Probabilistic Emulation of a Global Climate Model with Spherical DYffusion

Salva Rühling Cachay[1], Brian Henn[2], Oliver Watt-Meyer[2], Christopher S. Bretherton[2], Rose Yu[1]

[1]UC San Diego and [2]Allen Institute for AI (Ai2)

{sruhlingcachay, roseyu}@ucsd.edu

Paper

Code

Motivation and Key Idea

- Climate models are crucial for understanding long-term Earth system evolution.
  - ⋆ But, they are *computationally very expensive*.
  - ⇒ Limited exploration of mitigation and adaptation strategies.
  - ⇒ Large ensembles for reducing uncertainties are prohibitive to run.
- **Solution:** Train a *cheap-to-run* ML model to emulate (parts of) a climate model!
  - ⋆ Unfortunately, existing emulators are deterministic and using conventional diffusion models prohibitive.

Key contribution: We develop an *efficient* conditional generative model, integrating the DYffusion framework with a modified SFNO architecture, for *probabilistic emulation* of a realistic climate model.

Our Results

- Our method, Spherical DYffusion, lowers climate biases by 2× to 4× over relevant baselines.
- The climate biases are within $20 - 50\%$ of the lower-bound given by the reference model.
- Our method is effective for ensemble climate simulations, further reduces climate biases through ensemble-averaging, and produces consistent variability.
- Crucially, it's 25× faster to run than the physics-based climate model, and only less than 3× slower than deterministic emulators.
- Code: `https://github.com/Rose-STL-Lab/spherical-dyffusion`

Weather vs. Climate

Common: Given initial conditions, $\boldsymbol{x}_0$, of the Earth system, predict a sequence of H snapshots, $\boldsymbol{x}_{1:H}$.

Weather forecasting

- *Rapid progress:* ML models competitive or better than operational physics-based ones!
- *Short-term* time-specific focus (e.g. $H = 10$ days)
- Current ML models are often long-term *unstable*.
- Initial-value problem.

Climate Modeling

- Fairly little work, especially for fully data-driven temporal modeling.
- *Stable* & accurate reproduction of *long-term statistics* necessary (e.g. $H = 10 - 100$ years).
- Boundary-condition problem: Model is also conditioned on forcings $\boldsymbol{f}_{0:H}$.

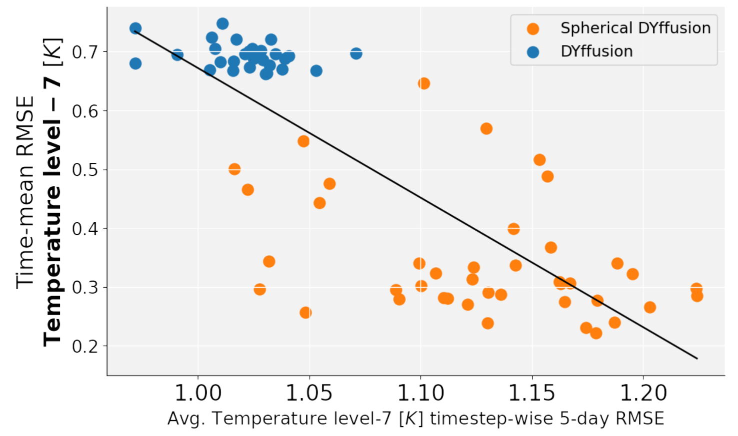

Figure 1. Weather performance (x-axis) is not a strong indicator of climate performance (y-axis).

Inference diagram of Spherical DYffusion

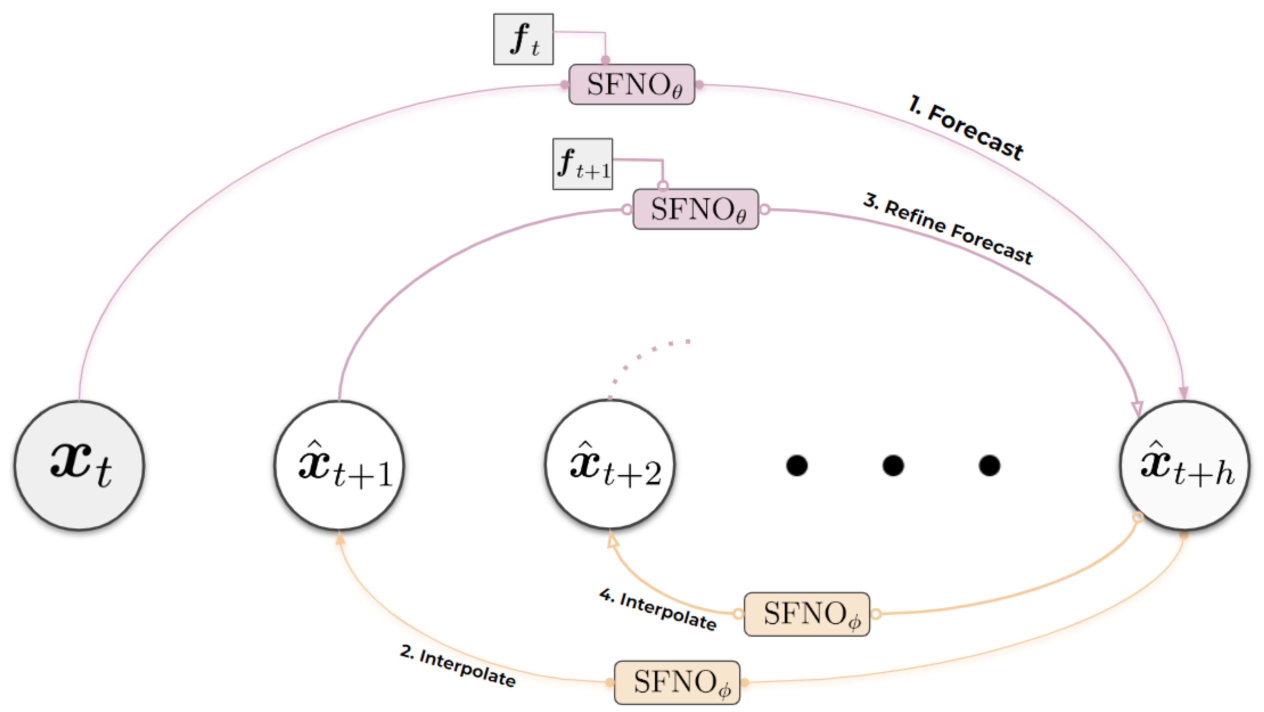

Figure 2. Given an initial condition $\boldsymbol{x}_t$ and forcings $\boldsymbol{f}_{t:t+h}$, our method uses the DYffusion framework, integrated with two SFNO backbone networks, to generate predictions for the next h time steps based on an alternation of refining the forecast of $\boldsymbol{x}_{t+h}$ and temporal interpolations. To forecast more time steps beyond $t + h$, our method is applied autoregressively.

Method

At its core, we integrate the Spherical Fourier Neural Operator (SFNO) [1] into the dynamics-informed diffusion model (DYffusion) framework [2]. This enables:

- Improved modeling of physical fields on a sphere versus U-Net from vanilla DYffusion.
- Ensemble climate simulations naturally.
- Guarantees high efficiency versus other diffusion models and only 3x more than ACE.

SFNO modifications. To achieve this, we need to modify the SFNO architecture to support *time-conditioning* and *inference stochasticity*. Our proposed, modified SFNO block is illustrated in Fig. 3.

Two-stage training. DYffusion requires two neural networks for temporal interpolation and direct multi-step forecasts. We propose to replace them with modified versions of the SFNO architecture, which we denote by SFNO_ϕ and SFNO_θ, respectively. We train them in two-stages such that:

$$\mathrm{SFNO}_\phi\left(\boldsymbol{x}_t, \boldsymbol{x}_{t+h}, \boldsymbol{f}_t, i \,|\xi\right) \approx \boldsymbol{x}_{t+i} \qquad i \in \{1, \dots, h-1\}$$
$$\mathrm{SFNO}_\theta(\mathrm{SFNO}_\phi\left(\boldsymbol{x}_t, \boldsymbol{x}_{t+h}, \boldsymbol{f}_t, j \,|\xi\right), \boldsymbol{f}_{t+j}, j) \approx \boldsymbol{x}_{t+h} \; j \in \{0, 1, \dots, h-1\},$$

where ξ refers to the random variable representing the interpolator network's inference stochasticity.

Inference. The inference process is illustrated in Fig. 2.

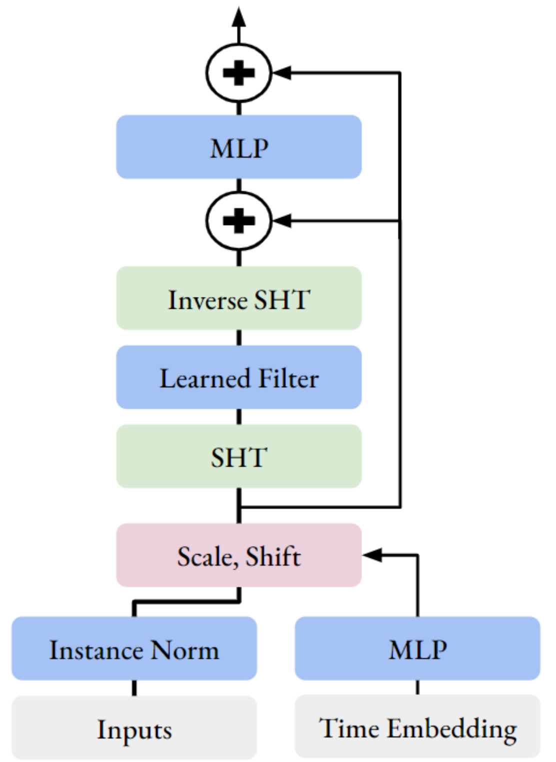

Figure 3. Diagram of one of the blocks of the modified SFNO architecture. Our newly introduced time-conditioning modules correspond to the Time Embedding, MLP on the right, and the scale-shift operation. Our method's interpolator stochasticity relies on dropout (part of the top two-layer MLP), and drop-path, which stochastically replaces all computations between the SHTs with the identity.

Experimental Setup

Dataset: Following ACE [3], we train on coarse-res simulated atmospheric data from FV3GFS, the US primary global forecast model:

- Variables include temperature, humidity, winds over 8 levels (34 variables in total).
- Forced by solar radiation and specified sea surface temperatures.
- Evaluate 10-year-long simulations at a 6-hourly interval (rollout of 14600 time steps!).

Metrics: The most crucial quality of a climate emulator is its ability to reproduce the *climatology*, i.e. the "time-mean" of weather states $\frac{1}{H}\sum_{t=1}^{H}\boldsymbol{x}_t$, of the reference model it aims to emulate. The similarity between the emulator means and the reference time-mean is evaluated using an area-weighted RMSE.

Baselines and Computational Complexity:

- ACE [3]: Deterministic SFNO-based. We also try a stochastic version, ACE-STO, via MC-Dropout.
- DYffusion [2]: Stochastic, U-Net-based.
- FV3GFS reference: We compare time-means from different FV3GFS trajectories with the validation time-mean to compute a "noise-floor" for the achievable climate biases due to natural variability.

Method	NFE	Runtime
ACE / SFNO	h	01:08
Standard diffusion	nh	N/A
Ours	$3(h-1)$	02:56
Physics-based FV3GFS	N/A	78:04
FV3GFS (2× coarser)	N/A	45:38

Table 1. Computational complexity of the different models in terms of: 1) the number of neural function evaluations (NFEs) needed to predict h time steps. and 2) Total inference runtime to simulate 10 years (in hours:minutes). n refers to the number of diffusion steps which usually ranges between 20 to 1000, making standard diffusion prohibitive to use for climate emulation.

Experiments

Quantitative climate biases.

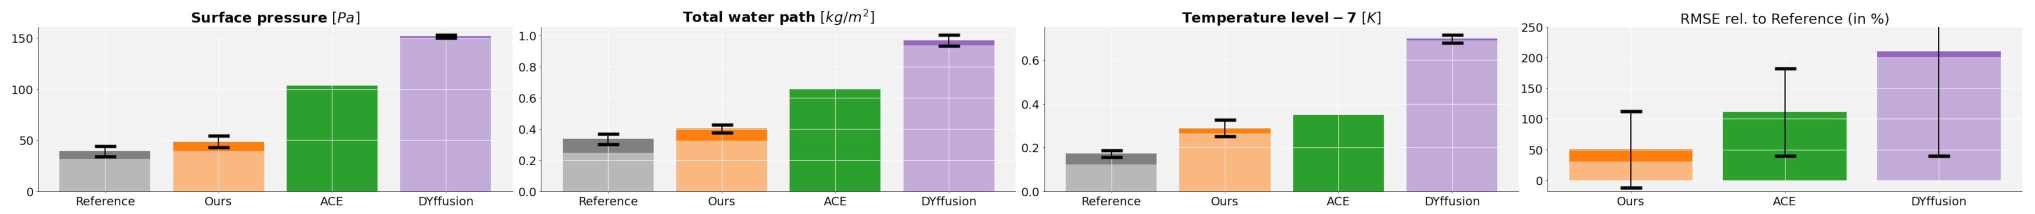

Figure 4. RMSE of 10-year time-means for 3 key variables (left panels) of our method are within $20 - 50\%$ of the reference and $2 - 4\times$ better than the baselines. The rightmost panel shows the average time-mean RMSE across all 34 variables, indicating 50% higher climate biases than the reference, which can be reduced to 30% by ensembling.

Experiments

Quantitative climate biases (global map).

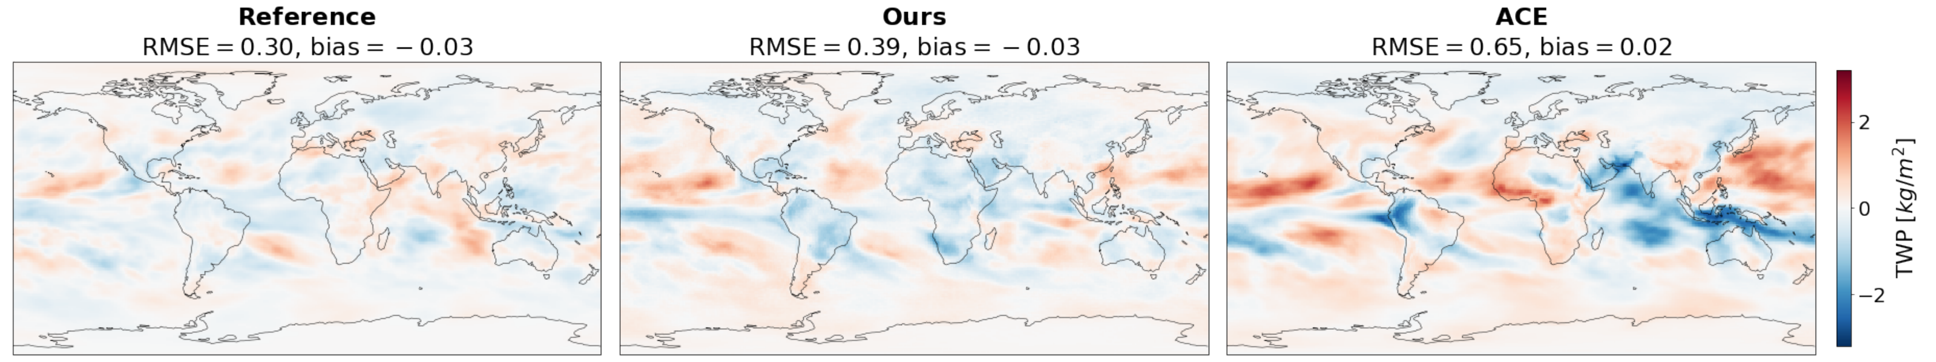

Figure 5. Global map of time-mean biases for total water path variable.

Qualitative video results. The emulated fields of Spherical DYffusion demonstrate high realism, closely mimicking the patterns and variability observed in actual climate model outputs. This showcases Spherical DYffusion's capability to generate plausible and physically consistent climate scenarios over long periods. Scan the QR code on the right to view the full 10-year-long video.

10-year rollout video:

Climate variability.

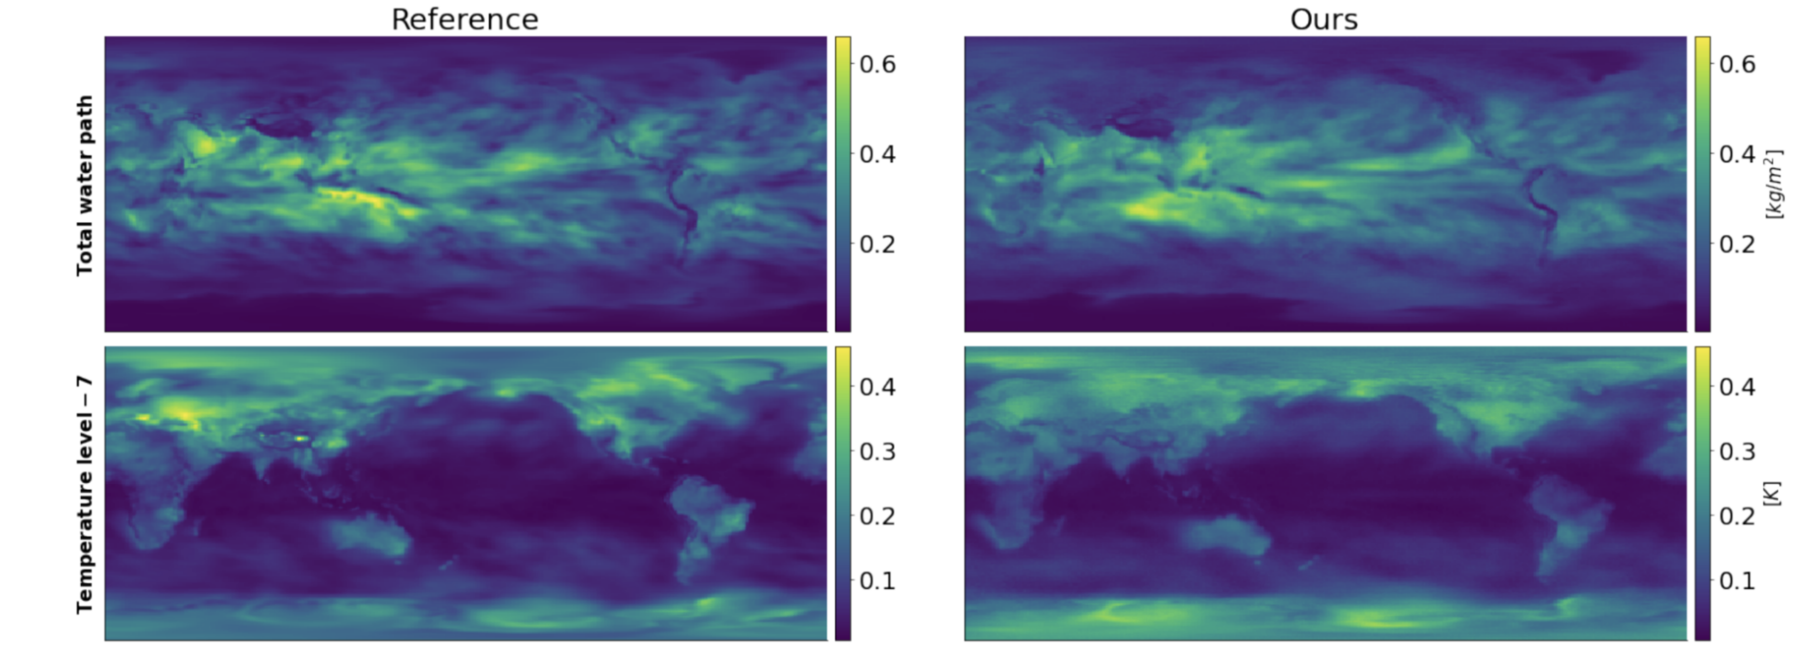

Figure 6. Our method shows a consistent ensemble variability in terms of the simulated climate that also largely reflects the spatial patterns and magnitudes of the reference ensemble.

100-year rollout.

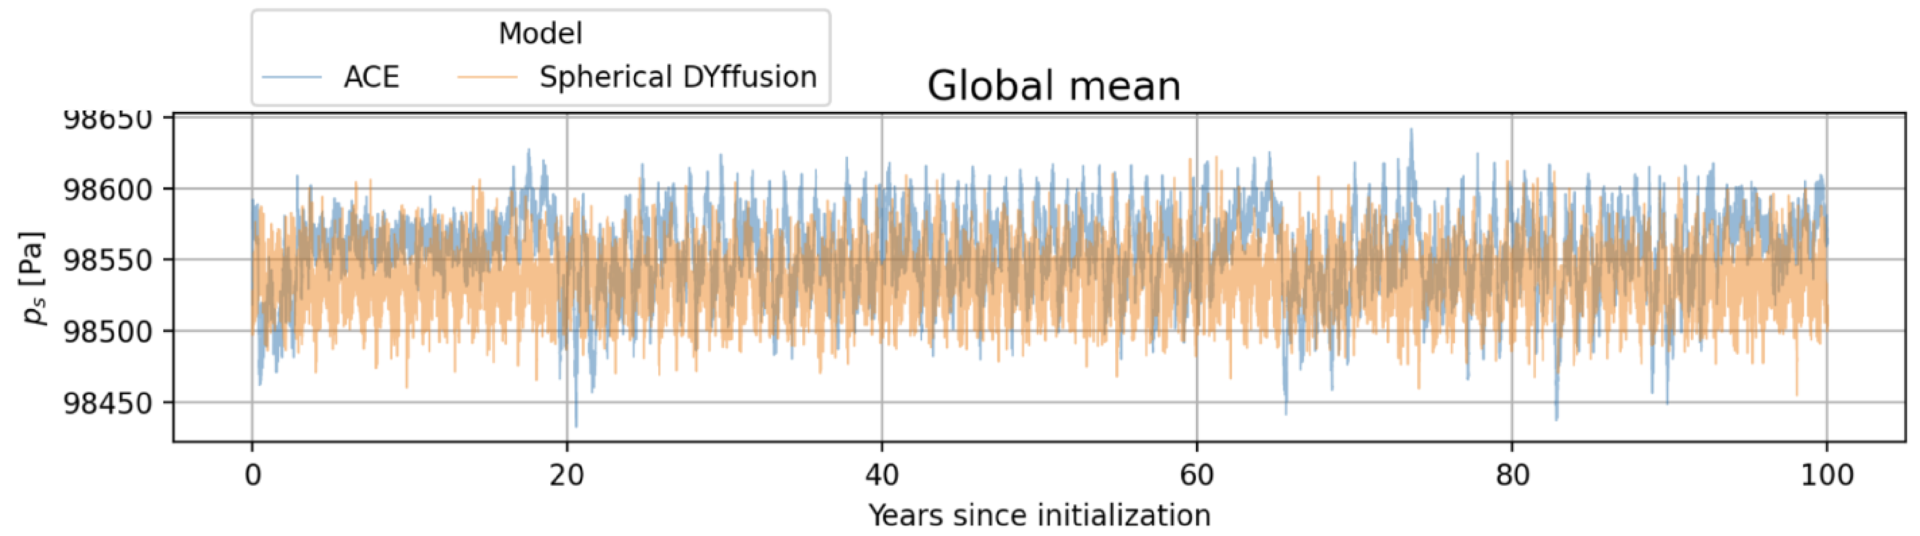

Figure 7. Spherical DYffusion remains stable for 100-year simulations–completing it in around 1 day of wall-clock time–and exhibits improved variability patterns compared to the baseline ACE model. Our model exhibits improved variability patterns compared to the baseline ACE, which suffers from unrealistic annual fluctuations.

References

[1] B. Bonev, T. Kurth, C. Hundt, J. Pathak, M. Baust, K. Kashinath, and A. Anandkumar. Spherical fourier neural operators: Learning stable dynamics on the sphere. *International Conference on Machine Learning*, 2023.

[2] S. Rühling Cachay, B. Zhao, H. Joren, and R. Yu. DYffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. *Advances in Neural Information Processing Systems*, 2023.

[3] O. Watt-Meyer, G. Dresdner, J. McGibbon, S. K. Clark, J. Duncan, B. Henn, M. Peters, N. D. Brenowitz, K. Kashinath, M. Pritchard, B. Bonev, and C. Bretherton. ACE: A fast, skillful learned global atmospheric model for climate prediction. *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*, 2023.

Acknowledgments This work was supported in part by the U.S. Army Research Office under Army-ECASE award W911NF-23-1-0231, the U.S. Department Of Energy, Office of Science, IARPA HAYSTAC Program, CDC-RFA-FT-23-0069, DARPA AIE FoundSci, DARPA YFA, NSF Grants #2205093, #2100237,#2146343, and #2134274.