Building the Foundation for Machine Learning-Based Mars Weather Forecasting

Mohammad Altiwainy[1], Cheng Li[2]
[1]Department of Computer Science, Wayne State University – Detroit, MI
[2]Department of Climate and Space Sciences and Engineering, University of Michigan – Ann Arbor, MI

ENGINEERING CLASP

Program in Climate and Space Science Observation

Motivation

- Mars is a leading candidate for future human exploration and possible habitation. However, its atmosphere remains difficult to predict, with hazards such as dust storms threatening mission safety.
- Earth weather forecasting has rapidly advanced with AI and machine learning models.
- Mars lacks a robust Machine Learning-based system for forecasting its weather.
- This project explores whether we can adapt a Earth-trained ML weather model to work on Martian data.

Model Selection

- Conducted Literature Review consisting of 12+ publications comparing different Machine Learning models for predicting climate on Earth.
- We proceeded with Microsoft's "Aurora" model for our experiment based off three factors:
    1. Data access/Feasibility
    2. Scalability
    3. Fine-Tuning Capacity

More specifically, we used Aurora's 0.25° Pretrained model because of its ability to be fine-tuned at any resolution

Problem Statement

The observed state of the atmosphere and surface at a discrete time t as a multidimensional array:

$$X^t = (S^t, A^t)$$

is split into surface (S^t) and atmospheric (A^t) components:

$$S^t \in \mathbb{R}^{V_S \times H \times W} \quad A^t \in \mathbb{R}^{V_A C \times H \times W}$$

- V_S and V_A are the number of surface-level and atmospheric variables, respectively.
- C represents the number of pressure levels.
- H and W are the number latitude and longitude coordinates, respectively.

To predict a future state at time $t' > t$, a simulator $\Phi(X^{t-1}, X^t) = \widehat{X}^{t+1}$ maps the observed states at the previous time X^{t-1} and current time X^t to a predicted state $\widehat{X}^{t+1}$ at the next time step.

For predictions at later time steps, an autoregressive roll-out is obtained by applying the function Φ iteratively, a total of k times:

$$\Phi(\widehat{X}^{t+k-2}, \widehat{X}^{t+k-1}) = \widehat{X}^{t+k}$$

Methodology

- All computations were performed on the Great Lakes High-Performance Computing (HPC) Cluster at the University of Michigan.
- We tested the full model pipeline by following the "Predictions for ERA5" example on Aurora's documentation to ensure we obtain the same results.
- We modified to run on cluster infrastructure using PyTorch, Xarray, and NetCDF4 on the Jupyter Notebook interface for easier visualization and debugging

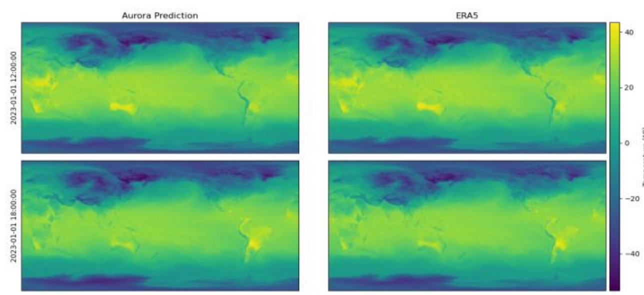

- *The heat maps on the left represent Aurora's predictions for the ERA5 at two different timesteps.*
- *The heat maps on the right represent the actual results from the ERA5 dataset at two different time steps.*
- *Comparing Aurora's predictions to the actual data, we can see they are nearly identical, concluding that Aurora gives accurate predictions for the ERA5*

- Transitioning to Mars, we then used the Ensemble Mars Atmosphere Reanalysis System (EMARS) dataset which uses the Local Ensemble Transform Kalman Filter (LETKF) for data assimilation with the GFDL/NASA Mars Global Climate Model (MGCM)
- Observations that are assimilated include the Thermal Emission Spectrometer (TES) and Mars Climate Sounder (MCS) temperature retrievals.
- Making predictions with the model involves three steps:
    1. prepare a batch of data,
    2. construct the model and load a checkpoint, and
    3. run the model on the batch.

In order to use EMARS data for our experiment, we must manipulate our data to fit the expected batch layout and formatting requirements for Aurora's pretrained model before fine-tuning:

1. To fit Aurora's dimensions of (721, 1440), we interpolated each variable in our EMARS netCDF file to have latitude and longitude dimensions of 721 and 1440, respectively
2. Aurora expects the latitude coordinates to be strictly decreasing(90 … -90) but our EMARS data is increasing(-90 …90) so we must sort them in descending order
3. We inspect our variables from EMARS and their shapes with Xarray and "ncdump" and match them to Aurora's variables, ensuring their dimensions and units match (for example, Aurora expects "t" temperature in Kelvin for one of its atmospheric variables, and we use EMARS' atmospheric temperature also named "t" with the same unit Kelvin

Results

- *The heat maps on the left represent Aurora's predictions for the EMARS at two different timesteps.*
- *The heat maps on the right represent the actual results from the EMARS dataset at two different time steps.*
- *Note: Due to missing variables to support Mars' surface and land structure, we will not be able to accurately plot the EMARS dataset without fine-tuning.*

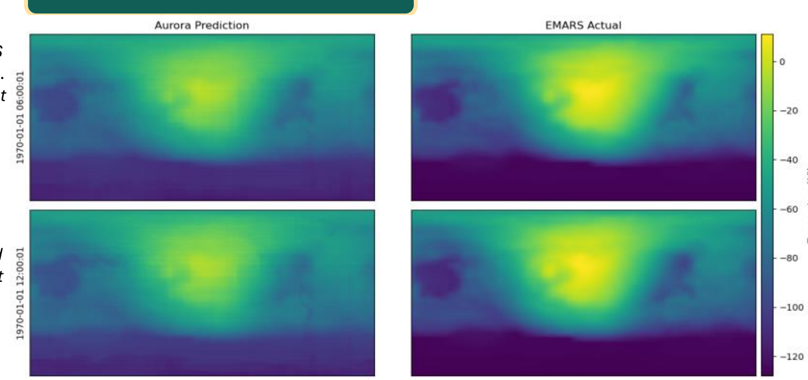

Conclusion

While we cannot make accurate predictions for Mars using rhe EMARS dataset due to missing variables and incorrect parameters, our results demonstrates successful model adaptation from Earth to Mars and serves as a critical step toward future ML-Based Mars Weather Forecasting.

Next Steps

- Our inference-ready pipeline may be used in the future for fine-tuning, by computing gradients and extending Aurora with new variable.
- Instructions on how to fine-tune is available on Aurora's documentation
- Our project will be open-source on GitHub:

References

- Bodnar, C., Bruinsma, W.P., Lucic, A. et al. A foundation model for the Earth system. Nature 641, 1180–1187 (2025). https://doi.org/10.1038/s41586-025-09005-y
- Greybush, S.J., Kalnay, E., Wilson, R.J., Hoffman, R.N., Nehrkorn, T., Leidner, M., et al. The Ensemble Mars Atmosphere Reanalysis System (EMARS) Version 1.0. Geoscience Data Journal 6, 137–150 (2019). https://doi.org/10.1002/gdj3.77
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., et al. The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society, 146(730), 1999–2049 (2020). https://doi.org/10.1002/qj.3803

Acknowledgements

This project was funded through a grant from the National Science Foundation's Research Experience for Undergraduates Program (Award Number 2348598)