# Deep learning to represent subgrid processes in climate models

Stephan Rasp[a,b,1], Michael S. Pritchard[b], and Pierre Gentine[c,d]

[a]Meteorological Institute, Ludwig-Maximilian-University, 80333 Munich, Germany; [b]Department of Earth System Science, University of California, Irvine, CA 92697; [c]Department of Earth and Environmental Engineering, Earth Institute, Columbia University, New York, NY 10027; and [d]Data Science Institute, Columbia University, New York, NY 10027

The representation of nonlinear subgrid processes, especially clouds, has been a major source of uncertainty in climate models for decades. Cloud-resolving models better represent many of these processes and can now be run globally but only for short-term simulations of at most a few years because of computational limitations. Here we demonstrate that deep learning can be used to capture many advantages of cloud-resolving modeling at a fraction of the computational cost. We train a deep neural network to represent all atmospheric subgrid processes in a climate model by learning from a multiscale model in which convection is treated explicitly. The trained neural network then replaces the traditional subgrid parameterizations in a global general circulation model in which it freely interacts with the resolved dynamics and the surface-flux scheme. The prognostic multiyear simulations are stable and closely reproduce not only the mean climate of the cloud-resolving simulation but also key aspects of variability, including precipitation extremes and the equatorial wave spectrum. Furthermore, the neural network approximately conserves energy despite not being explicitly instructed to. Finally, we show that the neural network parameterization generalizes to new surface forcing patterns but struggles to cope with temperatures far outside its training manifold. Our results show the feasibility of using deep learning for climate model parameterization. In a broader context, we anticipate that data-driven Earth system model development could play a key role in reducing climate prediction uncertainty in the coming decade.

climate modeling | deep learning | subgrid parameterization | convection

**M**any of the atmosphere's most important processes occur on scales smaller than the grid resolution of current climate models, around 50–100 km horizontally. Clouds, for example, can be as small as a few hundred meters; yet they play a crucial role in determining the Earth's climate by transporting heat and moisture, reflecting and absorbing radiation, and producing rain. Climate change simulations at such fine resolutions are still many decades away (1). To represent the effects of such subgrid processes on the resolved scales, physical approximations—called parameterizations—have been heuristically developed and tuned to observations over the last decades (2). However, owing to the sheer complexity of the underlying physical system, significant inaccuracies persist in the parameterization of clouds and their interaction with other processes, such as boundary-layer turbulence and radiation (1, 3, 4). These inaccuracies manifest themselves in stubborn model biases (5–7) and large uncertainties about how much the Earth will warm as a response to increased greenhouse gas concentrations (1, 8, 9). To improve climate predictions, therefore, novel, objective, and computationally efficient approaches to subgrid parameterization development are urgently needed.

Cloud-resolving models (CRMs) alleviate many of the issues related to parameterized convection. At horizontal resolutions of at least 4 km deep convection can be explicitly treated (10), which substantially improves the representation of land–atmosphere coupling (11, 12), convective organization (13), and weather

extremes. Further increasing the resolution to a few hundred meters allows for the direct representation of the most important boundary-layer eddies, which form shallow cumuli and stratocumuli. These low clouds are crucial for the Earth's energy balance and the cloud–radiation feedback (14). CRMs come with their own set of tuning and parameterization decisions but the advantages over coarser models are substantial. Unfortunately, global CRMs will be too computationally expensive for climate change simulations for many decades (1). Short-range simulations covering periods of months or even a few years, however, are beginning to be feasible and are in development at modeling centers around the world (15–18).

In this study, we explore whether deep learning can provide an objective, data-driven approach to using high-resolution modeling data for climate model parameterization. The paradigm shift from heuristic reasoning to machine learning has transformed computer vision and natural language processing over the last few years (19) and is starting to impact more traditional fields of science. The basic building blocks of deep learning are deep neural networks which consist of several interconnected layers of nonlinear nodes (20). They are capable of approximating arbitrary nonlinear functions (21) and can easily be adapted to novel problems. Furthermore, they can handle large datasets during training and provide fast predictions at inference time. All of

## Significance

Current climate models are too coarse to resolve many of the atmosphere's most important processes. Traditionally, these subgrid processes are heuristically approximated in so-called parameterizations. However, imperfections in these parameterizations, especially for clouds, have impeded progress toward more accurate climate predictions for decades. Cloud-resolving models alleviate many of the gravest issues of their coarse counterparts but will remain too computationally demanding for climate change predictions for the foreseeable future. Here we use deep learning to leverage the power of short-term cloud-resolving simulations for climate modeling. Our data-driven model is fast and accurate, thereby showing the potential of machine-learning–based approaches to climate model development.

these traits make deep learning an attractive approach for the problem of subgrid parameterization.

Extending on previous offline or single-column neural network cumulus parameterization studies (22–24), here we take the essential step of implementing the trained neural network in a global climate model and running a stable, prognostic multiyear simulation. To show the potential of this approach we compare key climate statistics between the deep learning-powered model and its training simulation. Furthermore, we tackle two crucial questions for a climate model implementation: First, does the neural network parameterization conserve energy? And second, to what degree can the network generalize outside of its training climate? We conclude by highlighting crucial challenges for future data-driven parameterization development.

## Climate Model and Neural Network Setup

Our base model is the superparameterized Community Atmosphere Model v3.0 (SPCAM) (25) in an aquaplanet setup (see *SI Appendix* for details). The sea surface temperatures (SSTs) are fixed and zonally invariant with a realistic equator-to-pole gradient (26). The model has a full diurnal cycle but no seasonal variation. The horizontal grid spacing of the global circulation model (GCM) is approximately 2° with 30 vertical levels. The GCM time step is 30 min. In superparameterization, a 2D CRM is embedded in each GCM grid column (27). This CRM explicitly resolves deep convective clouds and includes parameterizations for small-scale turbulence and cloud microphysics. In our setup, we use 84-km–wide columns with a CRM time step of 20 s, as in ref. 28. For comparison, we also run a control simulation with the traditional parameterization suite (CTRLCAM) that is based on an undilute plume parameterization of moist convection. CTRLCAM exhibits many typical problems associated with traditional subgrid cloud parameterizations: a double intertropical convergence zone (ITCZ) (5), too much drizzle and missing precipitation extremes, and an unrealistic equatorial wave spectrum with a missing Madden–Julian oscillation (MJO). In contrast, SPCAM captures the key benefits of full 3D CRMs in improving the realism all of these issues with respect to observations (29–31). In this context, a key test for a neural network parameterization is whether it learns sufficiently from the explicitly resolved convection in SPCAM to remedy such problems while being computationally more affordable.

Analogous to a traditional parameterization, the task of the neural network is to predict the subgrid tendencies as a function of the atmospheric state at every time step and grid column (*SI Appendix*, Table S1). Specifically, we selected the following input variables: the temperature $T(z)$, specific humidity $Q(z)$ and wind profiles $V(z)$, surface pressure $P_s$, incoming solar radiation $S_{in}$, and the sensible $H$ and latent heat fluxes $E$. These variables mirror the information received by the CRM and radiation scheme with a few omissions (*SI Appendix*). The output variables are the sum of the CRM and radiative heating rates $\Delta T_{phy}$, the CRM moistening rate $\Delta Q_{phy}$, the net radiative fluxes at the top of atmosphere and surface $F_{rad}$, and precipitation $P$. The input and output variables are stacked to vectors $\mathbf{x} = [T(z), Q(z), V(z), P_s, S_{in}, H, E]^T$ with length 94 and $\mathbf{y} = [\Delta T_{phy}(z), \Delta Q_{phy}(z), F_{rad}, P]^T$ with length 65 and normalized to have similar orders of magnitude (*SI Appendix*). We omit condensed water to reduce the complexity of the problem (*Discussion*). Furthermore, there is no momentum transport in our version of SPCAM. Informed by our previous sensitivity tests (24), we use 1 y of SPCAM simulation as training data for the neural network, amounting to around 140 million training samples.

The neural network itself $\hat{\mathbf{y}} = \mathcal{N}(\mathbf{x})$ is a nine-layer deep, fully connected network with 256 nodes in each layer. In total, the network has around 0.5 million parameters that are optimized to minimize the mean-squared error between the network's predictions $\hat{\mathbf{y}}$ and the training targets $\mathbf{y}$ (*SI Appendix*). This neural network architecture is informed by our previous sensitivity tests (24). Using deep rather than shallow networks has two main advantages: First, deeper, larger networks achieve lower training losses; and second, deep networks proved more stable in the prognostic simulations (for details see *SI Appendix* and *SI Appendix*, Fig. S1). Unstable modes and unrealistic artifacts have been the main issue in previous studies that used shallow architectures (22, 23).

Once trained, the neural network replaces the superparameterization's CRM as well as the radiation scheme in CAM. This neural network version of CAM is called NNCAM. In our prognostic global simulations, the neural network parameterization interacts freely with the resolved dynamics as well as with the surface flux scheme. The neural network parameterization speeds up the model significantly: NNCAM's physical parameterization is around 20 times faster than SPCAM's and even 8 times faster than NNCAM's, in which the radiation scheme is particularly expensive. The key fact to keep in mind is that the neural network does not become more expensive at prediction time even when trained with higher-resolution training data. The approach laid out here should, therefore, scale easily to neural networks trained with vastly more expensive 3D global CRM simulations.

The subsequent analyses are computed from 5-y prognostic simulations after a 1-y spin-up. All neural network, model, and analysis code is available in *SI Appendix*.

## Results

**Mean Climate.** To assess NNCAM's ability to reproduce SPCAM's climate we start by comparing the mean subgrid tendencies and the resulting mean state. The mean subgrid heating (Fig. 1*A*) and moistening rates (*SI Appendix*, Fig. S2) of SPCAM and NNCAM are in close agreement with a single latent heating tower at the ITCZ and secondary free-tropospheric heating maxima at the midlatitude storm tracks. The ITCZ peak, which is colocated with the maximum SSTs at 5° N, is slightly sharper in NNCAM compared with SPCAM. In contrast, CTRLCAM exhibits a double ITCZ signal, a common issue of traditional convection parameterizations (5). The resulting mean state in temperature (Fig. 1*B*), humidity, and wind (*SI Appendix*, Fig. S2 *B* and *C*) of NNCAM also closely resembles that of SPCAM throughout the troposphere. The only larger deviations are temperature biases in the stratosphere. Since the mean heating rate bias there is small, the temperature anomalies most likely have a secondary cause—for instance, differences in circulation or internal variability. In any case, these deviations are not of obvious concern because the upper atmosphere is poorly resolved in our setup and highly sensitive to changes in the model setup (*SI Appendix*, Fig. S5 *C* and *D*). In fact, CTRLCAM has even larger differences compared with SPCAM in the stratosphere but also throughout the troposphere for all variables.

The radiative fluxes predicted by the neural network parameterization also closely match those of SPCAM for most of the globe, whereas CTRLCAM has large differences in the tropics and subtropics caused by its double-ITCZ bias (Fig. 1*C* and *SI Appendix*, Fig. S2*D*). Toward the poles NNCAM's fluxes diverge slightly, the reasons for which are yet unclear. The mean precipitation of NNCAM and SPCAM follows the latent heating maxima with a peak at the ITCZ, which again is slightly sharper for NNCAM.

In general, the neural network parameterization, freely interacting with the resolved dynamics, reproduces the most important aspects of its training model's mean climate to a remarkable degree, especially compared with the standard parameterization.
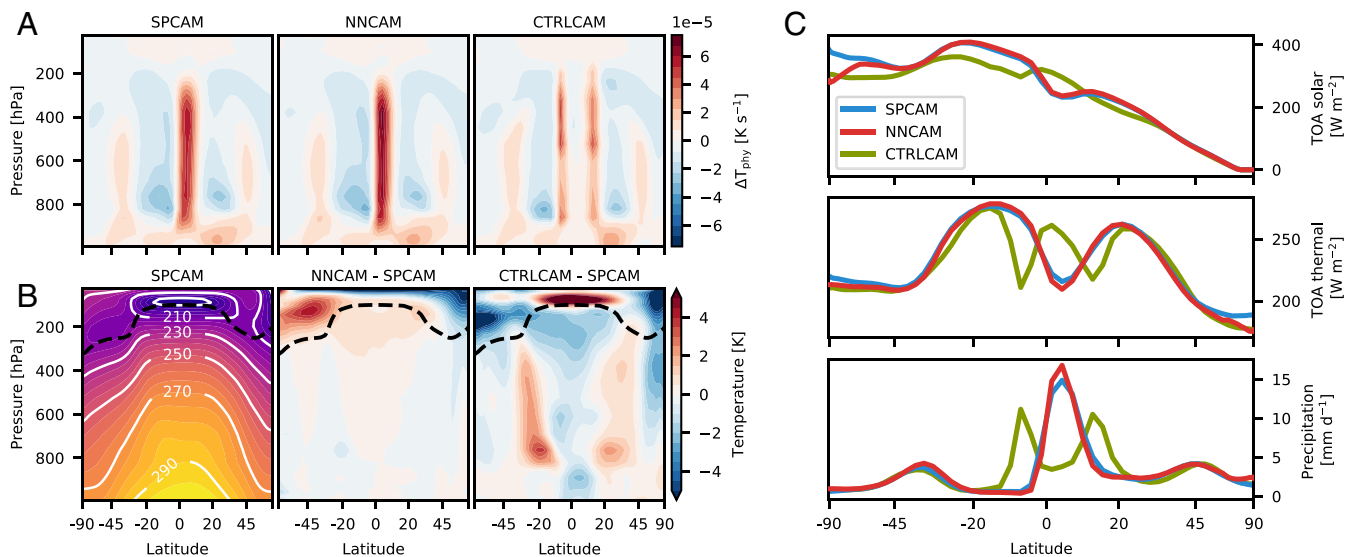
**Fig. 1.** (*A–C*) Longitudinal and 5-y temporal averages. (*A*) Mean convective and radiative subgrid heating rates $\Delta T_{phy}$. (*B*) Mean temperature *T* of SPCAM and biases of NNCAM and CTRLCAM relative to SPCAM. The dashed black line denotes the approximate position of the tropopause, determined by a $\partial p \theta$ contour. (*C*) Mean shortwave (solar) and longwave (thermal) net fluxes at the top of the atmosphere and precipitation. Note that the latitude axis is area weighted.

**Variability.** Next, we investigate NNCAM's ability to capture SPCAM's higher-order statistics—a crucial test since climate modeling is as much concerned about variability as it is about the mean. One of the key statistics for end users is the precipitation distribution (Fig. 2*A*). CTRLCAM shows the typical deficiencies of traditional convection parameterizations—too much drizzle and a lack of extremes. SPCAM remedies these biases and has been shown to better fit to observations (31). The precipitation distribution in NNCAM closely matches that of SPCAM, including the tail. The rarest events are slightly more common in NNCAM than in SPCAM, which is consistent with the narrower and stronger ITCZ (Fig. 1 *A* and *C*).

We now focus on the variability of the heating and moistening rates (Fig. 2*B* and *SI Appendix*, Fig. S3*A*). Here, NNCAM shows reduced variance compared with SPCAM and even CTRLCAM, mostly located at the shallow cloud level around 900 hPa and in the boundary layer. Snapshots of instantaneous heating and moistening rates (*SI Appendix*, Fig. S3 *B* and *C*) confirm that the neural network's predictions are much smoother; i.e., they lack the vertical and horizontal variability of SPCAM and CTRL-CAM. We hypothesize that this has two separate causes: First, low training skill in the boundary layer (24) suggests that much of SPCAM's variability in this region is chaotic and, therefore,

has limited inherent predictability. Faced with such seemingly random targets during training, the deterministic neural network will opt to make predictions that are close to the mean to lower its cost function across samples. Second, the omission of condensed water in our network inputs and outputs limits NNCAM's ability to produce sharp radiative heating gradients at the shallow cloud tops. Because the circulation is mostly driven by midtropospheric heating in tropical deep convection and midlatitude storms, however, the lack of low-tropospheric variability does not seem to negatively impact the mean state and precipitation predictions. This result is also of interest for climate prediction in general.

The tropical wave spectrum (32) depends vitally on the interplay between convective heating and large-scale dynamics. This makes it a demanding, indirect test of the neural network parameterization's ability to interact with the dynamical core. Current-generation climate models are still plagued by issues in representing tropical variability: In CTRLCAM, for instance, moist kelvin waves are too active and propagate too fast while the MJO is largely missing (Fig. 3). SPCAM drastically improves the realism of the wave spectrum (29), including in our aquaplanet setup (26). NNCAM captures the key improvements of SPCAM relative to CTRLCAM: a damped kelvin wave spectrum, albeit



**Fig. 2.** (*A*) Precipitation histogram of time-step (30 min) accumulation. The bin width is 3.9 mm·d$^{-1}$. Solid lines denote simulations for reference SSTs. Dashed lines denote simulations for +4-K SSTs (explanation in *Generalization*). The neural network in the +4-K case is NNCAM-ref + 4 K. (*B*) Zonally averaged temporal SD of convective and radiative subgrid heating rates $\Delta T_{phy}$.

**Fig. 3.** Space–time spectrum of the equatorially symmetric component of 15S–15N daily precipitation anomalies divided by background spectrum as in figure 3b in ref. 32. Negative (positive) values denote westward (eastward) traveling waves.

slightly weaker and faster in NNCAM, and an MJO-like intraseasonal, eastward traveling disturbance. The background spectra also agree well with these results (*SI Appendix*, Fig. S6*A*).

Overall, NNCAM's ability to capture key advantages of the cloud-resolving training model—representing precipitation extremes and producing realistic tropical waves—is to some extent unexpected and represents a major advantage compared with traditional parameterizations.

**Energy Conservation.** A necessary property of any climate model parameterization is that it conserves energy. In our setup, energy conservation is not prescribed during network training. Despite this, NNCAM conserves column moist static energy to a remarkable degree (Fig. 4*A*). Note that because of our omission of condensed water, the balance shown is only approximately true and exhibits some scatter even for SPCAM. The spread is slightly larger for NNCAM, but all points lie within a reasonable range, which shows that NNCAM never severely violates energy conservation. These results suggest that the neural network has approximately learned the physical relation between the input and output variables without being instructed to. This permits a simple postprocessing of the neural network's raw predictions to enforce exact energy conservation. We tested this correction without noticeable changes to the main results. Conservation of total moisture is equally as important but the lack of condensed water makes even an approximate version impossible.

The globally integrated total energy and moisture are also stable without noticeable drift or unreasonable scatter for multiyear simulations (Fig. 4*B*). This is still true for a 50-y NNCAM simulation that we ran as a test. The energy conservation properties of the neural network parameterization are promising and show that, to a certain degree, neural networks can learn higher-level concepts and physical laws from the underlying dataset.

**Generalization.** A key question for the prediction of future climates is whether such a neural network parameterization can generalize outside of its training manifold. To investigate this we run a set of sensitivity tests with perturbed SSTs. We begin by breaking the zonal symmetry of our reference state by adding a wavenumber one SST perturbation with 3-K amplitude (Fig. 5*A* and *SI Appendix*). Under such a perturbation SPCAM develops a thermally direct Walker circulation within the tropics with convective activity concentrated at the downwind sector of the warm pool. The neural network trained with the zonally invariant reference SSTs only (NNCAM) is able to generate a similar heating pattern even though the heating maximum is slightly weaker and more spread out. The resulting mean temperature state in the troposphere is also in close agreement, with biases of less than 1 K (*SI Appendix*, Fig. S4). Moreover, NNCAM runs stably despite the fact that the introduced SST perturbations exceed the training climate by as much as 3 K. CTRLCAM, for comparison, has a drastically damped heating maximum and a double ITCZ to the west of the warm pool.

Our next out-of-sample test is a global SST warming of up to 4 K in 1-K increments. We use the mass-weighted absolute temperature differences relative to the SPCAM reference solution at each SST increment as a proxy for the mean climate state difference (Fig. 5*B*). The neural network trained with the reference climate only (NNCAM) is unable to generalize to much warmer climates. A look at the mean heating rates for the +4-K SST simulation reveals that the ITCZ signal is washed out and unrealistic patterns develop in and above the boundary layer (*SI Appendix*, Fig. S5*B*). As a result the temperature bias is significant, particularly in the stratosphere (*SI Appendix*, Fig. S5*D*). This suggests that the neural network cannot handle temperatures that exceed the ones seen during training. To test the opposite case, we also trained a neural network with data from the +4-K SST SPCAM simulation only (NNCAM + 4 K). The respective prognostic simulation for the reference climate has a realistic heating rate and temperature structure at the equator but fails at the poles, where temperatures are lower than in the +4-K training dataset (*SI Appendix*, Fig. S5 *A* and *C*).

Finally, we train a neural network using 0.5 y of data from the reference and the +4-K simulations each, but not the intermediate increments (NNCAM-ref + 4 K). This version performs well for the extreme climates and also in between (Fig. 5*B* and *SI Appendix*, Fig. S5). Reassuringly, NNCAM-ref + 4 K is also able to capture important aspects of global warming: an increase in the precipitation extremes (Fig. 2*A*) and an amplification and acceleration of the MJO and kelvin waves (*SI Appendix*, Fig. S6*B*). These sensitivity tests suggest that the neural network is unable to extrapolate much beyond its training climate but can interpolate in between extremes.
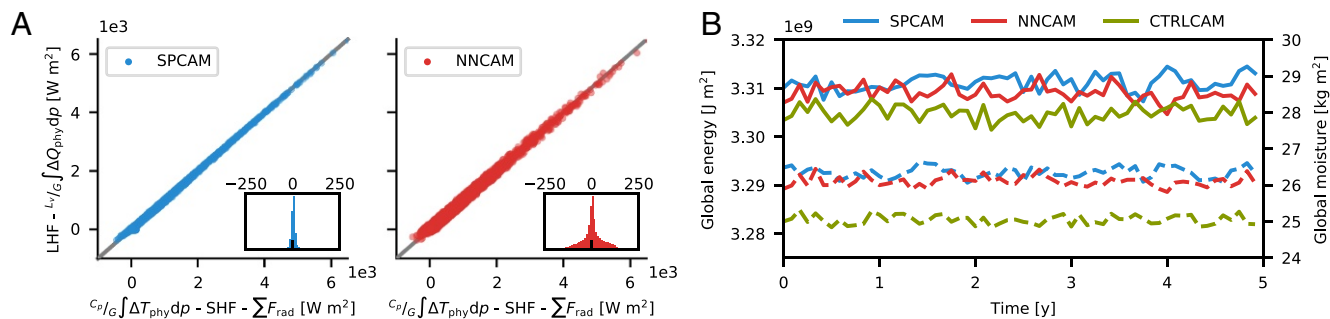


**Fig. 4.** (*A*) Scatter plots of vertically integrated column heating $^{C_p}/_G \int \Delta T_{phy} dp$ minus the sensible heat flux $H$ and the sum of the radiative fluxes at the boundaries $\sum F_{rad}$ against the vertically integrated column moistening $^{L_v}/_G \int \Delta T_{phy} dp$ minus the latent heat flux $H$. Each solid circle represents a single prediction at a single column. A total of 10 time steps are shown. *Inset* shows distribution of differences. (*B*) Globally integrated total energy (static, potential, and kinetic; solid lines) and moisture (dashed lines) for the 5-y simulations after 1 y of spin-up.

**Fig. 5.** (*A*) Vertically integrated mean heating rate $\frac{c_p}{G} \int \Delta T_{phy} dp$ for zonally perturbed SSTs. Contour lines show SST perturbation in 1-K intervals starting at 0.5 K. Dashed contours represent negative values. (*B*) Global mean mass-weighted absolute temperature difference relative to SPCAM reference at each SST increment. The different NNCAM experiments are explained in the key.

## Discussion

In this study we have demonstrated that a deep neural network can learn to represent subgrid processes in climate models from cloud-resolving model data at a fraction of the computational cost. Freely interacting with the resolved dynamics globally, our deep learning-powered model produces a stable mean climate that is close to its training climate, including precipitation extremes and tropical waves. Moreover, the neural network learned to approximately conserve energy without being told so explicitly. It manages to adapt to new surface forcing patterns but struggles with out-of-sample climates. The ability to interpolate between extremes suggests that short-term, high-resolution simulations which target the edges of the climate space can be used to build a comprehensive training dataset. Our study shows a potential way for data-driven development of climate and weather models. Opportunities but also challenges abound.

An immediate follow-up task is to extend this methodology to a less idealized model setup and incorporate more complexity in the neural network parameterization. This requires ensuring positive cloud water concentrations and stability which we found challenging in first tests. Predicting the condensation rate, which is not readily available in SPCAM, could provide a convenient way to ensure conservation properties. Another intriguing approach would be to predict subgrid fluxes instead of absolute tendencies. However, computing the flux divergence to obtain the tendencies amplifies any noise produced by the neural network. Additional complexities like topography, aerosols, and chemistry will present further challenges but none of those seem insurmountable from our current vantage point.

Limitations of our method when confronted with out-of-sample temperatures are related to the traditional problem of overfitting in machine learning—the inability to make accurate predictions for data unseen during training. Convolutional neural networks and regularization techniques are commonly used to fight overfitting. It may well be possible that a combination of these and novel techniques improves the out-of-sample predictions of a neural network parameterization. Note also that our idealized training climate is much more homogeneous than the real world climate, for instance a lack of the El Niño-Southern Oscillation, which probably exacerbated the generalization issues.

Convolutional and recurrent neural networks could be used to capture spatial and temporal dependencies, such as propagating mesoscale convective systems or convective memory across time steps. Furthermore, generative adversarial networks (20) could be one promising avenue toward creating a stochastic machine-learning parameterization that captures the variability of the training data. Random forests (33) have also recently been applied to learn and model subgrid convection in a global climate model (34). Compared with neural networks, they have the advantage that conservation properties are automatically obeyed but suffer from computational limitations.

Recently, it has been argued (35) that machine learning should be used to learn the parameters or parametric functions within a traditional parameterization framework rather than the full parameterization as we have done. Because the known physics are hard coded, this could lead to better generalization capabilities, a reduction of the required data amount, and the ability to isolate individual components of the climate system for process studies. On the flip side, it still leaves the burden of heuristically finding the framework equations, which requires splitting a coherent physical system into subprocesses. In this regard, our method of using a single network naturally unifies all subgrid processes without the need to prescribe interactions.

Regardless of the exact type of learned algorithm, once implemented in the prognostic model some biases will be unavoidable. In our current methodology there is no way of tuning after the training stage. We argue, therefore, that an online learning approach, where the machine-learning algorithm runs and learns in parallel with a CRM, is required for further development. Superparameterization presents a natural fit for such a technique. For full global CRMs this likely is more technically challenging.

A grand challenge is how to learn directly from observations—our closest knowledge of the truth—rather than high-resolution simulations which come with their own baggage of tuning and parameterization (turbulence and microphysics) (35). Complications arise because observations are sparse in time and space and often only of indirect quantities, for example satellite observations. Until data assimilation algorithms for parameter estimation advance, learning from high-resolution simulations seems the more promising route toward tangible progress in subgrid parameterization.

Our study presents a paradigm shift from the manual design of subgrid parameterizations to a data-driven approach that leverages the advantages of high-resolution modeling. This general methodology is not limited to the atmosphere but can equally as well be applied to other components of the Earth system and beyond. Challenges must still be overcome, but advances in computing capabilities and deep learning in recent years present novel opportunities that are just beginning to be investigated. We believe that machine-learning approaches offer great potential that should be explored in concert with traditional model development.

## Materials and Methods

Detailed explanations of the model and neural network setup can be found in *SI Appendix*. *SI Appendix* also contains links to the online code repositories. The raw model output data amount to several TB and are available from the authors upon request.

1. Schneider T, et al. (2017) Climate goals and computing the future of clouds. *Nat Clim Change* 7:3–5.
2. Hourdin F, et al. (2017) The art and science of climate model tuning. *Bull Am Meteorol Soc* 98:589–602.
3. Stevens B, Bony S, Ginoux P, Ming Y, Horowitz LW (2013) What are climate models missing? *Science* 340:1053–1054.
4. Bony S, et al. (2015) Clouds, circulation and climate sensitivity. *Nat Geosci* 8:261–268.
5. Oueslati B, Bellon G (2015) The double ITCZ bias in CMIP5 models: Interaction between SST, large-scale circulation and precipitation. *Clim Dyn* 44:585–607.
6. Arnold NP, Randall DA (2015) Global-scale convective aggregation: Implications for the Madden-Julian oscillation. *J Adv Model Earth Syst* 7:1499–1518.
7. Gentine P, et al. (2013) A probabilistic bulk model of coupled mixed layer and convection. Part I: Clear-sky case. *J Atmos Sci* 70:1543–1556.
8. Bony S, Dufresne J (2005) Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophys Res Lett* 32:L20806.
9. Sherwood SC, Bony S, Dufresne JL (2014) Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature* 505:37–42.
10. Weisman ML, Skamarock WC, Klemp JB (1997) The resolution dependence of explicitly modeled convective systems. *Mon Weather Rev* 125:527–548.
11. Sun J, Pritchard MS (2016) Effects of explicit convection on global land-atmosphere coupling in the superparameterized CAM. *J Adv Model Earth Syst* 8:1248–1269.
12. Leutwyler D, Lüthi D, Ban N, Fuhrer O, Schär C (2017) Evaluation of the convection-resolving climate modeling approach on continental scales. *J Geophys Res Atmos* 122:5237–5258.
13. Muller C, Bony S (2015) What favors convective aggregation and why? *Geophys Res Lett* 42:5626–5634.
14. Soden BJ, Vecchi GA (2011) The vertical distribution of cloud feedback in coupled ocean-atmosphere models. *Geophys Res Lett* 38:L12704.
15. Miyamoto Y, et al. (2013) Deep moist atmospheric convection in a subkilometer global simulation. *Geophys Res Lett* 40:4922–4926.
16. Bretherton CS, Khairoutdinov MF (2015) Convective self-aggregation feedbacks in near-global cloud-resolving simulations of an aquaplanet. *J Adv Model Earth Syst* 7:1765–1787.
17. Yashiro H, et al. (2016) Resolution dependence of the diurnal cycle of precipitation simulated by a global cloud-system resolving model. *SOLA* 12:272–276.
18. Klocke D, Brueck M, Hohenegger C, Stevens B (2017) Rediscovery of the doldrums in storm-resolving simulations over the tropical Atlantic. *Nat Geosci* 10:891–896.
19. LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444.
20. Goodfellow I, Bengio Y, Courville A (2016) *Deep Learning* (MIT Press, Cambridge, MA).
21. Nielsen MA (2015) *Neural Networks and Deep Learning*. Available at neuralnetworksanddeeplearning.com/. Accessed August 23, 2018.
22. Krasnopolsky VM, Fox-Rabinovitz MS, Belochitski AA (2013) Using ensemble of neural networks to learn stochastic convection parameterizations for climate and numerical weather prediction models from data simulated by a cloud resolving model. *Adv Artif Neural Syst* 2013:1–13.
23. Brenowitz ND, Bretherton CS (2018) Prognostic validation of a neural network unified physics parameterization. *Geophys Res Lett* 45:6289–6298.
24. Gentine P, Pritchard M, Rasp S, Reinaudi G, Yacalis G (2018) Could machine learning break the convection parameterization deadlock? *Geophys Res Lett* 45:5742–5751.
25. Collins WD, et al. (2006) The formulation and atmospheric simulation of the community atmosphere model version 3 (CAM3). *J Clim* 19:2144–2161.
26. Andersen JA, Kuang Z (2012) Moist static energy budget of MJO-like disturbances in the atmosphere of a zonally symmetric aquaplanet. *J Clim* 25:2782–2804.
27. Khairoutdinov MF, Randall DA (2001) A cloud resolving model as a cloud parameterization in the NCAR Community Climate System Model: Preliminary results. *Geophys Res Lett* 28:3617–3620.
28. Pritchard MS, Bretherton CS, DeMott CA (2014) Restricting 32-128 km horizontal scales hardly affects the MJO in the Superparameterized Community Atmosphere Model v.3.0 but the number of cloud-resolving grid columns constrains vertical mixing. *J Adv Model Earth Syst* 6:723–739.
29. Benedict JJ, Randall DA (2009) Structure of the Madden–Julian oscillation in the superparameterized CAM. *J Atmos Sci* 66:3277–3296.
30. Arnold NP, Randall DA (2015) Global-scale convective aggregation: Implications for the Madden-Julian oscillation. *J Adv Model Earth Syst* 7:1499–1518.
31. Kooperman GJ, Pritchard MS, O'Brien TA, Timmermans BW (2018) Rainfall from resolved rather than parameterized processes better represents the present-day and climate change response of moderate rates in the community atmosphere model. *J Adv Model Earth Syst* 10:971–988.
32. Wheeler M, Kiladis GN (1999) Convectively coupled equatorial waves: Analysis of clouds and temperature in the wavenumber–frequency domain. *J Atmos Sci* 56:374–399.
33. Breiman L (2001) Random forests. *Machine Learn* 45:5–32.
34. O'Gorman PA, Dwyer JG (2018) Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change and extreme events. arXiv:1806.11037.
35. Schneider T, Lan S, Stuart A, Teixeira J (2017) Earth system modeling 2.0: A blueprint for models that learn from observations and targeted high-resolution simulations. *Geophys Res Lett* 44:12,396–12,417.

EARTH, ATMOSPHERIC, AND PLANETARY SCIENCES