EGU24-10922, updated on 12 Jul 2024
https://doi.org/10.5194/egusphere-egu24-10922
EGU General Assembly 2024

Graph Neural Networks for Atmospheric Transport Modeling of CO_2

Vitus Benson[1,2,3], Ana Bastos[1,2], Christian Reimers[1,2], Alexander J. Winkler[1,2], Fanny Yang[3], and Markus Reichstein[1,2]
[1]Max Planck Institute for Biogeochemistry, Biogeochemical Integration, Jena, Germany (vbenson@bgc-jena.mpg.de)
[2]ELLIS Unit Jena
[3]ETH Zürich

Large deep neural network emulators are poised to revolutionize numerical weather prediction (NWP). Recent models like GraphCast or NeuralGCM can now compete and sometimes outperform traditional NWP systems, all at much lower computational cost. Yet to be explored is the applicability of large deep neural network emulators to other dense prediction tasks such as the modeling of 3D atmospheric composition. For instance the inverse modeling of carbon fluxes essential for estimating carbon budgets relies on fast CO_2 transport models.

Here, we present a novel approach to atmospheric transport modeling of CO_2 and other inert trace gases. Existing Eulerian transport modeling approaches rely on numerical solvers applied to the continuity equation, which are expensive: short time steps are required for numerical stability at the poles, and the loading of driving meteorological fields is IO-intensive. We learn high-fidelity transport in latent space by training graph neural networks, analogous to approaches used in weather forecasting, including an approach that conserves the CO_2 mass. For this, we prepare the CarbonBench dataset, a deep learning ready dataset based on Jena Carboscope CO_2 inversion data and NCEP NCAR meteorological reanalysis data together with ObsPack station observations for model evaluation.

Qualitative and quantitative experiments demonstrate the superior performance of our approach over a baseline U-Net for short-term (<40 days) atmospheric transport modeling of carbon dioxide. While the original GraphCast architecture achieves a similar speed to the TM3 transport model used to generate the training data, we show how various architectural changes introduced by us contribute to a reduced IO load (>4x) of our model, thereby speeding up forward runs. This is especially useful when applied multiple times with the same driving wind fields, e.g. in an inverse modeling framework. Thus, we pave the way towards integrating not only atmospheric observations (as is done in current CO_2 inversions), but also ecosystem surface fluxes (not yet done) into carbon cycle inversions. The latter requires backpropagating through a transport operator to optimize a flux model with many more parameters (e.g. a deep neural network) than those currently used in CO_2 inversions – which becomes feasible if the transport operator is fast enough. To the best of our knowledge, this work presents the first emulator of global Eulerian atmospheric transport, thereby providing an initial step towards next-gen inverse modeling of the carbon cycle with deep learning.