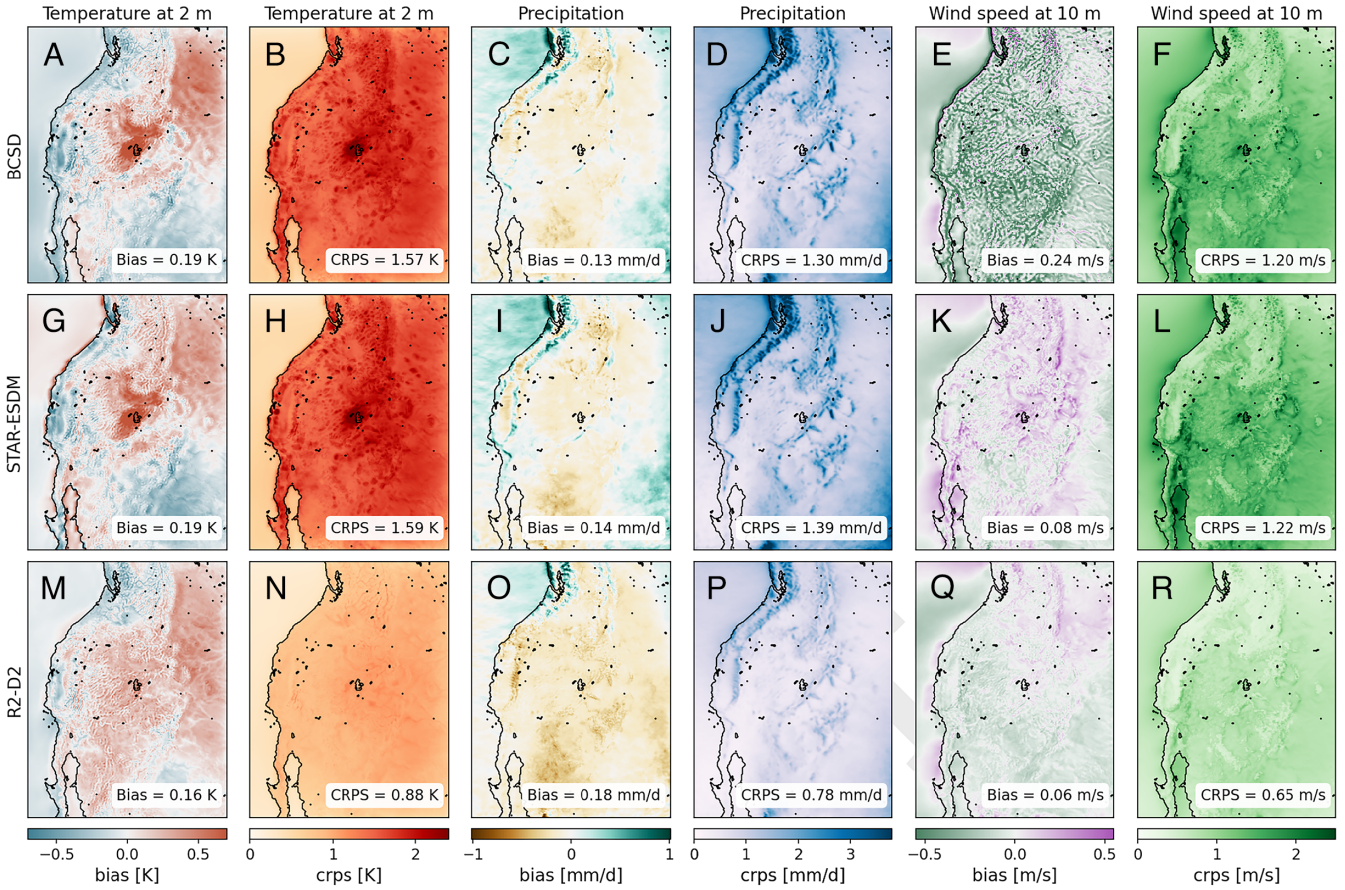

Fig. 3. Spatial distribution of downscaling bias (odd columns) and CRPS (even columns), shown for near-surface temperature, precipitation, and wind speed. Results are computed from 4-hourly data spanning years 2095-2097 of the multimodel SSP3-7.0 climate projection, and shown for BCSD (*A-F*), STAR-ESDM (*G-L*), and 32-member R2-D2 ensembles (M-R). The insets show the land-averaged absolute value of each metric.

affected by the inland extent of marine stratocumulus cloud decks (62). This makes Ensenada a location where downscaling can add important value in assessments of extreme heat risk. As shown in Fig. 4*D*, R2-D2 is able to track well the quantiles projected by the full dynamically downscaled ensemble, even at the tails of the distribution.

Downscaled projections are similarly essential for fire risk assessment in Southern California, where the hot and dry Santa Ana winds can fan wildfires in the fall, before the onset of the rainy season (7). The wildfire potential near Irvine, California, is explored in Fig. 4*E* in terms of the weather component of the Santa Ana Wildfire Threat Index [SAWTI; (63)], which is proportional to the kinetic energy of the near-surface winds and the dewpoint deficit. In this case, R2-D2 accurately captures the risk of the conditions most conducive to wildfires, characterized by high values of SAWTI. Both BCSD and STAR-ESDM underestimate the future wildfire risk. Finally, results are also shown for winter precipitation over Portland, Oregon, in Fig. 4*F*. In this case, too, the climatological distribution projected by the dynamical-generative framework closely follows that of the full dynamically downscaled climate ensemble, at a small fraction of the cost.

Regional Analysis of Compound Extremes. We further demonstrate the ability of dynamical-generative downscaling to capture regional compound extremes by analyzing the strongest Santa Ana wind event projected by the dynamically downscaled multimodel ensemble over the period September–November 2095. The event is selected as the 00 UTC snapshot with the highest average SAWTI over Southern California in the dynamically downscaled ensemble. This Santa Ana wind event occurred on November 13, 2095, of the SSP3-7.0 projection corresponding to the forcing model EC-Earth3-Veg (64). Fig. 5 *B* and *C* depict the 45-km resolution conditions during this event with respect to the fall climatology. The conditions on this date were characterized by stronger-than-usual northeasterly winds and anomalously dry air, particularly near the coast.

High spatial resolution is crucial to resolve the local acceleration of Santa Ana winds as they make their way from the Mojave Desert to the coast through the mountain passes, depicted in Fig. 5*A*. The wind flow and SAWTI values of the dynamically downscaled simulation are shown in Fig. 5*E*. The 9-km WRF simulations capture the deflection of the flow by the mountain ranges and the intensification of wildfire risk downstream, as the hot and dry desert air descends into the valleys and coastal areas of Southern California. R2-D2 captures these patterns as well, projecting similarly strong SAWTI conditions along the San Fernando and Santa Clara River valleys, and extending offshore into the Channel Islands. In contrast, STAR-ESDM fails to capture the offshore flow and high SAWTI conditions over the Channel Islands and tends to overestimate wildfire risk upstream of the mountain passes (Fig. 5*G*). BCSD is unable to capture the magnitude and spatial structure of wildfire risk, due to its inability to map coarse-resolution conditions to high-resolution climatological anomalies with realistic correlations (Fig. 5*H*).

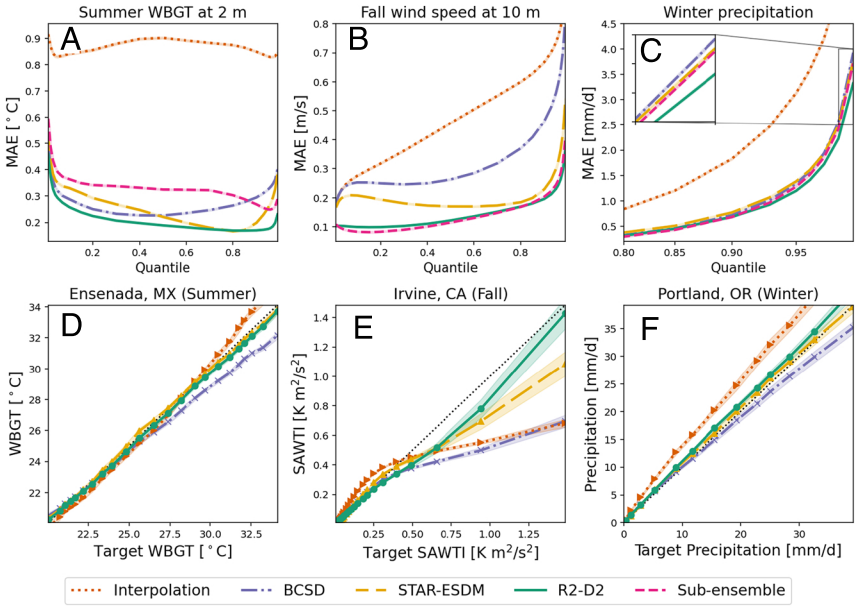

Fig. 4. Assessment of downscaled multimodel ensemble distribution fidelity. *Top*: Mean absolute error (MAE) of downscaled quantiles over land, with respect to the quantiles of the full dynamically downscaled ensemble. Results shown for summer WBGT (*A*), fall wind speed (*B*), and winter precipitation (*C*). *Bottom*: Quantile-quantile plots of summer WBGT (*D*), SAWTI (*E*), and precipitation (*F*) at specific locations with respect to the full dynamically downscaled ensemble. Quantiles from 0.01 to 0.99 are computed using daily snapshots at 00 UTC covering three-month seasons of 2095 to 2097. Results are shown for cubic interpolation, BCSD, STAR-ESDM, the generative model R2-D2, and for the average over 4-member dynamically downscaled subensembles. Uncertainty estimates represent the bootstrapped sample SD.

Wildfire risk in Southern California is most pronounced when strong Santa Ana winds blow over very dry vegetation, which acts as fuel. Since vegetation is at its driest before the start of the rainy season in late fall, projecting the timing of Santa Ana wind events is crucial to assess changes in wildfire risk over time. Fig. 5*D* illustrates the typical timing of the strongest

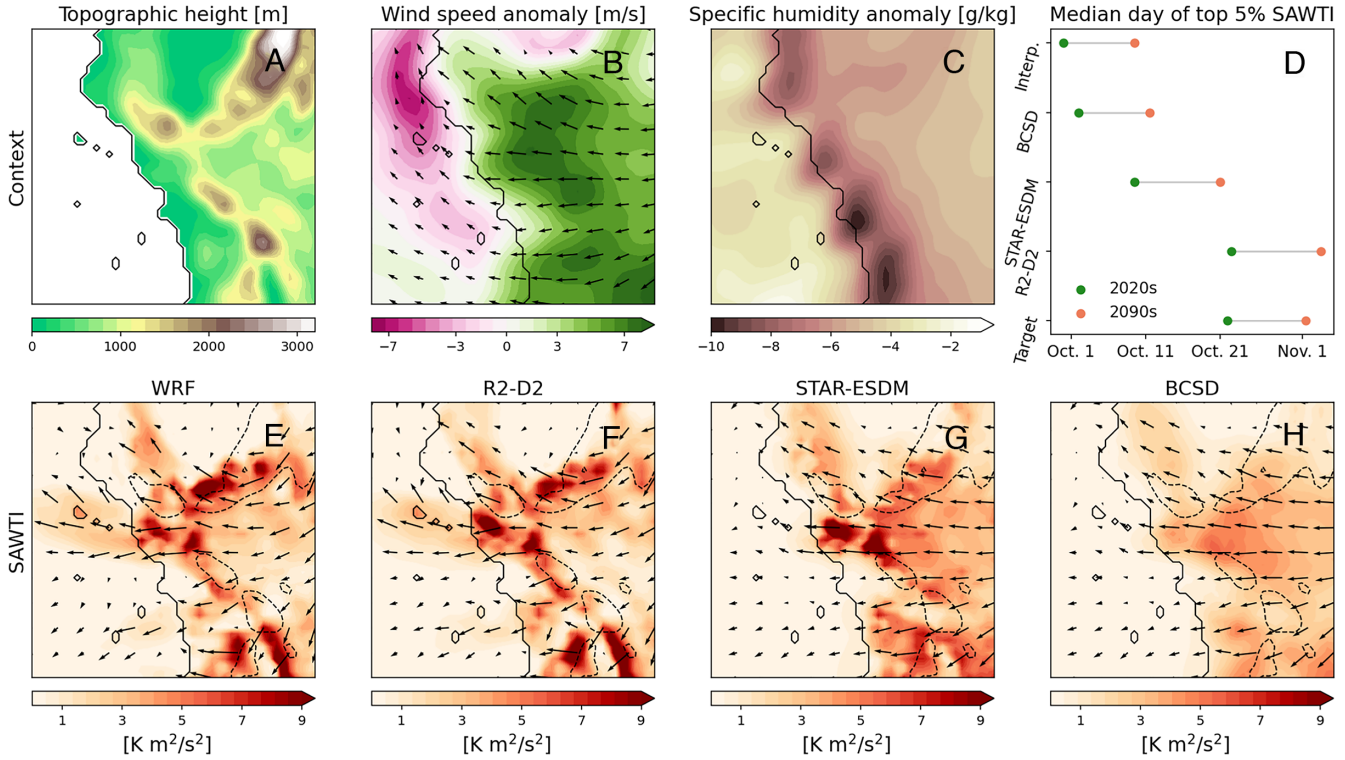

Fig. 5. Analysis of the strongest Santa Ana wind event in the multimodel projection for the period September–November 2095. (*A*) Topographic context. (*B* and *C*) 45-km resolution wind speed and specific humidity anomalies for the event date, relative to the September–November climatology (2095–2097). (*D*) Median day of the year on which the top 5% of Santa Ana wildfire threat index (SAWTI) conditions were observed over Southern California during fall, calculated using downscaled multimodel projections for 2025–2027 and 2095–2097. *Bottom*: Downscaled SAWTI from the target WRF (*E*), R2-D2 (*F*), STAR-ESDM (*G*), and BCSD (*H*). Quiver plots represent 10 m wind magnitude and direction, and dashed contours indicate the 1,200 m isohypse of the mountain ranges.

fall SAWTI conditions in Southern California, as projected by different downscaling methods. These conditions are shown for downscaled climate projections of the near future (2025–2027) and the late 21st century (2095–2097). Specifically, the timing is estimated by calculating the median day of the year for the top 5% of SAWTI conditions during the fall months (September–November) in each climate projection. The 45-km (Interp.) and 9-km (Target) WRF simulations both predict a similar shift of strong Santa Ana wind events toward later in the year with climate change, consistent with previous studies (65). However, the 9-km WRF simulations forecast peak Santa Ana wind conditions a few weeks later in the season, highlighting the impact of fine-scale processes on compound climate risks (4). These processes are not fully captured by BCSD and STAR-ESDM. In contrast, R2-D2 captures both the timing and the shift in Santa Ana wind events projected by dynamical downscaling.

Finally, another benefit of R2-D2 is its ability to assess the extent to which the strong wildfire conditions are determined by the coarse-resolution context. This can be analyzed in terms of the spread of the generative samples conditioned on the same coarse-resolution fields. In this particular case, the SAWTI conditions were tightly controlled by the coarse-resolution input, resulting in low variability in the generated samples and high confidence in severe wildfire risk. Additional samples supporting this tight large-scale control are included in *SI Appendix*, along with an analysis of the spread of the generative samples, which typically accounts for more than 70% of the downscaling error. More generally, the dynamical-generative framework provides estimates of the internal variability of downscaled fields given a large-scale context, which can be leveraged to construct counterfactuals and study the drivers of regional extremes (66, 67).

Discussion

Dynamical-generative downscaling combines the physical basis of dynamical downscaling with the sampling efficiency of diffusion models to provide regional climate projections that capture the full range of scenarios projected by existing ESMs. Dynamically downscaling coarse climate data from a wide range of ESMs to an intermediate resolution yields fields that are physically consistent with the RCM. This greatly simplifies the learning task of the generative model, by reducing the spectrum of conditioning inputs to those consistent with the dynamics of the RCM. The initial dynamical projection also creates a stable basis for incorporating high-resolution details, which facilitates residual learning.

These characteristics are exploited by the diffusion model in the generative stage, which can be trained on dynamically downscaled data from a single ESM to construct an efficient probabilistic sampler of high-resolution meteorological fields. The sampled fields display realistic spectra, extremal properties, and capture the multivariate uncertainty inherent in the downscaling process. Capturing spatial and field correlations maximizes the added value of the high-resolution projections by enabling downstream users to derive their own relevant climate indicators without loss of accuracy. Quantile-based methods such as BCSD do not capture these correlations, which can lead to the underestimation of compound risks (67, 68). Even in the case of univariate climate distributions, the generative R2-D2 model outperforms statistical downscaling methods such as BCSD and STAR-ESDM in downscaling skill. In addition, downscaled climate risk assessments with the proposed framework capture the uncertainty underlying climate projections better than dynamical downscaling of smaller ensembles, the current paradigm.

The dynamical-generative framework can be designed to be substantially more economical than pure dynamical downscaling, facilitating its application to very large climate model ensembles. In this study, the generative stage substitutes a component that consumes roughly 97.5% of the total computational budget in the original system. We document the use of the framework to downscale an ensemble of 8 ESMs. In this case, our method saves 85% of the dynamical downscaling cost, or 97.5% of the cost of 7 out of the 8 ESMs—a percentage that would increase for larger climate model ensembles. Inference with the diffusion model is relatively cheap: using a batch size of 32 samples on 16 NVIDIA A100 GPUs, the throughput exceeds 800 samples per hour. Moreover, the generative stage can sample downscaled fields without an initial spin-up time, which further boosts its efficiency compared to physics-based downscaling.

As previously demonstrated in weather forecasting (35), diffusion models can be used in conjunction with dynamical models of the atmosphere to reduce the cost of ensemble projections and augment the added value of the entire system. In the case of climate downscaling, this enables downscaling much larger climate model ensembles than those currently afforded by the dynamical downscaling paradigm (5, 25). This is essential to capture the regional impacts of climate change as projected by current state-of-the-art climate models.

Materials and Methods

Data for Learning and Evaluation. We derive our input and output data from the WUS-D3 dataset (41). WUS-D3 employs WRF to dynamically downscale global climate projections to an intermediate resolution of 45 km over a region covering the entire western United States, and extensions into the Midwest, the Pacific, western Canada, and northern Mexico, and shown in the center panel of Fig. 1. The 45 km WRF simulations are driven by 6-hourly lateral boundary conditions. They are also nudged toward the large-scale (>1,500 km) conditions of the forcing ESM with a relaxation timescale of 1.08 h. The 45 km grid is subsequently dynamically downscaled to 9 km resolution over the Western Electricity Coordinating Council US coverage area shown in the rightmost panel of Fig. 1, using 6-hourly lateral boundary conditions from the 45 km grid. In this study, we focus on the time-aligned hourly data available from both the coarse and high-resolution grids.

The data used for evaluation come from future climate projections under the SSP3-7.0 scenario using 8 different ESMs: CanESM5 (45), EC-Earth3-Veg (64), UKESM1-0-LL (69), MIROC6 (70), ACCESS-CM2 (71), MPI-ESM1-2-HR (72), NorESM2-MM (73), and TaiESM1 (74). The output of each ESM is debiased using the historical biases with respect to the ERA5 reanalysis (75) as a reference, prior to being used as input to WRF (17). In addition, sea surface temperatures in the Gulf of California are corrected to reflect its observed distribution, unresolved by most ESMs. This correction, as well as the resolution and specific climate projections used, are detailed in depth in the WUS-D3 description paper (41).

We train a probabilistic diffusion model to conditionally sample hourly snapshots of the meteorological field differences between the 9 km and 45 km resolution simulations. This difference is computed on the 9 km grid, after cubic interpolation of the 45 km resolution data to the 9 km grid. The generative model is conditioned on time-aligned and spatially interpolated 45 km meteorological fields, as well as static information about the 9 km grid. Therefore, all inputs and outputs cover the spatial grid with 340 × 270 degrees of freedom depicted in Fig. 1. The inputs and outputs of the model are shown in Table 1. Only spatial downscaling is considered, not time upsampling. To facilitate training, each input and output field is centered and normalized using its temporal and spatial mean and SD in the training set. Further details of our modeling framework are included in *SI Appendix*.

Given a forcing model used for training the generative model (CanESM5 in the main text), the period 2014–2094 is used for training, 2098–2100 for

validation, and 2095–2097 for testing. *SI Appendix* analyzes the climate change generalization of R2-D2 models trained on shorter time periods. Since the generative model only sees one ESM during training and evaluation, all other models used in the evaluation of the framework are part of the test set.

Generative Model Design and Training. The R2-D2 model is a conditional score-based diffusion model (76) for generative downscaling, with the training objective of denoising score matching (77). We adopt the denoising formulation introduced by Karras et al. (78) with a variance-exploding noise schedule. The denoiser is parameterized by a UNet-type convolutional architecture with about 180 million trainable parameters. The diffusion configuration and neural network architecture are described in detail in *SI Appendix*. Studies on the sensitivity of model performance to input selection, output selection, learning task, and climate change shifts between training and testing are also included therein.

We train and sample from the R2-D2 model using classifier-free guidance (79), with guidance strength of $g = 0.2$ and masking the conditioning input with probability $p_u = 0.1$ during training. We also employ a dropout probability $p_d = 0.1$ for the model weights during training. The diffusion model is trained using the Adam optimizer with a batch size of 128 for $2 \cdot 10^5$ steps, and exponential moving average decay of 0.9999. The learning rate schedule consists of cosine decay after an initial linear ramp-up, with a peak learning rate of $2 \cdot 10^{-4}$ and a terminal value of $1 \cdot 10^{-6}$. Training of the model takes about 5.5 d using 16 NVIDIA A100 GPUs.

Baselines. Dynamical-generative downscaling is evaluated against four baselines: cubic interpolation, BCSD, STAR-ESDM, and dynamically downscaled subensembles. Cubic interpolation outputs meteorological fields from the 45 km grid interpolated to the 9 km grid. In terms of the residual modeling approach used in our generative stage, cubic interpolation represents the null or zero residual output. BCSD follows a time-aligned implementation without temporal disaggregation, since inputs and outputs are aligned in our setting (48). The debiasing step is performed by quantile-mapping the coarse-resolution fields using as a reference the climatology of the low-pass filtered high-resolution fields over the dates and forcing model in the training dataset. Spatial disaggregation is then performed by retaining the debiased climatological anomalies and substituting the climatological mean by the high-resolution climatological mean. The retained climatological anomalies are additive for all variables, with the exception of precipitation, for which the multiplicative anomalies are retained (46).

STAR-ESDM decomposes input climate fields into three components with distinct timescales: a long-term trend, a slowly changing climatology, and a weather anomaly (49). Following this decomposition, the output of the method is the sum of three terms: the long-term third-order trend of the low-resolution data, debiased with respect to the high-resolution training data; the detrended climatological mean of the high-resolution training data, after a climate change adjustment inferred from changes in the coarse-resolution data climatology; and the quantile-mapped anomaly of the input data, also adjusted for climate change. The final baseline is a proxy for dynamical downscaling under computational constraints, which limits the number of ESMs that can be downscaled. We consider 4-member subensembles containing CanESM5 and three other ESMs. Four different subensembles are evaluated, such that all ESMs are used at least once in the analysis. This reduces the effect of model selection on the final results. The reported metrics for this baseline are the average metrics over the four subensembles.

Data, Materials, and Software Availability. Source code for our models, evaluation protocols, and tutorial notebooks are available on GitHub (https://github.com/google-research/swirl-dynamics/tree/main/swirl_dynamics/projects/probabilistic_diffusion/downscaling/gcm_wrf) (80). Pretrained model weights, as well as evaluation datasets, are available on Google Cloud (https://console.cloud.google.com/storage/browser/dynamical_generative_downscaling) (81). Previously published data were used for this work (WUS-D3 data were used for training and evaluation, described at https://doi.org/10.5194/gmd-17-2265-2024) (41).

ACKNOWLEDGMENTS. We thank Alex Hall and Stefan Rahimi for preliminary discussions and for providing information about the Western US Dynamically Downscaled Dataset dataset. We also thank Rob Carver for initial discussions about the data and Tyler Russell for technical program management. Finally, we thank Lizao Li, Stephan Hoyer, Michael Brenner, Peter Watson, Jorge Sebastián Moraga, and two anonymous reviewers for their insightful feedback.

Author affiliations: [a]Google Research, Mountain View, CA 94043; and [b]Department of Environmental Science and Engineering, California Institute of Technology, Pasadena, CA 91125

1. N. Pepin *et al.*, Elevation-dependent warming in mountain regions of the world. *Nat. Clim. Change* **5**, 424–430 (2015).
2. S. I. Seneviratne, D. Lüthi, M. Litschi, C. Schär, Land-atmosphere coupling and climate change in Europe. *Nature* **443**, 205–209 (2006).
3. E. J. Kendon *et al.*, Heavier summer downpours with climate change revealed by weather forecast resolution model. *Nat. Clim. Change* **4**, 570–576 (2014).
4. N. S. Diffenbaugh, J. S. Pal, R. J. Trapp, F. Giorgi, Fine-scale processes regulate the response of extreme events to global climate change. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 15774–15778 (2005).
5. W. J. Gutowski *et al.*, The ongoing need for high-resolution regional climate models: Process understanding and stakeholder information. *Bull. Am. Meteorol. Soc.* **101**, E664–E683 (2020).
6. G. Zängl, Dynamical aspects of wintertime cold-air pools in an alpine valley system. *Mon. Weather Rev.* **133**, 2721–2740 (2005).
7. M. Hughes, A. Hall, Local and synoptic mechanisms causing Southern California's Santa Ana winds. *Clim. Dyn.* **34**, 847–857 (2010).
8. C. J. Steele, S. R. Dorling, R. von Glasow, J. Bacon, Modelling sea-breeze climatologies and interactions on coasts in the southern North Sea: Implications for offshore wind energy. *Q. J. R. Meteorol. Soc.* **141**, 1821–1835 (2015).
9. B. Wang *et al.*, Sources of uncertainty for wheat yield projections under future climate are site-specific. *Nat. Food* **1**, 720–728 (2020).
10. C. Teutschbein, J. Seibert, Bias correction of regional climate model simulations for hydrological climate-change impact studies: Review and evaluation of different methods. *J. Hydrol.* **456–457**, 12–29 (2012).
11. Z. Khan *et al.*, Impacts of long-term temperature change and variability on electricity investments. *Nat. Commun.* **12**, 1643 (2021).
12. T. R. Knutson *et al.*, Dynamical downscaling projections of twenty-first-century Atlantic hurricane activity: CMIP3 and CMIP5 model-based scenarios. *J. Clim.* **26**, 6591–6617 (2013).
13. M. Goss *et al.*, Climate change is increasing the likelihood of extreme autumn wildfire conditions across California. *Environ. Res. Lett.* **15**, 094016 (2020).
14. F. Giorgi, Thirty years of regional climate modeling: Where are we and where are we going next? *J. Geophys. Res. Atmos.* **124**, 5696–5723 (2019).
15. A. Hall, Projecting regional change. *Science* **346**, 1461–1462 (2014).
16. C. L. Bruyère, J. M. Done, G. J. Holland, S. Fredrick, Bias corrections of global models for regional climate simulations of high-impact weather. *Clim. Dyn.* **43**, 1847–1856 (2014).
17. M. D. Risser *et al.*, Is bias correction in dynamical downscaling defensible?. *Geophys. Res. Lett.* **51**, e2023GL105979 (2024).
18. President's Council of Advisors on Science and Technology (PCAST), "Extreme weather risk in a changing climate: Enhancing prediction and protecting communities" (Report to the president, Washington, D.C., 2023).
19. F. Doblas-Reyes *et al.*, "Linking global to regional climate change" in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, V. Masson-Delmotte, Ed. (Cambridge University Press, 2021), pp. 1363–1512.
20. Y. Kurihara, R. E. Tuleya, M. A. Bender, The GFDL hurricane prediction system and its performance in the 1995 hurricane season. *Mon. Weather Rev.* **126**, 1306–1322 (1998).
21. X. Gao, J. S. Pal, F. Giorgi, Projected changes in mean and extreme precipitation over the Mediterranean region from a high resolution double nested RCM simulation. *Geophys. Res. Lett.* **33**, L03706 (2006).
22. K. Mahoney, M. A. Alexander, G. Thompson, J. J. Barsugli, J. D. Scott, Changes in hail and flood risk in high-resolution simulations over Colorado's mountains. *Nat. Clim. Change* **2**, 125–131 (2012).
23. X. Huang, D. L. Swain, A. D. Hall, Future precipitation increase from very high resolution ensemble downscaling of extreme atmospheric river storms in California. *Sci. Adv.* **6**, eaba1323 (2020).
24. D. W. Pierce, T. P. Barnett, B. D. Santer, P. J. Gleckler, Selecting global climate models for regional climate change studies. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 8441–8446 (2009).
25. N. Goldenson *et al.*, Use-inspired, process-oriented GCM selection: Prioritizing models for regional dynamical downscaling. *Bull. Am. Meteorol. Soc.* **104**, E1619–E1629 (2023).
26. K. Kornhuber *et al.*, Risks of synchronized low yields are underestimated in climate and crop model projections. *Nat. Commun.* **14**, 3528 (2023).
27. W. B. Anderson, R. Seager, W. Baethgen, M. Cane, L. You, Synchronous crop failures and climate-forced production variability. *Sci. Adv.* **5**, eaaw1976 (2024).
28. Y. Sha, D. J. Gagne, G. West, Deep-learning-based gridded downscaling of surface meteorological variables in complex terrain. Part I: Daily maximum and minimum 2-m temperature. *J. Appl. Meteorol. Climatol.* **59**, 2057–2073 (2020).
29. D. Quesada-Chacón, J. Baño-Medina, K. Barfus, C. Bernhofer, Downscaling CORDEX through deep learning to daily 1 km multivariate ensemble in complex terrain. *Earths Future* **11**, e2023EF003531 (2023).
30. P. M. M. Soares *et al.*, High-resolution downscaling of CMIP6 Earth system and global climate models using deep learning for Iberia. *Geosci. Model Dev.* **17**, 229–259 (2024).
31. A. Doury, S. Somot, S. Gadat, A. Ribes, L. Corre, Regional climate model emulator based on deep learning: Concept and first evaluation of a novel hybrid downscaling approach. *Clim. Dyn.* **60**, 1751–1779 (2023).

32. S. Hobeichi *et al*., Using machine learning to cut the cost of dynamical downscaling. *Earths Future* **11**, e2022EF003291 (2023).
33. M. van der Meer, S. de Roda Husman, S. Lhermitte, Deep learning regional climate model emulators: A comparison of two downscaling training frameworks. *J. Adv. Model. Earth Syst.* **15**, e2022MS003593 (2023).
34. D. Rastogi *et al*., Complementing dynamical downscaling with super-resolution convolutional neural networks. *Geophys. Res. Lett.* **52**, e2024GL111828 (2025).
35. L. Li, R. Carver, I. Lopez-Gomez, F. Sha, J. Anderson, Generative emulation of weather forecast ensembles with diffusion models. *Sci. Adv.* **10**, eadk4489 (2024).
36. M. Mardani *et al.*, Residual corrective diffusion modeling for km-scale atmospheric downscaling. *Commun Earth Environ* **6**, 124 (2025).
37. I. Price *et al*., Probabilistic weather forecasting with machine learning. *Nature* **637**, 84–90 (2025).
38. Z. Y. Wan *et al*., Debias coarsely, sample conditionally: Statistical downscaling through optimal transport and probabilistic diffusion models. *Adv. Neural Inf. Process. Syst.* **36**, 47749–47763 (2023).
39. L. Harris, A. T. T. McRae, M. Chantry, P. D. Dueben, T. N. Palmer, A generative deep learning approach to stochastic downscaling of precipitation forecasts. *J. Adv. Model. Earth Syst.* **14**, e2022MS003120 (2022).
40. F. Ling *et al*., Diffusion model-based probabilistic downscaling for 180-year east Asian climate reconstruction. *NPJ Clim. Atmos. Sci.* **7**, 131 (2024).
41. S. Rahimi *et al*., An overview of the western United States dynamically downscaled dataset (WUS-D3). *Geosci. Model Dev.* **17**, 2265–2286 (2024).
42. B. C. O'Neill *et al*., The scenario model intercomparison project (ScenarioMIP) for CMIP6. *Geosci. Model Dev.* **9**, 3461–3482 (2016).
43. K. Riahi *et al*., The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Glob. Environ. Chang.* **42**, 153–168 (2017).
44. W. Skamarock *et al*., "A description of the advanced research WRF model version 4" (Tech. Rep. NCAR/TN-556+STR, NCAR, 2021).
45. N. C. Swart *et al*., The Canadian Earth System Model version 5 (CanESM5.0.3). *Geosci. Model Dev.* **12**, 4823–4873 (2019).
46. A. W. Wood, E. P. Maurer, A. Kumar, D. P. Lettenmaier, Long-range experimental hydrologic forecasting for the eastern United States. *J. Geophys. Res. Atmos.* **107**, ACL 6–1–ACL 6–15 (2002).
47. A. W. Wood, L. R. Leung, V. Sridhar, D. P. Lettenmaier, Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Clim. Change* **62**, 189–216 (2004).
48. B. Thrasher, E. P. Maurer, C. McKellar, P. B. Duffy, Technical note: Bias correcting climate model simulated daily temperature extremes with quantile mapping. *Hydrol. Earth Syst. Sci.* **16**, 3309–3314 (2012).
49. K. Hayhoe, I. Scott-Fleming, A. Stoner, D. J. Wuebbles, STAR-ESDM: A generalizable approach to generating high-resolution climate projections through signal decomposition. *Earths Future* **12**, e2023EF004107 (2024).
50. S. Basile, A. R. Crimmins, C. W. Avery, B. D. Hamlington, K. E. Kunkel, "Appendix 3. Scenarios and datasets" in *Fifth National Climate Assessment*, A. Crimmins *et al*., Eds. (U.S. Global Change Research Program, 2023).
51. B. Barthel Sorensen et al., A probabilistic framework for learning non-intrusive corrections to long-time climate simulations from short-time training data. arXiv [Preprint] (2024). https://doi.org/10.48550/arXiv.2408.02688 (Accessed 1 September 2024).
52. H. Hersbach, Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather Forecast.* **15**, 559–570 (2000).
53. R. Schmidt, U. Stadtmüller, Non-parametric estimation of tail dependence. *Scand. J. Stat.* **33**, 307–335 (2006).
54. K. A. McKinnon, A. Poppick, I. R. Simpson, Hot extremes have become drier in the United States Southwest. *Nat. Clim. Change* **11**, 598–604 (2021).
55. G. Bürger, Y. Chen, Regression-based downscaling of spatial variability for hydrologic applications. *J. Hydrol.* **311**, 299–317 (2005).
56. D. Maraun *et al*., Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user. *Rev. Geophys.* **48**, 000314 (2010).
57. K. Emanuel, R. Sundararajan, J. Williams, Hurricanes and global warming: Results from downscaling IPCC AR4 simulations. *Bull. Am. Meteorol. Soc.* **89**, 347–368 (2008).
58. K. E. Trenberth, J. T. Fasullo, T. G. Shepherd, Attribution of climate extreme events. *Nat. Clim. Change* **5**, 725–730 (2015).
59. E. M. Fischer *et al*., Storylines for unprecedented heatwaves based on ensemble boosting. *Nat. Commun.* **14**, 4643 (2023).
60. X. Liu *et al*., How do climate model resolution and atmospheric moisture affect the simulation of unprecedented extreme events like the 2021 western North American heat wave?. *Geophys. Res. Lett.* **51**, e2024GL108160 (2024).
61. K. M. Willett, S. Sherwood, Exceedance of heat index thresholds for 15 regions under a warming climate using the wet-bulb globe temperature. *Int. J. Climatol.* **32**, 161–177 (2012).
62. R. E. S. Clemesha, A. Gershunov, S. F. Iacobellis, A. P. Williams, D. R. Cayan, The northward march of summer low cloudiness along the California coast. *Geophys. Res. Lett.* **43**, 1287–1295 (2016).
63. T. Rolinski *et al*., The Santa Ana wildfire threat index: Methodology and operational implementation. *Weather Forecast.* **31**, 1881–1897 (2016).
64. R. Döscher *et al*., The EC-Earth3 earth system model for the Coupled Model Intercomparison Project 6. *Geosci. Model Dev.* **15**, 2973–3020 (2022).
65. J. Guzman-Morales, A. Gershunov, Climate change suppresses Santa Ana winds of Southern California and sharpens their seasonality. *Geophys. Res. Lett.* **46**, 2772–2780 (2019).
66. J. Sillmann *et al*., Event-based storylines to address climate risk. *Earths Future* **9**, e2020EF001783 (2021).
67. J. Zscheischler *et al*., Future climate risk from compound events. *Nat. Clim. Change* **8**, 469–477 (2018).
68. V. S. Chandel, U. Bhatia, A. R. Ganguly, S. Ghosh, State-of-the-art bias correction of climate models misrepresent climate science and misinform adaptation. *Environ. Res. Lett.* **19**, 094052 (2024).
69. A. A. Sellar *et al*., UKESM1: Description and evaluation of the U.K. earth system model. *J. Adv. Model. Earth Syst.* **11**, 4513–4558 (2019).
70. H. Tatebe *et al*., Description and basic evaluation of simulated mean state, internal variability, and climate sensitivity in MIROC6. *Geosci. Model Dev.* **12**, 2727–2765 (2019).
71. D. Bi *et al*., Configuration and spin-up of ACCESS-CM2, the new generation Australian community climate and earth system simulator coupled model. *J. South. Hemisphere Earth Syst. Sci.* **70**, 225–251 (2020).
72. O. Gutjahr *et al*., Max Planck institute earth system model (MPI-ESM1.2) for the high-resolution model intercomparison project (HighResMIP). *Geosci. Model Dev.* **12**, 3241–3281 (2019).
73. Ø. Seland *et al*., Overview of the Norwegian earth system model (NorESM2) and key climate response of CMIP6 DECK, historical, and scenario simulations. *Geosci. Model Dev.* **13**, 6165–6200 (2020).
74. Y. C. Wang *et al*., Performance of the Taiwan earth system model in simulating climate variability compared with observations and CMIP6 model simulations. *J. Adv. Model. Earth Syst.* **13**, e2020MS002353 (2021).
75. H. Hersbach *et al*., The ERA5 global reanalysis. *Q. J. R. Meteorol. Soc.* **146**, 1999–2049 (2020).
76. J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* **33**, 6840–6851 (2020).
77. P. Vincent, A connection between score matching and denoising autoencoders. *Neural Comput.* **23**, 1661–1674 (2011).
78. T. Karras, M. Aittala, T. Aila, S. Laine, Elucidating the design space of diffusion-based generative models. *Adv. Neural Inf. Process. Syst.* **35**, 26565–26577 (2022).
79. J. Ho, T. Salimans, Classifier-free diffusion guidance. arXiv [Preprint] (2022). https://doi.org/10.48550/arXiv.2207.12598 (Accessed 7 August 2024).
80. I. Lopez-Gomez, Z. Y. Wan, Supporting software for "Dynamical-generative downscaling of climate model ensembles". GitHub. https://github.com/google-research/swirl-dynamics/tree/main/swirl_dynamics/projects/probabilistic_diffusion/downscaling/gcm_wrf. Deposited 20 February 2025.
81. I. Lopez-Gomez, Supporting data for "Dynamical-generative downscaling of climate model ensembles". Google Cloud. https://console.cloud.google.com/storage/browser/dynamical_generative_downscaling. Deposited 14 January 2025.