

An Online-Learned Neural Network Chemical Solver for Stable Long-Term Global Simulations of Atmospheric Chemistry

Makoto M. Kelp¹, Daniel J. Jacob^{1,2}, Haipeng Lin², Melissa P. Sulprizio²

¹Department of Earth and Planetary Sciences, Harvard University, Cambridge, MA 01238, USA

²School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 01238, USA

Corresponding author: Makoto M. Kelp (mkelp@g.harvard.edu)

Key Points:

- Application of ML to solve the kinetic equations in global 3-D atmospheric chemistry models requires online training of the ML solver
- Stable year-long global simulation of chemistry can be achieved with a ML solver applied to the Super-Fast mechanism in GEOS-Chem
- Computational speedup is five-fold relative to the reference Rosenbrock solver in GEOS-Chem

Abstract

A major computational barrier in global modeling of atmospheric chemistry is the numerical integration of the coupled kinetic equations describing the chemical mechanism. Machine-learned (ML) solvers can offer order-of-magnitude speedup relative to conventional implicit solvers, but past implementations have suffered from fast error growth and only run for short simulation times (<1 month). A successful ML solver for global models must avoid error growth over year-long simulations and allow for re-initialization of the chemical trajectory by transport at every time step. Here we explore the capability of a neural network solver equipped with an autoencoder to achieve stable full-year simulations of tropospheric oxidant chemistry in the global 3-D GEOS-Chem model, replacing its standard mechanism (228 species) by the Super-Fast mechanism (12 species) to avoid the curse of dimensionality. We find that online training of the ML solver within GEOS-Chem is essential for accuracy, whereas offline training from archived GEOS-Chem inputs/outputs produces large errors. After online training we achieve stable 1-year simulations with five-fold speedup compared to the standard implicit Rosenbrock solver with global tropospheric normalized mean biases of -0.3% for ozone, 1% for hydrogen oxide radicals, and -5% for nitrogen oxides. The ML solver captures the diurnal and synoptic variability of surface ozone at polluted and clean sites. There are however large regional biases for ozone and NO_x under remote conditions where chemical aging leads to error accumulation. These regional biases remain a major limitation for practical application, and ML emulation would be more difficult in a more complex mechanism.

Plain Language Summary

Global models of atmospheric chemistry are computationally expensive. A bottleneck is the chemical solver that integrates the large-dimensional coupled systems of kinetic equations describing the chemical mechanism. Machine learning (ML) could be transformative for reducing the cost of an atmospheric chemistry simulation by replacing the chemical solver with a faster emulator. However, past work found that ML chemical solvers experience rapid error growth and become unstable over time. Here we present results achieving for the first time a stable full-year global simulation of atmospheric chemistry with a ML solver and with five-fold speedup in computational performance over the reference simulation. We show that online training of the ML solver synchronously with the atmospheric chemistry model simulation produces considerably more stable results than offline training it from a static dataset of simulation results. Although our work represents an important step for using ML solvers in global atmospheric chemistry models, more work is needed to extend it to large chemical mechanisms and to reduce errors during long-term chemical aging.

1. Introduction

Global modeling of atmospheric chemistry is a grand computational challenge due to the large number of coupled chemical species, the non-linearity and numerical stiffness of chemical mechanisms, and the interactions with transport on all scales. The U.S. National Research Council's National Strategy for Advancing Climate Modeling identifies atmospheric chemistry as a priority frontier for Earth System Model (ESM) development (National Research Council, 2012). Current atmospheric chemistry models integrate the coupled chemical kinetic equations for the mechanism species over model time steps by using high-order implicit numerical solvers, but these solvers are expensive (Sandu et al., 1997) and often dominate the cost of an atmospheric simulation (Eastham et al., 2018). Here we explore the potential of machine learning (ML) neural network algorithms to dramatically reduce the computational intensity of atmospheric chemistry in global simulations.

Chemical solvers in atmospheric models compute the local evolution of species concentrations over a chemical time step that may range from minutes to hours depending on the model (Brasseur and Jacob, 2017). The chemical mechanism typically includes ~100 coupled species with lifetimes ranging from less than a second to much larger than the chemical time step. Rosenbrock and Gear high-order implicit solvers can integrate this system of stiff coupled differential equations with high accuracy, and fast implementations of these schemes are available for example through the Kinetic Pre-Processor (KPP) (Sandu and Sander, 2006) and SMVGEAR (Jacobson and Turco, 1994). They are still extremely costly for atmospheric models. Models combat that cost by decreasing the size of the chemical mechanism (Sportisse and Djouad, 2000), breaking down the stiffness of the problem (Young and Boris, 1977), or using lower-order approximations, as reviewed by Brasseur and Jacob (2017). But these methods rarely achieve a speed-up of more than a factor of two (Shen et al., 2020). ML methods could be transformative for reducing the cost.

ML methods would seem well-suited to chemical solvers in atmospheric models because the chemical computation is very repetitive, involving integration of similar conditions in neighboring grid cells and successive time steps. However, the large number of coupled species brings a curse of dimensionality to the problem. ML methods also have no check on error growth, unlike in standard chemical solvers where errors are dampened by the negative response to perturbations (Le Chatelier's principle). Keller and Evans (2019) created a prototype random forest integrator for the GEOS-Chem global 3-D chemical transport model (CTM) driven by archived meteorological data. They achieved successful short-term simulations but found large error growth after a few weeks. Kelp et al. (2020) trained a neural network integrator in a chemical box model, including an encoder/decoder (Lusch et al., 2018) to decrease dimensionality, and a recursive feedback loop over 24-h integration time to control error growth. They found that they could compress the 101-species dimension of their mechanism into fewer than 20 features without significant error penalty, and that they could avoid error growth over a 1-week integration time. Liu et al. (2021) developed a gas-phase neural network solver for the CMAQ regional CTM over China, combining a standard implicit solver for radicals and oxidants with a ML solver for volatile organic compounds (VOCs). They achieved an order-of-magnitude speedup over a one-month simulation but with error growth over remote ocean grid cells.

Error growth in a ML chemical solver may be tolerable for short-term simulations such as in chemical forecasts or in small-scale air quality applications. But global simulations of atmospheric chemistry need stability over longer-term horizons. For example, a global simulation of tropospheric oxidants (ozone and hydroxyl radical OH) with fixed concentration of methane has chemical modes of several months (Wild and Prather, 2000; Murray, 2016) and must typically be integrated over a year. Moreover, stability of the solution is required over the full range of tropospheric conditions from polluted to remote, and from the surface to the upper troposphere. Operator splitting between chemistry and transport resets initial conditions after each transport time step, meaning that one cannot project the solution along long-term chemical trajectories as with dedicated ML timeseries algorithms such as Recurrent Neural Networks (RNNs) (Rumelhart et al., 1986) and Long-Short-Term Memory networks (LSTMs) (Hochreiter and Schmidhuber, 1997). Success in applying ML solvers to box models, such as in Kelp et al. (2020), may not translate to a global CTM.

One possible cause of error growth in the above applications is the use of offline training. In offline training, the ML solver learns the chemical tendencies from an archived dataset of CTM inputs and outputs over chemical time steps. Training a ML solver offline is expedient, straightforward, and allows for easy manipulation of training data. However, it tends to overfit to the training data as the entire dataset is typically cycled multiple times to improve learning, and it may not properly represent the ensemble of conditions encountered by the CTM simulation in their temporal sequence (Rasp, 2020). An alternative is online training, in which the ML solver learns the chemical tendencies from the CTM simulation as it evolves with time. Online training is more expensive and difficult as it requires running the CTM and ML training in tandem at every chemical time step. It may suffer also from catastrophic forgetting where information from earlier training data is lost (McCloskey and Cohen, 1989) as each datapoint is used only once. However, it allows the ML solver to actually sample the conditions in the CTM as they evolve forward in time and learn from these chemical tendencies (Parisi et al., 2019; Rasp et al., 2020). To our knowledge, online training has not been used previously for atmospheric chemistry applications.

Here we demonstrate the capability of a neural network ML solver with online training to provide a stable representation of tropospheric chemistry in a global 3-D model environment over full-year simulations. We do so by emulating the 12-species ‘Super-Fast’ chemical mechanism (Cameron-Smith et al., 2006; Brown-Steiner et al., 2018) in the GEOS-Chem CTM. The Super-Fast mechanism is a reduced representation of tropospheric chemistry used in climate models (Lamarque et al., 2013). Although oversimplified in relation to the mechanisms used for atmospheric chemistry research, it is a useful prototype for our purpose because the limitation to 12 chemical variables alleviates the curse of dimensionality. This allows us to investigate other challenges in achieving stable and accurate ML solutions, thus providing a foundation for application of ML methods to more complicated mechanisms.

2. Materials and Methods

2.1 GEOS-Chem model and Super-Fast mechanism

We use the GEOS-Chem CTM version 12.0.0 (<https://doi.org/10.5281/zenodo.1343547>) driven by assimilated meteorological data from the NASA Global Modeling and Assimilation Office (GMAO) Goddard Earth Observing System (GEOS). GEOS-Chem computes the evolution of atmospheric composition by successive application over model time steps of operators simulating emissions, transport, chemistry, and deposition (Bey et al., 2001). The chemical operator computes the changes in concentrations over the time step by integrating the coupled system of ordinary differential equations describing chemical production and loss for the ensemble of species in the mechanism (Brasseur and Jacob, 2017). Here we conduct global simulations at $4^\circ \times 5^\circ$ degrees resolution and 47 vertical levels (25-37 in the troposphere) using the GEOS Modern-Era Retrospective analysis for Research and Applications, version 2 (MERRA-2) meteorological dataset with 3-hour temporal resolution (1-hour for surface variables). Time steps are 30 min for transport and 60 min for chemistry (Philip et al., 2016). Photolysis frequencies are calculated with the Fast-JX scheme (Bian and Prather, 2002), as implemented in GEOS-Chem by Mao et al. (2010).

The standard GEOS-Chem model includes a detailed oxidant-aerosol chemical mechanism for the troposphere and stratosphere with 228 species (Eastham et al., 2014; Wang et al., 2021), integrated with a 4th-order Rosenbrock Rodas3 chemical solver through the Kinetic Pre-Processor (KPP; Sandu et al., 1997). Here we replace the chemical mechanism in the troposphere with the Super-Fast mechanism for oxidant chemistry (Brown-Steiner et al., 2018), including 12 chemical species coupled through 21 thermal reactions and 6 photolysis reactions. We replace the chemical mechanism in the stratosphere with a simple linear relaxation to chemical equilibrium intended to provide reasonable flux boundary conditions at the tropopause (McLinden et al., 2000; Murray et al., 2012).

We integrate the Super-Fast mechanism with the KPP Rosenbrock solver in the same way as the standard mechanism and this defines the reference Super-Fast simulation to which our ML solver will be compared. The 12 coupled species in the Super-Fast mechanism include methane oxidation products (CH_3O_2 , CH_3OOH , CH_2O , CO), oxidants and related radical chemistry (OH , HO_2 , H_2O_2 , O_3 , NO , NO_2 , HNO_3), and biogenic isoprene (C_5H_8) which produces CH_3O_2 upon oxidation. The mechanism also includes CH_4 and O_2 with fixed concentrations, and H_2O with concentration specified by the meteorological data. The nitrogen oxide radicals ($\text{NO}_x \equiv \text{NO} + \text{NO}_2$) are oxidized to HNO_3 solely by OH , and HNO_3 is chemically inert and removed by deposition.

We use standard GEOS-Chem emission inventories for the years 2016 and 2017 including CEDS for NO_x and CO from fuel combustion (Hoesly et al., 2018), GFED4 for NO_x and CO from open fires (Randerson et al., 2015), MEGAN v2.1 for isoprene (Guenther et al., 2012), Murray et al. (2012) for lightning NO_x , and Hudman et al. (2012) for soil NO_x . We increase CO emissions by 19% for fuel combustion and 11% for open fires, following Fisher et al. (2017), to account for secondary production of CO from nonmethane volatile organic compounds (NMVOCs). The tropospheric methane concentration is imposed by latitude-dependent surface boundary

conditions (Murray, 2016). Standard GEOS-Chem modules for dry deposition (Bey et al., 2001) and wet deposition (Liu et al., 2001; Amos et al., 2012) are applied to CH_2O , H_2O_2 , O_3 , NO_2 , and HNO_3 .

2.2 Machine Learning (ML) neural network chemical solver

The ML chemical solver used here consists of three main components: an encoder, an integrator, and a decoder, each of which is a neural network. The encoder and decoder components (referred together as an autoencoder) are used for data compression and decompression, respectively (Kramer, 1991). The encoder learns to map chemical species to a compressed dimensional representation, the integrator learns to integrate the compressed representation forward in time, and the decoder learns to convert the compressed representation back to the original species. The encoder and decoder are shallow neural networks comprised of a single hidden layer with 16 nodes and linear activation. For the integrator, we use the ResNet residual neural network (He et al., 2016) with 1 block with 2 hidden layers and 128 nodes per layer, ReLu activation, and the Adam optimizer (Kingma and Ba, 2017). Each fully connected layer is preceded by a batch normalization operation (Ioffe and Szegedy, 2015), which normalizes the activations into the ResNet block to create a smoother optimization landscape for improved gradient flow (Santurkar et al., 2019). After each fully connected layer, a dropout rate of 0.5 is applied to prevent overfitting (Srivastava et al., 2014). We refer to the ML parameters as the coefficients of the regression algorithms.

Training of the ML solver involves fitting the chemical evolution in the reference Super-Fast simulation over model time steps. Each GEOS-Chem 1-h time step output constitutes one training sample, consisting of 20 input variables: concentrations of the 12 species at the beginning of the time step, 6 photolysis frequencies calculated by Fast-JX in the middle of the chemical time step (for O_3 , H_2O_2 , NO_2 , CH_2O by two branches, and CH_3OOH), temperature, and air density. The output variables are the concentrations of the 12 species at the end of the time step. Photolysis frequencies can themselves be emulated using neural networks (Krasnopolsky et al., 2005; Lagerquist et al., 2021; Sturm and Wexler, 2020) but their calculation is cheap compared to the chemical calculation.

All ML chemical solver code in this work is written using the Keras package Python ML routines. GEOS-Chem is written in Fortran and there are no programs to easily call Python ML algorithms from Fortran. To couple the ML chemical solver with GEOS-Chem, we use the C Foreign Function Interface for Python (CFFI; <https://cffi.readthedocs.io>) that calls the Python ML code from within Fortran. CFFI includes a process called embedding, which packages Python code into “dynamic libraries” that may be included and executed by a Fortran program.

Training of the ML solver to emulate the reference KPP Rosenbrock solver involves minimizing cost functions for the mean square error between output variables, i.e., the species concentrations computed at the end of the chemical time step. Kelp et al. (2020) found that a cost function that equally prioritizes all species is significantly less accurate than one that is specialized toward a single species of interest. Here, we create 12 separate ML solvers for each of the species. We use an encoder compression into 8 features and apply a log transformation for the concentrations of selected species to obtain more normal distributions that aid in neural network learning. For all

species, we rescale the inputs and outputs to $[0, 1]$ ranges using min-max normalization for ease of gradient optimization.

2.3 Offline and online training

We will present results from ML solvers trained in different ways offline and online. In standard offline training, we run GEOS-Chem to create a training dataset of input and output variables over individual 1-h chemical time steps. We use a training data batch size of 1024 and initial learning rate of 0.001, with learning-rate decay (You et al., 2019) occurring every time the validation set error plateaus for 10 epochs (an “epoch” is when an entire dataset is passed forward and backward through the neural network a single time). We use early stopping (Li et al., 2019) to halt ML solver training when the absolute error decreases less than 1×10^{-4} for 15 epochs.

Kelp et al. (2020) found that recursive training of 1-h chemical time steps over 24-h time horizons was critical in their box model application to capture slow chemical modes and prevent error growth. We implemented this recursive training here by mimicking the effects of operator splitting between chemistry and other operators in GEOS-Chem. This involved archiving the 24-h evolution of concentrations over 1-h time steps from the ensemble of non-chemistry operators and adding it to every hourly time step for recursive 24-h training of the ML solver. This recursive feedback is solely used for training; we archive the ML results only for the first 1-h time step and discard the remaining 23-h time steps. The expectation is that the first 1-h prediction will have learned from fitting the subsequent 23-h evolution.

In online training, we call the Python ML routines from Fortran as we run the GEOS-Chem model, sampling the same conditions as the offline training. We employ the same ML solver architecture as the offline-trained ML solver without the recursive time horizon and with a learning rate of 1×10^{-5} . At each chemical time step, we load the ML solver parameters from the previous training time step, fit the ML solver for one epoch given all the training data at the current chemical time step, then save the ML solver parameters to be loaded in the next chemical time step.

The online framework as described above trains ML solvers from scratch starting from randomly initialized parameters. Recent ML work has suggested starting online training from pretrained offline ML models (Rasp, 2020; Watt-Meyer et al., 2021) in order to have a better initialization of ML parameters. We also tried this approach and results will be presented below.

Keller and Evans (2019) previously used a random forest ML solver to emulate the GEOS-Chem mechanism, but here we employ a neural network ML solver for two reasons. First, random forest algorithms are much slower. Keller and Evans (2019) found that their random forest solver was 85% slower than the reference Rosenbrock solver, while neural networks should be much faster (Kelp et al., 2020; Liu et al., 2021). Second, random forests are not easily amenable to online training because the growing of the architecture to incorporate more trees and branches further slows performance (Lakshminarayanan et al., 2015), whereas online neural network training simply updates parameters. We did not consider the convolutional neural network

architectures commonly used in computer vision applications (Schmidhuber, 2015) because convolutional layers typically perform calculations slower than simple fully connected layers.

3. Results

3.1 Reference GEOS-Chem simulation with Super-Fast mechanism

We conducted the reference GEOS-Chem simulation using the Super-Fast mechanism integrated with the KPP Rosenbrock solver for three years (2015-2017). 2015 was used for initialization, 2016 for training the ML algorithms, and 2017 for testing them. Here we compare this reference Super-Fast simulation for 2016 with the standard full-chemistry GEOS-Chem simulation in GEOS-Chem 12.0.0 including 228 coupled species to represent oxidant-aerosol chemistry. The intent is to check that the Super-Fast mechanism, although crude, provides a sufficiently reasonable tropospheric simulation in GEOS-Chem to serve as useful reference for ML application.

Figure 1 compares zonal mean profiles of ozone and NO_x concentrations for December-February (DJF) and June-August (JJA) 2016. The overall patterns are consistent. The Super-Fast mechanism simulates excessive NO_x in the Northern Hemisphere winter because it does not include loss from the nighttime N_2O_5 hydrolysis pathway. This may also explain the higher ozone in winter since N_2O_5 hydrolysis is a loss of odd oxygen. An additional important ozone and NO_x sink in the standard mechanism in winter is halogen chemistry (Sherwen et al., 2016), which is not part of the Super-Fast mechanism. The Super-Fast biases are consistent with the previous evaluation by Brown-Steiner et al. (2018).

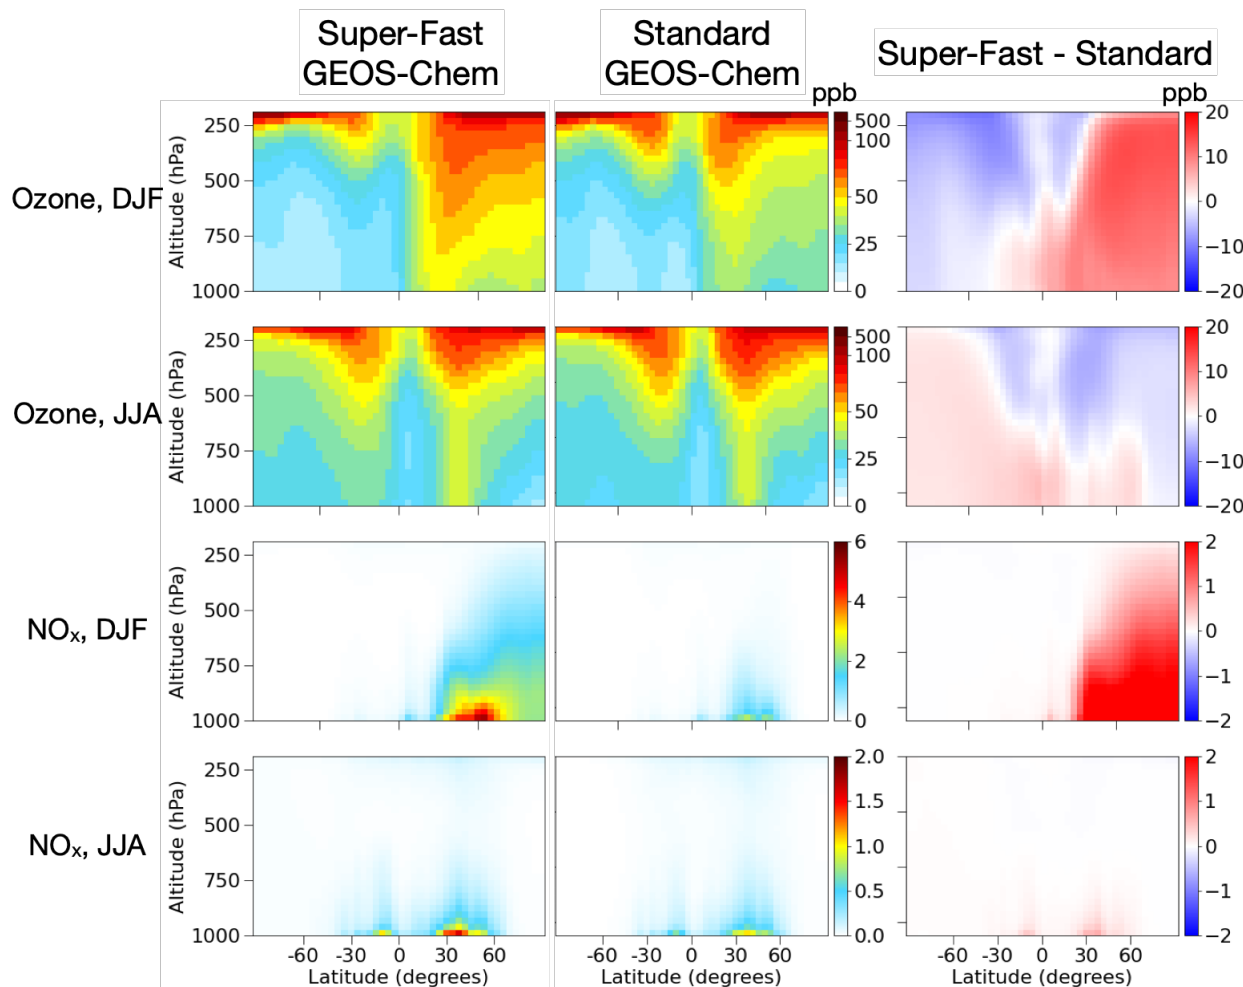


Figure 1. **Zonal mean concentrations of ozone and NO_x in GEOS-Chem.** The Figure compares GEOS-Chem simulations using the Super-Fast mechanism and the standard full-chemistry mechanism, both integrated with the KPP Rosenbrock solver. Results are for December-February (DJF) and June-August (JJA) 2016. Note different scales between panels.

Table 1 shows the global tropospheric ozone budgets in GEOS-Chem with the Super-Fast and standard mechanisms. The global tropospheric chemical production rate with Super-Fast is 10% lower, likely reflecting the crude treatment of NMVOCs (Wu et al., 2007). Ozone has a longer chemical lifetime with Super-Fast, likely due to absence of halogen chemistry and N₂O₅ hydrolysis. The global mean pressure-weighted tropospheric OH concentration is higher in Super-Fast (12.8×10^5 molecules cm⁻³) than in the standard mechanism (12.0×10^5 molecules cm⁻³), which can be explained by the higher NO_x.

Table 1. **Global budget of tropospheric ozone in GEOS-Chem^a**

	Super-Fast chemistry	Standard chemistry
Sources, Tg a⁻¹		
Chemical production	4480	4980
Cross-tropopause transport ^b	660	670
Total	5140	5650
Sinks, Tg a⁻¹		
Chemical loss	3980	4660
Dry deposition	1090	920
Wet deposition	69	70
Total	5140	5650
Ozone tropospheric mass, Tg	312	314
Ozone lifetime, days	22.2	20.3

^aFrom the GEOS-Chem version 12.0.0 simulation for year 2016 with the standard and Super-Fast chemical mechanisms. The budget is the annual mean for the odd oxygen (O_x) chemical family as defined in Hu et al., (2017) to account for rapid cycling between O_x components. In the Super-Fast mechanism this family is defined as $O_x \equiv O_3 + NO_2 + HNO_3$.

^bComputed from the vertical mass flux at 100 hPa.

Figure 2 shows the global distributions of ozone and NO_x concentrations at the surface and at 500 hPa simulated by GEOS-Chem with the standard and Super-Fast mechanisms for DJF and JJA. We explained above the higher wintertime ozone and NO_x in Super-Fast. Here we also see higher surface ozone and NO_x in Super-Fast over continental regions in the tropics and northern hemisphere summer, which may be due to the lack of peroxyacetylnitrate (PAN) as reservoir for NO_x . This may also explain the lower ozone in Super-Fast at 500 hPa in summer where decomposition of PAN provides an important source of NO_x in the standard mechanism. Super-Fast has very low ozone over eastern China in winter because of titration by NO and lack of radical production from HONO and NMVOC-produced formaldehyde photolysis, which sustain wintertime ozone production in the standard mechanism (Li et al., 2021).

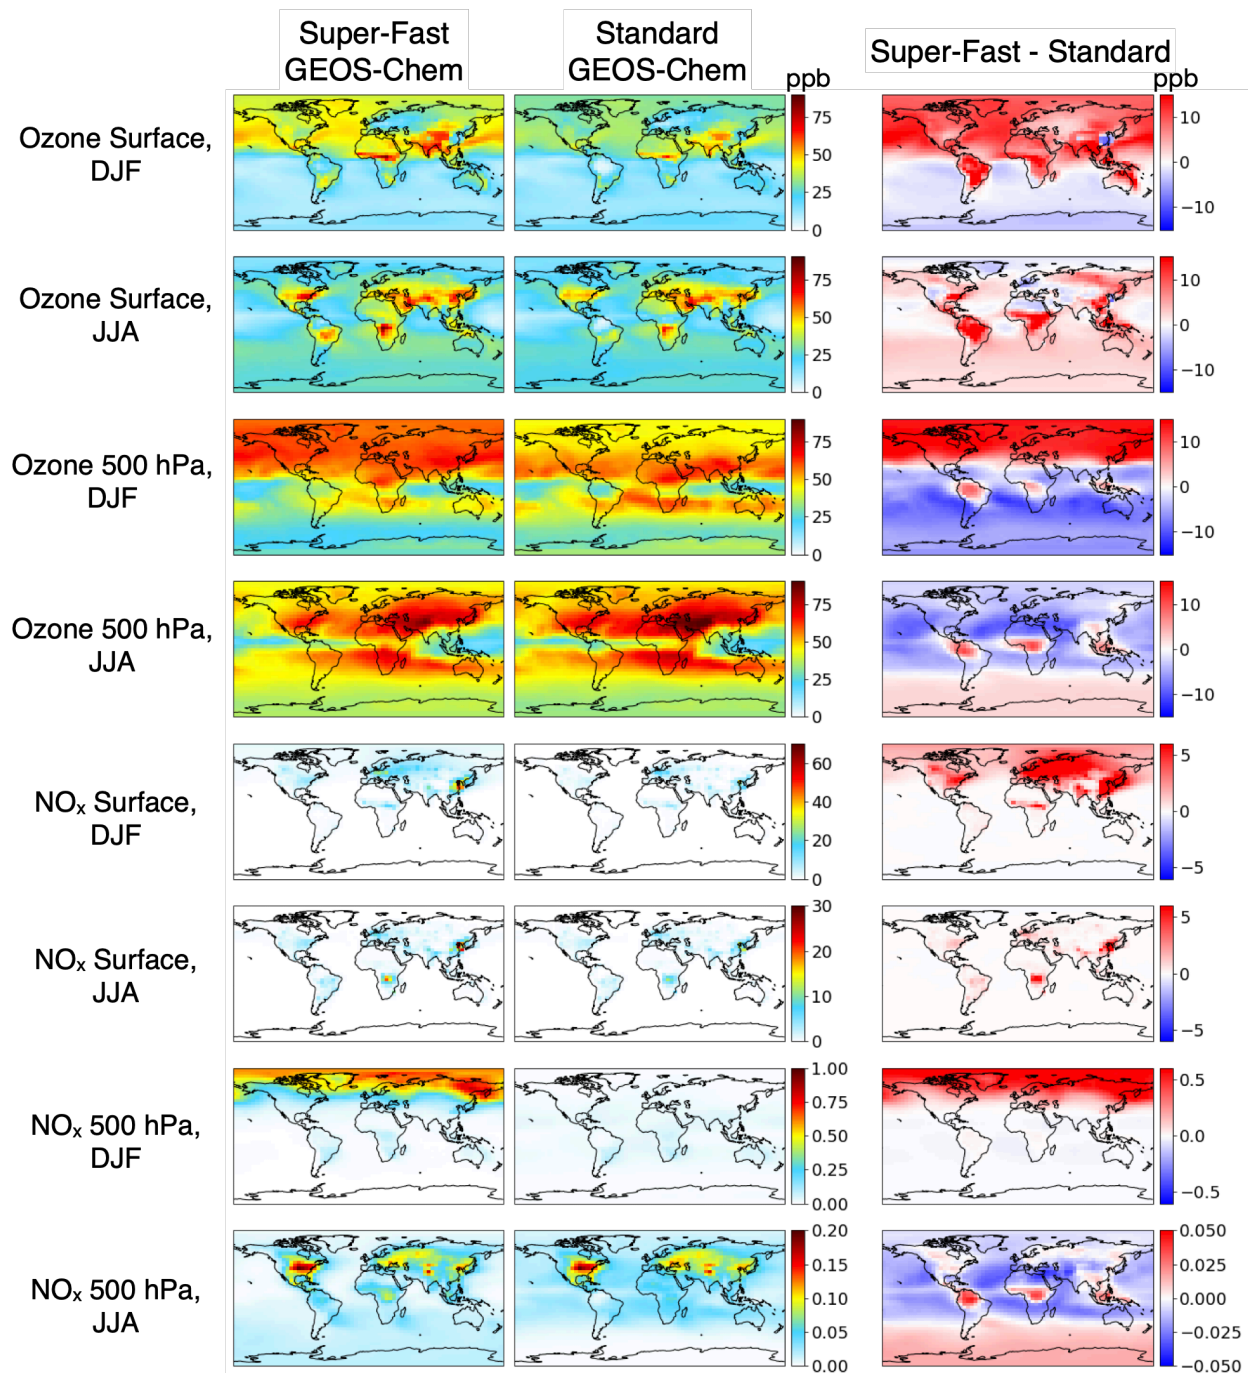


Figure 2. **Ozone and NO_x concentrations in surface air and at 500 hPa in GEOS-Chem.** The Figure compares GEOS-Chem simulations using the standard full-chemistry mechanism and the Super-Fast mechanism, both integrated with the KPP solver. Results are means for DJF and JJA 2016. Note different scales between panels.

3.2 Testing of offline and online ML solvers

We first test the accuracy and stability of the different offline and online ML solvers described in Section 2.3 by training a single-species ozone chemical solver, with all other species simulated with the Rosenbrock solver. The ML solver training is for June-August 2016 and the testing is for July 2017. Here and elsewhere, we will use four metrics to evaluate the ML solver (ML) relative to the reference Super-Fast simulation (R) for species $i \in [1, N]$ in a given grid cell:

$$\text{Eq1) } \textit{Normalized mean bias (NMB)} = \frac{\sum_{i=1}^N (ML_i - R_i)}{\sum_{i=1}^N (R_i)} \times 100$$

$$\text{Eq2) } \textit{Root mean square error (RMSE)} = \sqrt{\sum_{i=1}^N \frac{(ML_i - R_i)^2}{N}}$$

$$\text{Eq3) } \textit{Absolute error} = ML_i - R_i$$

$$\text{Eq4) } \textit{Fractional error} = \frac{2(ML_i - R_i)}{(ML_i + R_i)} \times 100$$

These metrics may be averaged spatially over the global domain and/or temporally over the period of interest.

Figure 3 shows the error statistics for surface ozone when using the different ML solvers. None of the ML solvers show runaway error growth, unlike in previous studies (Keller and Evans, 2019; Kelp et al., 2020), which we attribute to the relatively low dimensionality of the Super-Fast mechanism boosted by the use of the encoder/decoder to further reduce dimensionality.

The offline non-recursive ML solver shows large positive errors in remote regions, large negative errors in ozone production hotspots, and underestimate of temporal variability. We attribute this to the tendency of the ML solver trained on a randomly ordered ensemble of data to focus on simulating the mean. The offline-trained solver trained using a recursive 24-h feedback based on Kelp et al. (2020) improves the RMSE from 35.6 to 22.4 ppb, which is still very high. It features large positive land and negative ocean biases, as well as a small diurnal pattern in the RMSE. The reduction in error likely reflects better accounting of the predictable diurnal behavior of ozone concentrations.

When the offline 24-h recursive ML solver is retrained online within GEOS-Chem, we find that the RMSE decreases to 6.1 ppb. The online training on representative realizations in sequence (rather than random samples) prevents under/oversampling of specific chemical environments and captures better the temporal evolution of chemistry. But the patterns of biases learned from the offline training persist and are only partly corrected.

The ML solver trained online from scratch within GEOS-Chem performs the best by far and is the only viable solver for further consideration. It achieves a low RMSE of 1.3 ppb with fractional errors lower than 10%, which would be considered adequate for a global tropospheric ozone simulation (Hu et al., 2017). We attribute this success to the non-random order of the training, allowing the ML solver to emulate the temporal evolution within the CTM

environment. We find that training a ML solver online from scratch provides better performance than retraining an offline ML solver because the bias in the offline solver is difficult to unlearn.

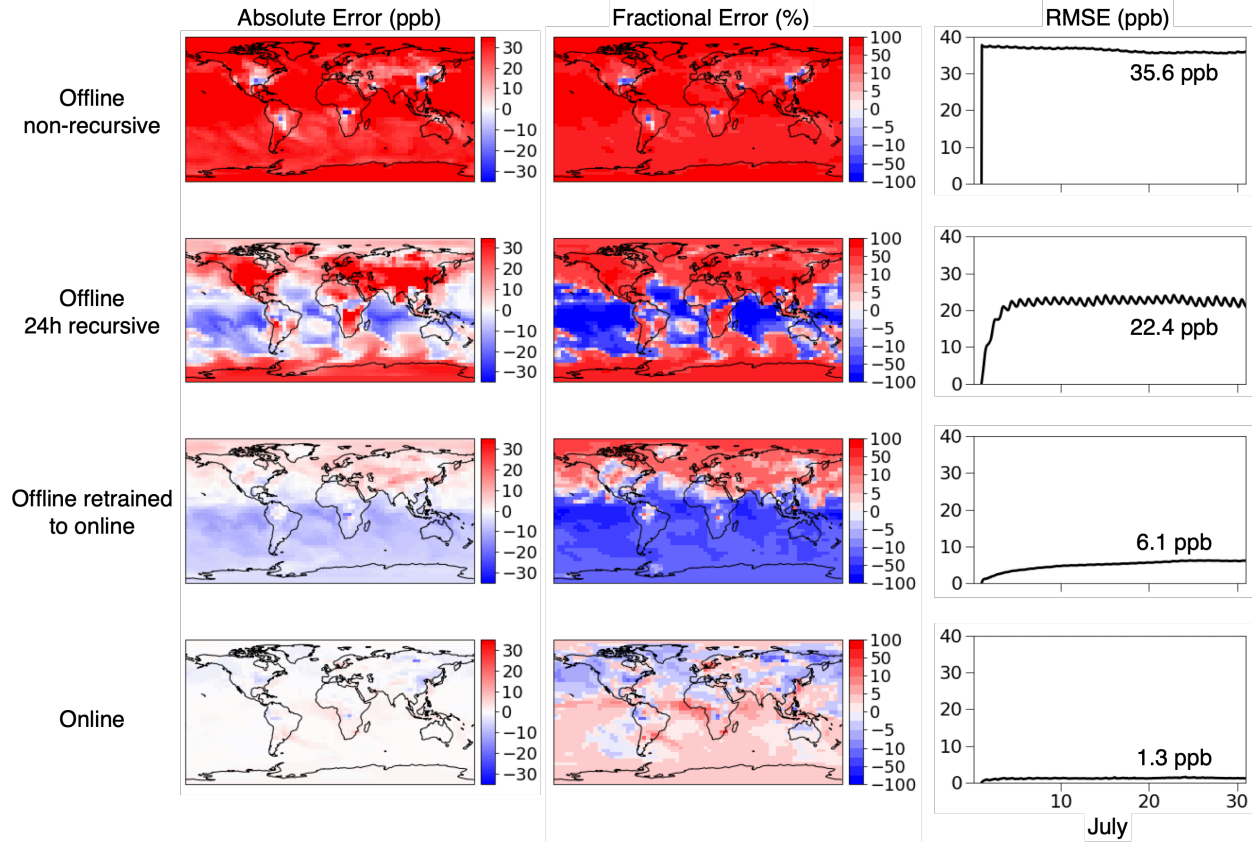


Figure 3. **Simulation of surface ozone by different ML solvers.** The left and middle panels show absolute and fractional errors at the end of a 31-day July 2017 simulation (24:00 UTC on July 31) relative to the reference Super-Fast baseline simulation using the KPP Rosenbrock solver. The right panel shows the temporal evolution of the global hourly root-mean-square error (RMSE) over the 31-day period. The mean RMSE for the last 10 days of July is given inset. See Section 2.3 for description of the different ML solvers. In this application the ML solver is applied to ozone only.

3.3 One-year simulation testing of online ML solver

We next apply the online ML solvers trained from scratch for all species to a 1-year GEOS-Chem simulation with the Super-Fast mechanism. We train the ML solvers on January-December 2016 and test them on January-December 2017. In early testing, we found that ML solvers for individual seasons (DJF, MAM, JJA, SON) outperformed ML solvers trained for the entire year. We also found that the ML solvers could not capture the discontinuity of hydrogen oxide radical ($\text{HO}_x \equiv \text{OH} + \text{HO}_2$) concentrations at sunrise/sunset because the ML training uses low-order continuous functions for its fits. Here we create separate ML solvers for OH and HO_2 at night, applying a log transformation to their concentrations in order to capture the fast nighttime decay. In the Super-Fast environment, it would alternatively be acceptable to set these concentrations to zero at night.

The online ML solver embedded within GEOS-Chem performs the chemical integration $5\times$ faster than the reference Super-Fast simulation (single Intel Broadwell CPU core; 2.10 GHz). This speedup is smaller than in Kelp et al. (2020) and Liu et al. (2021) because the Super-Fast mechanism is simpler and because of the overhead in accessing Python code at each time step. Further speedup could be achieved by reading the trained ML solver parameters through text files or by writing them in Fortran.

Figure 4 shows the daily evolution of the global normalized mean bias (NMB) for Super-Fast species over the full year. The global mean OH concentration computed with the ML solver ($13.2 \times 10^5 \text{ molecules cm}^{-3}$) reproduces that in the Super-Fast reference simulation ($13.1 \times 10^5 \text{ molecules cm}^{-3}$). Ozone has no significant bias averaged over the year (-0.3%) and remains within $\pm 9\%$ on a daily basis. HO_x is also successfully fitted, with an average bias of 1% and daily values within $\pm 6\%$. Other species except HNO_3 are also well fitted, and none shows error growth over the course of the simulation. The problem with HNO_3 is that – unlike other Super-Fast species – it does not have a chemical loss and there is therefore no first-order correction to growing biases in the ML solver. We also see in Figure 4 that the seasonal switch between solvers can rapidly erase the error from a poorly performing seasonal solver by switching to another solver in the next season. This suggests for future consideration that alternate application of separately trained solvers, or of the ML and reference solver, could significantly improve accuracy.

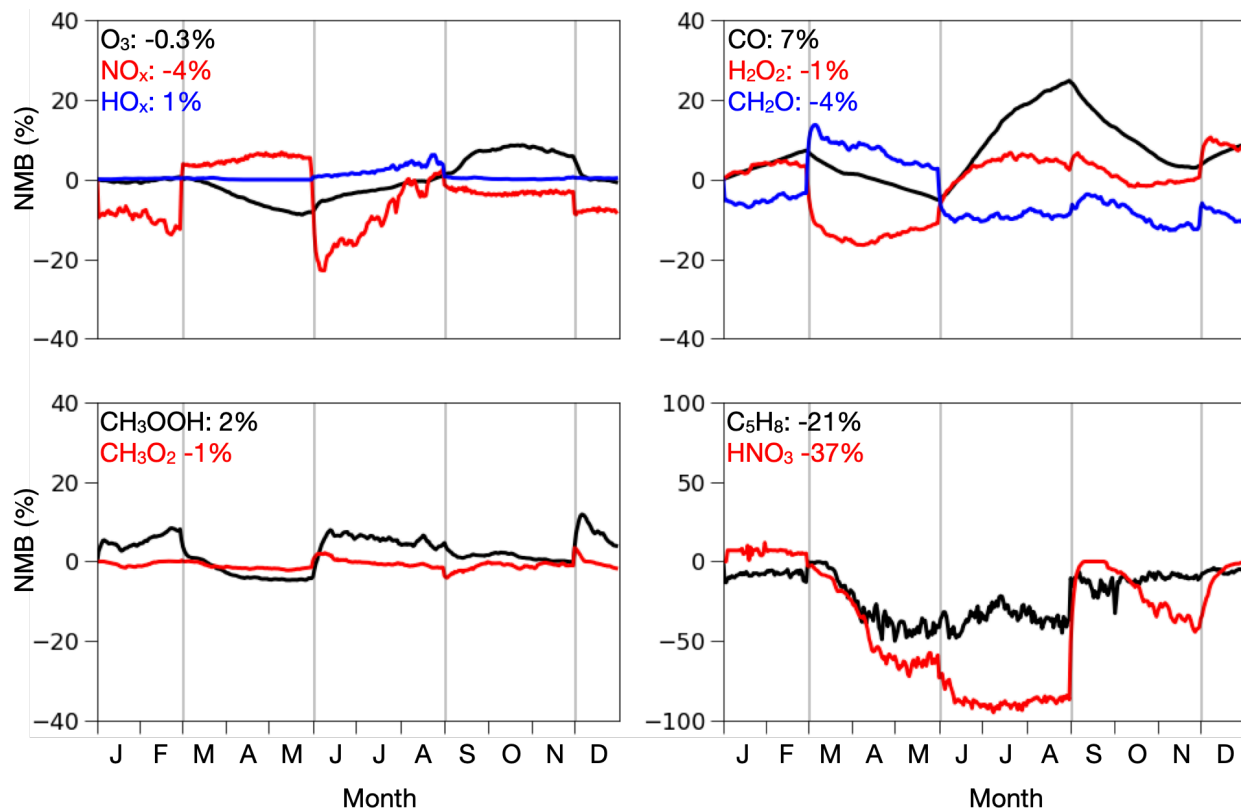


Figure 4. **Evolution of the global normalized mean bias (NMB) over a full-year simulation with the online ML solver.** The NMB is the global average calculated daily relative to the reference Super-Fast GEOS-Chem simulation. Results are shown for the 12 species of the Super-Fast mechanism. NO and NO₂ are grouped as NO_x, and OH and HO₂ are grouped as HO_x. Annual mean biases for each species are indicated in legend. Vertical lines indicate seasonal switches between different ML solvers.

Figure 5 compares zonal mean ozone and NO_x concentrations for DJF and JJA between simulations with the ML online and reference solvers, and Figure 6 compares surface and 500 hPa concentrations. Although the overall patterns are consistent, we find large errors in individual latitudinal bands, up to 20% for ozone and 100% for NO_x. The errors are largest at remote latitudes and high altitudes due to chemical error accumulation as air ages. The largest errors are in polar sunlit conditions where the effect of chemical aging during long-range transport is particularly important.

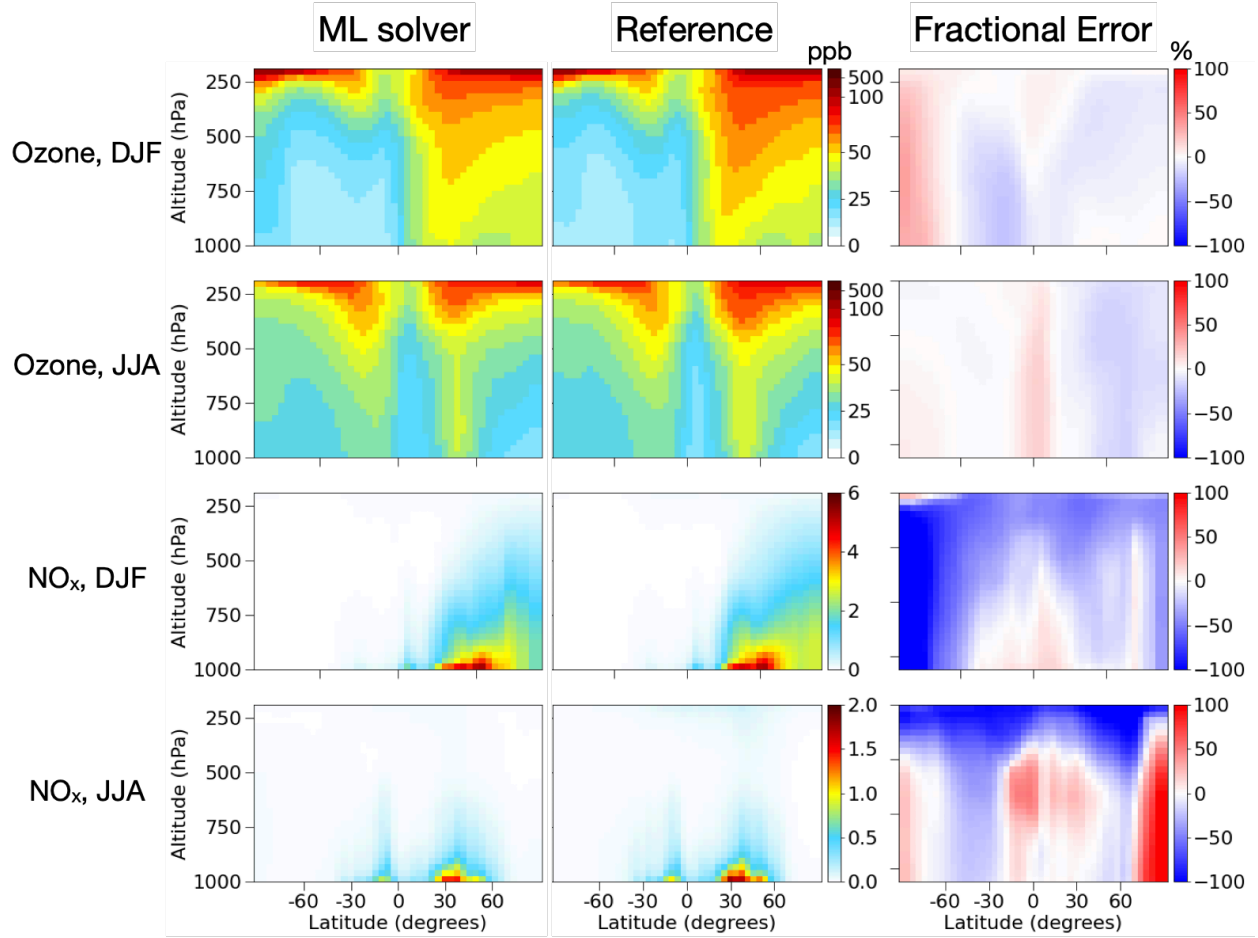


Figure 5. Zonal mean ozone and NO_x concentrations in GEOS-Chem using the online ML solver and compared to the reference simulation. Results are averages for DJF and JJA 2017. The Reference panels are the same as the Super-Fast panels in Figure 1.

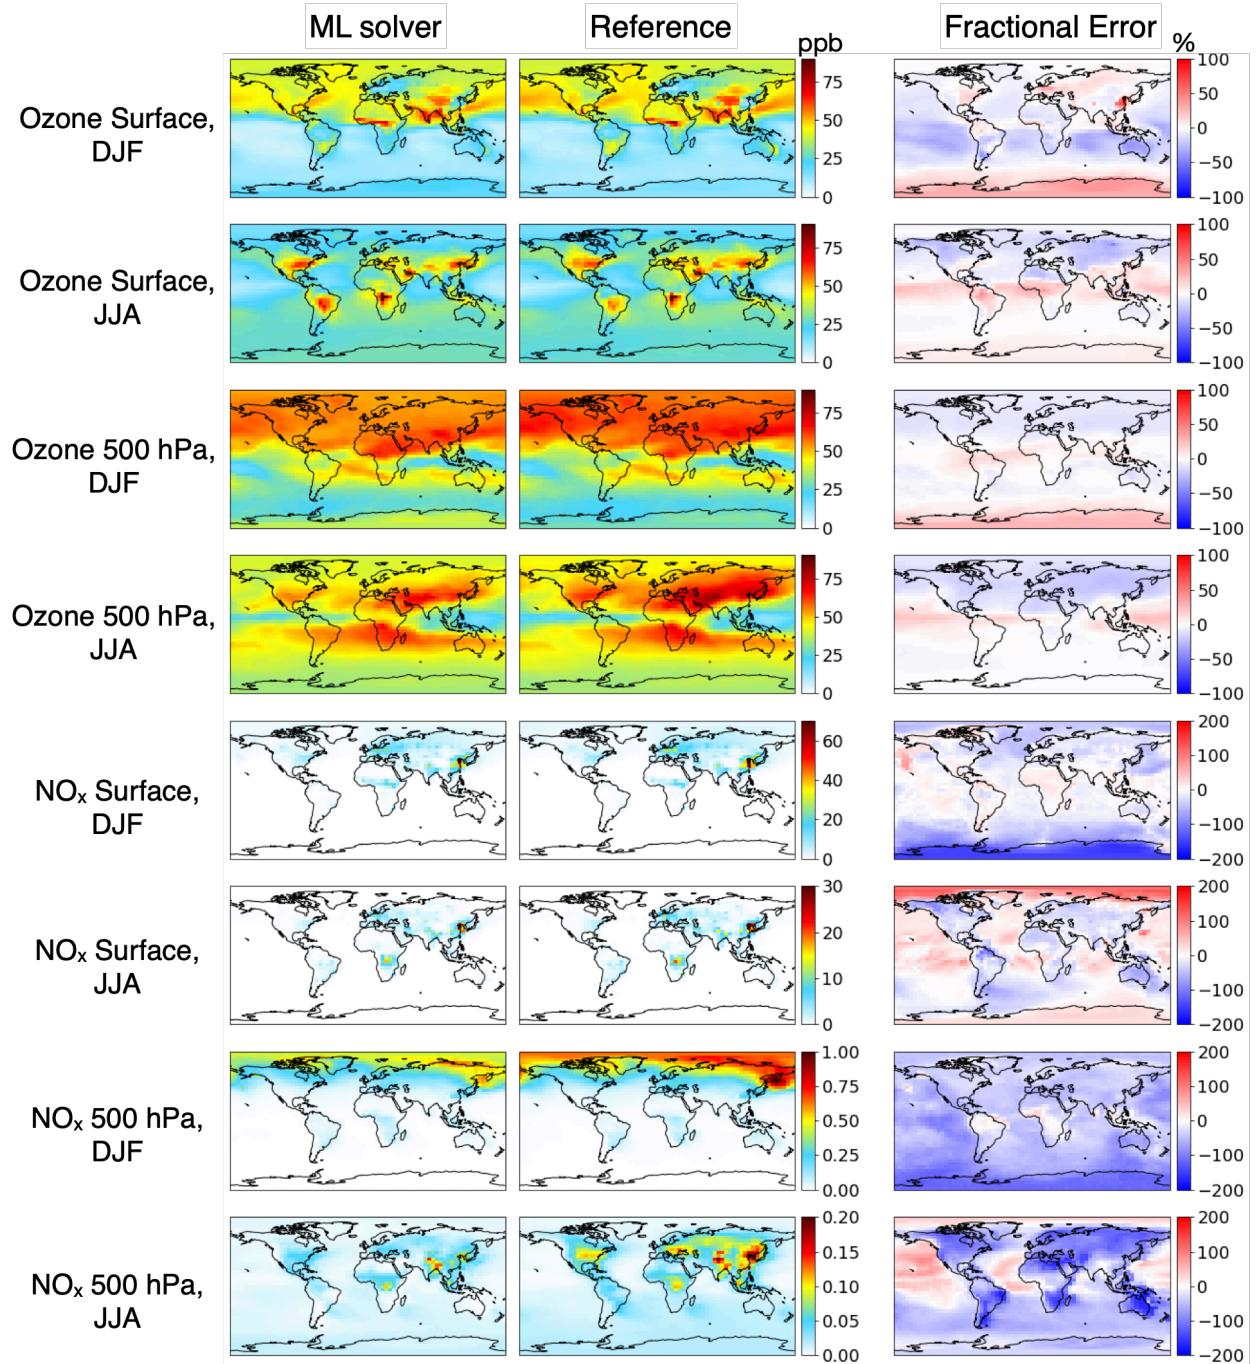


Figure 6. **Surface ozone and NO_x concentrations in GEOS-Chem using the online ML solver and compared to the reference simulation.** Results are averages for DJF and JJA 2017. The Reference panels are the same as the Super-Fast panels in Figure 2.

Figure 7 illustrates the ability of the ML solver to reproduce diurnal and synoptic variations of surface ozone for polluted (Beijing) and remote (Cape Verde) conditions. The Beijing time series shows large diurnal variation due to fast production in the daytime followed by fast loss at night from deposition and reaction with NO. Superimposed on this diurnal variation is synoptic (multi-day) variability with pollution episodes approaching 100 ppb. The ML solver reproduces these

features without systematic bias. Cape Verde shows by contrast much lower variability and no significant diurnal cycle, because of low NO_x and slow ozone deposition to the ocean, and again this is well reproduced by the ML solver.




Figure 7. **Hourly time series of surface ozone concentrations in Beijing (39.9 N, 116 E) and Cape Verde (16.5 N, 23.0 W) in July 2017.** The GEOS-Chem simulation using the online ML solver is compared to the reference simulation using the KPP Rosenbrock solver.

3.4 Alternative configurations

Although the online-trained ML solver as described here enables stable full-year global tropospheric chemistry simulations with reasonable accuracy for the main oxidants ozone and OH, there are large regional inaccuracies for species such as NO_x . We tried different approaches to overcome these inaccuracies but without success. Predicting the change in concentrations for longer-lived species rather than the concentrations themselves (Keller and Evans, 2019) worsened the fit. Grouping NO and NO_2 as NO_x in the prediction did not improve results and required a separate step to resolve the partitioning. Training separate ML solvers for different regions such as land, ocean, and upper troposphere did not improve results and led to high errors at the boundaries. Applying a log transform to all input species or using a L-1 norm (instead of least-squares) for fitting prevented successful simulation of polluted grid cells such as ozone in Beijing.

We found that fitting individual seasons with the ML solver, as opposed to the full year, led to significant improvement of results. Some of that improvement may relate to error correction in the switch between solvers from one season to the next (Figure 4). This suggests that alternating between independently trained ML solvers or between the ML solver and the Rosenbrock solver could help to reduce error. A similar approach would be to implement an online bias corrector (Ivatt and Evans, 2020) that either nudges the ML solver toward the Rosenbrock reference or learns to call the Rosenbrock solver when the ML solver starts to fail (Zheng et al., 2019). These could be directions for future research.

4. Conclusions

This work explored the capability of machine learning (ML) to speed up the kinetic integration of chemical mechanisms in full-year global simulations of atmospheric chemistry. The motivation was to remove a major computational bottleneck in global atmospheric chemistry models and for the inclusion of atmospheric chemistry in Earth system models (ESMs). A challenge was to avoid the runaway error growth that affected previous ML application to global models (Keller and Evans, 2019).

Chemical mechanisms in current-generation global atmospheric chemistry models such as GEOS-Chem typically include over 200 species (Shen et al., 2020) and kinetic integration is done with high-order implicit solvers (4th-order Rosenbrock in GEOS-Chem, implemented through KPP). The high dimensionality of the mechanism represents a major challenge for the application of ML solvers. As a first step and to avoid this complication, we implemented in GEOS-Chem the Super-Fast mechanism including only 12 coupled species to represent tropospheric oxidant chemistry (Lamarque et al., 2013; Brown-Steiner et al., 2018). We applied to that mechanism a neural network ML solver equipped with an autoencoder (Kelp et al., 2020), and compared the resulting simulation in GEOS-Chem to the reference simulation with the Rosenbrock solver.

We tried two approaches for training the ML solver over the global 3-D domain, offline using archived inputs/outputs from 1-h chemical integration time steps with the reference simulation, and online synchronously with the reference simulation. We found that the common practice of offline training resulted in large errors. We attributed these errors to training on a randomly ordered ensemble of data, and to overfitting caused by multiple passes through the data. Using a recursive algorithm over 24-h time horizons to capture diurnal and longer modes (Kelp et al., 2020) led to some improvement but errors were still large. The ML solver trained online had much better success, which we attribute to representative sampling of the GEOS-Chem simulation as it progresses in time. Online training from scratch performed much better than pre-training offline.

We applied the online ML solver to a 1-year GEOS-Chem simulation with the Super-Fast mechanism. The ML solver reduced the computational cost of the chemical integration five-fold. We found that training the ML solver for individual species and seasons led to best results. The ML solver achieved a stable simulation over the 1-year simulation period with no error growth. Global biases for ozone and OH were insignificant on an annual basis. Global daily biases for ozone were at most 9%. The ML solver was successful at reproducing the diurnal and synoptic variations of ozone at polluted and clean sites, including events of high concentrations. There were however systematic patterns of biases, worst in chemically aged air such as polar sunlit conditions and the middle/upper troposphere, and large biases for NO_x. Using different ML solver configurations did not readily solve that problem.

An important outcome of our work was to achieve for the first time a stable global simulation of atmospheric chemistry with a ML solver and with multi-fold improvement in computational performance. We found in the process that online training of the ML solver is considerably superior to offline training. Our application was limited to the oversimplified Super-Fast

mechanism, and even then regional biases could be large. Future work should focus on bias avoidance and there has been some studies to that effect (Zheng et al, 2019; Ivatt and Evans, 2020). The curse of dimensionality in full atmospheric chemistry mechanisms could be alleviated by the encoder/decoder processing to reduce dimensionality in the ML solver (Kelp et al., 2020).

Acknowledgements and Data

This work was supported by the US Environmental Protection Agency (EPA) Science To Achieve Results (STAR) program under grant number 840014 to Harvard University. The views and opinions expressed in this article are those of the authors and do not represent the official views of the U.S. EPA. The authors thank Kelvin Bates for assistance with KPP, Ben Brown-Steiner for Super-Fast troubleshooting, and Chris Tessum and Drew Pendergrass for valuable discussion.

The standard GEOS-Chem code is available through <https://doi.org/10.5281/zenodo.1343547> (version 12.0.0). The KPP implementation of the Super-Fast mechanism and the training and testing methods for the neural network ML solver can be found at Zenodo (<https://zenodo.org/record/5736524>).

REFERENCES

- Amos, H.M., Jacob, D.J., Holmes, C.D., Fisher, J.A., Wang, Q., Yantosca, R.M., Corbitt, E.S., Galarneau, E., Rutter, A.P., Gustin, M.S., Steffen, A., Schauer, J.J., Graydon, J.A., Louis, V.L.St., Talbot, R.W., Edgerton, E.S., Zhang, Y., Sunderland, E.M. (2012). Gas-particle partitioning of atmospheric Hg(II) and its effect on global mercury deposition. *Atmospheric Chem. Phys.* 12, 591–603. <https://doi.org/10.5194/acp-12-591-2012>
- Bey, I., Jacob, D.J., Yantosca, R.M., Logan, J.A., Field, B.D., Fiore, A.M., Li, Q., Liu, H.Y., Mickley, L.J., Schultz, M.G. (2001). Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation. *J. Geophys. Res.* *Atmospheres* 106, 23073–23095. <https://doi.org/10.1029/2001JD000807>
- Bian, H., Prather, M.J. (2002). Fast-J2: Accurate Simulation of Stratospheric Photolysis in Global Chemical Models. *J. Atmospheric Chem.* 41, 281–296. <https://doi.org/10.1023/A:1014980619462>
- Brasseur, G.P., Jacob, D.J. (2017). Modeling of Atmospheric Chemistry. Cambridge University Press, Cambridge. <https://doi.org/10.1017/9781316544754>
- Brown-Steiner, B., Selin, N.E., Prinn, R., Tilmes, S., Emmons, L., Lamarque, J.-F., Cameron-Smith, P. (2018). Evaluating simplified chemical mechanisms within present-day simulations of the Community Earth System Model version 1.2 with CAM4 (CESM1.2 CAM-chem): MOZART-4 vs. Reduced Hydrocarbon vs. Super-Fast chemistry. *Geosci. Model Dev.* 11, 4155–4174. <https://doi.org/10.5194/gmd-11-4155-2018>
- Cameron-Smith, P., Lamarque, J.-F., Connell, P., Chuang, C., Vitt, F. (2006). Toward an Earth system model: atmospheric chemistry, coupling, and petascale computing. *J. Phys. Conf. Ser.* 46, 343–350. <https://doi.org/10.1088/1742-6596/46/1/048>
- National Research Council (2012). A National Strategy for Advancing Climate Modeling. <https://doi.org/10.17226/13430>
- Eastham, S.D., Long, M.S., Keller, C.A., Lundgren, E., Yantosca, R.M., Zhuang, J., Li, C., Lee, C.J., Yannetti, M., Auer, B.M., Clune, T.L., Kouatchou, J., Putman, W.M., Thompson, M.A., Trayanov, A.L., Molod, A.M., Martin, R.V., Jacob, D.J. (2018). GEOS-Chem High Performance (GCHP v11-02c): a next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications. *Geosci. Model Dev.* 11, 2941–2953. <https://doi.org/10.5194/gmd-11-2941-2018>
- Eastham, S.D., Weisenstein, D.K., Barrett, S.R.H. (2014). Development and evaluation of the unified tropospheric–stratospheric chemistry extension (UCX) for the global chemistry-transport model GEOS-Chem. *Atmospheric Environment* 89, 52–63. <https://doi.org/10.1016/j.atmosenv.2014.02.001>
- Fisher, J.A., Murray, L.T., Jones, D.B.A., Deutscher, N.M. (2017). Improved method for linear carbon monoxide simulation and source attribution in atmospheric chemistry models illustrated using GEOS-Chem v9. *Geosci. Model Dev.* 10, 4129–4144. <https://doi.org/10.5194/gmd-10-4129-2017>
- Guenther, A.B., Jiang, X., Heald, C.L., Sakulyanontvittaya, T., Duhl, T., Emmons, L.K., Wang, X. (2012). The Model of Emissions of Gases and Aerosols from Nature version 2.1 (MEGAN2.1): an extended and updated framework for modeling biogenic emissions. *Geosci. Model Dev.* 5, 1471–1492. <https://doi.org/10.5194/gmd-5-1471-2012>
- He, K., Zhang, X., Ren, S., Sun, J. (2016). Deep residual learning for image recognition. *In Proc. 29th IEEE Conf. Comput. Vis. Pattern Recognit.* 770–778.

- Hochreiter, S., Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Comput.* 9, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Hoesly, R.M., Smith, S.J., Feng, L., Klimont, Z., Janssens-Maenhout, G., Pitkanen, T., Seibert, J.J., Vu, L., Andres, R.J., Bolt, R.M., Bond, T.C., Dawidowski, L., Kholod, N., Kurokawa, J., Li, M., Liu, L., Lu, Z., Moura, M.C.P., O'Rourke, P.R., Zhang, Q. (2018). Historical (1750–2014) anthropogenic emissions of reactive gases and aerosols from the Community Emissions Data System (CEDS). *Geosci. Model Dev.* 11, 369–408. <https://doi.org/10.5194/gmd-11-369-2018>
- Hudman, R.C., Moore, N.E., Mebust, A.K., Martin, R.V., Russell, A.R., Valin, L.C., Cohen, R.C. (2012). Steps towards a mechanistic model of global soil nitric oxide emissions: implementation and space based-constraints. *Atmospheric Chem. Phys.* 12, 7779–7795. <https://doi.org/10.5194/acp-12-7779-2012>
- Ioffe, S., Szegedy, C. (2015). Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. ArXiv:150203167 Cs.
- Ivatt, P.D., Evans, M.J. (2020). Improving the prediction of an atmospheric chemistry transport model using gradient-boosted regression trees. *Atmospheric Chem. Phys.* 20, 8063–8082. <https://doi.org/10.5194/acp-20-8063-2020>
- Keller, C.A., Evans, M.J. (2019). Application of random forest regression to the calculation of gas-phase chemistry within the GEOS-Chem chemistry model v10. *Geosci. Model Dev.* 12, 1209–1225. <https://doi.org/10.5194/gmd-12-1209-2019>
- Kelp, M.M., Jacob, D.J., Kutz, J.N., Marshall, J.D., Tessum, C.W. (2020). Toward Stable, General Machine-Learned Models of the Atmospheric Chemical System. *J. Geophys. Res. Atmospheres* 125, e2020JD032759. <https://doi.org/10.1029/2020JD032759>
- Kingma, D.P., Ba, J. (2017). Adam: A Method for Stochastic Optimization. ArXiv:14126980 Cs.
- Kramer, M.A. (1991). Nonlinear principal component analysis using autoassociative neural networks. *AIChE J.* 37, 233–243. <https://doi.org/10.1002/aic.690370209>
- Krasnopolsky, V.M., Fox-Rabinovitz, M.S., Chalikov, D.V. (2005). New Approach to Calculation of Atmospheric Model Physics: Accurate and Fast Neural Network Emulation of Longwave Radiation in a Climate Model. *Mon. Weather Rev.* 133, 1370–1383. <https://doi.org/10.1175/MWR2923.1>
- Lagerquist, R., Turner, D., Ebert-Uphoff, I., Stewart, J., Hagerty, V. (2021). Using deep learning to emulate and accelerate a radiative-transfer model. *J. Atmospheric Ocean. Technol.* 1. <https://doi.org/10.1175/JTECH-D-21-0007.1>
- Lakshminarayanan, B., Roy, D.M., Teh, Y.W. (2015). Mondrian Forests: Efficient Online Random Forests. ArXiv:14062673 Cs Stat.
- Lamarque, J.-F., Shindell, D.T., Josse, B., Young, P.J., Cionni, I., Eyring, V., Bergmann, D., Cameron-Smith, P., Collins, W.J., Doherty, R., Dalsoren, S., Faluvegi, G., Folberth, G., Ghan, S.J., Horowitz, L.W., Lee, Y.H., MacKenzie, I.A., Nagashima, T., Naik, V., Plummer, D., Righi, M., Rumbold, S.T., Schulz, M., Skeie, R.B., Stevenson, D.S., Strode, S., Sudo, K., Szopa, S., Voulgarakis, A., Zeng, G. (2013). The Atmospheric Chemistry and Climate Model Intercomparison Project (ACCMIP): overview and description of models, simulations and climate diagnostics. *Geosci. Model Dev.* 6, 179–206. <https://doi.org/10.5194/gmd-6-179-2013>
- Li, K., Jacob, D.J., Liao, H., Qiu, Y., Shen, L., Zhai, S., Bates, K.H., Sulprizio, M.P., Song, S., Lu, X., Zhang, Q., Zheng, B., Zhang, Y., Zhang, J., Lee, H.C., Kuk, S.K. (2021). Ozone

- pollution in the North China Plain spreading into the late-winter haze season. *Proc. Natl. Acad. Sci.* 118, e2015797118. <https://doi.org/10.1073/pnas.2015797118>
- Li, M., Soltanolkotabi, M., Oymak, S. (2019). Gradient Descent with Early Stopping is Provably Robust to Label Noise for Overparameterized Neural Networks. ArXiv:1903.11680 Cs Stat.
- Liu, C., Zhang, H., Cheng, Z., Shen, J., Zhao, J., Wang, Y., Wang, S., Cheng, Y. (2021). Emulation of an atmospheric gas-phase chemistry solver through deep learning: Case study of Chinese Mainland. *Atmospheric Pollut. Res.* 12, 101079. <https://doi.org/10.1016/j.apr.2021.101079>
- Liu, H., Jacob, D.J., Bey, I., Yantosca, R.M. (2001). Constraints from ^{210}Pb and ^7Be on wet deposition and transport in a global three-dimensional chemical tracer model driven by assimilated meteorological fields. *J. Geophys. Res. Atmospheres* 106, 12109–12128. <https://doi.org/10.1029/2000JD900839>
- Lusch, B., Kutz, J.N., Brunton, S.L. (2018). Deep learning for universal linear embeddings of nonlinear dynamics. *Nat. Commun.* 9, 4950. <https://doi.org/10.1038/s41467-018-07210-0>
- McCloskey, M., Cohen, N.J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem, in: Bower, G.H. (Ed.), *Psychology of Learning and Motivation*. Academic Press, pp. 109–165. [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)
- McLinden, C.A., Olsen, S.C., Hannegan, B., Wild, O., Prather, M.J., Sundet, J. (2000). Stratospheric ozone in 3-D models: A simple chemistry and the cross-tropopause flux. *J. Geophys. Res. Atmospheres* 105, 14653–14665. <https://doi.org/10.1029/2000JD900124>
- Murray, L.T. (2016). Lightning NO_x and Impacts on Air Quality. *Curr Pollution Rep* 2, 115–133. <https://doi.org/10.1007/s40726-016-0031-7>
- Murray, L.T., Jacob, D.J., Logan, J.A., Hudman, R.C., Koshak, W.J. (2012). Optimized regional and interannual variability of lightning in a global chemical transport model constrained by LIS/OTD satellite data: IAV OF LIGHTNING CONSTRAINED BY LIS/OTD. *J. Geophys. Res. Atmospheres* 117. <https://doi.org/10.1029/2012JD017934>
- Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Netw.* 113, 54–71. <https://doi.org/10.1016/j.neunet.2019.01.012>
- Philip, S., Martin, R.V., Keller, C.A. (2016). Sensitivity of chemistry-transport model simulations to the duration of chemical and transport operators: a case study with GEOS-Chem v10-01. *Geosci. Model Dev.* 9, 1683–1695. <https://doi.org/10.5194/gmd-9-1683-2016>
- Randerson, J.T., Van Der Werf, G.R., Giglio, L., Collatz, G.J., Kasibhatla, P.S. (2015). Global Fire Emissions Database, Version 4.1 (GFEDv4). *ORNL DAAC*. <https://doi.org/10.3334/ORNLDAAAC/1293>
- Rasp, S. (2020). Coupled online learning as a way to tackle instabilities and biases in neural network parameterizations: general algorithms and Lorenz 96 case study (v1.0). *Geosci. Model Dev.* 13, 2185–2196. <https://doi.org/10.5194/gmd-13-2185-2020>
- Rumelhart, D.E., Hinton, G.E., Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature* 323, 533–536. <https://doi.org/10.1038/323533a0>
- Sandu, A., Sander, R. (2006). Technical note: Simulating chemical systems in Fortran90 and Matlab with the Kinetic PreProcessor KPP-2.1. *Atmospheric Chem. Phys.* 6, 187–195. <https://doi.org/10.5194/acp-6-187-2006>

- Sandu, A., Verwer, J.G., Blom, J.G., Spee, E.J., Carmichael, G.R., Potra, F.A. (1997). Benchmarking stiff ode solvers for atmospheric chemistry problems II: Rosenbrock solvers. *Atmos. Environ.* 31, 3459–3472. [https://doi.org/10.1016/S1352-2310\(97\)83212-8](https://doi.org/10.1016/S1352-2310(97)83212-8)
- Santurkar, S., Tsipras, D., Ilyas, A., Madry, A. (2019). How Does Batch Normalization Help Optimization? ArXiv:180511604 Cs Stat.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Netw.* 61, 85–117. <https://doi.org/10.1016/j.neunet.2014.09.003>
- Shen, L., Jacob, D.J., Santillana, M., Wang, X., Chen, W. (2020). An adaptive method for speeding up the numerical integration of chemical mechanisms in atmospheric chemistry models: application to GEOS-Chem version 12.0.0. *Geosci. Model Dev.* 13, 2475–2486. <https://doi.org/10.5194/gmd-13-2475-2020>
- Sherwen, T., Schmidt, J.A., Evans, M.J., Carpenter, L.J., Großmann, K., Eastham, S.D., Jacob, D.J., Dix, B., Koenig, T.K., Sinreich, R., Ortega, I., Volkamer, R., Saiz-Lopez, A., Prados-Roman, C., Mahajan, A.S., Ordóñez, C. (2016). Global impacts of tropospheric halogens (Cl, Br, I) on oxidants and composition in GEOS-Chem. *Atmospheric Chemistry and Physics* 16, 12239–12271. <https://doi.org/10.5194/acp-16-12239-2016>
- Sportisse, B., Djouad, R. (2000). Reduction of Chemical Kinetics in Air Pollution Modeling. *J. Comput. Phys.* 164, 354–376. <https://doi.org/10.1006/jcph.2000.6601>
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R. (2014). Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.
- Sturm, P.O., Wexler, A.S. (2020). A mass- and energy-conserving framework for using machine learning to speed computations: a photochemistry example. *Geosci. Model Dev.* 13, 4435–4442. <https://doi.org/10.5194/gmd-13-4435-2020>
- Wang, X., Jacob, D.J., Downs, W., Zhai, S., Zhu, L., Shah, V., Holmes, C.D., Sherwen, T., Alexander, B., Evans, M.J., Eastham, S.D., Neuman, J.A., Veres, P., Koenig, T.K., Volkamer, R., Huey, L.G., Bannan, T.J., Percival, C.J., Lee, B.H., Thornton, J.A. (2021). Global tropospheric halogen (Cl, Br, I) chemistry and its impact on oxidants. *Atmospheric Chem. Phys. Discuss.* 1–34. <https://doi.org/10.5194/acp-2021-441>
- Watt-Meyer, O., Brenowitz, N.D., Clark, S.K., Henn, B., Kwa, A., McGibbon, J., Perkins, W.A., Bretherton, C.S. (2021). Correcting weather and climate models by machine learning nudged historical simulations. *Geophys. Res. Lett.*, e2021GL092555. <https://doi.org/10.1029/2021GL092555>
- Wild, O., Prather, M.J. (2000). Excitation of the primary tropospheric chemical mode in a global three-dimensional model. *Journal of Geophysical Research: Atmospheres* 105, 24647–24660. <https://doi.org/10.1029/2000JD900399>
- Wu, S., Mickley, L.J., Jacob, D.J., Logan, J.A., Yantosca, R.M., Rind, D. (2007). Why are there large differences between models in global budgets of tropospheric ozone? *J. Geophys. Res. Atmospheres* 112. <https://doi.org/10.1029/2006JD007801>
- You, K., Long, M., Wang, J., Jordan, M.I. (2019). How Does Learning Rate Decay Help Modern Neural Networks? ArXiv:190801878 Cs Stat.
- Young, T.R., Boris, J.P. (1977). A numerical technique for solving stiff ordinary differential equations associated with the chemical kinetics of reactive-flow problems. *J. Phys. Chem.* 81, 2424–2427. <https://doi.org/10.1021/j100540a018>

- Jacobson, M.Z., Turco, R.P. (1994). SMVGEAR: A sparse-matrix, vectorized gear code for atmospheric models. *Atmos. Environ.* 28, 273–284. [https://doi.org/10.1016/1352-2310\(94\)90102-3](https://doi.org/10.1016/1352-2310(94)90102-3)
- Zheng, S., Farahat, A., Gupta, C. (2019). Generative Adversarial Networks for Failure Prediction. ArXiv:191002034 Cs Stat.