9th International Workshop on Climate Informatics
October 02-04, 2019
Hosted by École Normale Supérieure, Paris, France

MACHINE LEARNING PARAMETERIZATIONS FOR OZONE: CLIMATE MODEL TRANSFERABILITY

Peer Nowack[1,2,3], Qing Yee Ellie Ong[4], Peter Braesicke[5], Joanna D. Haigh[1,2], Luke Abraham[6], John A. Pyle[6], and Apostolos Voulgarakis[2]

***Abstract*—Many climate modeling studies have demonstrated the importance of two-way interactions between ozone and atmospheric dynamics. However, atmospheric chemistry models needed for calculating changes in ozone are computationally expensive. Nowack et al. [1] highlighted the potential of machine learning-based ozone parameterizations in constant climate forcing simulations, with ozone being predicted as a function of the atmospheric temperature state. Here we investigate the role of additional time-lagged temperature information under preindustrial forcing conditions. In particular, we test if the use of Long Short-Term Memory (LSTM) neural networks can significantly improve the predictive skill of the parameterization. We then introduce a novel workflow to transfer the regression model to the new UK Earth System Model (UKESM). For this, we show for the first time how machine learning parameterizations could be transferred between climate models, a pivotal step to making any such parameterization widely applicable in climate science. Our results imply that ozone parameterizations could have much-extended scope as they are not bound to individual climate models but, once trained, could be used in a number of different models. We hope to stimulate similar transferability tests regarding machine learning parameterizations developed for other Earth system model components such as ocean eddy modeling, convection, clouds, or carbon cycle schemes.**

I. MOTIVATION

While being a greenhouse gas and air pollutant, ozone is also the only absorber of harmful solar UV-B radiation which would otherwise make life on Earth impossible [2]. Ozone's distribution in the atmosphere is constantly affected by anthropogenic and natural factors, from changes in the stratospheric circulation [3, 4] to chemical reactions [5–8]. Its importance for global radiative transfer in turn induces feedback effects on the Earth system by modulating temperature, dynamics and the biosphere [9–14]. A number of previous studies used machine learning methods to understand and to model factors influencing ozone, e.g. to forecast air quality [15, 16], to model ozone dry deposition [17] and iodide surface emissions [18], or to infer differences among chemistry models [19].

Here we explore the potential to use machine learning-based ozone parameterizations in constant forcing climate model simulations. Specifically, we focus on preindustrial simulations, which are core experiments in climate modeling intercomparsions [20, 21] and which are typically run for centennial to millennial time-scales. Atmospheric chemistry schemes add substantially to the overall high computational costs of such simulations, mainly because they require repeated numerical approximations to the transport equations of dozens of chemical tracers as well as to the large system of coupled chemical rate equations [22]. In preindustrial simulations, ozone's impact on climate is primarily determined by two-way interactions between the variability in ozone and climate. This is particularly true for the stratosphere, where changes in ozone have been found to modulate the Quasi-Biennial Oscillation [QBO, 23] or the polar vortices [24, 25].

Our paper is motivated by results presented in Nowack et al. [hereafter N2018, 1] who showed that, in certain simulations, the global ozone distribution could be well predicted by temperature-based machine learning regression functions. Here, we first revisit those results and test if further time-lagged temperature information improves the parameterization. We then investigate for the first time how such a parameterization could be transferred among climate models.

II. METHOD

As in N2018, we fit regression models that predict daily-mean ozone distributions at timestep t based on the previous day's temperature distribution. This time resolution is sufficient to capture the large-scale behaviour of the relatively slowly-moving stratosphere, where interactions between ozone and climate are par-

Corresponding author: Peer Nowack, p.nowack@imperial.ac.uk
[1]Grantham Institute, Imperial College London, UK. [2]Department of Physics, Imperial College London, UK. [3]Data Science Institute, Imperial College London, UK. [4]Department of Physics, University of Oxford, UK. [5]IMK-ASF, Karlsruhe Institute of Technology, Germany. [6]National Centre for Atmospheric Science and Department of Chemistry, University of Cambridge, UK.

ticularly important [e.g. 23]. More formally, we attempt to predict ozone mixing ratios $\mathbf{Y}$ in every model grid cell k of a global climate model as a function f of the global temperature state $\mathbf{X}$

$$Y_k^{(t)} = f(\mathbf{X}^{(t-1)}) \quad (1)$$

Below we further test the importance of lagged temperature information ($\mathbf{X}^{(t-2)}$, ..., $\mathbf{X}^{(t-\tau_{\max})}$).

Climate model data. The primary preindustrial climate model temperature and ozone data was produced using the HadGEM3-AO model from the UK Met Office [26], coupled to the atmospheric chemistry scheme UKCA [4, 27]. In these preindustrial simulations, atmospheric CO_2 is held at 285 ppmv. We use a continuous 50-year long time slice of daily mean ozone ($\mathbf{Y}_{\text{train}}$) and temperature ($\mathbf{X}_{\text{train}}$) data for training and cross-validation. Predictions (Figure 1) were then made on an independent 13-year long test set ($\mathbf{X}_{\text{test}}$, $\mathbf{Y}_{\text{test}}$). For the second part of this paper where we develop the model transferability workflow, we used 20-year long preindustrial temperature and ozone datasets ($\mathbf{X}_{\text{UKESM}}$, $\mathbf{Y}_{\text{UKESM}}$) produced by the new United Kingdom Earth System Model [UKESM, 28] for the Coupled Model Intercomparison Project Phase 6 [CMIP6, 21]. All UKESM data was bi-linearly interpolated to the HadGEM3-AO horizontal grid and we select identical vertical levels for the two models. All temperature data was pre-processed by removing each grid cell's training data mean $\mu_{\text{train},k}$ and by scaling to approximately unit variance

$$\mathrm{X}^{\text{norm}}_{\text{train,k}} = \frac{\mathrm{X}_{\text{train,k}} - \mu_{\text{train,k}}}{\sigma_{\text{train},k}} \quad (2)$$

$$\mathrm{X}^{\text{norm}}_{\text{test,k}} = \frac{\mathrm{X}_{\text{test,k}} - \mu_{\text{train,k}}}{\sigma_{\text{train},k}} \quad (3)$$

$$\mathrm{X}^{\text{norm}}_{\text{UKESM,k}} = \frac{\mathrm{X}_{\text{UKESM,k}} - \mu_{\text{train,k}}}{\sigma_{\text{train},k}} \quad (4)$$

Regression models and cross-validation. Following N2018, we use Ridge regression as the baseline approach for the temperature-ozone mapping. Ridge regression is a linear least squares regression augmented by L2-regularization to address the bias-variance tradeoff [29]. The cost function

$$J_k = \sum_{t=1}^{N} \left(Y_k^{(t)} - \sum_{j=1}^{p} c_{kj} X_{\text{pca},j}^{(t-1)} \right)^2 + \alpha \sum_{j=1}^{p} c_{kj}^2 \quad (5)$$

is minimized for each model grid cell k over N timesteps. We applied principal component analysis [PCA, 30] to $\mathbf{X}$ to speed up the training procedure. We here retain the first p=1000 components (equalling >95% represented variance; for extensive numerical tests see N2018 Supplementary Figure S1a) as a compromise between model complexity and model performance. Smaller (larger) values of α put weaker (stronger) constraints on the size of the coefficients, thus favoring overfitting (high bias). We use a standard time series cross validation method to find the best value for α, in which the time-ordered training data is split into five subsets of equal size. Preceding subsets are then sequentially used as training data for each subsequent subset (i.e. set 1 for 2, set 1+2 for 3 etc.). α is found according to the average generalization error.

We compare the performance of the Ridge regressions to Long Short-Term Memory (LSTM) neural networks, which can process time-lagged information highly effectively [31, 32]. We tested a range of bias and recurrent regularization parameter values in a non-stateful setting, varied the number of timesteps accounted for in the memory unit (up to 10 days), and also used different network architectures (one vs. two LSTM layers with up to 100 neurons per layer). All network architectures were fitted with ReLU activation functions for the LSTM layers and linear activations directed towards the output layer. In each case, we trained the network for 750 epochs using varying batch sizes (>256). The number of epochs was chosen according to their respective error learning curves as to avoid overfitting. To assert stability, we scaled all inputs/outputs to within (0,1)/(-1,1) range. Only the best settings after cross-validation on a 40-to-10 year training data split are discussed below. For all data pre-processing and regression tasks, we used Python's scikit-learn, keras and tensorflow packages [33, 34].

III. RESULTS

HadGEM3-AO test data results. Figure 1 shows four examples of ozone time series in different areas of the stratosphere[1]: in the tropical upper stratosphere (Figure 1a), where ozone concentrations are mainly determined by local photochemical reactions [5, 35], the tropical lower stratosphere where the longer time-scales of the Brewer-Dobson circulation and QBO pose the primary control mechanisms [Figure 1b; 3, 36], the mid-latitude

[1]There are on the order of 420,000 grid cells in HadGEM3-AO and more than 2,300,000 grid cells in UKESM. In the following, mainly to keep the computational expense of training the regression models in bay and to describe our approach intuitively, we focus our discussion on individual climate model grid cells which were found to be characteristic of the general method's performance in different atmospheric regimes. For global performance metrics using larger sets of grid cells see the extensive numerical results in the Supplementary Material of N2018.

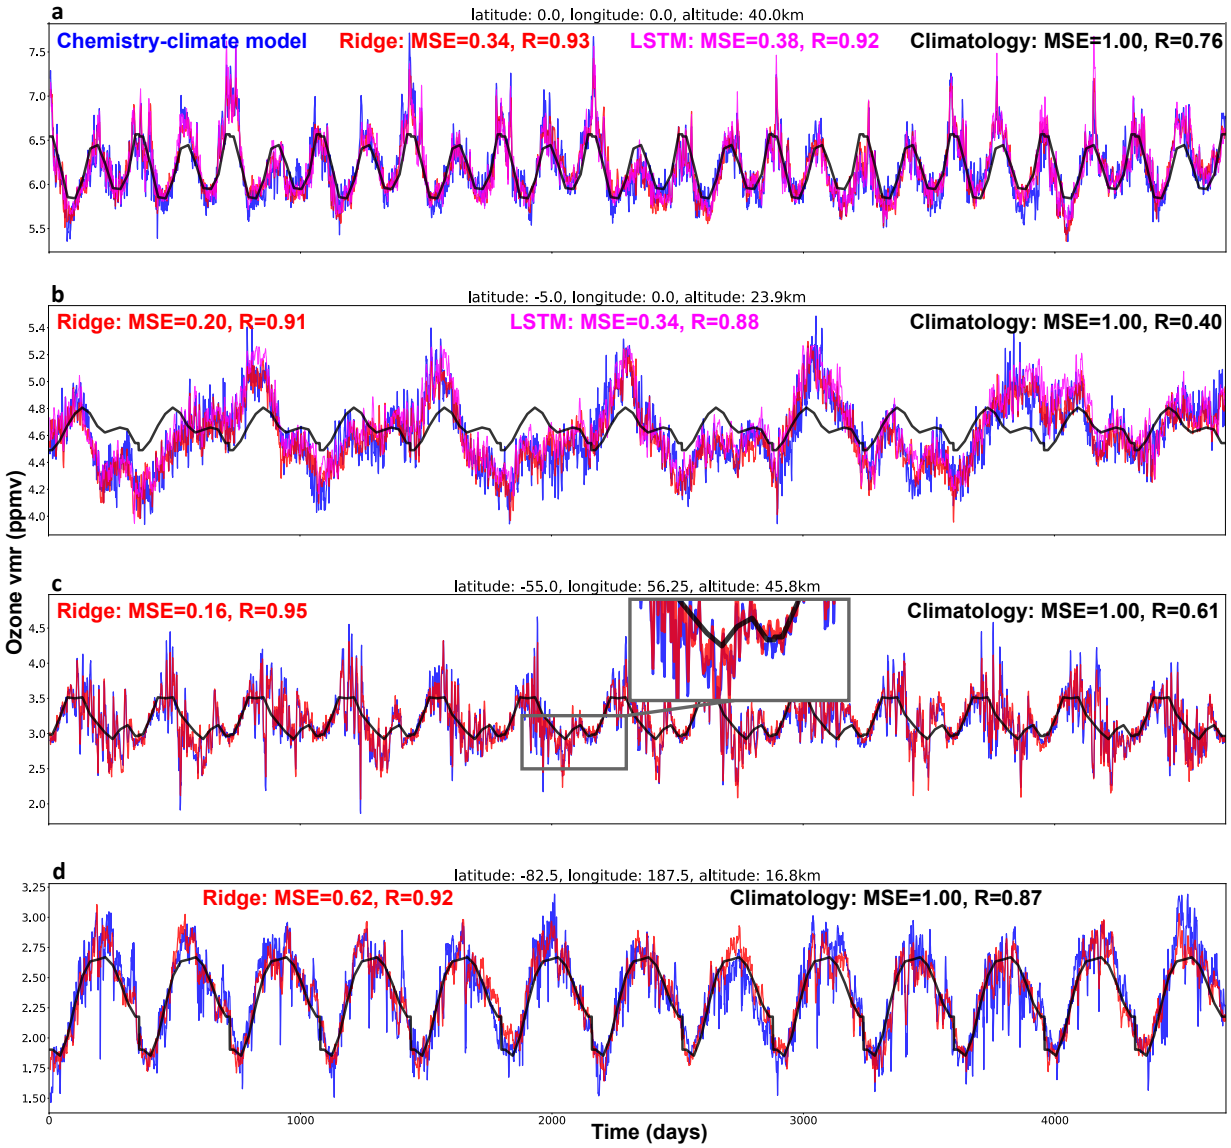

Fig. 1. Ozone time series for four climate model grid cells characteristic of different atmospheric regimes. (Black) Preindustrial climatology. (Blue) Ozone as simulated by the fully interactive chemistry-climate model. (Red) Ridge regression predictions. In the top two panels, we additionally show the results of LSTM predictions (magenta). Grid coordinates are as labeled. R is the Pearson correlation coefficient between the chemistry-climate model time series and the predictions. *MSE* is the ratio of mean squared errors of machine learning predictions relative to the chemistry-climate model data, divided by the same error for the climatologies (values <1 imply improvement).

upper stratosphere where seasonal wave-breaking is important [Figure 1c; 3] and the Southern Hemisphere lower stratosphere where the periodic break-down of the polar vortex is a key feature [Figure 1d; 8]. For each grid cell, we show the actual chemistry-climate model data (blue lines), the corresponding predictions using Ridge regressions (red) and seasonal ozone climatologies (black). For the first two grid cells we also show LSTM predictions (magenta), which here keep a temperature memory of up to five days. Intuitively, LSTMs could be useful as not only the present state of the atmosphere is important for ozone's distribution but also its history, e.g. in the lower stratosphere where ozone's lifetime is much longer than one day. Ozone climatologies (black) are frequently used in climate model simulations and here represent 50-year monthly-mean averages of $\mathbf{Y}_{\text{train}}$, which were subsequently linearly interpolated to daily time resolution. The climatology serves as a benchmark for a classic treatment of ozone in preindustrial simulations without interactive atmospheric chemistry. Due to the limited predictability of short-term ozone fluctuations one day in advance, see e.g. zoomed-in area in Figure 1c, we are interested in ozone predictions that broadly represent the state of ozone at any given time relative to this benchmark (see also discussion in N2018). We evaluate the predictive

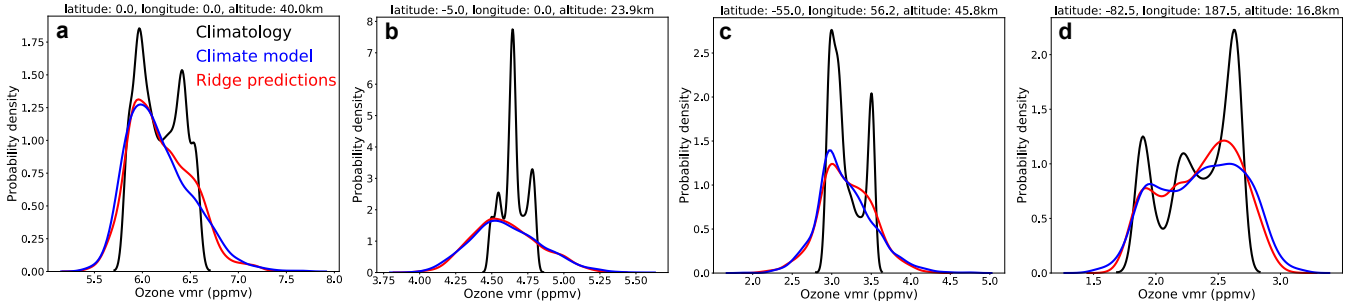

Fig. 2. Probability (kernel) density estimates for ozone volume mixing ratios in each of the four grid cells also shown in Figure 1.

skill relative to these climatologies through Pearson correlation coefficients (R) and mean squared error (MSE) ratios (details in caption of Figure 1), when matched against the actual chemistry-climate model data. These statistical results are shown inside each panel of Figure 1.

The Ridge regressions provide a good approximation to the true chemistry-climate model results, as is also evident from their respective ozone mixing ratio density estimates (Figure 2). Both MSE and R are improved in every grid cell relative to a model-consistent preindustrial climatology; in particular in the tropical mid-stratosphere where the QBO dominates and ozone variability is not captured well by an annual climatology (Figure 1b). Interestingly, the LSTM regressions do not provide a significant improvement over the Ridge regression approach despite their greater functional complexity and longer memory of the temperature history. Further performance gains might be achieved by more extensive parameter and network architecture tuning. However, the general result underlines previous findings in N2018, where other non-linear algorithms were also not found to perform better than Ridge regression. To further validate this result, we also carried out 5-day-lagged window Ridge regressions (not shown). We thus conclude that the single time-step Ridge regression is a solid method on this particular learning task and turn to the question of model transferability.

Model transferability. There are two main obstacles to transferring the parameterization from one climate model to another: differences in (a) the temperature fields and (b) the ozone fields. For example, even a constant background temperature difference, as will always occur among climate models [37], can seriously interfere with the regression task as the inputs may, for example, constantly take on extreme values relative to the original model's temperature distribution. This in turn renders the machine learning parameterization unable to make realistic ozone predictions.

To mitigate such effects, we found that the following intuitive re-calibration procedure led to good results: after standardization of the UKESM temperature data according to equation (4), the resulting $\mathbf{X}^{\text{norm}}_{\text{UKESM,k}}$ will still have average values in each grid cell significantly different from nil, simply due to the aforementioned time-average discrepancies in the background temperature state. We thus re-calibrated the temperature data to an approximately zero mean by subtracting the average offset over the first n number of years ($\overline{\mathbf{X}}^{\text{norm}}_{\text{UKESM,k}}$)

$$\mathbf{X}^{\text{adjusted}}_{\text{UKESM}} = \mathbf{X}^{\text{norm}}_{\text{UKESM}} - \overline{\mathbf{X}}^{\text{norm}}_{\text{UKESM}} \tag{6}$$

Empirically, we found that calibrations using the first n=5 or 10 years yielded almost identical results. In addition, we applied a re-calibration procedure to the predicted ozone field. We separate the ozone mixing ratios predicted by the regression (trained on HadGEM3-AO data) when provided with the re-calibrated temperature input from the UKESM model ($\mathbf{X}^{\text{adjusted}}_{\text{UKESM}}$) into a climatological plus a variability term for each cell

$$Y_{\text{HadGEM-consistent,k}} = Y_{\text{HadGEM-clim,k}} + Y_{\text{variability,k}} \tag{7}$$

We then use again n years of UKESM data to approximate a corresponding climatological term $Y_{\text{UKESM-clim}}$ which we use to replace the HadGEM climatology

$$Y_{\text{UKESM-consistent,k}} = Y_{\text{UKESM-clim,k}} + Y_{\text{variability,k}} \tag{8}$$

Figure 3 shows the results of the corresponding regressions (red) for three grid cells over the last ten years of the UKESM dataset, located in the tropical lower and mid-stratosphere as well as in the Northern Hemisphere polar stratospheric region. In addition, we show ozone climatologies from HadGEM3-AO (gray) and calculated from the first 5 years of the UKESM data (black). As before, we use the MSE ratio and correlation coefficients to compare the predictions to the actual UKESM interactive chemistry climate model results (blue). The quantitative results are given directly in each panel of Figure 3. The transferred Ridge regression performs typically far better than the climatologies

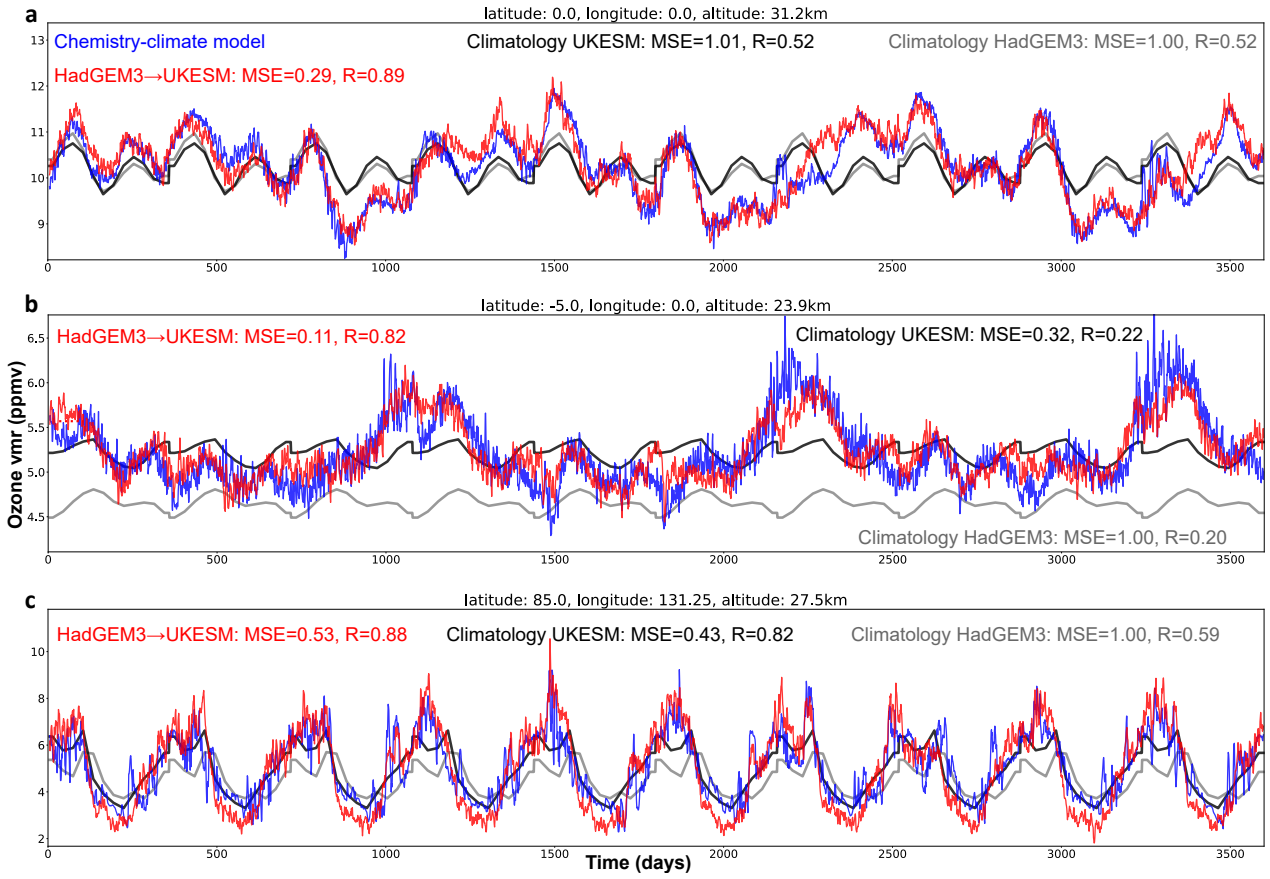

Fig. 3. As Figure 1, but for ten years of UKESM data using the re-calibrated Ridge regressions for predictions (red). To calculate the MSE ratio, we choose the errors when using the HadGEM3 climatology averaged over 50 years as the baseline (as in Figure 1).

both in terms of MSE and R. In terms of the MSE ratio, the HadGEM3-AO climatology performs particularly poorly in regions with major differences in background ozone levels between the two climate models (Figure 3b). Finally, note that a better than climatological performance is not guaranteed in this setting. For example in Figure 3c, where a UKESM-consistent climatology already achieves R-values greater than 0.8, the MSE is worse for the transferred regression, mainly due to constantly underestimated annual minimum values using the Ridge regression. Such a feature may even occur if we just used another atmospheric chemistry scheme. Its annually repeating consistency could also reflect changes in the underlying transport time-scales (i.e. dynamics) between the two models.

IV. DISCUSSION

We have presented two important extensions to temperature-based machine learning parameterizations:

1) Using LSTMs, which here took into account past temperature state information of up to 5-10 days, we tested for improved predictive skill. However, the results are comparable to those obtained with computationally cheaper Ridge regressions.

2) An intuitive re-calibration procedure to learn the ozone parameterization from data produced by one climate model (HadGEM3-AO) to then predict ozone values for another climate model (UKESM). The re-calibration is required for the temperature inputs, whereas the ozone correction can be chosen more flexibly, e.g. to adjust the field to a climatology of the modeler's choice. For example, in some cases there might be no ground truth ozone state (e.g. in models without interactive chemistry option) so that one might either choose ozone values consistent with HadGEM3-AO, or another background climatological field. Our method could therefore in principle be used to remove some persistent ozone model biases [see e.g. 38]. In conclusion, the method outlined here could be used to learn regression functions representing ozone variability from long existing chemistry-climate model datasets. Following re-calibration, the ozone parameterization could then be used in long climate model simulations such as under constant preindustrial forcing.

ACKNOWLEDGMENTS

Peer Nowack is supported through an Imperial College Research Fellowship. Qing Yee Ellie Ong's work on the transferability of the ozone parameterization was funded by the Laidlaw Scholars Undergraduate Research and Leadership Programme.

REFERENCES

[1] P. Nowack, P. Braesicke, J. Haigh, N. L. Abraham, J. Pyle, and A. Voulgarakis, "Using machine learning to build temperature-based ozone parameterizations for climate sensitivity simulations," *Environmental Research Letters*, vol. 13, no. 10, p. 104016, 2018.

[2] WMO, "World Meteorological Organization: Scientific assessment of ozone depletion: 2010, Global Ozone Research and Monitoring Project - Report No. 52, 516. pp, Geneva, Switzerland," 2011.

[3] SPARC, "SPARC CCMVal Report on the Evaluation of Chemistry-Climate Models," tech. rep., SPARC, 2010.

[4] P. J. Nowack, N. Luke Abraham, A. C. Maycock, P. Braesicke, J. M. Gregory, M. M. Joshi, A. Osprey, and J. A. Pyle, "A large ozone-circulation feedback and its implications for global warming assessments," *Nature Climate Change*, vol. 5, no. 1, pp. 41–45, 2015.

[5] J. D. Haigh and J. A. Pyle, "Ozone perturbation experiments in a two-dimensional circulation model," *Quart. J. R. Met. Soc.*, vol. 108, pp. 551–574, 1982.

[6] A. I. Jonsson, J. de Grandpré, V. I. Fomichev, J. C. McConnell, and S. R. Beagley, "Doubled CO2-induced cooling in the middle atmosphere: Photochemical analysis of the ozone radiative feedback," *Journal of Geophysical Research*, vol. 109, p. D24103, 2004.

[7] A. Voulgarakis, V. Naik, J. F. Lamarque, D. T. Shindell, P. J. Young, M. J. Prather, O. Wild, R. D. Field, D. Bergmann, P. Cameron-Smith, I. Cionni, W. J. Collins, S. B. Dalsøren, R. M. Doherty, V. Eyring, G. Faluvegi, G. A. Folberth, L. W. Horowitz, B. Josse, I. A. MacKenzie, T. Nagashima, D. A. Plummer, M. Righi, S. T. Rumbold, D. S. Stevenson, S. A. Strode, K. Sudo, S. Szopa, and G. Zeng, "Analysis of present day and future OH and methane lifetime in the ACCMIP simulations," *Atmospheric Chemistry and Physics*, vol. 13, no. 5, pp. 2563–2587, 2013.

[8] S. Solomon, D. J. Ivy, D. Kinnison, M. J. Mills, R. R. Neely, and A. Schmidt, "Emergence of healing in the Antarctic ozone layer," *Science*, vol. 353, no. 6296, pp. 269–274, 2016.

[9] S. Sitch, P. M. Cox, W. J. Collins, and C. Huntingford, "Indirect radiative forcing of climate change through ozone effects on the land-carbon sink.," *Nature*, vol. 448, no. 7155, pp. 791–794, 2007.

[10] C. Williamson and R. Zepp, "Solar ultraviolet radiation in a changing climate," *Nature Climate Change*, vol. 4, pp. 434–441, 2014.

[11] P. J. Nowack, "Cirrus and the Earth system," *Weather*, vol. 70, no. 11, p. 330, 2015.

[12] P. J. Nowack, N. L. Abraham, P. Braesicke, and J. A. Pyle, "Stratospheric ozone changes under solar geoengineering: implications for UV exposure and air quality," *Atmospheric Chemistry and Physics*, vol. 16, pp. 4191–4203, 2016.

[13] P. J. Nowack, P. Braesicke, N. L. Abraham, and J. A. Pyle, "On the role of ozone feedback in the ENSO amplitude response under global warming," *Geophysical Research Letters*, vol. 44, pp. 3858–3866, 2017.

[14] P. J. Nowack, N. L. Abraham, P. Braesicke, and J. A. Pyle, "The impact of stratospheric ozone feedbacks on climate sensitivity estimates," *Journal of Geophysical Research: Atmospheres*, vol. 123, pp. 4630–4641, 2018.

[15] V. Mallet and B. Sportisse, "Ensemble-based air quality forecasts: A multimodel approach applied to ozone," *Journal of Geophysical Research: Atmospheres*, vol. 111, p. D18302, 2006.

[16] C. A. Keller and M. J. Evans, "Application of random forest regression to the calculation of gas-phase chemistry within the GEOS-Chem chemistry model v10," *Geoscientific Model Development*, vol. 12, pp. 1209–1225, 2019.

[17] S. J. Silva, C. L. Heald, S. Ravela, I. Mammarella, and J. W. Munger, "A Deep Learning Parameterization for Ozone Dry Deposition Velocities," *Geophysical Research Letters*, vol. 46, no. 2, pp. 983–989, 2019.

[18] T. Sherwen, R. J. Chance, L. Tinel, D. Ellis, M. J. Evans, and L. J. Carpenter, "A machine learning based global sea-surface iodide distribution," *Earth System Science Data*, vol. 11, pp. 1239–1262, 2019.

[19] J. M. Nicely, R. J. Salawitch, T. Canty, D. C. Anderson, S. R. Arnold, M. P. Chipperfield, L. K. Emmons, J. Flemming, V. Huijnen, D. E. Kinnison, J. F. Lamarque, J. Mao, S. A. Monks, S. D. Steenrod, S. Tilmes, and S. Turquety, "Quantifying the causes of differences in tropospheric OH within global models," *Journal of Geophysical Research*, vol. 122, no. 3, pp. 1983–2007, 2017.

[20] B. Kravitz, A. Robock, O. Boucher, H. Schmidt, K. E. Taylor, G. Stenchikov, and M. Schulz, "The Geoengineering Model Intercomparison Project (GeoMIP)," *Atmospheric Science Letters*, vol. 12, no. 2, pp. 162–167, 2011.

[21] V. Eyring, S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor, "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization," *Geoscientific Model Development*, vol. 9, no. 5, pp. 1937–1958, 2016.

[22] E. Esentürk, L. Abraham, S. Archer-Nicholls, C. Mitsakou, P. Griffiths, and J. Pyle, "Quasi-Newton Methods for Atmospheric Chemistry Simulations : Implementation in UKCA UM Vn10 . 8," *Geoscientific Model Development*, vol. 11, no. February, pp. 3089–3108, 2018.

[23] D. Rind, J. Jonas, N. K. Balachandran, G. A. Schmidt, and J. Lean, "The QBO in two GISS global climate models: 1. Generation of the QBO," *Journal of Geophysical Research: Atmospheres*, vol. 119, no. 14, pp. 8798–8824, 2014.

[24] V. Silverman, N. Harnik, K. Matthes, S. W. Lubis, and S. Wahl, "Radiative effects of ozone waves on the Northern Hemisphere polar vortex and its modulation by the QBO," *Atmospheric Chemistry and Physics*, vol. 18, pp. 6637–6659, 2018.

[25] S. Haase and K. Matthes, "The importance of interactive chemistry for stratosphere-troposphere coupling," *Atmospheric Chemistry and Physics*, vol. 19, no. 5, pp. 3417–3432, 2019.

[26] H. T. Hewitt, D. Copsey, I. D. Culverwell, C. M. Harris, R. S. R. Hill, A. B. Keen, A. J. McLaren, and E. C. Hunke, "Design and implementation of the infrastructure of HadGEM3: The next-generation Met Office climate modelling system," *Geoscientific Model Development*, vol. 4, no. 2, pp. 223–253, 2011.

[27] O. Morgenstern, P. Braesicke, F. M. O'Connor, A. C. Bushell, C. E. Johnson, S. M. Osprey, and J. A. Pyle, "Evaluation of the new UKCA climate-composition model Part 1: The stratosphere," *Geoscientific Model Development*, vol. 2, no. 1, pp. 43–57, 2009.

[28] A. A. Sellar et al., "UKESM1: Description and evaluation of the UK Earth System Model," *Submitted to Journal of Advances in Modeling Earth Systems*, 2019.

[29] A. E. Hoerl and R. W. Kennard, "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, vol. 12, no. 1, pp. 55–67, 1970.

[30] C. M. Bishop, *Pattern recognition and machine learning*. Singapore: Springer Science+Business Media, 2006.

[31] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[32] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[34] A. Martin et al., "TensorFlow: A System for Large-Scale Machine Learning," 2015.

[35] J. A. Pyle, "A calculation of the possible depletion of ozone by chlorofluorocarbons using a two-dimensional model," *Pure and Applied Geophysics PAGEOPH*, vol. 118, pp. 355–377, 1980.

[36] R. A. Plumb, "Stratospheric Transport," *Journal of the Meteorological Society of Japan*, vol. 80, pp. 793–809, 2002.

[37] G. Flato et al., "Evaluation of Climate Models," in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ch. 9, pp. 741–866, Cambridge, UK and New York, USA: Cambridge University Press, 2013.

[38] S. C. Hardiman et al., "Processes controlling tropical tropopause temperature and stratospheric water vapor in climate models," *Journal of Climate*, vol. 28, pp. 6516–6535, 2015.