
Difference Learning for Air Quality Forecasting Transport Emulation

Reed R. Chen

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
reed.chen@jhuapl.edu

Christopher Ribaud

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
chris.ribaud@jhuapl.edu

Jennifer Sleeman

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
jennifer.sleeman@jhuapl.edu

Chace Ashcraft

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
chace.ashcraft@jhuapl.edu

Collin Kofroth

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
collin.kofroth@jhuapl.edu

Marisa Hughes

Johns Hopkins Applied Physics Laboratory
Laurel, MD 20723
marisa.hughes@jhuapl.edu

Ivanka Stajner

NOAA
College Park, MD 20740
ivanka.stajner@noaa.gov

Kevin Viner

NOAA
College Park, MD 20740
kevin.viner@noaa.gov

Kai Wang

NOAA
College Park, MD 20740
kai.wang@noaa.gov

Abstract

Human health is negatively impacted by poor air quality including increased risk for respiratory and cardiovascular disease. Due to a recent increase in extreme air quality events, both globally and locally in the United States, finer resolution air quality forecasting guidance is needed to effectively adapt to these events. The National Oceanic and Atmospheric Administration provides air quality forecasting guidance for the Continental United States. Their air quality forecasting model is based on a 15 km spatial resolution; however, the goal is to reach a three km spatial resolution. This is currently not feasible due in part to prohibitive computational requirements for modeling the transport of chemical species. In this work, we describe a deep learning transport emulator that is able to reduce computations while maintaining skill comparable with the existing numerical model. We show how this method maintains skill in the presence of extreme air quality events, making it a potential candidate for operational use. We also explore evaluating how well this model maintains the physical properties of the modeled transport for a given set of species.

1 Introduction

There has been a significant increase in high pollution air quality (AQ) events. These events are shown to have a strong sensitivity to extreme meteorological events such as heat waves [2]. Increased wildfire activity has specifically contributed to a sudden increase in fine particulate matter (PM_{2.5}) AQ pollution in the United States [3]. Studies show that increased emissions and climate change can

negatively impact air quality [9]. Since AQ has a direct correlation with increases in human-related illness and mortality [6, 8], it is important for AQ forecasting to address this changing environment. Operational AQ forecasting guidance is provided by the National Oceanic and Atmospheric Administration (NOAA) for the Continental United States (CONUS). The NOAA forecasting guidance system is computationally challenged by the transport of chemical species which involves solving a set of physical governing equations. This transport is a critical component of modeling AQ and presents a challenge for reaching finer spatial resolutions. In this study, we explore the feasibility of using deep learning as a transport emulator to provide a speed-up in overall computation, potentially enabling finer resolution modeling. We built our method to tolerate bursts of extreme AQ events without a loss in skill. This is an important factor, as extreme AQ events present a challenge for the existing NOAA AQ forecasting system.

1.1 Background

NOAA provides operational forecast guidance for AQ, including ozone and PM_{2.5}. NOAA uses the Unified Forecast System Air Quality (UFS-AQ) model for AQ forecasting [19]. The UFS-AQ transport of chemical tracers, which is used to calculate how species advect across the United States, is essential for AQ forecasts. Approximately 40% of the overall computation time is spent in the transport module, where, for each grid point across CONUS, for each of the 64 vertical levels that represent vertical atmospheric conditions, and at each timestep, calculations must be performed. The sheer number of computations contributes to the intractability of reaching a three km resolution. The transport of 183 chemical tracers is used to provide 72-hour forecasts. Each of the chemical species is sequentially passed into the transport module approximately every 30 minutes. After the transport module, each of the species is also processed in two other processing modules, the physics module and the chemistry module (which is applied to all species simultaneously) before the next transport timestep.

1.2 Related Work

Early Machine Learning (ML) methods which were used to speed up calculations by emulating parameterizations of atmospheric physics and chemistry include [4, 7, 15]. Work related to chemistry emulation in an effort to replace this component of a model [4] showed early promise of applying shallow neural networks to this class of problems. This work and the early promise of AI methods have been cautiously explored for weather forecasting. Data-driven AI methods have been shown to capture the physics of weather by using deep levels of abstraction across multiple layers of convolution. Recent research has shown impressive results when using data-driven AI methods for weather forecasting [11, 12, 14, 20] devoid of any explicit knowledge pertaining to the underlying physics. Many of the current state of the art weather forecasting AI models have shown excellent results forecasting a selection of weather variables and modeling a subset of vertical levels closest to the Earth’s surface [5, 10]. While models such as GraphCast [5] apply standard normalization to the residuals of each atmospheric variable, our model is species-agnostic, enabling a more flexible approach. Building on these ideas and our group’s previous work [17, 18] to emulate transport, we present a method that overcomes many issues present when working with atmospheric variables pertaining to chemical species such as highly skewed concentration distributions.

2 Air Quality Data

The NOAA AQ data was generated from the NOAA’s AQM.v7.0 UFS-AQ model. The model was run for one month, and data pairs were captured before and after transport. The data pairs consisted of chemical species concentrations and meteorological variables across CONUS. Data was collected for seven days between September 1, 2020 and October 1, 2020 every five days (approximately five TBs of data), a period of time when active wildfires were present. The resolution of the 183 species at each time point was $232 \times 396 \times 64$ pixels in latitude, longitude, and vertical levels. These data pairs were used to train the ML model.

Data was subset to only include the 87 of the 183 species which contribute to the Air Quality Index (AQI) [1]. These include the constituent species of PM_{2.5}, O₃, CO, NO₂, SO₂, and other contributors to the formation of ozone and PM_{2.5}. For each species, data across CONUS was evenly divided into a 4×6 grid of non-overlapping $58 \times 66 \times 16$ pixel patches, and only the lowest 16 vertical levels,

which most directly affect human health, were included. Since transport is advection-driven, the 3D wind velocity field and the altitude between vertical layers (which is not uniform) were included as features to the model. The surface geopotential, temperature, and pressure were also provided.

To assess the model’s performance during extreme AQ events, patches were designated as either extreme or non-extreme based on the U.S. EPA’s AQI thresholds for ozone and PM2.5. We considered AQI categories of *moderate* or above as extreme [1]. Thus, patches of ozone and ozone contributors were categorized as extreme if the maximum ozone concentration in the patch matched or exceeded 0.055 ppm. Similarly, patches of PM2.5 constituents were categorized as extreme if the maximum total PM2.5 concentration in that patch matched or exceeded 12.1 $\mu\text{g}/\text{m}^3$.

2.1 Advection

Modeling atmospheric variables and chemical species is difficult due to their highly skewed concentration distributions which vary greatly between variables and species. Traditionally, machine learning weather models normalize on a per-species basis, often with standard normalization. However, our model is species-agnostic and is not required to learn species-specific normalization parameters. Fortunately, the advection continuity equation, which the transport module solves for and our model aims to learn, is both species-agnostic and scale-invariant:

$$\frac{\partial c}{\partial t} + \nabla \cdot (\mathbf{v}c) = 0 \iff \frac{\partial(ac)}{\partial t} + \nabla \cdot (\mathbf{v}ac) = 0,$$

where c is the species concentration, \mathbf{v} is the velocity vector field, and a is the scaling factor. Furthermore, the advection equation is invariant under affine transformations if an incompressible flow is assumed due to the relatively low wind velocities:

$$\frac{\partial c}{\partial t} + \nabla \cdot (\mathbf{v}c) = 0 \iff \frac{\partial c}{\partial t} + \mathbf{v} \cdot \nabla c = 0 \iff \frac{\partial(ac + b)}{\partial t} + \mathbf{v} \cdot \nabla(ac + b) = 0,$$

where a and b parameterize an affine transformation. As a result of this invariance, we are able to apply linear and affine transformations such as min-max normalization on a per-species, per-batch, or even per-patch basis while largely preserving the underlying transport. This enables extreme species concentrations that are spatio-temporally localized to specific patches, e.g. high PM2.5 concentrations caused by wildfires, to be normalized independently from other patches. In longer-range forecasting applications, this invariance can be useful for data distribution shifts caused by climate change and applying data normalization in continual learning settings. Additionally, it opens the possibility of applying log transforms to the highly right-skewed species concentrations, i.e. $c' = \ln(ac + b)$. For example, data can be min-max normalized to a range of 1 to e before the log transform to avoid zero or negative values. The machine learning model can then learn a governing equation analogous to

$$\frac{\partial \exp(c')}{\partial t} + \mathbf{v} \cdot \nabla \exp(c') = 0$$

while maintaining invariance to species, batch, or patch-specific normalization parameters.

In our experiment, we demonstrated the potential of this approach by min-max normalizing the input and output data of the UFS-AQ model to a range of zero to one on a per-species basis. Since species concentrations decrease significantly as vertical level increases, we also min-max normalized by vertical layer. Although the advection equation is not invariant to this normalization, our approach aims to emulate rather than replicate the underlying transport.

3 Methodology

3.1 Difference Learning

Since the output of the UFS-AQ transport module is processed separately by the chemistry and physics modules at each timestep, our model must learn per-timestep advective transport. However, because the changes in species concentrations between the input and output of the UFS-AQ are relatively small as seen in Figure 1 (column 3), attempting to directly learn the translation between input and output can be akin to learning an autoencoding of the input data. Learning the small residuals between the output and input can also be challenging due to factors such as vanishing gradients. To overcome these challenges, we applied a cube root transformation to the difference.




Figure 1: Distributions of chemical species *anai*, a PM2.5 constituent, and ozone at vertical levels 1 (blue) and 16 (red). The *Input* and *Output* columns are concentration distributions from the UFS-AQ model after min-max normalization. Column 3 is the distribution of the difference between the *Input* and *Output* data. Column 4 is the distribution after taking the cube root of this difference. Note that the y-axis is log-scaled.

The cube root transformation is traditionally used to reduce the skewness of a distribution [16]. For our data, it serves to increase the spread of the concentration distribution (Figure 1). The cube root transformation has several immediate advantages; it does not require species-specific normalization parameters, can be applied to zero and negative values, and reshapes the target distribution to a wider range of -1 to 1. Since n 'th root transformations are less sensitive to changes at the extremities of the range from -1 to 1, higher root transformations may yield more accurate predictions when the residuals are small, but may negatively impact prediction accuracy and contrast when the residuals are closer to -1 or 1 (Appendix, Figure 4).

3.2 Model and Training

The deep learning model used to emulate the transport is a 3D U-Net with four downsampling and upsampling blocks [13]. The U-Net difference learning approach is illustrated in Figure 3 (Appendix). The mean squared error (MSE) loss function and the Adam optimizer with a learning rate of 0.001, $\beta_1 = 0.9$, and $\beta_2 = 0.999$ were used. During training, the first 4 days of the data were divided into an 80/20 train/validation split, giving 394,214 training patches and 98,554 validation patches. The test dataset consists of the last three days of data, or 369,576 patches. Of these 369,576 patches, 233,950 patches were classified as extreme. The U-Net has a total of 90,310,657 trainable parameters and was trained on a single Tesla V100 for 20 epochs over two days and 19 hours.

4 Results

Figure 2 demonstrates that the U-Net's predictions align well with the ground truth across multiple species, vertical layers, and non-extreme/extreme patches. The RMSE across the entire test dataset, calculated in the min-max normalized space, is 0.0115. The model successfully predicts concentration changes during non-extreme and extreme events with RMSEs of 0.00838 and 0.0129 respectively (Table 1). Inference on a Tesla V100 takes 4.74 ms for a batch of 32 patches. Extrapolating this time to the entirety of CONUS for all 183 species and 64 layers, the ML approach can produce a prediction in only 2.6 seconds per timestep.

Table 1: U-Net RMSEs calculated in the min-max normalized space for non-extreme patches, extreme patches, and all patches in the test dataset.

Non-Extreme	Extreme	All Data
0.00838	0.0129	0.0115

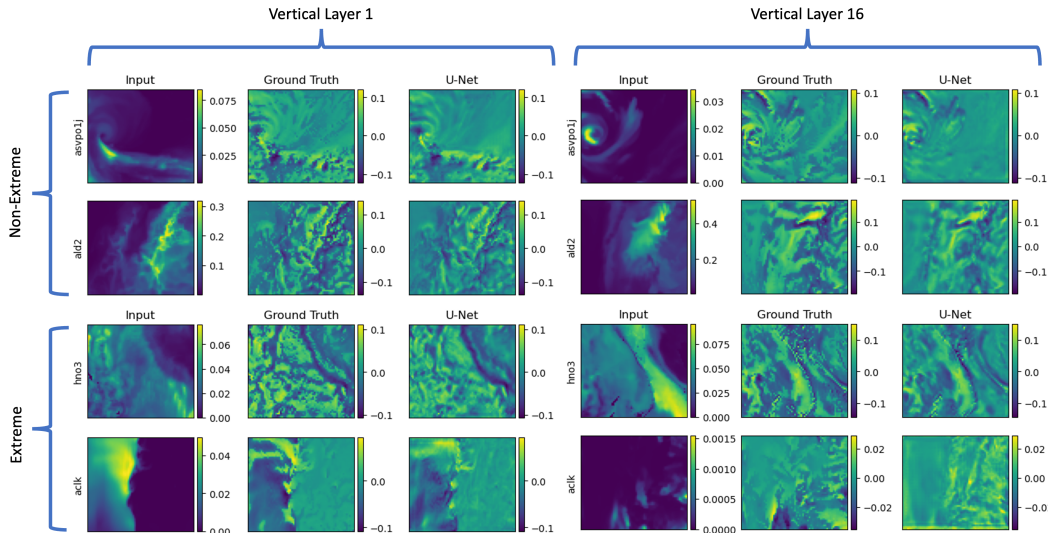


Figure 2: U-Net predictions, where $Ground\ Truth = (Output - Input)^{1/3}$. The first three columns show patches from vertical layer 1, and the last three columns show patches from vertical layer 16. The first two rows show non-extreme patches, and the last two rows show extreme patches. Species *asvp01j* and *acik* are PM2.5 constituents, and *ald2* and *hno3* are contributors to ozone and PM2.5 formation respectively.

In an attempt to quantify the model’s efficacy in learning the underlying physics, the mass of *asvp01j* (a PM2.5 constituent) was computed as a preliminary physics-based evaluation metric. The mean percent difference in *asvp01j* mass, calculated between the U-Net prediction and the UFS-AQ transport module, is 0.0741%. This demonstrates the potential of the U-Net model in preserving the mass advected by the UFS-AQ model.

5 Conclusions

Our model emulates the per-timestep advective transport of atmospheric chemical species. With an overall RMSE of 0.0115, good performance during both extreme and non-extreme AQ events (Table 1), and estimated prediction time of 2.6 seconds on a single GPU, this ML method exhibits significant potential for integration into the NOAA operational AQ environment. To achieve these results, we utilize data transformations that the underlying transport, governed by the advection equation, is largely invariant under.

5.1 Future Work

In future work, we plan to further explore data transformations, applied on a per-patch basis, which preserve the underlying transport. These include log transformations on the input and output data of the UFS-AQ model, as well as n ’th root transformations on the residuals. To eliminate boundary artifacts, we will train on larger, overlapping patches. We will also explore adding physics-informed regularization terms to the MSE loss function and further evaluate mass conservation. Ultimately, we aim to develop an ML model which efficiently emulates advective transport over large time-scales for implementation in the UFS-AQ model.

Acknowledgments and Disclosure of Funding

This work has been funded by NOAA NA21OAR4310383, SUBAWD003728.

References

- [1] Office of Air Quality Planning and U.S. Environmental Protection Agency Standards. *Technical Assistance Document for the Reporting of Daily Air Quality – the Air Quality Index (AQI)*. 2018. URL: <https://www.airnow.gov/sites/default/files/2020-05/aqi-technical-assistance-document-sept2018.pdf>.
- [2] Pei Hou and Shiliang Wu. “Long-term changes in extreme air pollution meteorology and the implications for air quality”. In: *Scientific reports* 6.1 (2016), pp. 1–9.
- [3] Daniel A Jaffe et al. “Wildfire and prescribed burning impacts on air quality in the United States”. In: *Journal of the Air & Waste Management Association* 70.6 (2020), pp. 583–615.
- [4] Makoto M Kelp, Christopher W Tessum, and Julian D Marshall. “Orders-of-magnitude speedup in atmospheric chemistry modeling through neural network-based emulation”. In: *arXiv preprint arXiv:1808.03874* (2018).
- [5] Remi Lam et al. “Learning skillful medium-range global weather forecasting”. In: *Science* 0.0 (), eadi2336. DOI: 10.1126/science.adi2336. eprint: <https://www.science.org/doi/pdf/10.1126/science.adi2336>. URL: <https://www.science.org/doi/abs/10.1126/science.adi2336>.
- [6] Philip J Landrigan. “Air pollution and health”. In: *The Lancet Public Health* 2.1 (2017), e4–e5.
- [7] Qi Liao et al. “Deep learning for air quality forecasts: a review”. In: *Current Pollution Reports* 6 (2020), pp. 399–409.
- [8] Julia Ling, Reese Jones, and Jeremy Templeton. “Global impact of landscape fire emissions on surface level PM_{2.5} concentrations, air quality exposure and population mortality”. In: *Journal of Computational Physics* 318 (2016), pp. 22–35.
- [9] Mojtaba Moghani and Cristina L Archer. “The impact of emissions and climate change on future ozone concentrations in the USA”. In: *Air Quality, Atmosphere & Health* 13 (2020), pp. 1465–1476.
- [10] Jaideep Pathak, Shashank Subramanian, et al. “Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators”. In: *arXiv preprint arXiv:2202.11214* (2022).
- [11] Stephan Rasp and Nils Thuerey. “Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench”. In: *Journal of Advances in Modeling Earth Systems* 13.2 (2021), e2020MS002405.
- [12] Stephan Rasp and Nils Thuerey. “Purely data-driven medium-range weather forecasting achieves comparable skill to physical models at similar resolution”. In: *arXiv preprint arXiv:2008.08626* (2020).
- [13] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. “U-net: Convolutional networks for biomedical image segmentation”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer. 2015, pp. 234–241.
- [14] Sebastian Scher. “Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning”. In: *Geophysical Research Letters* 45.22 (2018), pp. 12–616.
- [15] Lu Shen et al. “A machine-learning-guided adaptive algorithm to reduce the computational cost of integrating kinetics in global atmospheric chemistry models: application to GEOS-Chem versions 12.0.0 and 12.9.1”. In: *Geoscientific Model Development* 15.4 (2022), pp. 1677–1687.
- [16] Maria Sinimaa et al. “Feature Engineering of Weather Data for Short-Term Energy Consumption Forecast”. In: *2021 IEEE Madrid PowerTech*. 2021, pp. 1–6. DOI: 10.1109/PowerTech46648.2021.9494920.
- [17] Jennifer Sleeman et al. “Artificial Intelligence Air Quality Forecast Emulation of Atmospheric Tracers”. In: *103rd AMS Annual Meeting*. AMS. 2023.
- [18] Jennifer Sleeman et al. “The Integration of Artificial Intelligence for Improved Operational Air Quality Forecasting”. In: *AGU Fall Meeting Abstracts*. Vol. 2021. 2021, A15E–1680.
- [19] Ivanka Stajner et al. “Development of Next-Generation Air Quality Predictions for the United States in the Unified Forecast System”. In: *103rd AMS Annual Meeting*. AMS. 2023.

- [20] Jonathan A Weyn, Dale R Durran, and Rich Caruana. “Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere”. In: *Journal of Advances in Modeling Earth Systems* 12.9 (2020), e2020MS002109.

A Appendix

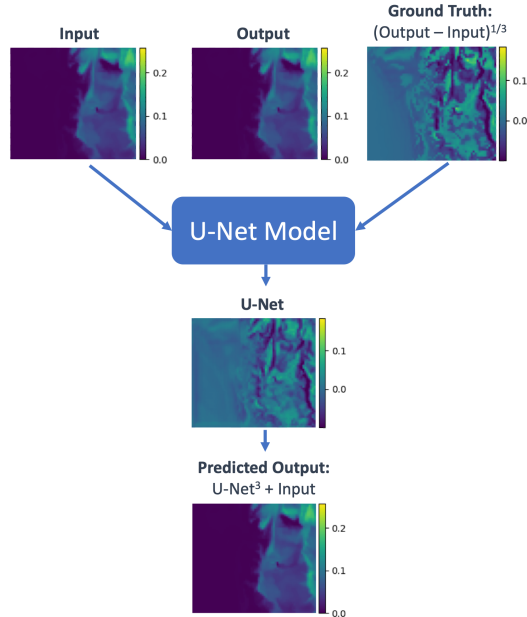


Figure 3: Illustration of the difference learning approach. The U-Net learns the mapping between *Input* and $(\text{Output} - \text{Input})^{1/3}$. The *Predicted Output* is the U-Net’s prediction in the min-max normalized space.

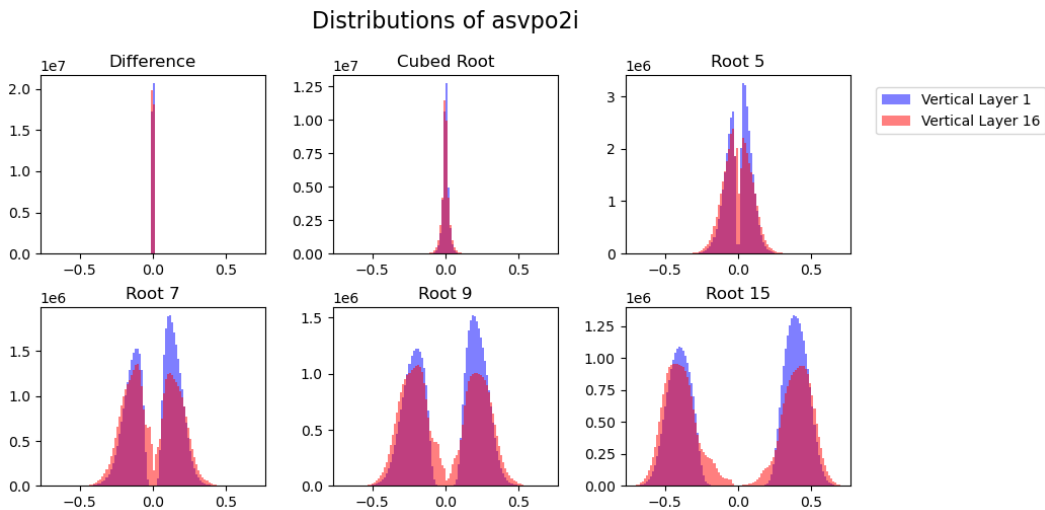


Figure 4: Effects of $(\text{Output} - \text{Input})^{1/n}$ transformations on the distribution of *asvpo2i*, a PM2.5 constituent, at vertical levels 1 (blue) and 16 (red). From left to right, top to bottom, $n = 1, 3, 5, 7, 9, 15$.