



RESEARCH ARTICLE

10.1029/2025MS005431

Key Points:

- CERA improves generalization across climates via reducing distribution shift, outperforming raw-input and physics-informed baselines
- Latent alignment using Earth Mover's Distance improves both predictive accuracy and robustness across random seeds
- The framework is tested on a parameterization of moist processes, but has potential for other climate applications such as downscaling

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

S. Liu,
shuchang.liu@hotmail.com

Citation:

Liu, S., & O'Gorman, P. A. (2026). CERA: A framework for improved generalization of machine learning models to changed climates. *Journal of Advances in Modeling Earth Systems*, 18, e2025MS005431. <https://doi.org/10.1029/2025MS005431>

Received 21 AUG 2025

Accepted 12 JUN 2026

Author Contributions:

Conceptualization: Shuchang Liu, Paul A. O'Gorman

Data curation: Shuchang Liu

Formal analysis: Shuchang Liu

Investigation: Shuchang Liu

Methodology: Shuchang Liu

Software: Shuchang Liu

Validation: Shuchang Liu

Visualization: Shuchang Liu

Writing – original draft: Shuchang Liu

Writing – review & editing:

Shuchang Liu, Paul A. O'Gorman

© 2026 The Author(s). Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

CERA: A Framework for Improved Generalization of Machine Learning Models to Changed Climates

Shuchang Liu¹  and Paul A. O'Gorman¹ 

¹Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract Robust generalization under climate change remains a major challenge for machine learning applications in climate science. Most existing approaches struggle to extrapolate beyond the climate they were trained on, leading to a strong dependence on training data from model simulations of warm climates. Use of climate-invariant inputs improves generalization but requires challenging manual feature engineering. Here, we present CERA (Climate-invariant Encoding with Representation Alignment), a machine learning framework consisting of an autoencoder with explicit latent-space alignment, followed by a predictor for downstream process estimation. We test CERA on the problem of parameterizing moist-physics processes. Without training on labeled data from a +4 K climate, CERA leverages labeled control-climate data and unlabeled warmer-climate inputs to improve generalization to the warmer climate, outperforming both raw-input and physically informed baselines in predicting key moisture and energy tendencies. It captures not only the vertical and meridional structures of the moisture tendencies, but also shifts in the intensity distribution of precipitation including extremes. Ablation experiments show that latent alignment improves both accuracy and the robustness across random seeds used in training. While some reduced skill remains in the boundary layer, the framework offers a data-driven alternative to manual feature engineering of climate invariant inputs. Beyond parameterizations used in hybrid ML-physics systems, the approach holds promise for other climate applications such as statistical downscaling.

Plain Language Summary Predicting how the atmosphere will behave in a warmer world is one of the biggest challenges in climate science. While machine learning has shown promise in improving weather and climate models, most approaches struggle to work reliably outside the climate conditions they were trained on. In this study, we introduce CERA (Climate-invariant Encoding with Representation Alignment), a machine learning method designed to improve performance under climate change. CERA first learns a shared internal representation of the atmosphere using data from both present-day (control) and warmer (+4°C) climates, without training on any outputs from the warmer climate. A second stage of the model then learns to predict key atmospheric processes using only outputs from the control climate. CERA makes more accurate predictions under warming than models that rely on raw inputs or manually engineered physical features. CERA offers a flexible, data-driven alternative to previous approaches to climate-change generalization. We apply CERA to moist physics processes in the atmosphere (including deep convection), but it holds potential for broader uses in climate science such as for estimating higher-resolution climate fields from coarse model output.

1. Introduction

Machine learning (ML) has emerged as a powerful tool for advancing weather and climate science (Bracco et al., 2025; Chantry et al., 2021; Eyring et al., 2024; Nguyen et al., 2023; Schneider et al., 2022). It has been applied across a wide range of tasks, such as weather forecasting (Bi et al., 2022; Kochkov et al., 2024; Kurth et al., 2023; Lam et al., 2023; McNally et al., 2024; Price et al., 2025; Weyn et al., 2021), parameterization of subgrid processes (Brenowitz & Bretherton, 2018; Gentile et al., 2018; O'Gorman & Dwyer, 2018; Yuval & O'Gorman, 2020), emulation of global climate models (Cachay et al., 2024; Chapman et al., 2025; Watt-Meyer et al., 2025), bias-correction of weather and climate models (Bora et al., 2023; Bretherton et al., 2022; Chapman & Berner, 2025; Gregory et al., 2025; Watt-Meyer et al., 2021) and statistical downscaling (Flora & Potvin, 2025; Hobeichi et al., 2023; Rampal et al., 2022). In particular, ML-based parameterizations have shown promise in representing the effects of complex subgrid processes such as moist convection and boundary-layer turbulence by training on high-resolution model data (Brenowitz & Bretherton, 2018; Han et al., 2023; Yuval & O'Gorman, 2020) or observations (McCandless et al., 2022; Zhao et al., 2019).

Despite this promise, a persistent challenge for ML models in climate applications is their limited generalization to out-of-distribution conditions. This problem is especially acute under climate change, where shifts in temperature, moisture, and circulation patterns can push the system into regimes not well represented in the training data. For instance, ML parameterizations trained on present-day conditions often perform poorly when evaluated in warmer climates, leading to inaccurate predictions of precipitation and subgrid tendencies (O’Gorman & Dwyer, 2018; Rasp et al., 2018; Scher & Messori, 2019). This generalization gap limits the reliability and robustness of ML-based tools in climate change scenarios, precisely where improved predictions are most needed.

Several approaches have been proposed to improve performance in different climates. The most straightforward approach is to train ML models on high-resolution model output across multiple climates, so that the learned representation is implicitly exposed to climatic variability (Bodnar et al., 2025; Clark et al., 2022; O’Gorman & Dwyer, 2018; Rasp et al., 2018; Y. Q. Sun et al., 2024), although this implies a strong reliance on the quality of the high-resolution model output in different climates which is difficult to validate using observations. Also there is a need to ensure the training data covers a sufficiently wide range of climates for the application at hand. Another strategy is to incorporate physics-based constraints, such as conservation of energy and moisture, non-negativity of precipitation, or constraints based on known thermodynamic relationships which may improve generalization (Brenowitz et al., 2020; Kashinath et al., 2021; Perezhogin et al., 2025; Yuval et al., 2021). In addition, Beucler et al. (2024) proposed a climate-invariant feature engineering framework, in which input variables are transformed (e.g., from specific humidity and temperature to relative humidity (RH) and plume buoyancy) to reduce their distribution shift under climate forcing. While such physically motivated transformations have shown promise, they rely on hand-crafted features derived from expert knowledge. These transformations often require time-consuming trial-and-error tuning and may not yield optimal solutions, as they could sacrifice important information in pursuit of approximate climate invariance. Moreover, some approaches assume access to labeled outputs from future climates, which may not be available in practice, especially when using real-world observations. Note that labeled data in the context of machine learning refers to the ground truth output data that is used for supervised learning.

In this study, we introduce a third path: learning representations that are approximately climate invariant (i.e., have smaller distribution shift between climates) directly from data, without relying on expert-designed features or supervision from warmer-climate labels. Our model consists of an autoencoder (AE) and a multilayer perceptron (MLP) trained jointly in an end-to-end fashion. The AE learns a shared latent representation of the input profiles from both climates, while the MLP predicts subgrid thermodynamic tendencies from this latent code. Inputs from both present-day and +4 K warmer climates are used to train the AE, but the MLP is supervised only using control-climate outputs. The AE ensures that the latent representation preserves the information necessary to reconstruct the original temperature and humidity profiles for both climates. To encourage the latent space to remain aligned across climates, we introduce a distributional regularization term based on the Earth Mover’s Distance (EMD) (Rubner et al., 2000), which penalizes divergence between the latent encodings of control and +4 K inputs. We recognize that not all physical information can be climate invariant while still reconstructing both climates, and thus we reserve part of the latent space that is not aligned by the EMD loss and is not used for prediction of outputs.

The results show that our method achieves comparable or better performance than approaches based on hand-crafted physical inputs, without requiring any outputs from warmer climate conditions. We do still need inputs from the warmer climate, and in practice these could be obtained from climate-model simulations at coarse resolution or potentially a pseudo-global warming approach applied to observations (Schär et al., 1996). By not requiring outputs, and in particular the relationship between inputs and outputs, in the warmer climate, we are reducing dependence on high-resolution model output in a warmer climate. This highlights the potential of our method as a data-driven strategy for improving generalization in climate ML.

We test our approach on the problem of parameterization of subgrid moist processes in the atmosphere using the neural-network parameterization of Yuval et al. (2021). We consider a +4 K warming as a stringent test of generalization, but we emphasize that a useful application in practice could be for a much smaller magnitude of warming. For example, one could train on recent decades and potentially use climate-invariance to improve generalization for a modest warming over the subsequent few decades or to improve robustness to extremes of internal variability that are not present in the training data.

We describe the data, architecture and baseline models in Section 2. We then describe the offline performance of CERA (not coupled to a physical model) in Section 3 with an additional emphasis on precipitation which is important for climate impacts. We present online results from coarse-resolution simulations run with the subgrid parameterization trained with CERA in Section 4. Finally we discuss our results and their implications in Section 5.

2. Methods

2.1. Input–Output Specification and Experimental Setup

We use the same high-resolution model output and coarse-graining procedure as in Yuval et al. (2021) supplemented by a warmer simulation from O’Gorman et al. (2021) to test generalization across climates. Both simulations were performed using the System for Atmospheric Modeling (SAM, version 6.3) (Khairoutdinov & Randall, 2003) in a quasi-global aquaplanet setup on an equatorial beta-plane with 48 vertical levels and a zonally symmetric SST distribution. A hypohydrostatic rescaling factor of 4 was used to enable convection-resolving simulations at 12 km horizontal resolution. The high-resolution simulation output was coarse-grained by a factor of 8 (to 96 km) using spatial averaging. For a range of thermodynamic and moisture variables, subgrid fluxes were calculated by differencing coarse-grained and resolved fluxes, and coarse-grained tendencies were calculated by coarse-graining sources/sinks. The control simulation (0 K) has baseline SSTs while the +4 K warmer simulation has SSTs uniformly increased by 4 K. These are the same SAM simulations as used in Beucler et al. (2024) but in that study they were referred to as −4 and 0 K. Each high-resolution simulation was run for 600 days, with the final 500 days used for training and evaluation. The data from this period are randomly split into 95% for training and 5% for testing.

The inputs to the moist physics parameterization consist of vertical profiles of temperature T and total non-precipitating water mixing ratio q_T , evaluated at the lowest 30 full model levels. This yields a total of $30 \times 2 = 60$ input features. The outputs include five vertically resolved quantities: the subgrid contribution to the flux of non-precipitating liquid/ice static energy H_L due to vertical advection ($H_{L,adv}$) and the total tendency of H_L from freezing and melting of precipitation ($H_{L,phase}$); the subgrid contribution to the flux of q_T due to vertical advection ($q_{T,adv}$) and cloud ice sedimentation ($q_{T,sed}$), and the total tendency of q_T due to microphysical conversions between q_T and precipitating water ($q_{T,mic}$). These outputs are defined on different vertical grids: vertical advective fluxes are predicted at 29 “half” model levels above the surface, ice sedimentation fluxes at 30 “half” model levels, and total tendencies at 30 “full” model levels. This results in a total of $29 \times 2 + 30 \times 3 = 148$ output features.

Following Yuval et al. (2021), we predict subgrid fluxes rather than tendencies for advection and cloud ice sedimentation. Furthermore, the NN does not directly predict energy changes due to cloud microphysics and ice sedimentation. Instead, these are diagnosed physically from the corresponding moisture tendencies. This overall approach ensures good conservation of energy and water by the subgrid parameterization. Since CERA does not alter the outputs that are predicted, it maintains the energy and water conservation properties of the original subgrid parameterization.

For full details of the simulation setup, coarse-graining methodology, and the energy and moisture variables used we refer the reader to Yuval et al. (2021). One deviation from that setup is that we omit radiative heating from the parameterization, as it is not clear a priori if it should have the same climate invariant inputs as the moist physics processes that are our main focus. Consequently, we exclude from the inputs the absolute meridional distance from the equator, which was previously used as a proxy for insolation, and we omit radiative heating from the outputs. Another difference is that our machine learning model uses input data that are normalized per column, whereas Yuval et al. (2021) apply normalization per vertical level. Output variables are normalized per column, but for simplicity we do not apply the output reweighting for different types of physical variables used by Yuval et al. (2021).

However, we did include radiative heating in an alternative version of our analysis (results shown in Figure 3) which was designed to be directly comparable to Beucler et al. (2024). The inputs in this case include the absolute meridional distance from the equator, $|y|$, which serves as a proxy for insolation. The five vertically resolved outputs listed earlier plus radiative heating are combined into the subgrid tendencies of H_L and q_T using equations S2 and S5 of Yuval et al. (2021). The model was retrained to directly predict these two combined tendencies.




Figure 1. Schematic of CERA. The top half illustrates the autoencoder stage, where inputs from control (X_{control}) and warm (X_{warm}) climates are encoded into latent representations (Z_{control} and Z_{warm}) using shared 1D convolutional layers. Reconstruction losses $MSE(X, \hat{X})$ are applied to both climates, while an Earth Mover's Distance loss $EMD(Z_{\text{control}}, Z_{\text{warm}})$ encourages alignment between their latent spaces. The bottom half depicts the predictor stage, where the latent representation from the control climate is reshaped and passed through a predictor network to estimate target outputs $\hat{Y}$, with an $MSE(Y, \hat{Y})$ loss applied for supervision. The framework enables learning a latent space with reduced distribution shift between climates, supporting improved generalization to warmer climate conditions.

We compare all our results with the climate invariant method as introduced in Beucler et al. (2024), in which the inputs are transformed to a plume buoyancy variable (B) and RH. These transformations were designed to improve the generalization of learned parameterizations across different climate states.

2.2. Model Architecture, Training, and Evaluation

We propose CERA (Climate-invariant Encoding with Representation Alignment), a self-supervised model that learns structure with reduced distribution shift across climate regimes directly from raw inputs, without any feature engineering or labels from warmer climates. As illustrated in Figure 1, the model consists of two components: an autoencoder (AE) that processes multi-level input features (i.e., vertical profiles), and a MLP that predicts target outputs from the learned latent representations.

The AE operates on a 60-dimensional input vector comprising vertical profiles of temperature and total non-precipitating water. The encoder uses a stack of eight one-dimensional convolutional layers. Each convolution has a kernel size of one and 64 channels, except for the final encoder layer, which outputs a three channel latent tensor $Z \in \mathbb{R}^{B \times 3 \times L}$, where B is the batch size and L is the number of vertical levels, which is 30 in our case. A kernel size of one means that each vertical level is transformed independently of its neighbors. This preserves vertical locality and can be viewed as applying the same dense transformation at each vertical level, which makes the latent alignment straightforward and simplifies interpretation of the mapping as a local variable transformation. With this configuration, the encoder does not represent vertical interactions. Any non-locality enters only through the predictor, which maps the latent space from all levels to all outputs. The decoder mirrors the encoder with transposed 1D convolutions, known as deconvolutions (see Zeiler et al. (2010)) to reconstruct the original input from the full latent code Z , enabling self-supervised learning via reconstruction loss.

To recognize that not all input information can be reliably aligned while still reconstructing the inputs in different climates in the decoder, the latent space is partitioned. One channel of Z at each vertical level is reserved for such non-aligned information for the AE and does not participate in either the distribution alignment or the downstream predictor.




Figure 2. Offline R^2 comparison across models for outputs in control and +4 K climates. Blue dots indicate the coefficient of determination (R^2) for each method (Baseline, RH + B baseline, CERA and CERA-noAlign) in the control climate, evaluated across five subgrid variables. Red triangles show the corresponding R^2 scores under the +4 K warming scenario. Results are displayed over five individual random seeds used in training per method. CERA achieves the highest overall performance, especially under warming, while RH + B exhibits greater robustness. CERA-noAlign improves upon the raw baseline but not to the same extent as CERA and with increased variability across random seeds, underscoring the importance of latent alignment. Note the different vertical axes for the different panels.

The MLP predictor takes only the aligned subset of Z as input and maps it to a 148-dimensional output vector representing vertically resolved subgrid fluxes and tendencies. The MLP predictor is a five-layer fully connected neural network with 128 neurons per hidden layer and leaky ReLU activations.

We adapted the model for the alternative version of our analysis (results shown in Figure 3) to ensure a direct comparison with Beucler et al. (2024), accounting for differences in input and output variables described in the previous section. Since the distance to the equator $|y|$ is assumed to affect only the radiative heating component, we process $|y|$ through a fully-connected two-layer MLP with 8 neurons in the hidden layer and 30 neurons in the

output layer, and add its output as a residual correction exclusively to the energy tendency. This design choice minimizes the influence of $|y|$ on the rest of the network.

The AE is trained using input data from both the control and +4 K warmer climate states, while the MLP is trained using only control-climate outputs as supervision. To enforce a shared latent structure across climate regimes, we introduced a loss term based on the Earth Mover's Distance (EMD) (Rubner et al., 2000) between the latent distributions of control and +4 K climate inputs. In one dimension, the EMD simplifies to the distance between sorted samples, allowing for an efficient closed-form computation (Levina & Bickel, 2001). In our implementation, the EMD is computed separately for each latent channel at each vertical level and then averaged across all dimensions. This alignment mechanism forces the encoder to extract features that are common across both climates, effectively learning a representation space with reduced distribution shift.

The model is trained end-to-end using a combined loss that includes reconstruction loss for the autoencoder, EMD loss to encourage latent alignment between climates, and prediction (pred) loss between predicted and true subgrid fluxes and tendencies in the control climate. The total training loss is defined as:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{pred}} - \lambda_{\text{EMD}}) \cdot \mathcal{L}_{\text{reconstruction}} + \lambda_{\text{pred}} \cdot \mathcal{L}_{\text{pred}} + \lambda_{\text{EMD}} \cdot \text{EMD}(Z_0, Z_{+4K}), \quad (1)$$

where:

- $\mathcal{L}_{\text{reconstruction}}$ is the mean squared error between inputs and their reconstructions for both control and +4 K climates,
- $\text{EMD}(Z_0, Z_{+4K})$ is the Earth Mover's Distance between latent spaces from the two climate states,
- $\mathcal{L}_{\text{pred}}$ is the supervised loss (mean squared error) between predicted and ground-truth subgrid fluxes and tendencies, using control-climate labels, and
- $\lambda_{\text{EMD}}, \lambda_{\text{pred}}$ are tunable weights balancing the three objectives.

Hyperparameters controlling latent alignment and prediction were selected to balance robustness and predictive accuracy. In particular, the weight on the latent alignment loss, λ_{EMD} , was chosen as the largest value that avoided latent space collapse, defined as a collapse of latent variability (average standard deviation of latent variables < 0.1). The prediction loss weight, λ_{pred} , was then selected to maintain predictive skill while preserving stable latent alignment. Further training details and hyperparameters are provided in Section S1 in Supporting Information S1.

We compare CERA against four alternative models: (a) a baseline MLP trained on raw input profiles from the control climate (Baseline) (Yuval et al., 2021); (b) a physically informed MLP trained on hand-crafted, climate-invariant features of RH and plume buoyancy, following Beucier et al. (2024) (RH + B baseline); (c) an ablation variant of CERA trained without the EMD alignment loss (CERA-noAlign); and (d) CERA-noAlign trained on control data only. To assess robustness, all models were trained using five independent random seeds (Haynes et al., 2023).

3. Offline Results

3.1. Autoencoding and Latent Alignment Enable Representations With Improved Generalization

Figure 2 shows the predictive performance of Baseline, RH + B Baseline, CERA, and CERA-noAlign for each of the five subgrid outputs. Averaging R^2 over all output variables, CERA achieves the highest overall accuracy, with a mean R^2 of 0.75 in the control climate and 0.53 in the +4 K climate. This indicates both strong in-distribution performance and substantial generalization capacity under warming.

In comparison, the Baseline model performs similarly to the RH + B baseline in the control climate (both have $R^2 = 0.72$) and slightly worse than CERA ($R^2 = 0.75$). That CERA outperforms the Baseline model in the control climate suggests that restricting the inputs to features with reduced distribution shift across climate regimes actually helps generalization beyond the training data even for test data in the control climate. The performance of the Baseline model deteriorates sharply under warming, with the mean R^2 falling to 0.26, highlighting a lack of ability to extrapolate beyond its training distribution as might be expected given distribution shifts in the input variables. The RH + B baseline maintains its performance to a much greater extent in the warm climate ($R^2 = 0.46$), outperforming the Baseline model but performing slightly worse than CERA ($R^2 = 0.53$).




Figure 3. Same as Figure 2, but making a direct comparison to Beucler et al. (2024) by retraining to include radiative heating in the parameterization and assessing skill for the tendencies of non-precipitating liquid/ice static energy H_L and total non-precipitating water q_T .

These results underscore the effectiveness of the proposed latent-alignment strategy in enabling generalization across climates. Ablation experiments further emphasize its importance. The autoencoder-only variant (CERA-noAlign) improves upon the raw Baseline across all random seeds, indicating that learning a reversible, shared latent representation across both climates helps the model extract more informative features for downstream prediction, even without the alignment. The addition of the EMD loss further improves the performance. As for the model using CERA-noAlign trained on control data only, under the +4 K climate its R^2 is overall much lower than that of CERA-noAlign trained on inputs from both climates (Figure S1 in Supporting Information S1). This indicates that including inputs from both climate states during autoencoder training is important for achieving robust performance.

In addition to improving mean accuracy, CERA also enhances robustness across random seeds relative to the raw Baseline. The variance in R^2 among trained models is notably reduced for CERA compared to the Baseline and CERA-noAlign in the +4 K climate. However, the RH + B baseline demonstrates the highest robustness overall, in particular for the sedimentation output. These results suggest that CERA may converge to different solutions across random seeds, leading to variability in performance. In contrast, the RH + B baseline represents a deterministic, hand-crafted solution that yields more consistent results across trained models. There may exist alternative hand-crafted input formulations capable of achieving higher accuracy. Identifying such configurations, however, would require extensive manual tuning and domain expertise, suggesting that self-supervised learning approaches like CERA may offer a more scalable alternative.

To further assess performance, we make a direct comparison to Beucler et al. (2024) by retraining to include radiative heating and assessing skill for the total subgrid liquid/ice static energy (H_L) and moistening (q_T) tendencies. As shown in Figure 3, CERA achieves the highest R^2 scores in the control climate, outperforming both the raw-input Baseline and the RH + B baseline. In the +4 K climate, CERA and RH + B produce comparable performance, both substantially outperforming the Baseline model. These results show that CERA is robust to changes in the outputs, and that again CERA not only generalizes effectively to a warmer climate, but it maintains high predictive skill in the training climate.

Results from further sensitivity experiments are described in Section S2 in Supporting Information S1 including increasing kernel size to include vertical non-locality in the autoencoder, aligning all channels in the latent space, varying hyperparameters in the training loss, and training CERA on outputs from both the control and +4 K climates.

3.2. Evaluation of Performance for Precipitation

Given the importance of precipitation and associated latent heating for climate-change and its impacts, we next evaluate how well the ML models capture precipitation-related processes. In particular, we analyze the vertical profile of prediction skill for q_{T_mic} , the tendency of q_T due to microphysical conversion to precipitation, and the




Figure 4. Offline R^2 for $q_{T,mic}$ versus latitude and pressure across methods and climates. Shown are the coefficient of determination (R^2) values for predicted $q_{T,mic}$ (moisture tendency due to microphysical conversion to precipitating water) in the control (top) and +4 K (bottom) climates, averaged over five random seeds. Each row compares three models: Baseline, RH + B baseline, and CERA. Regions with variance less than 1% of the mean variance across levels and latitudes were masked to remove near-constant areas. In the control climate, all models exhibit moderate-to-high skill. Under warming, performance differences become more pronounced. The Baseline model degrades sharply in the tropics, the loss in skill is much less in RH + B, while CERA retains high skill aloft though with some loss of skill in the boundary layer.

resulting offline surface precipitation distributions. Details for computing the instantaneous surface precipitation rate can be found in Section S3 in Supporting Information S1. All results were averaged across five random seeds per model.

Figure 4 shows R^2 for $q_{T,mic}$ versus pressure and latitude in the control and +4 K climates. In the control climate, all three models have a similar skill structure with moderate-to-high skill in the mid-to-upper troposphere. In the +4 K climate, performance disparities between methods becomes pronounced. The Baseline model suffers a substantial drop in predictive skill, with negative R^2 values in the tropics where it fails to extrapolate beyond its training climate. The RH + B baseline does not degrade to the same extent, but its skill is still substantially lower than in the control climate. In contrast, CERA maintains strong performance in the mid and upper troposphere but shows reduced skill in the boundary layer, particularly in the tropics. The lower R^2 values in the boundary layer may arise from our choice of model architecture, as there is no communication among vertical levels in the autoencoder, or they may reflect a trade-off that favors overall performance across vertical levels.

Figure 5 shows the normalized frequency distributions of instantaneous precipitation rates under control and +4 K warming scenarios. Precipitation rates are binned using 100 logarithmically spaced intervals starting at 1 mm day^{-1} . Frequencies are normalized by the number of samples exceeding this threshold, such that the sum across all bins equals one, yielding a unitless vertical axis. These distributions provide a sensitive diagnostic of each model's ability to reproduce the full range of precipitation intensities, from drizzle to extremes.

In the control climate (Figure 5a), all models exhibit similar performance and closely follow the reference distribution from the high-resolution simulation across most of the range. However, two systematic discrepancies are apparent. First, all models overestimate the frequency of light precipitation events in the $1\text{--}2 \text{ mm day}^{-1}$ range. Second, all models slightly underestimate the frequency of extreme events above $\sim 100 \text{ mm day}^{-1}$. Despite these deviations, the overall agreement with the reference remains strong.

Under +4 K warming (Figure 5b), more pronounced differences emerge. The Baseline model exhibits substantial overestimation across a wide range of intensities. Its bias for weak intensities becomes more pronounced, with an overestimation of precipitation below 4 mm day^{-1} , and it substantially overestimates the frequency of extreme




Figure 5. Offline frequency distributions of instantaneous precipitation rates in the control and +4 K climates. Panel (a) shows distributions in the control climate, while panel (b) shows distributions under +4 K warming. Solid lines show model predictions averaged across five random seeds: purple for the Baseline, orange for RH + B, and green for CERA. The high-resolution reference simulation is shown in dashed black. Precipitation rates are derived from offline-integrated q_{T_mic} and binned using 100 logarithmically spaced intervals starting at 1 mm day^{-1} . All models are compared against the high-resolution simulation reference (ground truth). In the control climate, models generally reproduce the overall distribution shape but tend to overestimate drizzle ($1\text{--}2 \text{ mm day}^{-1}$) and slightly underestimate extremes. Under warming, model performances diverge more substantially. The baseline model overestimates across a broad range, RH + B underestimates throughout, and CERA closely tracks the reference in the tail but underpredicts the $40\text{--}300 \text{ mm day}^{-1}$ range.

events above 200 mm day^{-1} . In contrast, the RH + B baseline systematically underestimates the frequency of precipitation across almost all intensities. CERA more accurately reproduces the high-resolution reference, particularly at the extreme tail (above 300 mm day^{-1}), where it tracks the enhanced frequency of intense events associated with warming. However, it tends to underestimate the frequency of precipitation in the range of $40\text{--}300 \text{ mm day}^{-1}$. The performance of CERA-noAlign for precipitation was found to be comparable to that of CERA. To maintain visual clarity, we have omitted these results from the figures.

Further insight into model performance is provided by considering precipitation versus latitude as shown in Figure 6. In the control climate (Figures 6a and 6c), all models reproduce the general meridional structures of the mean and extreme (99.9th percentile) precipitation. There are minor deviations in the extremes, with all models slightly underestimating the peak intensity.

In the +4 K climate (Figures 6b and 6d), model differences become pronounced in the tropics. CERA successfully captures both the intensification and the meridional structure of tropical precipitation, even if it exhibits a slight underestimation of both mean and extreme precipitation near the peak. In contrast, the Baseline model substantially overestimates the magnitude of tropical precipitation, particularly for extreme precipitation, while the RH + B baseline systematically underestimates both mean and extreme values across the tropics. The poor performance of RH + B in the tropics arises because it makes large vertically coherent mean errors in q_{T_mic} which then leads to large biases in surface precipitation, whereas the mean errors for CERA are smaller and less coherent vertically (Figure S3 in Supporting Information S1).

Lastly, we evaluate the Root Mean Square Error for the instantaneous precipitation across the evaluated models as compared to the high-resolution reference (Table 1). For the control climate, the CERA model exhibits the lowest error (28 mm day^{-1}), whereas the Baseline model shows the greatest deviation from the high-resolution reference (33 mm day^{-1}). RMSE increases substantially in the +4 K climate, and the discrepancy between models widens substantially; the CERA model maintains the lowest error (67 mm day^{-1}), while the Baseline model exhibits the highest, reaching 156 mm day^{-1} .




Figure 6. Offline mean and extreme precipitation versus latitude in control and +4 K climates. Panels (a and b) show the mean precipitation for the control and +4 K climates, respectively, while panels (c and d) display the corresponding 99.9th percentile (extreme) precipitation. All values are derived from offline-integrated q_{T_mic} . Solid lines represent model predictions averaged over five random seeds (purple for baseline, orange for RH + B, green for CERA), and the dashed black line shows the high-resolution reference simulation. Extreme precipitation curves have been lightly smoothed using a 1-2-1 filter for readability.

Together, these results demonstrate that CERA generalizes robustly across climates, effectively capturing the vertical and zonal structure of moist processes, although its performance degrades in the tropical boundary layer. The Baseline model shows a sharp decline in skill for +4 K, particularly in the deep tropics, while the RH + B baseline has much less of a decline in skill for q_{T_mic} but consistently underestimates precipitation throughout the tropics. In Section S4 in Supporting Information S1, we describe additional results in which CERA is applied to improve generalization from the control climate to a hemispherically asymmetric climate, demonstrating that CERA can be useful for climate changes beyond a uniform SST warming.

4. Online Results: Simulations With a CERA Trained Parameterization

Although we primarily focus on offline results which are straightforward to interpret, we also conducted a series of online experiments. We compare the performance of the CERA trained parameterization against the Baseline MLP parameterization when the parameterizations are coupled to SAM in coarse-resolution simulations. The CERA parameterization leads to stable simulations at coarse resolution that generalize better than baseline to a warmer climate.

SAM was configured at a coarse resolution of 96 km, utilizing the Fortran-implemented parameterization that includes trained CERA or baseline MLP with weights loaded via a NetCDF file. We maintained the same SST distributions, timestep, and vertical grid as the high-resolution simulations. Implementation details for the parameterization follow Yuval et al. (2021). For simplicity, and unlike Yuval et al. (2021), we did not include a second neural network for predicting eddy diffusivity and surface flux corrections. Additionally, we do not predict radiative tendencies which are instead handled by the coarse SAM model. Simulations were run for 500 days, with the first 100 days discarded as spin-up. The remaining 400 days were used for analysis. The high-resolution simulations were treated in the same way for consistency.

For each of CERA and Baseline MLP, we performed three runs using parameterizations trained with (three) different random seeds. Due to the higher computational cost of online simulations, our evaluation does not include as many configurations and random seeds as for the offline results, but it is

Table 1
Offline Root Mean Square Error of Instantaneous Precipitation (mm day^{-1})
Across Different Models for Control and +4 K Climates

Climate	CERA	RH + B baseline	Baseline
Control	28	31	33
+4 K	67	107	156




Figure 7. Mean precipitation versus latitude for coarse-resolution simulations with the ML subgrid parameterizations. Panels (a and b) show the zonal- and time-mean precipitation for the control and +4 K climates, respectively. Solid lines represent simulations for three random seeds used in training (purple for Baseline and green for CERA), and the dashed black line shows the high-resolution reference simulation.

sufficient to get a sense of the potential of CERA. All simulations were found to run stably, and the resulting zonal- and time-mean precipitation distributions versus latitude are shown in Figure 7. We focus on mean precipitation rather than extreme precipitation to reduce sampling error. For both CERA and Baseline MLP, one out of the three seeds resulted in a simulation that transitioned to a single ITCZ in the +4 K climate, whereas the other seeds and the high-resolution simulation remained in a double ITCZ configuration. We have previously found that with the SST distribution used, the climate is close to a dynamical transition to a single ITCZ, and evidently there are sufficient errors in generalizing to a warm climate for a transition to a single ITCZ for some seeds. Aside from the seeds that transition to a single ITCZ, CERA and the baseline MLP both do a reasonably good job in capturing the mean precipitation distribution in both the control and +4 K climates.

Interestingly, for the baseline MLP runs that didn't transition to a single ITCZ in the +4 K climate, the mean precipitation error is much smaller than in the offline tests (compare Figures 6b and 7b). This occurs because, although the baseline MLP generates biased moisture tendencies that are too large, the additional latent heating overstabilizes the tropical troposphere which tends to suppress precipitation (Figure 8). The peak in mean warming in the tropical upper troposphere from control to +4 K climate is 11 K for the CERA simulations which is similar to 13 K for the high-resolution simulations, and much smaller than 20 K for the Baseline MLP simulations. In other words, there is error compensation in which biased static stability offsets the precipitation overestimation observed offline. Such a compensation of errors is unsurprising given that mean precipitation is energetically constrained for a given SST distribution and CO₂ concentration (O'Gorman et al., 2012), and the ML parameterization we are using is energy conserving by design. These findings underscore the need for caution when evaluating parameterizations, including ML-based parameterizations, as models can produce plausible results for specific metrics through internal error compensation.

Overall, while the online evaluation is more complex to interpret because of ITCZ transitions for some random seeds, we nonetheless find that the CERA parameterization leads to stable coarse-resolution simulations that improve on the performance of the MLP baseline for the +4 K climate. The clear improvement in performance shows up in smaller temperature biases in the tropical troposphere rather than in precipitation because mean precipitation is energetically constrained.

5. Discussion

CERA demonstrates that combining autoencoding with latent space alignment offers a powerful and flexible approach for improved generalization across climates of subgrid processes. Unlike models that rely on hand-crafted input features, CERA learns directly from raw inputs, eliminating the need for manual feature engineering. This approach reduces development effort and avoids restrictive assumptions that could limit model skill or generalizability. Notably, CERA does not degrade performance in the control climate and in some cases even




Figure 8. Mean temperature versus latitude and pressure for control climate, +4 K climate, and the climate change response (+4 K minus control) for coarse-resolution simulations with the ML subgrid parameterizations. Panels (a–c) display results from the baseline multilayer perceptron (MLP) coarse-resolution simulation, while panels (d–f) show the CERA coarse-resolution simulation. Panels (g–i) provide the high-resolution simulation as a reference. For both the baseline MLP and CERA, we present the averaged results from the two seeds that did not transition to a single ITCZ in the +4 K simulations.

improves it, suggesting that latent alignment can enhance generalization without sacrificing in-distribution accuracy. The CERA trained parameterization was successfully implemented in the SAM model and led to stable simulations at coarse-resolution. Although the online evaluation was complicated by ITCZ transitions for some random seeds, CERA outperformed the baseline MLP by having smaller temperature biases in the +4 K climate.

While this study focuses on moist-physics parameterization, the framework is broadly applicable and could be extended to other parameterized processes, or even to applications beyond parameterization. For example, CERA may benefit statistical downscaling tasks that demand robust extrapolation across climate regimes.

Nevertheless, the framework has limitations. While the latent alignment strategy enhances robustness, not all physical processes can be represented in a climate-invariant-like manner. This imposes inherent limits on the generalizability of this approach. Future work could address these challenges by incorporating additional physical constraints or adapting the framework to account for climate-dependent processes.

Additionally, the autoencoder in our framework employs one-dimensional convolutional layers with a kernel size of one, resulting in vertically local transformations that act independently at each level. While this design preserves interpretability, it precludes modeling vertical interactions. Future work could explore vertically non-local architectures to capture cross-level dependencies. Moreover, further analysis of the learned latent representations, such as applying symbolic regression, may also help uncover interpretable, low-dimensional expressions of key physical relationships. Other domain adaptation methods have been developed in the field of machine learning (Ganin et al., 2016; B. Sun et al., 2017), and future work could investigate their effectiveness regarding the problem of generalization to different climates.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Availability Statement

The data and code used for the figures in this study are available at Zenodo (Liu, 2026).

Acknowledgments

Shuchang Liu was supported by SNSF Postdoc.Mobility Grant P500PN_222180 and the MIT Climate Grand Challenge on Weather and Climate Extremes. Support was also provided by Schmidt Sciences. We thank BingXin Ke for fruitful discussions. We thank Bill Boos and Alexey Fedorov for providing the output from the hemispherically symmetric high-resolution simulations. We also thank Janni Yuval, Tom Beucler and the M²LInES team. We thank the reviewers for their constructive comments and suggestions, which helped improve the manuscript.

References

- Beucler, T., Gentine, P., Yuval, J., Gupta, A., Peng, L., Lin, J., et al. (2024). Climate-invariant machine learning. *Science Advances*, 10(6), eadj7250. <https://doi.org/10.1126/sciadv.adj7250>
- Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2022). Pangu-weather: A 3D high-resolution model for fast and accurate global weather forecast. *arXiv preprint arXiv:2211.02556*.
- Bodnar, C., Bruinsma, W. P., Lucic, A., Stanley, M., Allen, A., Brandstetter, J., et al. (2025). A foundation model for the Earth system. *Nature*, 641(8065), 1180–1187. <https://doi.org/10.1038/s41586-025-09005-y>
- Bora, A., Shukla, K., Zhang, S., Harrop, B., Leung, R., & Karniadakis, G. E. (2023). Learning bias corrections for climate models using deep neural operators. *arXiv preprint arXiv:2302.03173*.
- Bracco, A., Brajard, J., Dijkstra, H. A., Hassanzadeh, P., Lessig, C., & Monteleoni, C. (2025). Machine learning for the physics of climate. *Nature Reviews Physics*, 7(1), 6–20. <https://doi.org/10.1038/s42254-024-00776-3>
- Brenowitz, N. D., Beucler, T., Pritchard, M., & Bretherton, C. S. (2020). Interpreting and stabilizing machine-learning parametrizations of convection. *Journal of the Atmospheric Sciences*, 77(12), 4357–4375. <https://doi.org/10.1175/JAS-D-20-0082.1>
- Brenowitz, N. D., & Bretherton, C. S. (2018). Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12), 6289–6298. <https://doi.org/10.1029/2018gl078510>
- Bretherton, C. S., Henn, B., Kwa, A., Brenowitz, N. D., Watt-Meyer, O., McGibbon, J., et al. (2022). Correcting coarse-grid weather and climate models by machine learning from global storm-resolving simulations. *Journal of Advances in Modeling Earth Systems*, 14(2), e2021MS002794. <https://doi.org/10.1029/2021ms002794>
- Cachay, S. R., Henn, B., Watt-Meyer, O., Bretherton, C. S., & Yu, R. (2024). Probabilistic Emulation of a Global Climate Model with Spherical DYffusion. *arXiv preprint arXiv:2406.14798*.
- Chantry, M., Christensen, H., Dueben, P., & Palmer, T. (2021). Opportunities and challenges for machine learning in weather and climate modelling: Hard, medium and soft AI. *Philosophical Transactions of the Royal Society A*, 379(2194), 20200083. <https://doi.org/10.1098/rsta.2020.0083>
- Chapman, W. E., & Berner, J. (2025). Improving climate bias and variability via CNN-based state-dependent model-error corrections. *Geophysical Research Letters*, 52(6), e2024GL114106. <https://doi.org/10.1029/2024GL114106>
- Chapman, W. E., Schreck, J. S., Sha, Y., Gagne, D. J., II, Kimpara, D., Zanna, L., et al. (2025). CAMulator: Fast emulation of the community atmosphere model. *arXiv preprint arXiv:2504.06007*.
- Clark, S. K., Brenowitz, N. D., Henn, B., Kwa, A., McGibbon, J., Perkins, W. A., et al. (2022). Correcting a 200 km resolution climate model in multiple climates by machine learning from 25 km resolution simulations. *Journal of Advances in Modeling Earth Systems*, 14(9), e2022MS003219. <https://doi.org/10.1029/2022MS003219>
- Eyring, V., Collins, W. D., Gentine, P., Barnes, E. A., Barreiro, M., Beucler, T., et al. (2024). Pushing the frontiers in climate modelling and analysis with machine learning. *Nature Climate Change*, 14(9), 916–928. <https://doi.org/10.1038/s41558-024-02095-y>
- Flora, M. L., & Potvin, C. (2025). WoFSCast: A machine learning model for predicting thunderstorms at watch-to-warning scales. *Geophysical Research Letters*, 52(10), e2024GL112383. <https://doi.org/10.1029/2024GL112383>
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., et al. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59), 1–35.
- Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., & Yacalis, G. (2018). Could machine learning break the convection parameterization deadlock? *Geophysical Research Letters*, 45(11), 5742–5751. <https://doi.org/10.1029/2018gl078202>
- Gregory, W., Bushuk, M., Zhang, Y.-F., Adcroft, A., Zanna, L., McHugh, C., & Jia, L. (2025). Advancing global sea ice prediction capabilities using a fully-coupled climate model with integrated machine learning. *arXiv preprint arXiv:2505.18328*.
- Han, Y., Zhang, G. J., & Wang, Y. (2023). An ensemble of neural networks for moist physics processes, its generalizability and stable integration. *Journal of Advances in Modeling Earth Systems*, 15(10), e2022MS003508. <https://doi.org/10.1029/2022MS003508>
- Haynes, K., Lagerquist, R., McGraw, M., Musgrave, K., & Ebert-Uphoff, I. (2023). Creating and evaluating uncertainty estimates with neural networks for environmental-science applications. *Artificial Intelligence for the Earth Systems*, 2(2), 220061. <https://doi.org/10.1175/ai-es-d-22-0061.1>
- Hobeichi, S., Nishant, N., Shao, Y., Abramowitz, G., Pitman, A., Sherwood, S., et al. (2023). Using machine learning to cut the cost of dynamical downscaling. *Earth's Future*, 11(3), e2022EF003291. <https://doi.org/10.1029/2022EF003291>
- Kashinath, K., Mustafa, M., Albert, A., Wu, J. L., Jiang, C., Esmailzadeh, S., et al. (2021). Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A*, 379(2194), 20200093. <https://doi.org/10.1098/rsta.2020.0093>
- Khairoutdinov, M. F., & Randall, D. A. (2003). Cloud resolving modeling of the ARM summer 1997 IOP: Model formulation, results, uncertainties, and sensitivities. *Journal of the Atmospheric Sciences*, 60(4), 607–625. [https://doi.org/10.1175/1520-0469\(2003\)060<0607:crmota>2.0.co;2](https://doi.org/10.1175/1520-0469(2003)060<0607:crmota>2.0.co;2)
- Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., et al. (2024). Neural general circulation models for weather and climate. *Nature*, 632(8027), 1060–1066. <https://doi.org/10.1038/s41586-024-07744-y>
- Kurth, T., Subramanian, S., Harrington, P., Pathak, J., Mardani, M., Hall, D., et al. (2023). Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive Fourier neural operators. In *Proceedings of the platform for advanced scientific computing conference* (pp. 1–11).
- Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirsberger, P., Fortunato, M., Alet, F., et al. (2023). Learning skillful medium-range global weather forecasting. *Science*, 382(6677), 1416–1421. <https://doi.org/10.1126/science.adi2336>
- Levina, E., & Bickel, P. (2001). The Earth Mover's distance is the Mallows distance: Some insights from statistics. In *Proceedings of the 8th IEEE International Conference on Computer Vision (ICCV 2001)* (Vol. 2, pp. 251–256). <https://doi.org/10.1109/ICCV.2001.937632>
- Liu, S. (2026). Data and code for 'CERA: A Framework for Improved Generalization of Machine Learning Models to Changed Climates' [Dataset]. *Zenodo*. <https://doi.org/10.5281/zenodo.21216188>
- McCandless, T., Gagne, D. J., Kosović, B., Haupt, S. E., Yang, B., Becker, C., & Schreck, J. (2022). Machine learning for improving surface-layer-flux estimates. *Boundary-Layer Meteorology*, 185(2), 199–228. <https://doi.org/10.1007/s10546-022-00727-4>

- McNally, A., Lessig, C., Lean, P., Boucher, E., Alexe, M., Pinnington, E., et al. (2024). Data driven weather forecasts trained and initialised directly from observations. *arXiv preprint arXiv:2407.15586*.
- Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K., & Grover, A. (2023). Climax: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*.
- O'Gorman, P. A., Allan, R. P., Byrne, M. P., & Previdi, M. (2012). Energetic constraints on precipitation under climate change. *Surveys in Geophysics*, 33(3–4), 585–608. <https://doi.org/10.1007/s10712-011-9159-6>
- O'Gorman, P. A., & Dwyer, J. G. (2018). Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events. *Journal of Advances in Modeling Earth Systems*, 10(10), 2548–2563. <https://doi.org/10.1029/2018ms001351>
- O'Gorman, P. A., Li, Z., Boos, W. R., & Yuval, J. (2021). Response of extreme precipitation to uniform surface warming in quasi-global aquaplanet simulations at high resolution. *Philosophical Transactions of the Royal Society A: Mathematical, Physical & Engineering Sciences*, 379(2195), 20190543. <https://doi.org/10.1098/rsta.2019.0543>
- Perezhogin, P., Adcroft, A., & Zanna, L. (2025). Generalizable neural-network parameterization of mesoscale eddies in idealized and global ocean models. *arXiv preprint arXiv:2505.08900*.
- Price, L., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., et al. (2025). Probabilistic weather forecasting with machine learning. *Nature*, 637(8044), 84–90. <https://doi.org/10.1038/s41586-024-08252-9>
- Rampal, N., Gibson, P. B., Sood, A., Stuart, S., Fauchereau, N. C., Brandolino, C., et al. (2022). High-resolution downscaling with interpretable deep learning: Rainfall extremes over New Zealand. *Weather and Climate Extremes*, 38, 100525. <https://doi.org/10.1016/j.wace.2022.100525>
- Rasp, S., Pritchard, M. S., & Gentile, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689. <https://doi.org/10.1073/pnas.1810286115>
- Rubner, Y., Tomasi, C., & Guibas, L. J. (2000). The Earth Mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2), 99–121. <https://doi.org/10.1023/A:1026543900054>
- Schär, C., Frei, C., Lüthi, D., & Davies, H. C. (1996). Surrogate climate-change scenarios for regional climate models. *Geophysical Research Letters*, 23(6), 669–672. <https://doi.org/10.1029/96GL00265>
- Scher, S., & Messori, G. (2019). Generalization properties of feed-forward neural networks trained on Lorenz systems. *Nonlinear Processes in Geophysics*, 26(4), 381–399. <https://doi.org/10.5194/npg-26-381-2019>
- Schneider, R., Bonavita, M., Geer, A., Arcucci, R., Dueben, P., Vitolo, C., et al. (2022). ESA-ECMWF Report on recent progress and research directions in machine learning for Earth System observation and prediction. *npj Climate and Atmospheric Science*, 5(1), 51. <https://doi.org/10.1038/s41612-022-00269-z>
- Sun, B., Feng, J., & Saenko, K. (2017). Correlation alignment for unsupervised domain adaptation. In *Domain adaptation in computer vision applications* (pp. 153–171). Springer.
- Sun, Y. Q., Pahlavan, H. A., Chattopadhyay, A., Hassanzadeh, P., Lubis, S. W., Alexander, M. J., et al. (2024). Data imbalance, uncertainty quantification, and transfer learning in data-driven parameterizations: Lessons from the emulation of gravity wave momentum transport in WACCM. *Journal of Advances in Modeling Earth Systems*, 16(7), e2023MS004145. <https://doi.org/10.1029/2023ms004145>
- Watt-Meyer, O., Brenowitz, N. D., Clark, S. K., Henn, B., Kwa, A., McGibbon, J., et al. (2021). Correcting weather and climate models by machine learning nudged historical simulations. *Geophysical Research Letters*, 48(15), e2021GL092555. <https://doi.org/10.1029/2021gl092555>
- Watt-Meyer, O., Henn, B., McGibbon, J., Clark, S. K., Kwa, A., Perkins, W. A., et al. (2025). ACE2: Accurately learning subseasonal to decadal atmospheric variability and forced responses. *npj Climate and Atmospheric Science*, 8(1), 205. <https://doi.org/10.1038/s41612-025-01090-0>
- Weyn, J. A., Durran, D. R., Caruana, R., & Cresswell-Clay, N. (2021). Sub-seasonal forecasting with a large ensemble of deep-learning weather prediction models. *Journal of Advances in Modeling Earth Systems*, 13(7), e2021MS002502. <https://doi.org/10.1029/2021ms002502>
- Yuval, J., & O'Gorman, P. A. (2020). Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions. *Nature Communications*, 11(1), 3295. <https://doi.org/10.1038/s41467-020-17142-3>
- Yuval, J., O'Gorman, P. A., & Hill, C. N. (2021). Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision. *Geophysical Research Letters*, 48(6), e2020GL091363. <https://doi.org/10.1029/2020gl091363>
- Zeiler, M. D., Krishnan, D., Taylor, G. W., & Fergus, R. (2010). Deconvolutional networks. In *2010 IEEE computer Society conference on computer vision and pattern recognition* (pp. 2528–2535). IEEE.
- Zhao, W. L., Gentile, P., Reichstein, M., Zhang, Y., Zhou, S., Wen, Y., et al. (2019). Physics-constrained machine learning of evapotranspiration. *Geophysical Research Letters*, 46(24), 14496–14507. <https://doi.org/10.1029/2019gl085291>