

PROTOTYPE-ENHANCED PREDICTION IN GRAPH NEURAL NETWORKS FOR CLIMATE APPLICATIONS

**Nawid Keshtmand, Elena Fillola, Jeffrey Nicholas Clark,
Raul Santos-Rodriguez & Matthew Rigby**
University of Bristol
Bristol, England
y118410@bristol.ac.uk

ABSTRACT

Data-driven emulators are increasingly being used to learn and emulate physics-based simulations, reducing computational expense and run time. Here, we present a structured way to improve the quality of these high-dimensional emulated outputs, through the use of prototypes: an approximation of the emulator’s output passed as an input, which informs the model and leads to better predictions. We demonstrate our approach to emulate atmospheric dispersion, key for greenhouse gas emissions monitoring, by comparing a baseline model to models trained using prototypes as an additional input. The prototype models achieve better performance, even with few prototypes and even if they are chosen at random, but we show that choosing the prototypes through data-driven methods (k-means) can lead to almost 10% increased performance in some metrics.

1 INTRODUCTION

Earth system modeling often relies on complex physics-based simulations, which, are computationally expensive and can limit their applicability for real-time decision-making or large-scale datasets. Machine learning (ML) models are increasingly being used to accelerate, or fully replace, these expensive simulations (Fillola et al., 2023a; Vinuesa & Brunton, 2022; Bhowmik et al., 2022). However, ML-driven emulators are often required to predict a high-dimensional output state from a high-dimensional input state, and exploring techniques to improve their performance is key to ensuring they can serve as reliable surrogates for physics-based simulations.

Here, we introduce the concept of ‘prototypes’: an approximation of the emulator’s output provided as an input, serving as a prior or guide to help the emulator produce higher-quality predictions.

We demonstrate the utility of prototypes in monitoring greenhouse gas (GHG) emissions. Traditionally, countries report their emissions through national greenhouse gas inventories (NGHGs), compiled through a bottom-up approach based on socioeconomic activity data and emissions factors (Olague, 2016). An alternative is the top-down inverse modeling approach, which uses atmospheric observations to estimate GHG emissions. This method employs atmospheric dispersion models to solve for atmospheric transport, with Lagrangian Particle Dispersion Models (LPDMs) being among the most popular for this application Jacob et al. (2022). When used in inverse modelling systems, LPDMs simulate gas ‘particles’ moving backward in time from a GHG measurement location, generating a 2D histogram - or ‘footprint’ - that identifies the contribution of upwind fluxes to the observed GHG concentrations (Hegarty et al., 2013) (see examples in Fig. 1). With one LPDM run required for each measurement and taking ~ 20 CPU-minutes, these methods face challenges with the vast increase in GHG data from new satellite-based instruments, now totalling thousands of measurements every day. The computational cost of generating footprints creates a bottleneck in emission estimation. Here, we aim to improve an existing LPDM-emulator by Fillola et al. (2023b), which uses an architecture based on Graph Neural Networks (GNNs)(Battaglia et al., 2018; Bacciu et al., 2020) and can predict satellite GHG measurement footprints based on meteorological and topographical features $\sim 1000\times$ faster than the physics-based dispersion model.

Our contributions include: 1) the development of a prototype-based approach for GNN model predictions and, 2) experimental validation of this method in the context of GHG emissions footprint

prediction. We show that incorporating prototypes improves the model’s performance across several metrics compared to the same model without prototypes.

2 METHOD

Architecture We use the same Encoder-Processor-Decoder architecture as Fillola et al. (2023b), which is similar to the approach used in meteorological applications (Keisler, 2022; Lam et al., 2022; Pfaff et al., 2020). The inputs, including meteorological and topographical data, are arranged in a latitude-longitude grid that represents the native data space. In the Encoder this data gets encoded into a triangular, lower-resolution mesh, and in the Processor each mesh node is updated iteratively with a GNN block (Battaglia et al., 2018), using information from neighbouring nodes. The Decoder then maps the mesh to the original data space, emulating the footprint value at each grid node. For a visual representation of the architecture and more details, see Fillola et al. (2023b).

Data We use footprints produced by NAME, the UK’s Met Office LPDM (Jones et al., 2007) for GOSAT measurements (Greenhouse gases Observing SATellite (Parker et al., 2020)) over Brazil for 2014-16, used originally by Tunnicliffe et al. (2020). The footprints have a resolution of $0.352^\circ \times 0.234^\circ$ ($\sim 25 \times 35$ km in mid-latitudes) and have been cropped to a size of 50×50 gridcells ($\sim 1000 \times 1500$ km), centered around the coordinates of the satellite measurement. They aggregate gas dispersion over 30 days prior to each observation. The inputs are meteorological features at a range of heights, plus information about the topography and coordinates (see Fillola et al. (2023b)). We train on data for 2014-2015 (~ 11000 data points after sampling by using every third sample in the dataset as samples close in time would be similar to one another), and test on data for 2016 (~ 7100 data points after sampling).

Here, we compare the performance of the original model (Fillola et al., 2023b), with the same model with an additional input feature: a prototype.

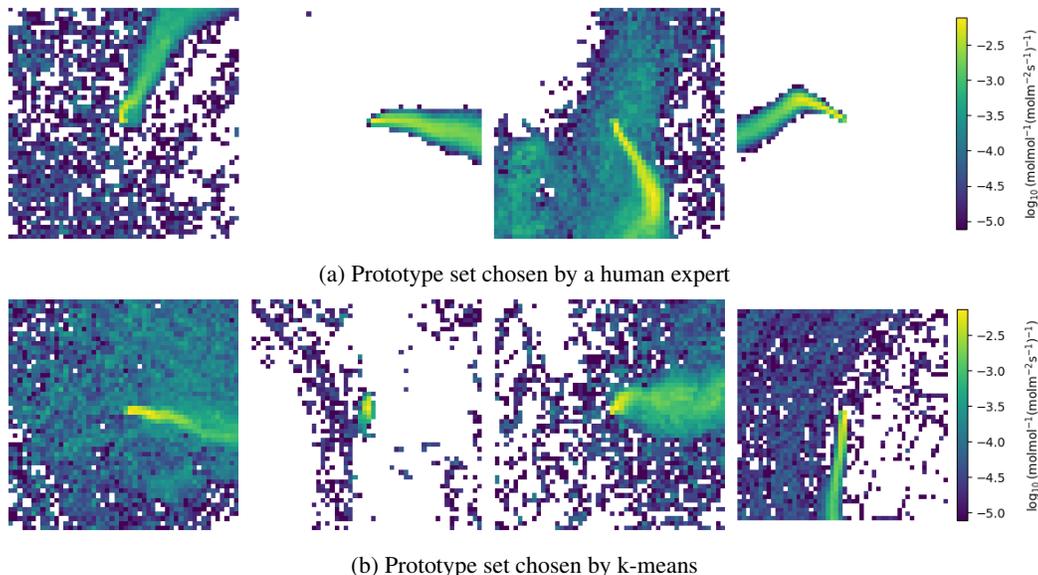


Figure 1: Examples of two prototype sets, with size $n = 4$. Prototypes are chosen from the true footprints in the training dataset. During training and testing, the footprint from the set closest to the true label is passed as an input to improve prediction. A human expert (a) selects footprints that have distinct wind directions, whereas k-means (b) chooses footprints based on the data distribution.

Prototype Design The prototype, i.e. an initial inexpensive approximation of the output, can be produced in several ways: low-cost calculations based on the inputs, using artificial samples, or using selected examples from the training data. He et al. (2023) designed a CNN-based emulator for an LPDM, but they operate at a smaller domain and higher resolution (400×400 km domain, km-scale resolution). At this scale, Gaussian plumes can be used as prototypes: a simple mathematical

approximation of dispersion generated using mainly wind direction and speed (He et al., 2023). They demonstrate how the addition of this input improves the performance of their model. However, Gaussian plumes do not generalize to large domains (e.g. continental scales) and therefore we choose an alternative approach: we select n footprints from the training data to be prototypes. Each training sample then gets assigned one of these prototypes as an additional input, to aid prediction. There are two main steps:

1. Choosing the set of prototypes: the chosen samples should be representative of the range of potential outputs. Here, we demonstrate two methods to curate the prototype set: a human expert, and a data-driven approach. An atmospheric dispersion expert chooses manually n footprints, aiming to cover a wide range of different conditions, such as where the upwind areas of the footprint are one of the four main cardinal directions (Fig. 1a). For comparison, we also train a model where $n = 4$ prototype footprints are chosen at random. For the data-driven approach, we implement k-means clustering with n clusters, and use the footprint closest to the center of each group as prototype (Fig. 1b). If there is a bias in the data distribution (for example, in our dataset for South America many footprints have a East direction due to prevailing winds), the data-driven prototypes will reflect it (e.g. see Fig. 1b).
2. Assigning each sample the best-fitting prototype from the prototype set: to demonstrate the impact of prototypes alone, we show an ‘oracle case’ where we have access to the true footprints to assign prototypes. Each footprint in the dataset gets assigned the prototype with the lowest L_2 distance in a lower dimensional (64) PCA space. In deployment, the true footprint to be predicted is unknown, so the prototype for a particular data point would need to be chosen using the inputs (e.g. using a classifier).

Training details All models are trained for 100 epochs with a batch size of 5, using the Adam optimizer with a learning rate of $5e^{-5}$. We train and run the model on a 32GB NVIDIA V100 GPU which takes approximately 4 hours. Each model is trained for three different random seeds.

Performance metrics Test set footprints were evaluated quantitatively and qualitatively. Quantitatively, we calculated Mean Squared Error (MSE) on a pixel-by-pixel basis, and intersection over union (IoU), to measure spatial overlap between the true and the predicted footprints:
$$\text{IoU} = \frac{\text{Area of Overlap between predicted and true images}}{\text{Area of Union of the predicted and true images}}$$
 (Rahman & Wang, 2016).

3 RESULTS AND DISCUSSION

Quantitative analysis We compare the performance of a baseline model (0 prototypes, same as Fillola et al. (2023b)) with our prototype approach for a range of n prototypes, using a) prototype sets chosen by an expert and b) prototype sets chosen through k-means (Fig. 2). Using at least two prototypes, through any method, improves the performance on the baseline.

We find that the data-driven approach, utilising prototypes selected with k-means, achieves better improvements on the baseline than the expert-driven method. This improvement becomes more significant with a higher number of prototypes, with the k-means model with $n = 20$ prototypes producing a mean IoU score 8% higher than the baseline, and 4% higher than $n = 2$ expert prototypes. We chose to compare the $n = 2$ expert prototypes and $n = 20$ k-means prototypes as increasing the number of expert prototypes further or increasing the number of k-means prototypes further does not lead to a significant enhancement in the performance as shown by the small difference between $n = 20$ and $n = 30$, indicating that $n \simeq 20$ prototypes are able to sufficiently represent the data space. Therefore $n = 2$ expert prototypes and $n = 20$ k-means prototypes would be a good option for a practitioner aiming to maximise performance whilst minimizing computational cost.

The random experiment, where $n = 4$ footprints are chosen arbitrarily to be the prototype set, achieves comparable performance to the expert-chosen set, showing that manually curating the set might not provide the expected advantage.

Qualitative analysis The predictions made by the baseline model without prototypes (Fig 3, second column) capture the main upwind direction of the true footprint (Fig 3, first column). However, the predictions tend to be smoothed compared to the “true” footprints, and do not always represent well the higher values. Predictions by models that use prototypes produce sharper footprints that better capture the wind direction and high values, particularly in complex meteorological scenarios.

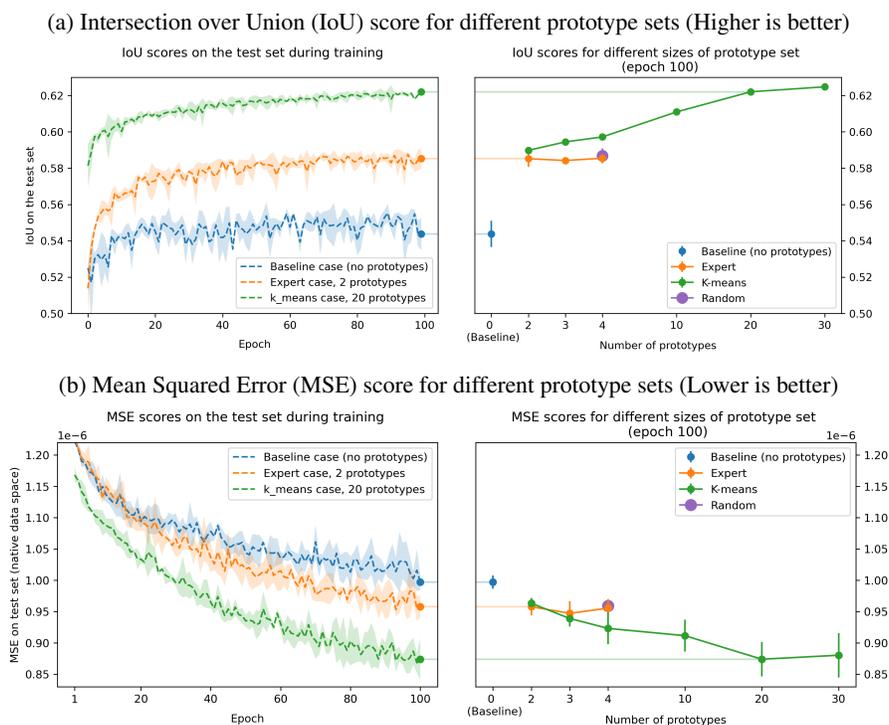


Figure 2: Comparison of different models, showing metrics (a) Intersection over Union and (b) Mean Squared Error. The left panels show training curves for three selected models averaged over three seeds (shaded area shows standard distribution). The right panels show the mean score at epoch 100 for different prototype sets and sizes. Error bars show standard deviation across three seeds.

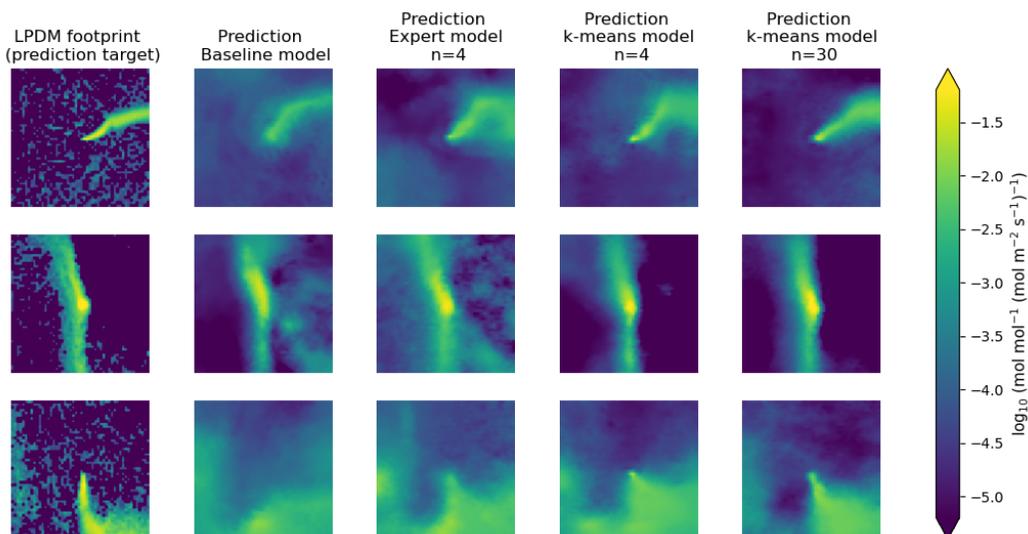


Figure 3: For three random samples from the test set (each row), comparison of the true footprint (first column) with the baseline output not utilising prototypes (second column) and three prototype models. Models are shown ordered by increasing performance, from left to right.

4 CONCLUSION AND FUTURE WORK

Here, we propose using prototypes as an additional input to emulators of physical processes. We demonstrate a structured, easy-to-implement approach to improve GNN-based footprint predictions: footprints predicted using prototypes achieve better IoU and MSE scores, capture the upwind direction and shape better, and are less prone to over-smoothing (Fig 3). This work enhances efficient GHG atmospheric transport simulators with machine learning, making them practical for real-world use and enabling analysis of GHG sources and sinks with vastly more data. We demonstrate an oracle set-up, where we assign the optimal prototypes. In practical applications, prototype assignments would need to utilize input features, rather than the target footprint itself. This could be achieved by designing a classifier that chooses a prototype for each sample based on meteorological features.

Future work will focus on developing a deployment pipeline, developing prototype classifiers and exploring how they impact the performance. It is also key to explore other applications of prototypes, collaborating with domain experts and those developing emulators, since it is likely that other physical emulators could benefit from incorporating prototypes into their emulation pipeline. Increasing the performance and reliability of these emulators will in turn lead to wider and faster adoption of systems that are key for climate modeling, such as near-real time greenhouse gas emissions monitoring.

ACKNOWLEDGMENTS

NK, JNC, and MR are funded by the Self-Learning Digital Twins for Sustainable Land Management grant [grant number EP/Y00597X/1]. EFM is funded by a Google PhD Fellowship. RSR is funded by the UKRI Turing AI Fellowship [grant number EP/V024817/1].

REFERENCES

- Davide Bacciu, Federico Errica, Alessio Micheli, and Marco Podda. A gentle introduction to deep learning for graphs. *Neural Networks*, 129:203–221, 2020.
- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Moumita Bhowmik, Manmeet Singh, Suryachandra Rao, and Souvik Paul. Deepclouds. ai: Deep learning enabled computationally cheap direct numerical simulations. *arXiv preprint arXiv:2208.08956*, 2022.
- Elena Fillola, Raul Santos-Rodriguez, Alistair Manning, Simon O’Doherty, and Matt Rigby. A machine learning emulator for lagrangian particle dispersion model footprints: a case study using name. *Geoscientific Model Development*, 16(7):1997–2009, 2023a.
- Elena Fillola, Raul Santos Rodriguez, and Matt Rigby. Accelerating ghg emissions inference: A lagrangian particle dispersion model emulator using graph neural networks. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*, 2023b. URL <https://www.climatechange.ai/papers/neurips2023/74>.
- Tai-Long He, Nikhil Dadheech, Tammy M Thompson, and Alexander J Turner. Footnet: Development of a machine learning emulator of atmospheric transport. 2023.
- Jennifer Hegarty, Roland R Draxler, Ariel F Stein, Jerome Brioude, Marikate Mountain, Janusz Eluszkiewicz, Thomas Nehrkorn, Fong Ngan, and Arlyn Andrews. Evaluation of lagrangian particle dispersion models with measurements from controlled tracer releases. *Journal of Applied Meteorology and Climatology*, 52(12):2623–2637, 2013.
- Daniel J Jacob, Daniel J Varon, Daniel H Cusworth, Philip E Dennison, Christian Frankenberg, Ritesh Gautam, Luis Guanter, John Kelley, Jason McKeever, Lesley E Ott, et al. Quantifying methane emissions from the global scale down to point sources using satellite observations of atmospheric methane. *Atmospheric Chemistry and Physics*, 22(14):9617–9646, 2022.

- Andrew Jones, David Thomson, Matthew Hort, and Ben Devenish. The uk met office's next-generation atmospheric dispersion model, name iii. In *Air pollution modeling and its application XVII*, pp. 580–589. Springer, 2007.
- Ryan Keisler. Forecasting global weather with graph neural networks. *arXiv preprint arXiv:2202.07575*, 2022.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Graphcast: Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*, 2022.
- Eduardo P Olaguer. *Atmospheric impacts of the oil and gas industry*. Academic Press, 2016.
- Robert J Parker, Alex Webb, Hartmut Boesch, Peter Somkuti, Rocio Barrio Guillo, Antonio Di Noia, Nikoleta Kalaitzi, Jasdeep S Anand, Peter Bergamaschi, Frederic Chevallier, et al. A decade of gosat proxy satellite ch 4 observations. *Earth System Science Data*, 12(4):3383–3412, 2020.
- Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *arXiv preprint arXiv:2010.03409*, 2020.
- Md Atiqur Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation. In *International symposium on visual computing*, pp. 234–244. Springer, 2016.
- Rachel L Tunnicliffe, Anita L Ganesan, Robert J Parker, Hartmut Boesch, Nicola Gedney, Benjamin Poulter, Zhen Zhang, Jošt V Lavrič, David Walter, Matthew Rigby, et al. Quantifying sources of brazil's ch 4 emissions between 2010 and 2018 from satellite data. *Atmospheric Chemistry and Physics*, 20(21):13041–13067, 2020.
- Ricardo Vinuesa and Steven L Brunton. Enhancing computational fluid dynamics with machine learning. *Nature Computational Science*, 2(6):358–366, 2022.