Using machine learning to build temperature-based ozone parameterizations for climate sensitivity simulations

Environmental Research Letters

LETTER

Using machine learning to build temperature-based ozone parameterizations for climate sensitivity simulations

Peer Nowack[1,2,3], **Peter Braesicke**[4], **Joanna Haigh**[1,2], **Nathan Luke Abraham**[5,6], **John Pyle**[5,6] and **Apostolos Voulgarakis**[2]

[1] Grantham Institute for Climate Change, Imperial College London, United Kingdom
[2] Department of Physics, Blackett Laboratory, Imperial College London, United Kingdom
[3] Data Science Institute, Imperial College London, United Kingdom
[4] IMK-ASF, Karlsruhe Institute of Technology, Germany
[5] National Centre for Atmospheric Science, United Kingdom
[6] Department of Chemistry, University of Cambridge, United Kingdom

E-mail: p.nowack@imperial.ac.uk

Keywords: climate change, climate sensitivity, ozone, parameterization, machine learning, big data, climate modeling

Supplementary material for this article is available online

OPEN ACCESS

RECEIVED
29 May 2018

REVISED
28 August 2018

ACCEPTED FOR PUBLICATION
14 September 2018

PUBLISHED
9 October 2018

Abstract

A number of studies have demonstrated the importance of ozone in climate change simulations, for example concerning global warming projections and atmospheric dynamics. However, fully interactive atmospheric chemistry schemes needed for calculating changes in ozone are computationally expensive. Climate modelers therefore often use climatological ozone fields, which are typically neither consistent with the actual climate state simulated by each model nor with the specific climate change scenario. This limitation applies in particular to standard modeling experiments such as preindustrial control or abrupt $4xCO_2$ climate sensitivity simulations. Here we suggest a novel method using a simple linear machine learning regression algorithm to predict ozone distributions for preindustrial and abrupt $4xCO_2$ simulations. Using the atmospheric temperature field as the only input, the regression reliably predicts three-dimensional ozone distributions at monthly to daily time intervals. In particular, the representation of stratospheric ozone variability is much improved compared with a fixed climatology, which is important for interactions with dynamical phenomena such as the polar vortices and the Quasi-Biennial Oscillation. Our method requires training data covering only a fraction of the usual length of simulations and thus promises to be an important stepping stone towards a range of new computationally efficient methods to consider ozone changes in long climate simulations. We highlight key development steps to further improve and extend the scope of machine learning-based ozone parameterizations.

1. Introduction

The trace gas ozone plays multiple roles in the Earth system. Besides being an important greenhouse gas, it is the only absorber of harmful solar UV-B radiation which would otherwise make life on Earth impossible (WMO 2011). However, ozone's distribution in the atmosphere is subject to change. Anthropogenic and natural factors force variability and trends in its concentrations, mainly related to catalytic ozone depletion cycles (Hunt 1966, Crutzen 1970, Molina and Rowland 1974, Solomon *et al* 2016), the stratospheric circulation (SPARC 2010, Eyring *et al* 2013), background temperature (Fels *et al* 1980, Haigh and Pyle 1982), the solar cycle (Haigh 1996, Ball *et al* 2016) as well as changes in two-way interactions with other constituents (Young *et al* 2013, Voulgarakis *et al* 2013). Ozone's importance for global radiative transfer, in turn, implies feedback effects on the Earth system by modulating temperature, dynamics and the biosphere (Lacis *et al* 1990, Thompson and Solomon 2002, Son *et al* 2008, Williamson *et al* 2014). Here we introduce a novel method for representing the time evolution of ozone (or *ozone feedbacks*) in state-of-the-art

Earth system models. For this *ozone parameterization* we focus specifically on preindustrial control and abrupt 4xCO_2 climate simulations where ozone feedbacks are typically not included. Our work is timely due to a number of recent studies that have demonstrated the importance of ozone in this context.

Preindustrial control and abrupt 4xCO_2 simulations are core experiments in many climate modeling intercomparsion projects (Taylor *et al* 2012, Kravitz *et al* 2013, Eyring *et al* 2016). Preindustrial simulations help us understand climate variability in the absence of human influences (e.g. Bellenger *et al* 2013). Abrupt 4xCO_2 simulations, in which atmospheric carbon dioxide (CO_2) mixing ratios are quadrupled starting from preindustrial conditions, are most frequently used to study climate sensitivity (e.g. Ceppi and Gregory 2017, Knutti *et al* 2017); in the broadest sense the global warming response to increased CO_2. They also provide insights on how climate dynamics might change in high-CO_2 climates, for example in the form of major climate modes (Guilyardi *et al* 2012, Rashid *et al* 2016).

Many studies have highlighted that the representation of ozone can impact the results of such climate sensitivity simulations. The use of non-adaptive ozone climatologies in abrupt 4xCO_2 simulations can affect climate sensitivity estimates (Li *et al* 2013, Muthers *et al* 2014, Dietmüller *et al* 2014, Nowack *et al* 2015, 2018), the position of the jet streams (Chiodo and Polvani 2017, Nowack *et al* 2018) as well as the response of the Walker circulation and ENSO (Nowack *et al* 2017). Similar effects have been found for paleo-climates and solar forcing scenarios (Haigh 1996, Heinemann 2009, Chiodo and Polvani 2016, Muthers *et al* 2016, Noda *et al* 2017). Already the use of zonally averaged ozone climatologies, even if otherwise consistent with the climate state, can affect atmospheric variability due to the lack of ozone-dynamics interactions (Gabriel *et al* 2007, Crook *et al* 2008, Gillett *et al* 2009, McCormack *et al* 2011, Albers and Nathan 2012, Rind *et al* 2014, Nowack *et al* 2018, Silverman *et al* 2018).

A key issue is the high computational cost associated with calculating ozone feedbacks. For example, the climate model used here is slowed down by a factor of around three by including a comprehensive atmospheric chemistry component. While this number will depend on the chemistry scheme itself and the underlying climate model, the slow-down is expected to be substantial for any similar case. There are two main contributors to this slow-down (Esentürk *et al* 2018). Firstly, atmospheric chemistry is represented by a large system (>100) of coupled partial differential chemical rate equations. Secondly, chemical species are subject to transport as part of the atmospheric circulation. Representing either requires significant computing power to numerically solve the respective chemical and dynamical equations at each timestep.

The chemistry-climate modeling community has thus provided an ozone dataset for the Representative Concentration Pathways scenarios defined for the Coupled Model Intercomparison Project Phase 5 (CMIP5; Cionni *et al* 2011, Eyring *et al* 2013), which has been widely used in climate models that lack fully interactive chemistry schemes. While being an improvement over neglecting changes in ozone altogether (Son *et al* 2008, 2010), these climatologies were neither consistent with individual model responses nor the actual forcing scenario. In comparison, there has been no coordinated effort to provide similar, or even better datasets, for preindustrial control or abrupt 4xCO_2 simulations. Consequently, it is imperative to test alternative, computationally less expensive approaches with the specific application to preindustrial and 4xCO_2 simulations in mind.

Here, we show that ozone parameterizations using machine learning (ML) could be a powerful alternative to fully interactive chemistry schemes for climate studies that do not specifically focus on atmospheric chemical mechanisms. ML is a subset of artificial intelligence and the two terms are colloquially often used interchangeably. More specifically, ML describes a collection of statistical methods that allow computers to learn relationships from data without being explicitly programmed (e.g. Samuel 1959). As such, ML has the potential to advance scientific understanding and modeling efforts in climate science (Monteleoni *et al* 2013). For example, ML has been used to study coupling mechanisms in the climate system (Boers *et al* 2014, Runge *et al* 2015), for statistical downscaling in regional climate studies (Vandal *et al* 2017, Anderson and Lucas 2018) and has been suggested as a new tool to parameterize convection and clouds (Schneider *et al* 2017, Gentine *et al* 2018).

To model ozone, the resulting parameterization requires atmospheric temperature as the only input variable, which will be available from any climate model simulation, thus putting virtually no constraint on the general applicability of the method. In section 2, we describe our method, the climate model and its data output used to fit the ML model. In section 3, we first demonstrate that our parameterization can well predict the general ozone climatology and trends. We then show that it does also capture large parts of variability in the ozone distribution, in particular in the stratosphere. Finally, we discuss our results as well as future challenges in section 4.

2. Methods

2.1. General approach

Our goal is to predict the distribution of ozone without calculating the underlying system of coupled partial differential chemical rate equations (Pyle 1980, Jonsson *et al* 2004). Ideally, we aim to only use variables in this parameterization that are readily available in climate models without atmospheric chemistry component. As demonstrated below and discussed in section 4, variations and trends in atmospheric

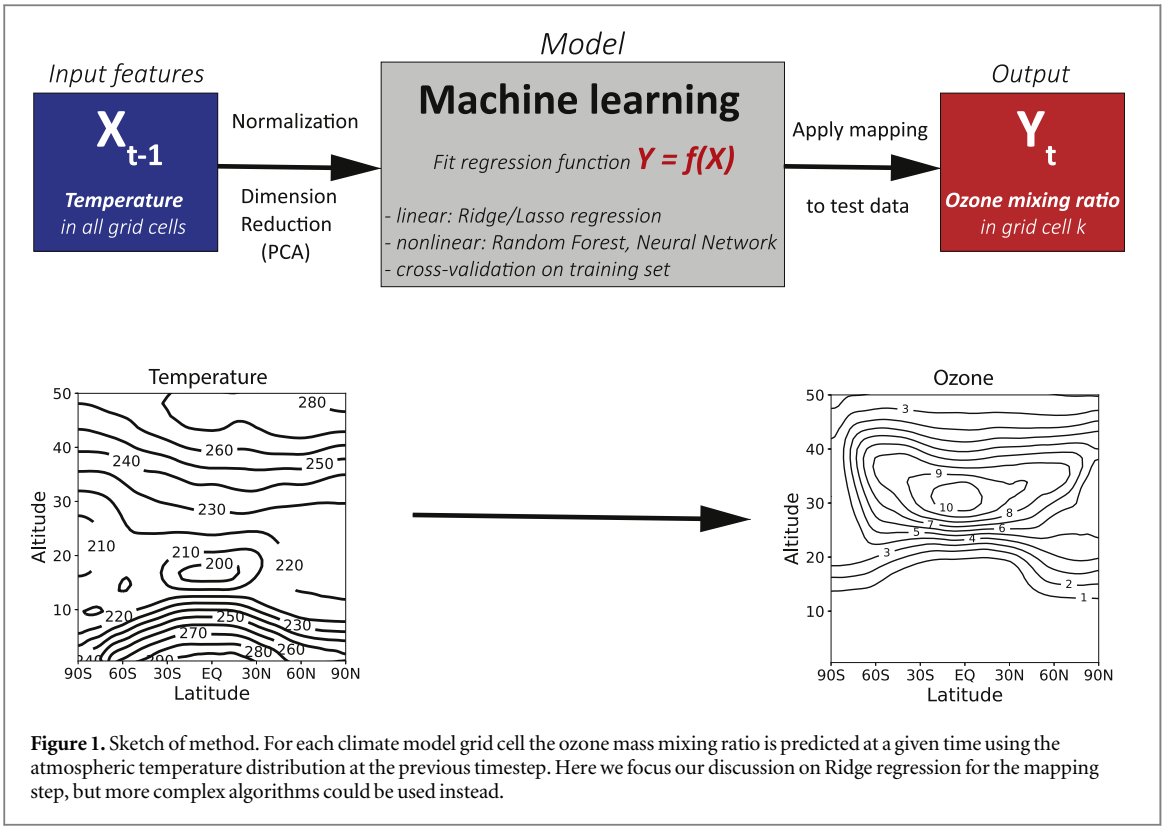

Figure 1. Sketch of method. For each climate model grid cell the ozone mass mixing ratio is predicted at a given time using the atmospheric temperature distribution at the previous timestep. Here we focus our discussion on Ridge regression for the mapping step, but more complex algorithms could be used instead.

temperatures are reflective of processes affecting the production, depletion and transport of ozone and can thus be used as a good predictor for ozone trends and variability, in particular in the slowly-evolving stratosphere.

Figure 1 shows the general work-flow of the method suggested here. We approach this task as a typical multi-output regression problem, with the predicted variables being ozone mass mixing ratios at timestep t in *each* grid cell and the input features being the modeled temperature at the previous timestep in *all* grid cells, giving rise to individual regressions for each grid cell. The time resolution (t versus $t-1$) is a question of choice; we tested climatologically relevant time-scales of daily-mean to monthly-mean ozone. For the regression itself, we tested a number of linear and non-linear ML methods from Ridge regression to neural networks (see supplementary materials stacks.iop.org/erl/13/104016/mmedia). We found that the relationship between temperatures and ozone is highly linear, thus making the linear Ridge regression a computationally and conceptually favorable approach. As shown below, this applies in particular to the 4xCO_2 scenario, where a linear method allows for extrapolation outside the training domain. In order to speed up the algorithm, we apply principal component analysis (PCA, Bishop 2006) to the temperature input, thereby reducing the number of input regression variables under the condition of maintaining >98% of variance. The PCA dimension reduction significantly reduces the number of input variables on which the algorithm has to be trained, which was essential in order for us being able to carry out the large number of tests presented in this paper and its supplementary materials. We found the loss in predictive skill due to using PCA to be very small as compared to using the entire temperature field as input, whereas the time needed to train the algorithm decreased by several orders of magnitude. However, once the model is fit to the data, the prediction of ozone fields for new temperature input is fast with or without prior PCA step. A final version of the algorithm presented here, for example to be implemented in a climate model, might as well be trained directly on the entire temperature field.

2.2. Climate model and data

The data for this study was produced using the atmosphere-ocean coupled configuration of the Hadley Centre Global Environment Model version 3 (HadGEM3-AO) from the UK Met Office (Hewitt *et al* 2011), coupled to the United Kingdom Chemistry and Aerosols atmospheric chemistry scheme (Morgenstern *et al* 2009, Nowack *et al* 2015). The model, which has a self-contained Quasi-Biennial-Oscillation (QBO), is described in detail in supplementary section 1.

We use data from two types of simulations: (1) a preindustrial simulation with atmospheric CO_2 at 285 ppmv that has been spun up for >1000 years. (2) an abrupt 4xCO_2 simulation with CO_2 quadrupled to 1140 ppmv, which was initialized from this preindustrial run. As shown in figure 2(a), we observe the characteristic rapid increase in global mean surface

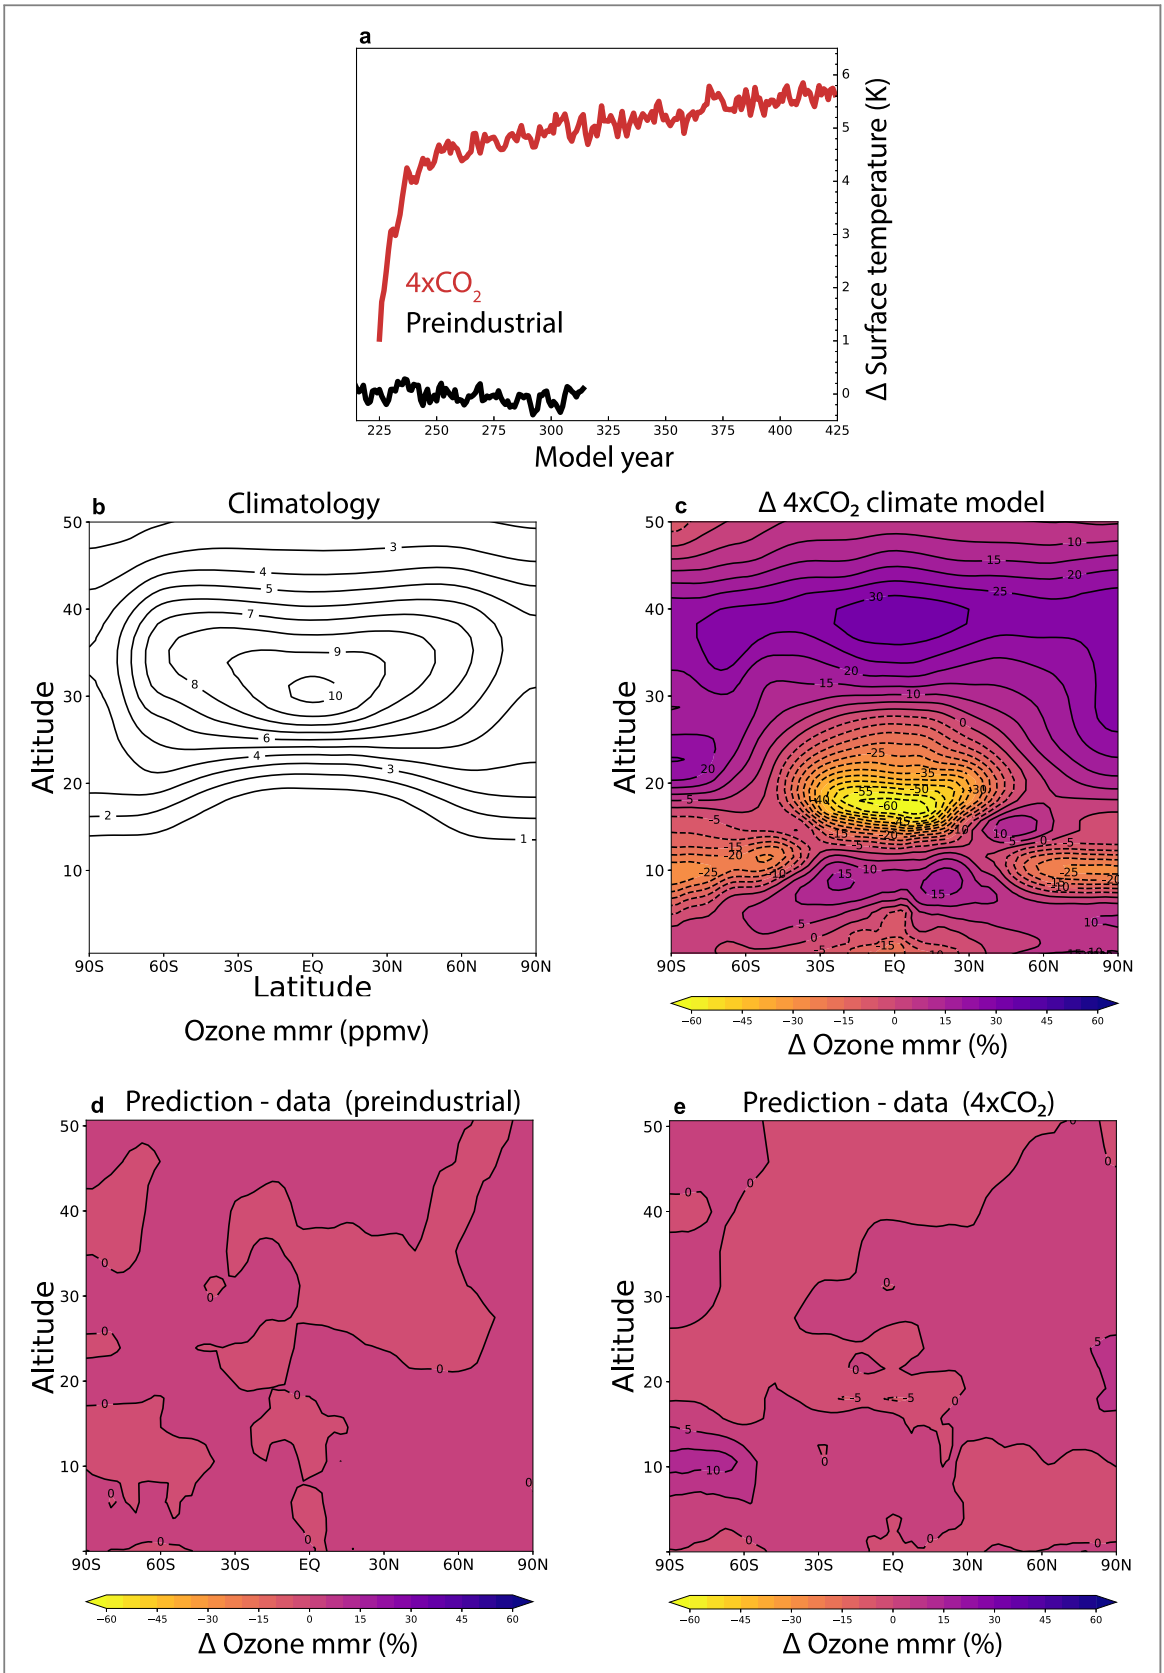

Figure 2. (a) Global mean surface temperature anomalies for the preindustrial control simulation and following an abrupt quadrupling of atmospheric CO_2. The time axis is chosen to conform with two previous publications (Nowack *et al* 2015, 2018). (b) 50-year average zonal mean ozone distribution (ppmv) for the preindustrial control run. (c) Percentage changes in ozone under $4xCO_2$, showing circulation-driven decreases in the tropical lower stratosphere and temperature-induced increases in the upper stratosphere. (d), (e) Percentage differences between the ozone climatology predicted by the ML model and simulated by the chemistry-climate model under preindustrial and $4xCO_2$ conditions, respectively.

temperature following the abrupt increase in CO_2 (Andrews *et al* 2012, Nowack *et al* 2015, 2016, 2018).

In order to train the ML model, the data was split into training and test datasets (Bishop 2006). The training datasets were used to optimize the model concerning bias and overfitting (section 2.3), whereas the independent test datasets served to evaluate the learned model. The 4xCO_2 simulation was run for 200 years in total (figure 2(a)). In section 3, we explore various ways of splitting this dataset into training and test datasets. For the preindustrial run, we randomly picked a 50-year-long interval for training and another 13-year-long independent extract for testing. For all regressions we used daily-mean model output for temperature and ozone mass mixing ratios, for the former on pressure levels, for the latter on model altitude levels (measured in kilometres above the surface). We selected a subset of the actual model vertical resolution (14 pressure, 22 altitude levels) covering the entire troposphere and stratosphere. Choosing vertical subsets allowed us to run a larger number of tests, without impacting our general results or conclusions.

2.3. Regression model

We use Ridge regression, also referred to as Tikhonov regularization, for the mapping between the temperature state of the atmosphere and ozone's distribution. Ridge regression is a linear least squares regression augmented by L2-regularization to address the bias-variance trade-off (Hoerl and Kennard 1970, James *et al* 2013). For the regression, we minimize the cost function

$$J_{\mathrm{Ridge,k}} = \sum_{t=1}^{N}\left(Y_k^{(t)} - \sum_{j=1}^{p} c_{kj} X_j^{(t-1)}\right)^2 + \alpha \sum_{j=1}^{p} c_{kj}^2 \quad (1)$$

for each model grid cell k over N timesteps t given time series training data for ozone (Y) and the first p principal component scores of temperature variability (X), with the latter preceding the former by one timestep in order to predict ozone. The ozone prediction for each grid cell is thus (potentially) a function of temperature variability in all grid cells, as characterized by the time series weights of the leading p principal components of temperature variability. The cost function determines the optimization goal. Its first term is the ordinary least squares regression error, the second penalty term avoids overfitting by nudging the regression towards small regression coefficients c_{kj}. In general, 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), which is referred to as validation set. α is found according to the average generalization error on these validation sets. Prior to the regression and PCA step, both X and Y were preprocessed by removing each grid cell's mean and by scaling to unit variance. Therefore, the final ozone predictions for Y on a separate testset t'

$$Y_k^{(t')} = \sum_{j=1}^{p} c_{kj} X_j^{(t'-1)} \quad (2)$$

have to be re-scaled to mass mixing ratios. The number of components chosen for X and the optimized regularization parameters α for each fit are documented in supplementary section 2. For all regression, normalization and dimension reduction tasks, we used Python's scikit-learn package (Pedregosa *et al* 2011).

For the ML models presented in this main part of the paper, we choose a prediction timestep of one day, that is, we predict daily average ozone distributions using the temperature state of the atmosphere one day earlier. We also implemented models for 5-day-mean (predicting 5-day-average ozone fields five days in advance) and monthly-mean data (predicting one month in advance). Increasing the lead time results in the expected decline in the quality of the predictions (supplementary figure S1).

For the preindustrial and the 4xCO_2 scenario, we used the first 2000 and 2500 PCA temperature components as input, respectively, as compared to around 100000 spatial grid cells. The components explain >99% of variance for the preindustrial training set and >98% for the 25-year-long 4xCO_2 training set (introduced below). For daily-mean data, these choices represent a good compromise between minimizing the number of input variables and the prediction error (figures S1a–d). For the 4xCO_2 case, we also dimension-reduced the ozone field prior to the regression fit, which, after the necessary additional transformations, yields essentially identical results to grid-cell-wise regressions (see supplementary section 2).

3. Results

3.1. Average ozone fields

Any ozone parameterization should be able to reproduce the average ozone climatology of an interactive chemistry model with high fidelity. For 4xCO_2 simulations, this requires that the method also captures ozone trends under forcing (Nowack *et al* 2015). Figure 2(c) shows percentage differences in ozone mixing ratios under 4xCO_2 relative to the preindustrial climatology (figure 2(b)), averaged over years 50–200 after the CO_2 forcing. We find ozone decreases in the lower stratosphere, where ozone is long-lived enough to be influenced by transport, most notably in the tropical lower stratosphere, which is an ubiquitous feature in chemistry-climate model simulations (SPARC 2010). In the upper stratosphere, we find ozone increases, which are also well understood based on the temperature-dependency of catalytic ozone depletion cycles (Haigh and Pyle 1982, Meul *et al* 2014). All of these changes are not included

in simulations with fixed ozone climatologies (Nowack *et al* 2015).

After minimizing the cost function according to equation (1), the ML model predicts the time-mean ozone climatology for both preindustrial and 4xCO_2 conditions well (figures 2(d), (e) and S2/S3). Specifically, the prediction error on the preindustrial testset for the zonal mean ozone climatology is smaller than 5% everywhere. A model trained separately on the initial 25 years of the 4xCO_2 run can also predict the average ozone field for the remaining 175 years to within 5% almost everywhere (figure 2(e)). For this, the regression necessarily captures long-term trends, which are key to the evolution of ozone in the dynamically-driven lower stratosphere where there are significant changes over the entire period (Dietmüller *et al* 2014, Nowack *et al* 2015, figures S4/S5). We found that the same extrapolation would not work using the model fit on preindustrial data; using the preindustrial regression model leads to a catastrophic prediction failure for the 4xCO_2 scenario, implying that the co-variability in temperature and ozone for the preindustrial state is not a valid predictor of ozone's response to an external forcing.

However, further tests indicate that the length of the training dataset could be reduced for most purposes. The bulk of the prediction error vanishes using 10–15 years of training data (figures S1e,f and S5). In comparison, taking 50 years of training data would lead to mean squared errors (MSE) almost as small as for the preindustrial testset (figure 2(d)), despite the additional complexity of underlying trends under 4xCO_2 (figures S1d,f and S5). Here we focus on the model trained on 25 years of data as an intermediate case. Next, we show that the resulting ML model also poses a significant improvement over a fixed ozone climatology in terms of resolving temporal and spatial ozone variability.

3.2. Representation of variability

Figure 3 shows six representative cases of ozone times series data produced by the chemistry-climate model (blue lines) and the corresponding predictions of the ML model (red lines) on the 13-year-long preindustrial testset for specific grid cells. Figure 4 shows similar data for the 4xCO_2 testset. Ideally, the time series for the climate model and the ML model should be in perfect sync. In each figure, the ozone field as given by a 'standard' fixed ozone climatology (gray lines) is indicated. Such a field varies on a seasonal basis, but not interannually and does also not adapt to the CO_2 forcing. Extensive samples for other grid cells and for other time resolutions can be found in supplementary figures S8-S12.

For the preindustrial testset, the ML model generally captures key aspects of ozone variability across the entire atmosphere. At all levels the MSE is reduced, reaching improvement factors greater than seven in the tropical stratosphere (figure S6). This is also evident from figures 3(a)–(c) where the predictions and climate model match well from locations in the lower to the upper tropical stratosphere. In comparison, a fixed ozone climatology misses many undulations and sometimes even the periodicity is completely out of sync. Figure 3(b) exhibits just one prominent example for the tropical mid-stratosphere, where the QBO is a key driver of ozone variability (figure 3(b)). The good representation of ozone variability there is crucial because ozone feedbacks are well-known for two-way interactions with the QBO as well as its remote impacts on the extratropics (e.g. Rind *et al* 2014, Silverman *et al* 2018). Major performance gains are also found around the positions of the jet streams and the polar vortices (figure S6). Our method also represents the zonal structure in the ozone distribution at any given time well (figures S12.1–S12.5), both in the troposphere (e.g. differences between land and oceanic regions) and in the stratosphere (e.g. due to wavy polar vortices).

In some areas the MSE improvement over a fixed climatology is not as large. For the example shown in figure 3(d), the climatology is already a very good approximation due to little interannual variability. More pronounced are the differences in the troposphere (e.g. Figure 3(f)), but even in the least favorable cases, the ML model still performs better than the climatology and captures the seasonal cycle. We argue that this should be sufficient for typical climate sensitivity simulations, where tropospheric ozone abundances are particularly low (Chiodo *et al* 2018) so that any variability around their mean state will have little impact on the modeled climate. Instead, we argue that it is more important to capture tropospheric ozone trends under 4xCO_2 (figure 2(c)), which as we showed are modeled well by the parameterization (figure 2(e)). As discussed in section 4, improving the representation of tropospheric ozone variability could be an exciting aspect of future work, especially with regard to other forcing scenarios.

Results for the abrupt 4xCO_2 scenario are similar (figure 4). However, there the ML model does not only reproduce well the internal ozone variability but also ozone trends, again with limited ability to pick up on very short-term tropospheric fluctuations. The ability to reproduce trends is crucial for the 4xCO_2 case, because the climate continues to warm over the entire 175 years prediction period following the initial 25-year training period (figure 2(a)). As a result, temperatures will increase beyond those encountered in the training dataset. Leaving the training domain can lead to catastrophic failure of ML models, because they usually perform well on interpolation tasks, but perform poorly for extrapolations. This is particularly true for highly non-linear regression functions such as neural networks. In contrast, the linear ML algorithm used here remains stable and robust in its predictions. This is probably not surprising given that it simply

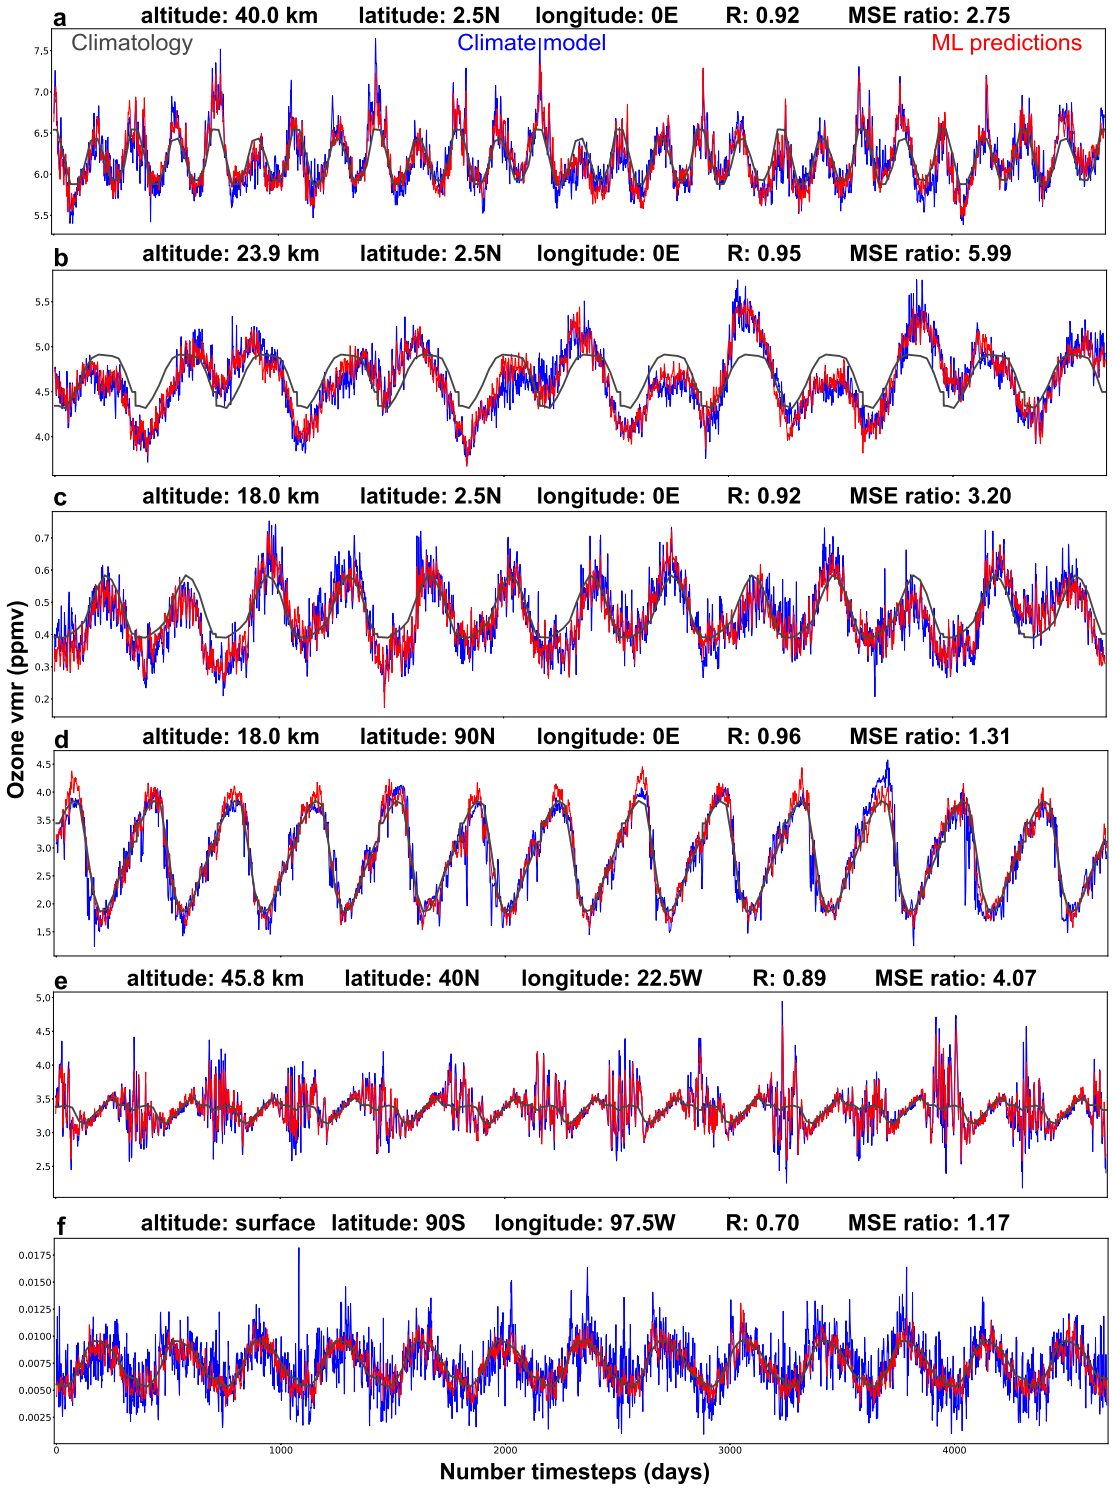

Figure 3. Ozone (ppmv) time series data for various model grid cells under preindustrial conditions. (Gray) Fixed preindustrial climatology. (Blue) Ozone as simulated by the fully interactive chemistry-climate model. (Red) Predictions on the testdata using the ML model. Coordinates of each cell as labeled. *R* is the Pearson correlation coefficient between the chemistry-climate model time series and the predictions. *MSE ratio* is the ratio of mean squared errors of the climatology and the ML predictions relative to the chemistry-climate model data (values >1 imply an improvement using the ML model).

assumes the linear relationships between temperature inputs and ozone outputs to hold even outside the training domain, therefore acting as a better extrapolation tool than many popular non-linear algorithms.

The time series of the predictions in figure 4 remain highly correlated with the ozone field produced by the chemistry-climate model, reaching almost perfect correlations in the stratosphere. Improvements compared to a fixed climatology in terms of the MSE are much larger, which is not surprising given the large differences not just in variability but also in the mean values for ozone; an important factor that is missed in models using climatologies. As a result, improvement factors >150 are attained in the

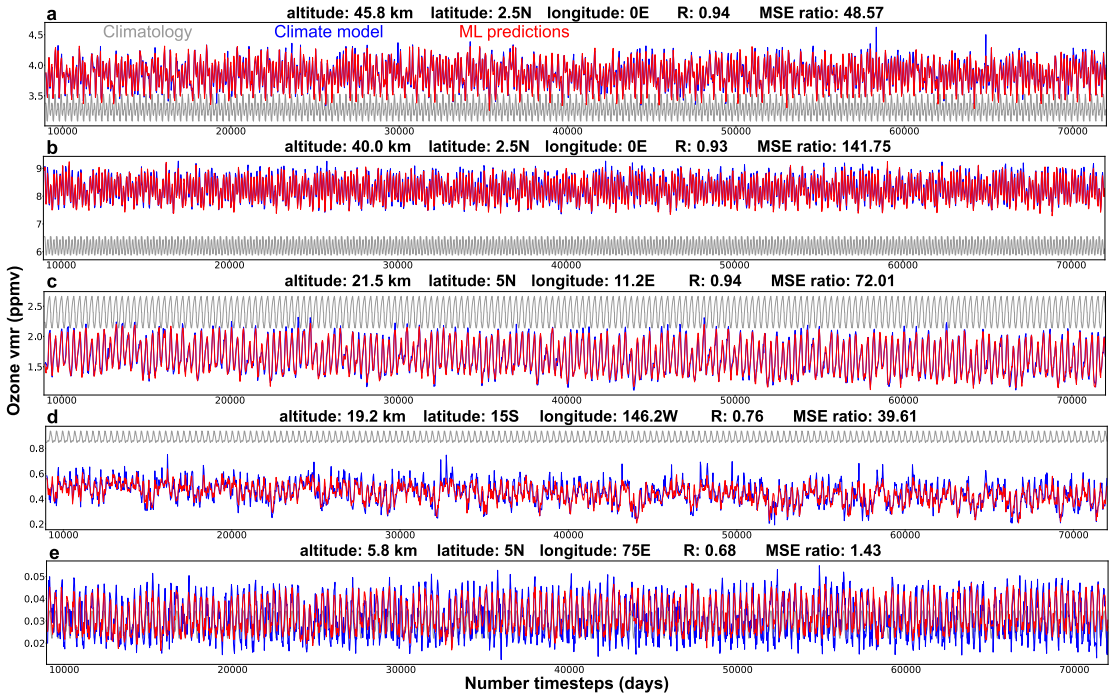

Figure 4. Ozone (ppmv) time series data for the 175-year-long 4xCO_2 testset (timesteps 9001–72000 based on a 360-day model calendar). Color code and labels as in figure 3; however, the daily-mean predictions and model output were smoothed by calculating their monthly-means to aid visualization. The time series remain in sync long after the training period has ended, despite remaining trends in the system, see (c) and (d). Actual daily-mean time series are documented in supplementary figures S8-S11. The correlation coefficients and MSE ratios were calculated using daily-mean data.

stratosphere (figure S7). Clearly, the ML model outperforms the use of a preindustrial climatology by far. We further note that climatologies taken over the training period would only be able to capture mean changes in the quickly responding upper stratosphere (figure 4(a)), but not the slower, surface-temperature correlated trends in the climatologically so important lower stratosphere (figures 4(c) and S4, Dietmüller *et al* 2014, Nowack *et al* 2015), or the interannual variability in ozone.

4. Discussion and conclusions

4.1. Physical intuition

Our parameterization represents well the climatologically important variability and trends in ozone for preindustrial and 4xCO_2 scenarios. The only input variable is the atmospheric temperature state at the previous timestep. While temperature does affect ozone production and depletion, it is not the only driver of changes in ozone. This raises an obvious question: why does the regression work so well?

The simple answer is that the regression does not just pick up on direct causal relationships, but can also gain predictive skill from systematic indirect correlations. For example, a stable polar vortex characterized by cold, isolated polar stratospheric air will be tantamount to hindered poleward transport of ozone across the vortex edge. A faster stratospheric Brewer-Dobson circulation will be reflected in a colder tropical lower stratosphere and increased poleward heat transport. The resulting temperature structure in turn is indicative of enhanced poleward transport of ozone. Long-term changes in the circulation, which are key for ozone trends under 4xCO_2 forcing, are also strongly correlated with the degree of tropospheric warming found in climate models (Shepherd and McLandress 2011, Chiodo *et al* 2018). The seasonal cycle in ozone, which is primarily driven by changes in transport and photochemical ozone production, will equally be reflected in the atmospheric temperature field. Atmospheric humidity, an important factor in tropospheric ozone production and loss mechanisms and future ozone trends, will also be correlated with temperature (Held and Soden 2006, Young *et al* 2018). In the upper stratosphere, where ozone is short-lived, catalytic ozone depletion mechanisms are directly dependent on temperature (Haigh and Pyle 1982, Jonsson *et al* 2004).

All these factors will interact in complex ways at any time and altogether give rise to a temperature state of the atmosphere that can be related to a corresponding, representative ozone distribution. ML can identify these relationships automatically where each process (for example the vortex state) could be characterized by the time-dependent weights of a large number of principal components in temperature (or alternatively, temperatures in specific grid cells). Each temperature component in turn can be important for predicting a number of different processes needed to

model ozone trends and variability. For example, a temperature component reflecting the seasonal cycle might characterize certain seasonal aspects of both the atmospheric circulation and regional changes in solar radiation.

In summary, the main processes driving ozone variability and trends will in some way be correlated with atmospheric temperature. Even though this does not not always represent direct cause-effect relationships, it implies that once a certain temperature state of the atmosphere has been associated with a specific ozone distribution, such information can be used to predict future ozone distributions.

4.2. Outlook

Despite promising results, we have only conducted a first feasibility study, which we hope will motivate further work in this direction. While our method could immediately reduce the necessary length of interactive chemistry simulations to periods as short as 10–15 years (whereas preindustrial and 4xCO_2 simulations are usually run on centennial to millennial time-scales), it still asks for significant effort on the model developer side. However, since we find that linear models provide good approximations, there is hope that our scheme could become transferable across models by using scaling arguments involving the strength of the circulation response, surface temperature change etc, without the need for repeated training episodes. If used across a number of climate models, such a parameterization could result in another significant cut in necessary computing and model development effort.

The next logical step is the actual implementation of our approach in a (fully coupled) climate model. In particular, this would allow us to evaluate its performance in terms of climate statistics. Here, we demonstrated for the first time that ozone can be modeled solely as a function of atmospheric temperatures. However, the fit is necessarily an approximation so that the climate model would see a slightly different ozone field compared to using a fully interactive chemistry scheme (see our validation), even though the error would probably be within the range produced by different state-of-the-art chemistry schemes. As a result, there might also be small differences in simulated temperatures, which in turn would feedback onto the ozone again. With our approach, climatic differences should be significantly smaller compared to using a fixed ozone climatology instead of interactive chemistry. We hope future studies will investigate these two-way interactions in detail.

Additional opportunities lie in the inclusion of other input variables. In particular, a better representation of short-term tropospheric ozone variability should be aimed for, which is not yet captured in great detail by the current parameterization (figures 3(f) and 4e). This is not surprising given our prediction lead time of one day, which is long in terms of tropospheric time-scales and processes such as convection or turbulent transport. Increasing the time resolution might help, but also the inclusion of emissions of ozone precursors, for example of nitrogen oxides by lightning. Ultimately, similar methods could be applied to other, more complex forcing scenarios. The resulting advanced regression models could be a useful tool for comparing chemistry-climate models to observations in terms of their representation of co-variability in temperature, or any other possible input features, and ozone. Such a tool could open up new pathways for detecting and removing significant model biases or for identifying missing processes. More immediately, we see scope for using similar methods in solar forcing or paleoclimate scenarios. We also highlight the potential significance of taking lagged variables into account, ideally in combination with other regression methods (Runge *et al* 2015, Kretschmer *et al* 2017).

Having actual implementations of ozone parameterizations in climate models in mind, we also point towards the intrinsic limitations of semi-statistical models. For example, we found that in tropospheric regions of relatively high ozone variability, the parameterization could sometimes 'overshoot' slightly, resulting in rare cases of small negative predictions for ozone mass mixing ratios (figures S8-S11). While not necessarily of great importance for radiative transfer, such unphysical negative values should be prevented by setting specific boundary conditions (Schneider *et al* 2017), e.g. by imposing the value from the previous timestep. Finally, we note that trace gases other than ozone, in the form of methane and nitrous oxide, are also typically prescribed as climatologies.

In conclusion, we have demonstrated that machine learning regression models could be a useful tool for representing ozone more consistently in climate simulations. We hope that our research will be a starting point for an informed discussion on possible ML parameterizations of ozone, moving beyond the current binary framework of fixed ozone climatologies and interactive chemistry schemes. In this context, we wish to acknowledge relaxation schemes as a simplified option to implement interactive ozone chemistry (e.g. Cariolle and Déqué 1986, McLinden *et al* 2000, Cariolle and Teyssèdre 2007). These calculate purely chemical tendencies and therefore still require advection of ozone. They are ideal to evaluate interannual variability and small forcings (Braesicke and Pyle 2003), but have limited scope under extreme climate change scenarios. In comparison, our ML method directly accounts for both local and non-local (e.g. circulation) effects and is easily extendible to extreme scenarios. Another difference is that relaxation schemes linearize the chemical equation system, whereas we use a regularized linear regression approach to model ozone as a function of temperature. Therefore, ozone can still vary non-linearly in

time given non-linear variations in temperature. We invite data collaborations for climate sensitivity simulations with interactive chemistry schemes, which will facilitate the process of building robust and transferable ML-based ozone parameterizations.

Acknowledgments

PN is supported through an Imperial College Research Fellowship. Data production using the chemistry-climate model was funded through the ACCI project number 267760 of the European Research Council. The corresponding climate model development was part of the QESM-ESM project supported by the UK Natural Environment Research Council (NERC) under contract numbers RH/H10/19 and R8/H12/124. We further acknowledge use of the MONSooN system, a collaborative facility supplied under the Joint Weather and Climate Research Programme, which is a strategic partnership between the UK Met Office and NERC. We used the JASMIN postprocessing system (Lawrence *et al* 2013) provided through the Centre for Environmental Data Analysis (CEDA).

ORCID iDs

Peer Nowack https://orcid.org/0000-0003-4588-7832
Peter Braesicke https://orcid.org/0000-0003-1423-0619
Joanna Haigh https://orcid.org/0000-0001-5504-4754
Nathan Luke Abraham https://orcid.org/0000-0003-3750-3544
John Pyle https://orcid.org/0000-0003-3629-9916
Apostolos Voulgarakis https://orcid.org/0000-0002-6656-4437

References

Albers J R and Nathan T R 2012 Pathways for communicating the effects of stratospheric ozone to the polar vortex: Role of zonally asymmetric ozone *J. Atmos. Sci.* **69** 785–801
Anderson G J and Lucas D D 2018 Machine learning predictions of a multiresolution climate model ensemble *Geophys. Res. Lett.* **45** 4273–80
Andrews T, Gregory J M, Webb M J and Taylor K E 2012 Forcing, feedbacks and climate sensitivity in CMIP5 coupled atmosphere-ocean climate models *Geophys. Res. Lett.* **39** L09712
Ball W T, Haigh J D, Rozanov E V, Kuchar A, Sukhodolov T, Tummon F, Shapiro A V and Schmutz W 2016 High solar cycle spectral variations inconsistent with stratospheric ozone observations *Nat. Geosci.* **9** 206–11
Bellenger H, Guilyardi E, Leloup J, Lengaigne M and Vialard J 2013 ENSO representation in climate models: from CMIP3 to CMIP5 *Clim. Dyn.* **42** 1999–2018
Bishop C M 2006 *Pattern Recognition and Machine Learning* (Berlin: Springer)
Boers N, Bookhagen B, Barbosa H M J, Marwan N, Kurths J and Marengo J A 2014 Prediction of extreme floods in the eastern Central Andes based on a complex networks approach *Nat. Commun.* **5** 5199
Braesicke P and Pyle J A 2003 Changing ozone and changing circulation in northern mid-latitudes: possible feedbacks? *Geophys. Res. Lett.* **30** 1059
Cariolle D and Déqué M 1986 Southern hemisphere medium-scale waves and total ozone disturbances in a spectral general circulation model *J. Geophys. Res.* **91** 10825–46
Cariolle D and Teyssèdre H 2007 A revised linear ozone photochemistry parameterization for use in transport and general circulation models: Multi-annual simulations *Atmos. Chem. Phys.* **7** 2183–96
Ceppi P and Gregory J M 2017 Relationship of tropospheric stability to climate sensitivity and Earth's observed radiation budget *Proc. Natl Acad. Sci.* **114** 13126–31
Chiodo G and Polvani L M 2016 Reduction of climate sensitivity to solar forcing due to stratospheric ozone feedback *J. Clim.* **29** 4651–63
Chiodo G and Polvani L M 2017 Reduced Southern Hemispheric circulation response to quadrupled CO_2 due to stratospheric ozone feedback *Geophys. Res. Lett.* **44** 465–74
Chiodo G, Polvani L, Marsh D, Stenke A, Ball W, Rozanov E, Muthers S and Tsigaridis K 2018 The response of the ozone layer to quadrupled CO_2 concentrations *J. Clim.* **31** 3893–907
Cionni I, Eyring V, Lamarque J F, Randel W J, Stevenson D S, Wu F, Bodeker G E, Shepherd T G, Shindell D T and Waugh D W 2011 Ozone database in support of CMIP5 simulations: Results and corresponding radiative forcing *Atmos. Chem. Phys.* **11** 11267–92
Crook J A, Gillett N P and Keeley S P E 2008 Sensitivity of Southern Hemisphere climate to zonal asymmetry in ozone *Geophys. Res. Lett.* **35** L07806
Crutzen P J 1970 The influence of nitrogen oxides on the atmospheric ozone content *Q. J. R. Meteorol. Soc.* **96** 320–5
Dietmüller S, Ponater M and Sausen R 2014 Interactive ozone induces a negative feedback in CO_2-driven climate change simulations *J. Geophys. Res.: Atmos.* **119** 1796–805
Esentürk E, Abraham L, Archer-Nicholls S, Mitsakou C, Griffiths P and Pyle J 2018 Quasi-Newton Methods for Atmospheric Chemistry Simulations : Implementation in UKCA UM Vn10 . 8 *Geosci. Mod. Dev.* **11** 3089–108
Eyring V *et al* 2013 Long-term ozone changes and associated climate impacts in CMIP5 simulations *J. Geophys. Res.: Atmos.* **118** 5029–60
Eyring V, Bony S, Meehl G A, Senior C A, Stevens B, Stouffer R J and Taylor K E 2016 Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization *Geosci. Mod. Dev.* **9** 1937–58
Fels S B, Mahlman J D, Schwarzkopf M D and Sinclair R W 1980 Stratospheric sensitivity to perturbations in ozone and carbon dioxide: radiative and dynamical response *J. Atmos. Sci.* **37** 2265–97
Gabriel A, Peters D, Kirchner I and Graf H F 2007 Effect of zonally asymmetric ozone on stratospheric temperature and planetary wave propogation *Geophys. Res. Lett.* **34** L06807
Gentine P, Pritchard M, Rasp S, Reinaudi G and Yacalis G 2018 Could machine learning break the convection parameterization deadlock? *Geophys. Res. Lett.* **45** 5742–51
Gillett N P, Scinocca J F, Plummer D A and Reader M C 2009 Sensitivity of climate to dynamically-consistent zonal asymmetries in ozone *Geophys. Res. Lett.* **36** L10809
Guilyardi E, Bellenger H, Collins M, Ferrett S, Cai W and Wittenberg A 2012 A first look at ENSO in CMIP5 *Clivar Exchanges* **17** 29–32
Haigh J D 1996 The impact of solar variability on climate *Science* **272** 981–4
Haigh J D and Pyle J A 1982 Ozone perturbation experiments in a two-dimensional circulation model *Quart. J. R. Met. Soc.* **108** 551–74
Heinemann M 2009 *Warm and Sensitive Paleocene-Eocene Climate* (Hamburg: Max Planck Institute for Meteorology)
Held I M and Soden B J 2006 Robust responses of the hydrological cycle to global warming *J. Clim.* **19** 5686–99
Hewitt H T, Copsey D, Culverwell I D, Harris C M, Hill R S R, Keen A B, McLaren A J and Hunke E C 2011 Design and

implementation of the infrastructure of HadGEM3: the next-generation Met Office climate modelling system *Geosci. Mod. Dev.* **4** 223–53
Hoerl A E and Kennard R W 1970 Ridge Regression: Biased Estimation for Nonorthogonal Problems *Technometrics* **12** 55–67
Hunt B G 1966 The need for a modified photochemical theory of the ozonosphere *J. Atmos. Sci.* **23** 88–95
James G, Witten D, Hastie T and Tibshirani R 2013 *An Introduction to Statistical Learning* (New York: Springer)
Jonsson A I, de Grandpré J, Fomichev V I, McConnell J C and Beagley S R 2004 Doubled CO_2-induced cooling in the middle atmosphere: Photochemical analysis of the ozone radiative feedback *J. Geophys. Res.* **109** D24103
Knutti R, Rugenstein M A A and Hegerl G C 2017 Beyond equilibrium climate sensitivity *Nat. Geosci.* **10** 727–36
Kravitz B, Robock A, Forster P M, Haywood J M, Lawrence M G and Schmidt H 2013 An overview of the Geoengineering Model Intercomparison Project (GeoMIP) *J. Geophys. Res.: Atmos.* **118** 13103–7
Kretschmer M, Runge J and Coumou D 2017 Early prediction of extreme stratospheric polar vortex states based on causal precursors *Geophys. Res. Lett.* **44** 8592–600
Lacis A A, Wuebbles D J and Logan J A 1990 Radiative forcing of climate by changes in the vertical distribution of ozone *J. Geophys. Res.* **95** 9971–81
Lawrence B N, Bennett V L, Churchill J, Juckes M, Kershaw P, Pascoe S, Pepler S, Pritchard M and Stephens A 2013 Storing and manipulating environmental big data with JASMIN *Proceedings —2013 IEEE Int. Conf. on Big Data, Big Data 2013* pp 68–75
Li C, von Storch J-S and Marotzke J 2013 Deep-ocean heat uptake and equilibrium climate response *Clim. Dyn.* **40** 1071–86
McCormack J P, Nathan T R and Cordero E C 2011 The effect of zonally asymmetric ozone heating on the Northern Hemisphere winter polar stratosphere *Geophys. Res. Lett.* **38** L03802
McLinden C A, Olsen S C, Hannegan B, Wild O, Prather M J and Sundet J 2000 Stratospheric ozone in 3-D models: a simple chemistry and the cross-tropopause flux *J. Geophys. Res.: Atmos.* **105** 14653–65
Meul S, Langematz U, Oberländer S, Garny H and Jöckel P 2014 Chemical contribution to future tropical ozone change in the lower stratosphere *Atmos. Chem. Phys.* **14** 2959–71
Molina M J and Rowland F S 1974 Stratospheric sink for chlorofluoromethanes: chlorine atomc-atalysed destruction of ozone *Nature* **249** 810–2
Monteleoni C *et al* 2013 Climate informatics *Computational Intelligent Data Analysis for Sustainable Development; Data Mining and Knowledge Discovery Series* pp 81–126 ch 4
Morgenstern O, Braesicke P, O'Connor F M, Bushell A C, Johnson C E, Osprey S M and Pyle J A 2009 Evaluation of the new UKCA climate-composition model: I. The stratosphere *Geosci. Mod. Dev.* **2** 43–57
Muthers S *et al* 2014 The coupled atmosphere-chemistry-ocean model SOCOL-MPIOM *Geosci. Mod. Dev.* **7** 2157–79
Muthers S, Raible C C, Rozanov E and Stocker T F 2016 Response of the AMOC to reduced solar radiation—the modulating role of atmospheric chemistry *Earth Syst. Dynam.* **7** 877–92
Noda S, Kodera K, Adachi Y, Deushi M, Kitoh A, Mizuta R, Murakami S, Yoshida K and Yoden S 2017 Impact of interactive chemistry of stratospheric ozone on southern hemisphere paleoclimate simulation *J. Geophys. Res.: Atmos.* **122** 878–95
Nowack P J, Abraham N L, Braesicke P and Pyle J A 2016 Stratospheric ozone changes under solar geoengineering: implications for UV exposure and air quality *Atmos. Chem. Phys.* **16** 4191–203
Nowack P J, Abraham N L, Braesicke P and Pyle J A 2018 The impact of stratospheric ozone feedbacks on climate sensitivity estimates *J. Geophys. Res.: Atmos.* **123** 4630–41
Nowack P J, Braesicke P, Abraham N L and Pyle J A 2017 On the role of ozone feedback in the ENSO amplitude response under global warming *Geophys. Res. Lett.* **44** 3858–66
Nowack P J, Luke Abraham N, Maycock A C, Braesicke P, Gregory J M, Joshi M M, Osprey A and Pyle J A 2015 A large ozone-circulation feedback and its implications for global warming assessments *Nat. Clim. Change* **5** 41–5
Pedregosa F *et al* 2011 Scikit-learn: Machine learning in Python *J. Mach. Learn. Res.* **12** 2825–30
Pyle J A 1980 A calculation of the possible depletion of ozone by chlorofluorocarbons using a two-dimensional model *Pure Appl. Geophys. PAGEOPH* **118** 355–77
Rashid H A, Hirst A C and Marsland S J 2016 An atmospheric mechanism for ENSO amplitude changes under an abrupt quadrupling of CO_2 concentration in CMIP5 models *Geophys. Res. Lett.* **43** 1687–94
Rind D, Jonas J, Balachandran N K, Schmidt G a and Lean J 2014 The QBO in two GISS global climate models: 1. Generation of the QBO *J. Geophys. Res.: Atmos.* **119** 8798–824
Runge J, Petoukhov V, Donges J F, Hlinka J, Jajcay N, Vejmelka M, Hartman D, Marwan N, Paluš M and Kurths J 2015 Identifying causal gateways and mediators in complex spatio-temporal systems *Nat. Commun.* **6** 8502
Samuel A L 1959 Some studies in machine learning using the game of checkers *IBM J. Res. Dev.* **3** 210–29
Schneider T, Lan S, Stuart A and 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** 12396–417
Shepherd T G and McLandress C 2011 A robust mechanism for strengthening of the Brewer-Dobson circulation in response to climate change: critical-layer control of subtropical wave breaking *J. Atmos. Sci.* **68** 784–97
Silverman V, Harnik N, Matthes K, Lubis S W and Wahl S 2018 Radiative effects of ozone waves on the Northern Hemisphere polar vortex and its modulation by the QBO *Atmos. Chem. Phys.* **18** 6637–59
Solomon S, Ivy D J, Kinnison D, Mills M J, Neely III RR and Schmidt A 2016 Antarctic ozone layer *Science* **353** 269–74
Son S-W *et al* 2010 Impact of stratospheric ozone on Southern Hemisphere circulation change: a multimodel assessment *J. Geophys. Res.: Atmos.* **115** 1–18
Son S-W, Polvani L M, Waugh D W, Akiyoshi H, Garcia R, Kinnison D, Pawson S, Rozanov E, Shepherd T G and Shibata K 2008 The impact of stratospheric ozone recovery on the Southern Hemisphere westerly jet *Science* **320** 1486–9
SPARC 2010 *CCMVal Report on the Evaluation of Chemistry-Climate Models, Technical Report, SPARC* http://sparc-climate.org/publications/sparc-reports/
Taylor K E, Stouffer R J and Meehl G A 2012 An overview of CMIP5 and the experiment design *Bull. Am. Meteorol. Soc.* **93** 485–98
Thompson D W J and Solomon S 2002 Interpretation of recent Southern Hemisphere climate change *Science* **296** 895–9
Vandal T, Kodra E and Ganguly A R 2017 *Intercomparison of Machine Learning Methods for Statistical Downscaling: The Case of Daily and Extreme Precipitation* 1–20 http://arxiv.org/abs/1702.04018
Voulgarakis A *et al* 2013 Analysis of present day and future OH and methane lifetime in the ACCMIP simulations *Atmos. Chem. Phys.* **13** 2563–87
Williamson C E *et al* 2014 Solar ultraviolet radiation in a changing climate *Nat. Clim. Change* **4** 434–41
WMO 2011 *(World Meteorological Organization): Scientific Assessment of Ozone Depletion: 2010, Global Ozone Research and Monitoring Project—Report No. 52* Geneva, Switzerland
Young P J *et al* 2013 Pre-industrial to end 21st century projections of tropospheric ozone from the Atmospheric Chemistry and Climate Model Intercomparison Project (ACCMIP) *Atmos. Chem. Phys.* **13** 2063–90
Young P J *et al* 2018 Tropospheric Ozone Assessment Report (TOAR): Assessment of global-scale model performance for global and regional ozone distributions, variability, and trends *Elem. Sci. Anth.* **6**