CAMBRIDGE
UNIVERSITY PRESS

RESEARCH ARTICLE

Separation of climate models and observations based on daily output using machine learning

Lukas Brunner[1,2] [ID]* and Sebastian Sippel[2] [ID]

[1]Department of Meteorology and Geophysics, University of Vienna, Vienna, Austria
[2]Institute for Atmospheric and Climate Science, ETH Zurich, Zurich, Switzerland
*Corresponding author. Email: l.brunner@univie.ac.at

NOTE: This manuscript has been submitted to Environmental Data Science and has not yet undergone peer-review.

Abstract

Climate models are primary tools to investigate processes in the climate system, to project future changes, and to inform decision makers. The latest generation of models provides simulations of unprecedented complexity and resolution, resulting in increasingly realistic representations of the real climate system. At the same time, there is growing awareness that not all models produce equally independent and plausible simulations. Therefore, methods are being developed to evaluate how model dependence affects output similarity and how well models match observed climate. Such methods typically draw on climatological averages over several decades to minimize internal variability. Here we show that models can robustly be identified as different from observations based on temperature output from only a single day using two statistical and machine learning classifiers. The first classifier, a logistic regression, achieves almost perfect accuracy in identifying models never seen in training. An analysis of the derived regression coefficients reveals that this skill is founded in several climatological model biases emerging on daily time-scales already. The second classifier, a convolutional neural network, is even able to correctly identify the majority of 43 different models based on spatial dependencies between daily global weather patterns alone, i.e., after removing the climatological model biases. Misclassifications appear mostly between models developed at the same institution, indicating that the effects of shared code also emerge already on daily times scales. These results build the basis for further research on the usage of statistical and machine learning for evaluating models without relying on climatological time-scales.

Impact Statement

Climate models are our main tool to investigate changes in the climate system and to provide information for decision makers. While the models become increasingly better in representing climate system processes, some differences compared to observations remain. Such differences are typically investigated on climatological time-scales of several decades. Here we show that statistical and machine learning classifiers can reliably distinguish models from observations based on the temperature map from only a single day. We find that several areas of climatological model biases emerge already on a daily basis and can, hence, be used for separation. These results provide new avenues for the evaluation and interpretation of model-observation differences, on short time scales, and targeting pattern-based model features.

1. Introduction

Multi-model ensembles such as the latest sixth Coupled Model Intercomparison Project (CMIP6; Eyring et al. [2016]) are frequently used to investigate physical climate mechanisms, to attribute past and project future changes, and inform political decisions. However, there is growing awareness in the climate community that not all models contained in CMIP6 provide equally plausible and independent simulations of the climate system [Tebaldi and Knutti, 2007, Knutti, 2010, Bishop and Abramowitz, 2013, Annan and Hargreaves, 2017, Eyring et al., 2019]. Model evaluation methods, typically based on climatological means over multiple decades, have identified persistent biases across models compared to observations and highlighted the effect of model dependence on model output similarity [Masson and Knutti, 2011, Knutti et al., 2013, Boé, 2018, Brunner et al., 2020, Bock et al., 2020]. Both, model bias and dependence, are ultimately due to the parametrizations used to represent processes which cannot explicitly be resolved in simulations of the observed climate in numerical models.

Work is ongoing to further reduce these parametrizations in a new generation of kilometer-scale, storm-resolving global models with the ultimate aim of creating a digital twin of Earth [Bauer et al., 2021, Rackow et al., 2021]. These developments continue to blur the border between weather prediction, climate predictions, and climate projections [Meehl et al., 2021] and call for new, innovative evaluation methods that allow to contrast models with each other and with observations on weather time-scales. This includes investigating if model biases and dependencies can be identified already on such scales.

Here we present initial results from a setup drawing on different statistical and machine learning classifiers that allows such an evaluation even on daily data. Similar approaches based on statistical and machine learning have been used previously to identify forced signals from spatial patterns of temperature, precipitation, and humidity [Barnes et al., 2019, Sippel et al., 2020], to explore the role of single forcing agents [Labe and Barnes, 2021], and to predict modes of atmosphere-ocean variability [Gordon et al., 2021]. In this exploratory study we showcase and discuss the potential of our approach based on several examples drawing on reanalysis data (used to represent observations), CMIP6 and experimental storm resolving simulations from the NextGEMS project[1].

The work is guided by two main questions that will be further detailed in sections 3 and 4: based on their daily temperature output, can models reliably be identified (1) as different from observational and (2) as different from each other?

2. Data and Methods

2.1. *Model and reanalysis data*

We use all available models from CMIP6 which provide daily data for surface air temperature in the historical period. In total this sums up to 43 models which can broadly be grouped into 22 families by developing institution (see table S1 in the supplement for a full list). Daily reanalysis data, taken as pseudo-observations, are derived from the fifth generation of the European Centre for Medium-Range Weather Forecasts (ECMWF) Retrospective Analysis (ERA5; Hersbach et al. [2020]). The period from 1950 to 2014 is used to match the first availability of ERA5 and the end of the historical forcing period in CMIP6, respectively and all daily temperature fields are regridded to 2.5°×2.5°. In addition, we use temperature projections for the end of the century (2091-2100) driven by the high emission scenario SSP5-8.5 [Meinshausen et al., 2020] if available to test the robustness of our approach under severe warming (see table S1). Finally, we also draw on one year (February 2020 to January 2021) of prototype data from the NextGEMS project and use global temperature fields from an ICON run with an atmospheric resolution of 5km.

---

[1]https://nextgems-h2020.eu

2.2. Statistical and machine learning methods

We use two different statistical and machine learning methods to separate models from observations, and from each other: First, logistic regression, which allows insights of learned coefficients but has the limitation of being a linear method. Second, a convolutional neural network (CNN) which represents rather the other end of the complexity spectrum, allowing to learn non-linear spatial relations between features but lacking the easy interpretability logistic regression allows. Both methods are briefly described here and we provide Jupyter notebooks in the supplement showcasing their exact implementation.

Logistic regression is linear in its parameters and takes an $M \times N$ matrix as input, where N is the number of days and M the number of grid cells. For training an additional vector of length N is provided containing the labels (i.e., 0 for 'model' and 1 for 'observation') for each day. A continuous probability is then assigned to each test sample with 1 indicating the highest certainty that a given day belongs into the observation category and 0.5 representing the decision threshold below which it is assigned into the model category. The logistic regression classifier is constrained by a L_2 regularization (i.e., a penalty on the squared sum of regression coefficients) to avoid overfitting and to ensure smooth coefficients in space. This regularization is optimized using 5-fold cross-validation. More general details on regularized logistic regression can be found, for instance, in Hastie et al. [2009] and an example implementation is provided in the supplement.

To complement the logistic regression classifier we use a CNN as a second method. Deep neural networks such as CNNs are considerably more complex and less interpretable than traditional methods but their performance is superior in exchange. Therefore, their use in different scientific disciplines, not least in climate sciences, has been rapidly increasing in recent years [Hsieh, 2022, Kashinath et al., 2021]. Due to their complexity many different design choices in the exact layout of the network are possible, here we use an out-of-the-box setup consisting of four convolutional and maximum pooling layers resulting in about 122'000 trainable parameters (see supplement for the exact architecture). The 2-dimensional temperature field from each day can directly be interpreted by the CNN equivalent to a picture classification task, therefore, the input layer takes a $K \times L \times N$ matrix where N is the number of days, K the number of latitudes, and L the number of longitudes (hence, $K * L = M$, the number of grid cells). In contrast to the logistic regression the CNN is set up to allow distinguishing more than two output categories and a given input is identified as belonging to the class which is assigned the highest probability. For the case with only the model and observation classes, discussed in the first part of the manuscript, this means the threshold for assignment is 0.5 as for the logistic regression.

2.3. Data preprocessing and conventions

Each daily 2-dimensional temperature field is used as a sample and each of the 10'368 grid cells (resulting from 72 latitudes and 144 longitudes) as a feature in machine learning terminology. The 2-dimensional daily fields are either used directly in the case of the CNN classifier or reduced to a 1-dimensional feature array in the case of logistic regression. Note that the changing grid cell area with latitude is not explicitly accounted for in this study but might implicitly be learned by the classifiers. Training and validation data are drawn from the 50-year period 1950-1999 and test data from the slightly offset 10-year period 2005-2014 to ensure that they are temporarily out-of-sample without leakage in the temporal domain. For reasons of computational performance only subsets of days are used for training as specified in the corresponding sections. We also use a setup with individual models being withheld from the training sample altogether and only used in the test dataset (out-of-sample in the model domain in addition to the time domain).

Two pre-processing steps are applied to investigate their influence on the classification accuracy: (1) removing the global mean from each day and from the entire spatial pattern, and (2) removing the mean seasonal cycle from each grid cell on a model-by-model basis in addition to (1). The mean seasonal cycle is calculated on a day-of-the-year basis for each model and the observations separately using the
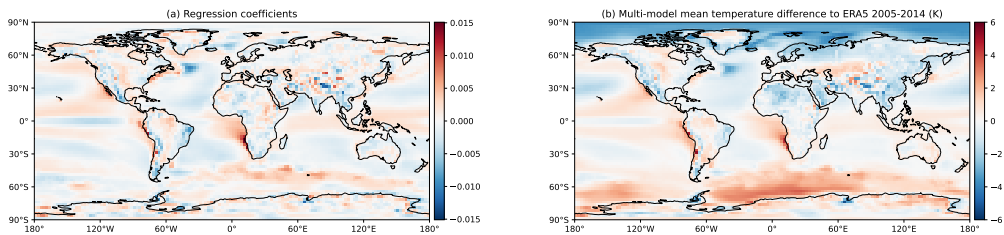
Figure 1. *(a) Logistic regression coefficients learned from daily data in the period 1950-1999 and used to separate daily temperature maps from models and observations. (b) Multi-model mean temperature difference to ERA5 based on climatologies in the period 2005-2014. The climatologies are calculated from daily data with the global mean removed.*

training period 1950-1999 and a 31 day running window. A physical interpretation of (1) is that the classifiers train on persistent regional biases as far as they emerge from internal variability on daily scales. In (2) such regional biases are explicitly removed together with any biases in the equator-pole gradient or between the hemispheres. This means that the classifiers can only train on the spatial relations of internal variability making the problem considerably harder. Finally, each sample is associated with one of two labels categories to be predicted by the machine learning classifiers: either 'model'/'observation' or the name of the dataset (i.e., 'ERA5' or one of 43 model names). Figures S6 and S7 in the supplement show an example test day from each model and ERA5 for the two pre-processing cases.

3. Separating models and observations

3.1. Logistic regression classifier

First, we train a single logistic regression classifier with the aim of distinguishing models from observations based only on the temperature maps from a single day. For this case we only remove the global mean (i.e., the mean seasonal cycle is still included; see figure S6 for an example). We use 200 different, randomly drawn training days from each of the 43 models resulting in 8'600 training samples labeled 'model' which is matched by random 8'600 days labeled 'observation' (out of 18'250 total days available from ERA5). Sensitivity tests show that using less samples for training leads to strongly decreasing test accuracy while using more samples generally leads to slightly better results as one would expect, but also greatly increases computing resource usage.

For this setup the accuracy for the test sample (i.e., individual days from 2005-2014) is almost 100 % with only individual samples (less than 10 from over 150'000) being misclassified. Since temperatures between different grid cells used as features are not independent the cross-validation yields a strong L_2 regularization parameter of about $C = 0.005$. Overall, this very high accuracy shows that there are persistent distinguishing features that can be used to reliably separate climate models and reanalysis data even in the presence of large internal variability on daily time scales. The spatial patterns in these features are reflected in the regression coefficients learned by the classifier and shown in figure 1. The distribution of coefficients identifies areas important for the separation of climate models and observations. In addition, the sign of the coefficients can be physically interpreted with positive values indicating regions where models tend to be warmer than the observations and vice-versa.

Over the ocean the most prominent region of negative coefficients is found in the North Atlantic, near the so-called North Atlantic warming hole [Chemke et al., 2020, Keil et al., 2020]. Here models seem to systematically underestimate temperatures on a daily basis compared with observations and relative to the global mean. In contrast, there are regions of high coefficients on the eastern edges of the Pacific and Atlantic ocean basins. These regions correspond to persistent model biases in the representation of clouds and their radiative effect which are known to appear on all times scales from daily to decadal

[Williams et al., 2013, Ma et al., 2014, Brient et al., 2019, Bock et al., 2020, Chen et al., 2022, figure 1b]. In the equatorial Pacific, known as a region with notorious climate model biases that typically show too cold and too narrow equatorial cold tongues, negative coefficients are also found by the logistic regression classifier. This is accompanied by warm biases to the north and south (shown as positive regression coefficients) connected to the models representation of the intertropical convergence zone [e.g., Hirota et al., 2011, Li and Xie, 2014, Tian and Dong, 2020]. Overall, the consistency in the pattern between the logistic regression coefficients identified for classification on daily time scales (figure 1a) and the long-term climatological model-biases (figure 1b) is remarkable. A notable exception are the northern high latitudes where the clear cold-bias in the climatological multi-model mean is not reflected in a corresponding pattern in the regression coefficients. Similarly for the southern part of the southern ocean where the climatological warm-bias in models is not reflected in positive regression coefficients. Daily temperatures from these regions do, hence, not carry information relevant for the separation of models and observations. This seems to be caused by the large model spread in these regions with models both considerably over- and underestimating ERA5 temperatures on climatological time scales (see figure S8).

Focusing on land grid cells, it becomes clear that the corresponding regression coefficients are considerably more prone to (over-)fitting to rather small-scale structures in particular in regions with complex terrain such as the Andes or the Himalayan. Given the finer native resolution of ERA5 compared to models it is also possible that the classifier learns information from the different representations of topography in observations and models. To avoid this behavior we continue to use land masked data in the rest of the manuscript. A map of regression coefficients from a classifier trained on such land-masked data can be found in figure S1 in the supplement and shows very similar patterns over ocean compared to figure 1a.

Now, we look more closely into how confident the classifier is in separating observations and models split by model variant. Figure 2a shows the probability for test samples from 44 datasets (43 different model variants and ERA5) to be in either category. Since the probabilities for the categories are complementary, a predicted probability of 0 % to be an observation is equivalent to 100 % to be a model ($p(0|X) = 1 - p(1|X)$ in logistic regression). For the vast majority of cases the logistic regression classifier assigns the correct category with close to 100 % probability. This is quite remarkable considering that only 200 random days per model variant were used for training, meaning that for an individual model the seasonal cycle is not sampled completely in training, and decadal- and multi-decadal variability can be sampled only superficially in those 200 days. Common features, therefore, seem to get generalised between models. However, for several models a fraction of days from the 10 years of test data is also predicted with less certainty (whiskers in figure 2a). The model families most prone to get confused with observations are CMCC, CNRM, EC-Earth3, GFDL, HadGEM3, and IPSL. An intuitive interpretation of this behavior might be that these models provide the 'best' representation of 'true' temperatures on a daily basis. However, this is not conclusive and would require additional evidence. First, there are only very few samples with a non-zero probability of belonging into the observation category for each given model, therefore the confusion could be based on chance (i.e., a model sample could be classified as observations, but for the wrong reasons so that no conclusion on the overall model performance should be drawn). A second important consideration is the possible codebase overlap also between models and ERA5 that might get picked up by the classifier. For EC-Earth3 (for which all four variants show a tendency to get confused), for example, there are documented dependencies on the ECMWF atmosphere [Döscher et al., 2022].

We continue to investigate if there are more fundamental features that distinguish models from observations and if a model that has never been used in training can still be correctly identified. For each tested model variant a separate classifier which was not trained on that particular model nor on any model from the same model family is used in figure 2b (see table S1 for a list of model families). The exclusion of models from the same family ensures that the tested model is not merely identified as similar to another model from the same family (this will be discussed further in the second part of the manuscript). Note that observations are not tested for this case as they have to be included in
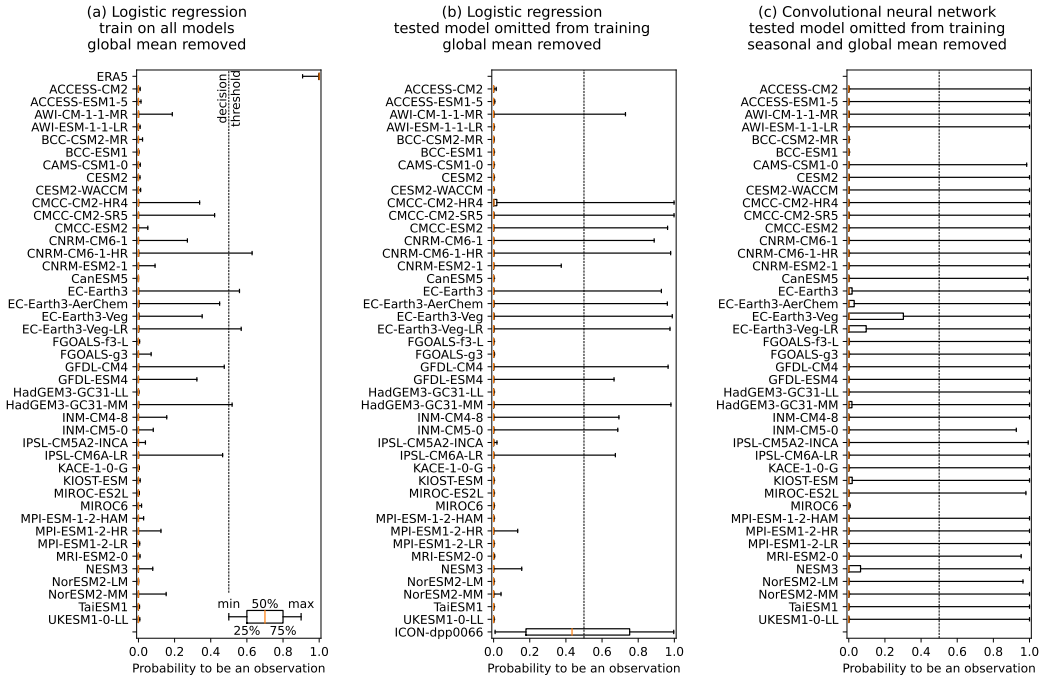
Figure 2. *Distribution of predicted probabilities for the testset split by dataset. The vertical dotted line at 50 % marks the decision threshold between the two categories. (a) Results for the logistic regression classifier using landmasked data with the global mean removed and training on all models. (b) Same as (a) but models from the same family as tested model are withheld from the training sample in turn, i.e., using 22 separately trained classifiers based on the model families given in table S1. The last model, ICON-dpp0066, is never used in training and has only 1 year of test samples available. (c) Results for the convolutional neural network and the same settings as (b) with the seasonal cycle removed in addition.*

training. The more strict out-of-sample interpretation leads to more test samples getting confused with the observations, but overall accuracy remains close to 100 %. This means that most samples can clearly be identified as belonging to the model category even if the classifier has not been trained on any models from the same family. This could be due to remaining dependencies across families, which have been quite loosely defined based on institutions in this study. However, it could also be an indication for the existence of shared patterns in they daily temperature fields persisting across all CMIP6 models.

Focusing on model families with high and low resolution variants (e.g., AWI, CMCC, CNRM, HadGEM, MPI) one can speculate about a certain resolution dependence in the probability to be mis-classified (figures 2a and 2b). Several higher resolution model variants seem to have a higher probability to be misclassified compared to their lower resolution siblings. Such a behavior would be consistent with the findings of Bock et al. [2020] who note that long-standing regional model biases are smaller in higher resolution versions of the same model. To investigate this behavior further we include pre-liminary results from the NextGEMS project. We apply the classifier trained on all models (used in figure 2a) to predict the category of one year of preliminary data from a global, 5km native resolu-tion simulation using ICON that has never been used in training. The results are shown alongside the model-out-of-sample test in figure 2b (bottommost box). Indeed, the daily temperature data from the 5 km resolution ICON can no longer be identified as clearly either CMIP model or observation. Since only one year of data is available it is not possible to investigate full climatological differences to ERA5

(as in figure 1b) but figure S2 shows the difference between the mean over one year in ICON and the year 2014 in ERA5. Most strikingly, the model warm bias at the eastern edges of the ocean basins is greatly reduced in ICON. In other regions differences are increased compared to figure 1b (note that this comparison is very disadvantageous for ICON as figure 1b shows a mean over 10 years and over 43 models).

To summarize, a linear and interpretable logistic regression classifier is found to be able to identify biases that can be used to reliably separate models from observations based on only a single day of temperatures. Patterns in the estimated regression coefficients reveal that known regions of climatological model-observation differences are also important for the identification of climate models on a daily basis. Based on this result it can be hypothesized that also the inverse is true: models which can no longer be distinguished from the observations on a daily basis might have reduced climatological biases in one or several regions. Should such a relation hold in future research this will enable innovative, new ways of model evaluation based on considerably shorter time scales than the typically used 20+ years as pioneered for the case of a 5 km ICON simulation.

3.2. Convolutional neural network

Based on the results in the last section, we here test if models and observations can still be separated in the absence of any climatological biases. To do so we remove the mean seasonal cycle from each grid cell and model separately in addition to removing the daily global mean as in the last section (see figure S3 for the resulting multi-model mean bias compared to ERA5 and figure S7 for a breakdown by individual models). The logistic regression classifier no longer has any skill for this case as the only remaining sources of information are non-linear relations between the spatial structures of daily global weather (see probabilities assigned to the test samples in figure S4a).

We, therefore, move on to use a more complex, yet less interpretable CNN in the rest of the study. First, we repeat the same case as for the logistic regression with only the global mean removed and tested models omitted from training also achieving almost perfect accuracy (see figure S4b). In contrast to the logistic regression classifier the CNN is also able to correctly identify the 5 km ICON runs as a model, indicating that it learns some more fundamental model properties that persist also at high resolutions. This becomes even more obvious if training and testing on samples with the seasonal cycle removed: while the logistic regression did not have any skill for this case the CNN achieves an overall accuracy of about 98 % showcasing the power of this non-linear method. Figure 2c shows the corresponding breakdown of predicted probabilities revealing that now almost all models get confused for observations a number of times while the vast majority of samples are still correctly assigned to the model category.

The skill is remarkable if we remind ourselves that subtracting the mean seasonal cycle for each individual model effectively removes the entire model specif regional bias used in the first part of the manuscript. This means that the only remaining source of information to learn from are amplitude and spatial dependencies of the remaining daily temperature variability.

4. Identifying models by name

In the first part of the manuscript we showed that there are common features across all models that enable us to reliably separate them from observations. Complementary, we now ask if there are also separating features that allow to distinguish models from each other. Past research, based on climatological time scales, has shown that models can be separated as well as clustered into families based only on their output [e.g., Masson and Knutti, 2011, Knutti et al., 2013, Boé, 2018, Brunner et al., 2020, Merrifield et al., 2020]. Here we investigate if models still possess unique features that allow such a separation even on daily time scales.
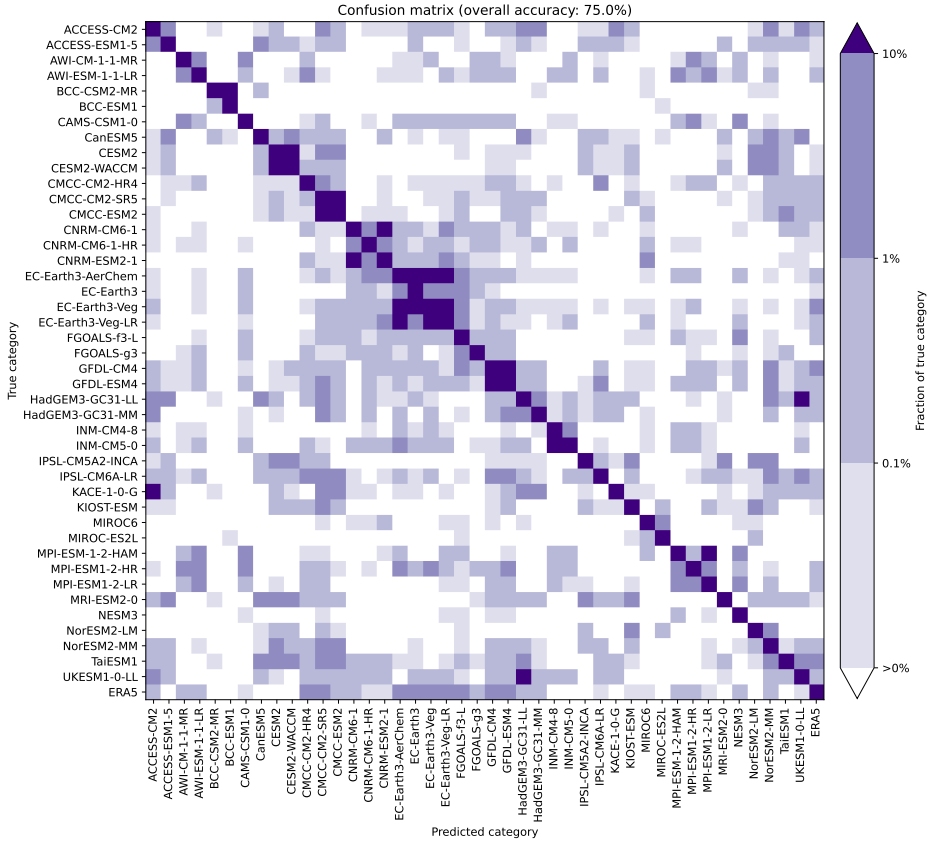
Figure 3. *Confusion matrix of true and predicted labels for classifying daily temperature fields using a convolutional neural network. The vertical axis gives the true, the horizontal axis the predicted label and the shading the fraction of cases relative to the total number in the true category (3'650 samples). Note the logarithmic color bar.*

We use the CNN on the deseasonalized data and train it to recognise each of the 43 models as well as the observations. We use 800 training samples per model and from the observations for this case corresponding to an overall doubling of samples compared to the first part of the manuscript. Assigning the correct label (i.e., model name) to each of the 44 datasets yields an accuracy of 75%. To put this into context, note that compared to the last section (with 2 categories) we now aim to separate 44 categories which considerably increases the difficulty of the classification. The CNN is thus able to pick up patterns unique to each model in order to separate it from all other models. Note that an out-of-sample test for completely unknown models is obviously no longer possible in this classification-by-name setup.

Figure 3 shows the confusion matrix of true sample names versus the predicted sample names. As expected from the overall accuracy the majority of samples is assigned correctly with misclassifications exceeding 10% almost exclusively found within model families (i.e. models from the same institution). The only notable exception is the Korean KACE model being classified as the Australian ACCESS-CM2 for more than 10% of samples. Closer investigation reveals that both models are related to the United Kingdom's HadGEM model family, with models from this larger group all being prone to be confused with each other also in general [Andrews et al., 2019, Bi et al., 2020, Lee et al., 2020].

Finally, we ask how the differences between models relate to the differences due to global warming and check if the classification remains robust in a changing climate. We draw test samples from the period 2091-2100 under SSP8-5.8 rather than from 2005-2014. For this case only 33 models and no observations are available. As figure S5 in the supplement shows the patterns learned by the CNN in the historical period still persist even after extreme climate change and allow the correct identification of about half of the test samples from the end of the century.

Since we use only 800 samples per model for training in this section there is vast potential for improvement (depending on the computational resources spent for training). For example, one could use the entire period 1950-2014 (about 24'000 days per model) for training. This would probably allow for an identification of the model name corresponding to a randomly selected daily temperature field under any SSP with an accuracy considerably higher than 50 %. Ultimately such an approach could be used, for example, in model development to investigate the impact of changes in parametrizations, resolution, or model components on the model output without the need to run the model for several decades. Conversely, if a model is no longer recognised in a future run this might also be an indication for potential problems warranting a closer investigation.

5. Summary and Conclusions

We have shown that machine learning classifiers can distinguish CMIP6 generation climate models from the ERA5 reanalysis based on only a single day of global temperature. Both investigated approaches (logistic regression and convolutional neural network - CNN) are found to perform very well in separating surface air temperature fields where only the global mean has been removed. If the mean seasonal cycle is removed from each grid cell in addition, the logistic regression does no longer show any skill while the CNN still achieves an accuracy of about 98 %.

This performance is quite remarkable given that removing the mean seasonal cycle from each model means that the classifier cannot learn from regions that show absolute biases in the models, and hence the classifier relies only on the spatial dependency structures (and, possibly, the amplitudes of variations) of daily global weather. The properties of the CNN as a highly non-linear, deep learning method, however, do not allow a straightforward extraction of the features used to separate both categories, but a further investigation in the future might be able to reveal more details.

In addition to the CMIP6 models we also tested the classifiers on one year of prototype data from a storm resolving simulation run with ICON at a resolution of 5km. The logistic regression classifier was not able to clearly assign the samples into either category, indicating that this high-resolution case has different or potentially reduced biases compared to the CMIP6 models. The CNN, in turn, manages to correctly predict that the ICON simulations belong to the model category for most samples. This case raises a range of interesting questions, that we will follow up upon once more data is available, for example: If training would be based also/only on ICON output, would it be possible to correctly classify out-of-sample data from ICON and CMIP6 using logistic regression? Does the CNN have skill also for the case where the mean seasonal cycle is removed from ICON?

In the second part of the study we investigated if the CNN is able to separate between 43 model variants, which considerably increases the difficulty of the classification task. Again, we used daily temperature fields with the mean seasonal cycle and the global mean removed. We found an overall accuracy of 75 %, which is about 40 times better than the baseline of a random choice. These results show that the CNN is clearly able to pick up relations between features that reliably separate models, including very similar variants. It is also important to keep in mind that we only used 800 samples per model variant for training due to computational performance restrictions. In theory, orders of magnitude more training samples would be available leaving ample room for improving the classifier. Finally, we showed that the CNN is able to correctly identify about half of the test samples even after 100 years of strong climate change using training samples from the period 1950-2000 and test samples from

2091-2100 in the high-emission SSP5-8.5 scenario. The features identifying climate models are, hence, state-invariant remaining robust even under a warming exceeding several degrees centigrade.

Future applications could build on the approaches illustrated here in several ways: First, the approach could add to the model evaluation toolbox [Eyring et al., 2019], and could target, for example, the classification of individual model components (e.g., atmospheric or ocean component), model generations (e.g., CMIP6 versus the about 10 years older CMIP5), perturbed parameter ensembles, or strains of model development in general. This would add to model evaluation, as the classifiers employed here target spatial patterns, thus likely related to dynamical characteristics, rather than model mean climatology, seasonal cycles, or other error metrics compared to observations. The clustering could also be applied in this context to infer model performance directly (e.g., predicting an estimate of model error), potentially enabling to establish a model's ability to simulate observed climate based on considerably shorter periods than has been possible so far and potentially directly incorporating physical processes [Kashinath et al., 2021]. Second, the classification approaches could be used to pinpoint model-model or model-observation differences. This could be done by analyzing the spatial scales of separability, that is, whether models on regional domains are less separable than globally, and/or whether this may depend on specific regions. Additionally, illustrating and understanding the patterns of separability, by using explainable neural network techniques [e.g., Toms et al., 2020] revealing how the neural network has learnt to distinguish models and observations, could provide additional insights. Lastly, and most speculatively, recent progress in machine learning towards image-to-image translation, using for instance techniques such as generative adversarial networks [Stengel et al., 2020], could provide an avenue towards bias correcting model output towards observations, such that a hypothetically 'bias-corrected' spatial pattern would not be distinguishable from observations.

References

Timothy Andrews, Martin B. Andrews, Alejandro Bodas-Salcedo, Gareth S. Jones, Till Kuhlbrodt, James Manners, Matthew B. Menary, Jeff Ridley, Mark A. Ringer, Alistair A. Sellar, Catherine A. Senior, and Yongming Tang. Forcings, Feedbacks, and Climate Sensitivity in HadGEM3-GC3.1

and UKESM1. *Journal of Advances in Modeling Earth Systems*, 11(12):4377–4394, 2019. ISSN 19422466. doi: 10.1029/2019MS001866.

James D. Annan and Julia C. Hargreaves. On the meaning of independence in climate science. *Earth System Dynamics*, 8(1):211–224, 2017. ISSN 21904987. doi: 10.5194/esd-8-211-2017.

Elizabeth A. Barnes, James W. Hurrell, Imme Ebert-Uphoff, Chuck Anderson, and David Anderson. Viewing Forced Climate Patterns Through an AI Lens. *Geophysical Research Letters*, 46(22):13389–13398, 2019. ISSN 19448007. doi: 10.1029/2019GL084944.

Peter Bauer, Bjorn Stevens, and Wilco Hazeleger. A digital twin of Earth for the green transition. *Nature Climate Change*, 11(2):80–83, 2021. ISSN 17586798. doi: 10.1038/s41558-021-00986-y. URL http://dx.doi.org/10.1038/s41558-021-00986-y.

Daohua Bi, Martin Dix, Simon Marsland, Siobhan O'farrell, Arnold Sullivan, Roger Bodman, Rachel Law, Ian Harman, Jhan Srbinovsky, Harun A. Rashid, Peter Dobrohotoff, Chloe Mackallah, Hailin Yan, Anthony Hirst, Abhishek Savita, Fabio Boeira Dias, Matthew Woodhouse, Russell Fiedler, and Aidan Heerdegen. Configuration and spin-up of ACCESS-CM2, the new generation Australian Community Climate and Earth System Simulator Coupled Model. *Journal of Southern Hemisphere Earth Systems Science*, 70(1):225–251, 2020. ISSN 22065865. doi: 10.1071/ES19040.

Craig H. Bishop and Gab Abramowitz. Climate model dependence and the replicate Earth paradigm. *Climate Dynamics*, 41(3-4):885–900, 2013. ISSN 09307575. doi: 10.1007/s00382-012-1610-y.

L. Bock, A. Lauer, M. Schlund, M. Barreiro, N. Bellouin, C. Jones, G. A. Meehl, V. Predoi, M. J. Roberts, and V. Eyring. Quantifying Progress Across Different CMIP Phases With the ESMValTool. *Journal of Geophysical Research: Atmospheres*, 125(21):1–28, 2020. ISSN 21698996. doi: 10.1029/2019JD032321.

Julien Boé. Interdependency in Multimodel Climate Projections: Component Replication and Result Similarity. *Geophysical Research Letters*, 45(6):2771–2779, 3 2018. ISSN 00948276. doi: 10.1002/2017GL076829. URL http://doi.wiley.com/10.1002/2017GL076829.

Florent Brient, Romain Roehrig, and Aurore Voldoire. Evaluating Marine Stratocumulus Clouds in the CNRM-CM6-1 Model Using Short-Term Hindcasts. *Journal of Advances in Modeling Earth Systems*, 11(1):127–148, 2019. ISSN 19422466. doi: 10.1029/2018MS001461.

Lukas Brunner, Angeline G. Pendergrass, Flavio Lehner, Anna L. Merrifield, Ruth Lorenz, and Reto Knutti. Reduced global warming from CMIP6 projections when weighting models by performance and independence. *Earth System Dynamics*, 11(4):995–1012, 11 2020. ISSN 2190-4987. doi: 10.5194/esd-11-995-2020. URL https://esd.copernicus.org/articles/11/995/2020/.

Rei Chemke, Laure Zanna, and Lorenzo M. Polvani. Identifying a human signal in the North Atlantic warming hole. *Nature Communications*, 11(1):1–7, 2020. ISSN 20411723. doi: 10.1038/s41467-020-15285-x. URL http://dx.doi.org/10.1038/s41467-020-15285-x.

Guoxing Chen, Wei-Chyung Wang, Qing Bao, and Jiandong Li. Evaluation of Simulated Cloud Diurnal Variation in CMIP6 Climate Models. *Journal of Geophysical Research: Atmospheres*, 127(6):1–14, 2022. ISSN 2169-897X. doi: 10.1029/2021jd036422.

Ralf Döscher, Mario Acosta, Andrea Alessandri, Peter Anthoni, Thomas Arsouze, Tommi Bergman, Raffaele Bernardello, Souhail Boussetta, Louis-philippe Caron, Glenn Carver, Miguel Castrillo, Franco Catalano, Ivana Cvijanovic, Paolo Davini, Evelien Dekker, Francisco J. Doblas-Reyes, David Docquier, Pablo Echevarria, Uwe Fladrich, Ramon Fuentes-Franco, Matthias Gröger, Jost v. Hardenberg, Jenny Hieronymus, M. Pasha Karami, Jukka-Pekka Keskinen, Torben Koenigk, Risto Makkonen, François Massonnet, Martin Ménégoz, Paul A. Miller, Eduardo Moreno-Chamarro, Lars Nieradzik, Twan van Noije, Paul Nolan, Declan O'Donnell, Pirkka Ollinaho, Gijs van den Oord, Pablo Ortega, Oriol Tintó Prims, Arthur Ramos, Thomas Reerink, Clement Rousset, Yohan Ruprich-Robert, Philippe Le Sager, Torben Schmith, Roland Schrödner, Federico Serva, Valentina Sicardi, Marianne Sloth Madsen, Benjamin Smith, Tian Tian, Etienne Tourigny, Petteri Uotila, Martin Vancoppenolle, Shiyu Wang, David Wårlind, Ulrika Willén, Klaus Wyser, Shuting Yang, Xavier Yepes-Arbós, and Qiong Zhang. The EC-Earth3 Earth system model for the Coupled Model Intercomparison Project 6. *Geoscientific Model Development*, 15(7):2973–3020, 4 2022. ISSN

1991-9603.   doi: 10.5194/gmd-15-2973-2022.   URL https://gmd.copernicus.org/articles/15/2973/2022/.

Veronika Eyring, Sandrine Bony, Gerald A. Meehl, Catherine A. Senior, Bjorn Stevens, Ronald J. Stouffer, and Karl E. Taylor.   Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5):1937–1958, 5 2016. ISSN 1991-9603. doi: 10.5194/gmd-9-1937-2016. URL https://www.geosci-model-dev.net/9/1937/2016/.

Veronika Eyring, Peter M. Cox, Gregory M. Flato, Peter J. Gleckler, Gab Abramowitz, Peter Caldwell, William D. Collins, Bettina K. Gier, Alex D. Hall, Forrest M. Hoffman, George C. Hurtt, Alexandra Jahn, Chris D. Jones, Stephen A. Klein, John P. Krasting, Lester Kwiatkowski, Ruth Lorenz, Eric Maloney, Gerald A. Meehl, Angeline G. Pendergrass, Robert Pincus, Alex C. Ruane, Joellen L. Russell, Benjamin M. Sanderson, Benjamin D. Santer, Steven C. Sherwood, Isla R. Simpson, Ronald J. Stouffer, and Mark S. Williamson.   Taking climate model evaluation to the next level. *Nature Climate Change*, 9(2):102–110, 2019. ISSN 17586798. doi: 10.1038/s41558-018-0355-y. URL http://dx.doi.org/10.1038/s41558-018-0355-y.

Emily M. Gordon, Elizabeth A. Barnes, and James W. Hurrell. Oceanic Harbingers of Pacific Decadal Oscillation Predictability in CESM2 Detected by Neural Networks. *Geophysical Research Letters*, 48(21), 2021. ISSN 0094-8276. doi: 10.1029/2021gl095392.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman.   *The Elements of Statistical Learning*.   Springer, 2 edition, 3 2009.   ISBN 9780387848570.   doi: 10.1080/01443610062940. URL http://www.springerlink.com/index/D7X7KX6772HQ2135.pdfhttp://www.ncbi.nlm.nih.gov/pubmed/15512507.

Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnoti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sebastien Villaume, and Jean Noël Thépaut. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. ISSN 1477870X. doi: 10.1002/qj.3803.

Nagio Hirota, Yukari N. Takayabu, Masahiro Watanabe, and Masahide Kimoto. Precipitation reproducibility over tropical oceans and its relationship to the double ITCZ problem in CMIP3 and MIROC5 climate models. *Journal of Climate*, 24(18):4859–4873, 2011. ISSN 08948755. doi: 10.1175/2011JCLI4156.1.

William W. Hsieh. Evolution of machine learning in environmental science—A perspective. *Environmental Data Science*, 1:1–8, 2022. doi: 10.1017/eds.2022.2.

K. Kashinath, M. Mustafa, A. Albert, J. L. Wu, C. Jiang, S. Esmaeilzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, A. Manepalli, D. Chirila, R. Yu, R. Walters, B. White, H. Xiao, H. A. Tchelepi, P. Marcus, A. Anandkumar, P. Hassanzadeh, and Prabhat. Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194), 2021. ISSN 1364503X. doi: 10.1098/rsta.2020.0093.

Paul Keil, Thorsten Mauritsen, Johann Jungclaus, Christopher Hedemann, Dirk Olonscheck, and Rohit Ghosh. Multiple drivers of the North Atlantic warming hole. *Nature Climate Change*, 10(7):667–671, 2020. ISSN 17586798. doi: 10.1038/s41558-020-0819-8. URL http://dx.doi.org/10.1038/s41558-020-0819-8.

Reto Knutti. The end of model democracy? *Climatic Change*, 102(3-4):395–404, 10 2010. ISSN 0165-0009. doi: 10.1007/s10584-010-9800-2. URL http://link.springer.com/10.1007/s10584-010-9800-2.

Reto Knutti, David Masson, and Andrew Gettelman. Climate model genealogy: Generation CMIP5 and how we got there. *Geophysical Research Letters*, 40(6):1194–1199, 3 2013. ISSN 00948276.

doi: 10.1002/grl.50256. URL http://doi.wiley.com/10.1002/grl.50256.

Zachary M. Labe and Elizabeth A. Barnes. Detecting Climate Signals Using Explainable AI With Single-Forcing Large Ensembles. *Journal of Advances in Modeling Earth Systems*, 13(6), 2021. ISSN 19422466. doi: 10.1029/2021MS002464.

Johan Lee, Jinwon Jisun Kim, Min Ah Sun, Byeong Hyeon Kim, Hyejin Moon, Hyun Min Sung, Jinwon Jisun Kim, and Young Hwa Byun. Evaluation of the Korea Meteorological Administration Advanced Community Earth-System model (K-ACE). *Asia-Pacific Journal of Atmospheric Sciences*, 56(3):381–395, 2020. ISSN 19767951. doi: 10.1007/s13143-019-00144-7.

Gen Li and Shang Ping Xie. Tropical biases in CMIP5 multimodel ensemble: The excessive equatorial pacific cold tongue and double ITCZ problems. *Journal of Climate*, 27(4):1765–1780, 2014. ISSN 08948755. doi: 10.1175/JCLI-D-13-00337.1.

Hsi Yen Ma, S. Xie, S. A. Klein, K. D. Williams, J. S. Boyle, S. Bony, H. Douville, S. Fermepin, B. Medeiros, S. Tyteca, M. Watanabe, and D. Williamson. On the correspondence between mean forecast errors and climate errors in CMIP5 models. *Journal of Climate*, 27(4):1781–1798, 2014. ISSN 08948755. doi: 10.1175/JCLI-D-13-00474.1.

D. Masson and R. Knutti. Climate model genealogy. *Geophysical Research Letters*, 38(8):1–4, 2011. ISSN 00948276. doi: 10.1029/2011GL046864.

Gerald A. Meehl, Jadwiga H. Richter, Haiyan Teng, Antonietta Capotondi, Kim Cobb, Francisco Doblas-Reyes, Markus G. Donat, Matthew H. England, John C. Fyfe, Weiqing Han, Hyemi Kim, Ben P. Kirtman, Yochanan Kushnir, Nicole S. Lovenduski, Michael E. Mann, William J. Merryfield, Veronica Nieves, Kathy Pegion, Nan Rosenbloom, Sara C. Sanchez, Adam A. Scaife, Doug Smith, Aneesh C. Subramanian, Lantao Sun, Diane Thompson, Caroline C. Ummenhofer, and Shang Ping Xie. Initialized Earth System prediction from subseasonal to decadal timescales. *Nature Reviews Earth and Environment*, 2(5):340–357, 2021. ISSN 2662138X. doi: 10.1038/s43017-021-00155-x.

Malte Meinshausen, Zebedee R. J. Nicholls, Jared Lewis, Matthew J. Gidden, Elisabeth Vogel, Mandy Freund, Urs Beyerle, Claudia Gessner, Alexander Nauels, Nico Bauer, Josep G. Canadell, John S. Daniel, Andrew John, Paul B. Krummel, Gunnar Luderer, Nicolai Meinshausen, Stephen A. Montzka, Peter J. Rayner, Stefan Reimann, Steven J. Smith, Marten van den Berg, Guus J. M. Velders, Martin K. Vollmer, and Ray H. J. Wang. The shared socio-economic pathway (SSP) greenhouse gas concentrations and their extensions to 2500. *Geoscientific Model Development*, 13 (8):3571–3605, 8 2020. ISSN 1991-9603. doi: 10.5194/gmd-13-3571-2020. URL https://gmd. copernicus.org/articles/13/3571/2020/.

Anna Louise Merrifield, Lukas Brunner, Ruth Lorenz, Iselin Medhaug, and Reto Knutti. An investigation of weighting schemes suitable for incorporating large ensembles into multi-model ensembles. *Earth System Dynamics*, 11(3):807–834, 9 2020. ISSN 2190-4987. doi: 10.5194/esd-11-807-2020. URL https://esd.copernicus.org/articles/11/807/2020/.

Thomas Rackow, Nils Wedi, Kristian Mogensen, Peter Dueben, Helge F Goessling, Christian Kühnlein, Lorenzo Zampieri, and Thomas Jung. DYAMOND-II simulations with IFS-FESOM2, 2021.

Sebastian Sippel, Nicolai Meinshausen, Erich M Fischer, Enikő Székely, and Reto Knutti. Climate change now detectable from any single day of weather at global scale. *Nature Climate Change*, 10 (1):35–41, 1 2020. ISSN 1758-678X. doi: 10.1038/s41558-019-0666-7. URL http://dx.doi.org/10. 1038/s41558-019-0666-7http://www.nature.com/articles/s41558-019-0666-7.

Karen Stengel, Andrew Glaws, Dylan Hettinger, and Ryan N. King. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences of the United States of America*, 117(29):16805–16815, 2020. ISSN 10916490. doi: 10.1073/pnas.1918964117.

Claudia Tebaldi and Reto Knutti. The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1857):2053–2075, 2007. ISSN 1364-503X. doi: 10.1098/rsta.2007.2076. URL http://rsta.royalsocietypublishing.org/cgi/doi/10.1098/rsta.2007.2076.

Baijun Tian and Xinyu Dong. The Double-ITCZ Bias in CMIP3, CMIP5, and CMIP6 Models Based on Annual Mean Precipitation. *Geophysical Research Letters*, 47(8):1–11, 2020. ISSN 19448007.

doi: 10.1029/2020GL087232.

Benjamin A. Toms, Elizabeth A. Barnes, and Imme Ebert-Uphoff. Physically Interpretable Neural Networks for the Geosciences: Applications to Earth System Variability. *Journal of Advances in Modeling Earth Systems*, 12(9):1–20, 2020. ISSN 19422466. doi: 10.1029/2019MS002002.

Keith D. Williams, A. Bodas-Salcedo, M. Déqué, S. Fermepin, B. Medeiros, M. Watanabe, C. Jakob, S. A. Klein, C. A. Senior, and D. L. Williamson. The transpose-AMIP II experiment and its application to the understanding of southern ocean cloud biases in climate models. *Journal of Climate*, 26 (10):3258–3274, 2013. ISSN 08948755. doi: 10.1175/JCLI-D-12-00429.1.