DiffObs: Generative Diffusion for Global Forecasting of Satellite Observations

Jason Stock[1,2], Jaideep Pathak[1], Yair Cohen[1], Mike Pritchard[1],
Piyush Garg[1], Dale Durran[1], Morteza Mardani[1] & Noah Brenowitz[1]

[1] *NVIDIA Corporation and* [2] *Colorado State University, 2024*

Introduction

Presenting a computationally ambitious, autoregressive generative diffusion model (DiffObs) to predict the high-resolution global evolution of daily precipitation from a satellite observational product.

Question: can we generate realistic convectively coupled tropical disturbances across daily to multi-week simulations?

Motivation: tropical atmospheric variability regulates subseasonal predictability, but **(a)** is challenging to capture realistically in physics-based models and is incompletely understood, and **(b)** most machine learning approaches use complete state information at smaller spatial/temporal scales; GPU advances enable more ambition today.

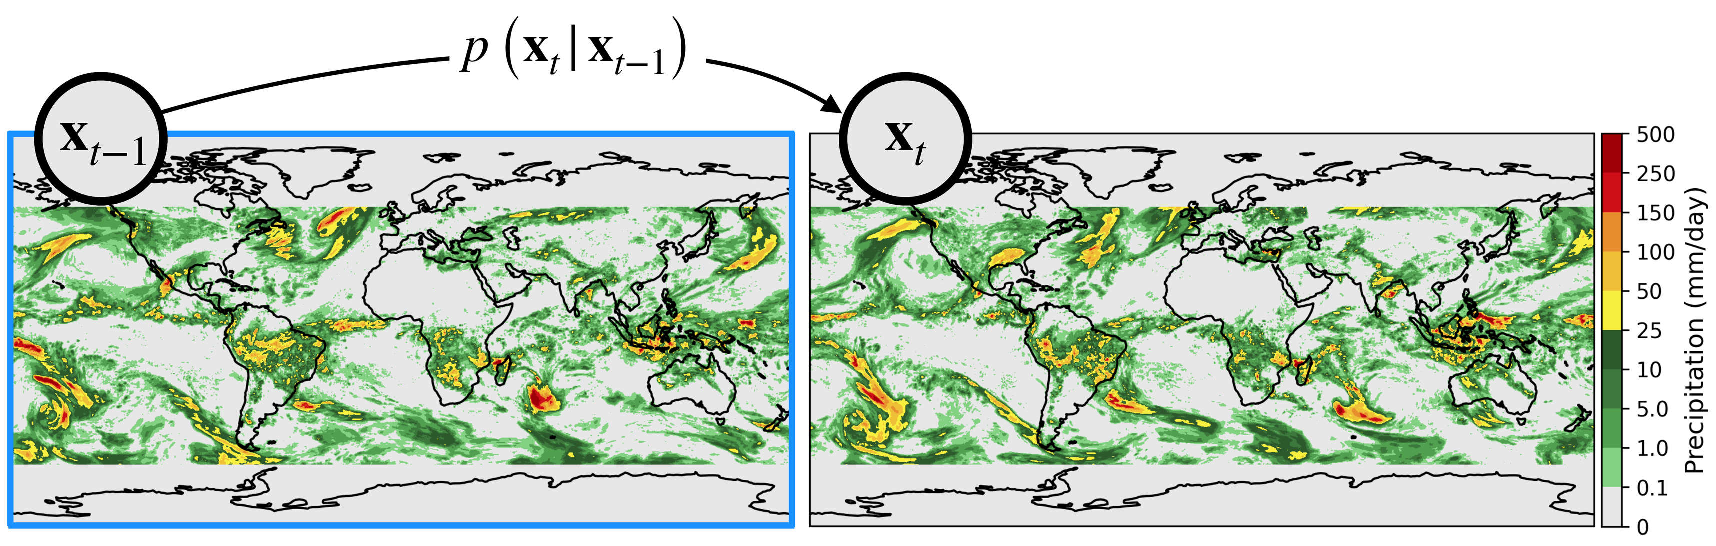

Figure 1: Single step output $\mathbf{x}_t$ from DiffObs when conditioned on the previous state $\mathbf{x}_{t-1}$

Dataset Details

Final precipitation, half hourly Integrated Multi-satellitE Retrievals for Global Perception Measurements (IMERG) L3 Version 06B data (Figure 1).

- **Aggregation**: all half hour samples are daily-accumulated (in mm/d).
- **Spatial Coverage**: grid coarsening from $0.1° \rightarrow 0.4°$ with cropping in the meridional direction (at poles) between $56.2°$N and $61.8°$S.
- **Temporal Partitioning**: separated by years, with 2000–2016 (6,041) for training and 2017–2022 (1,729) for testing, total samples in parentheses.

Experimental Results

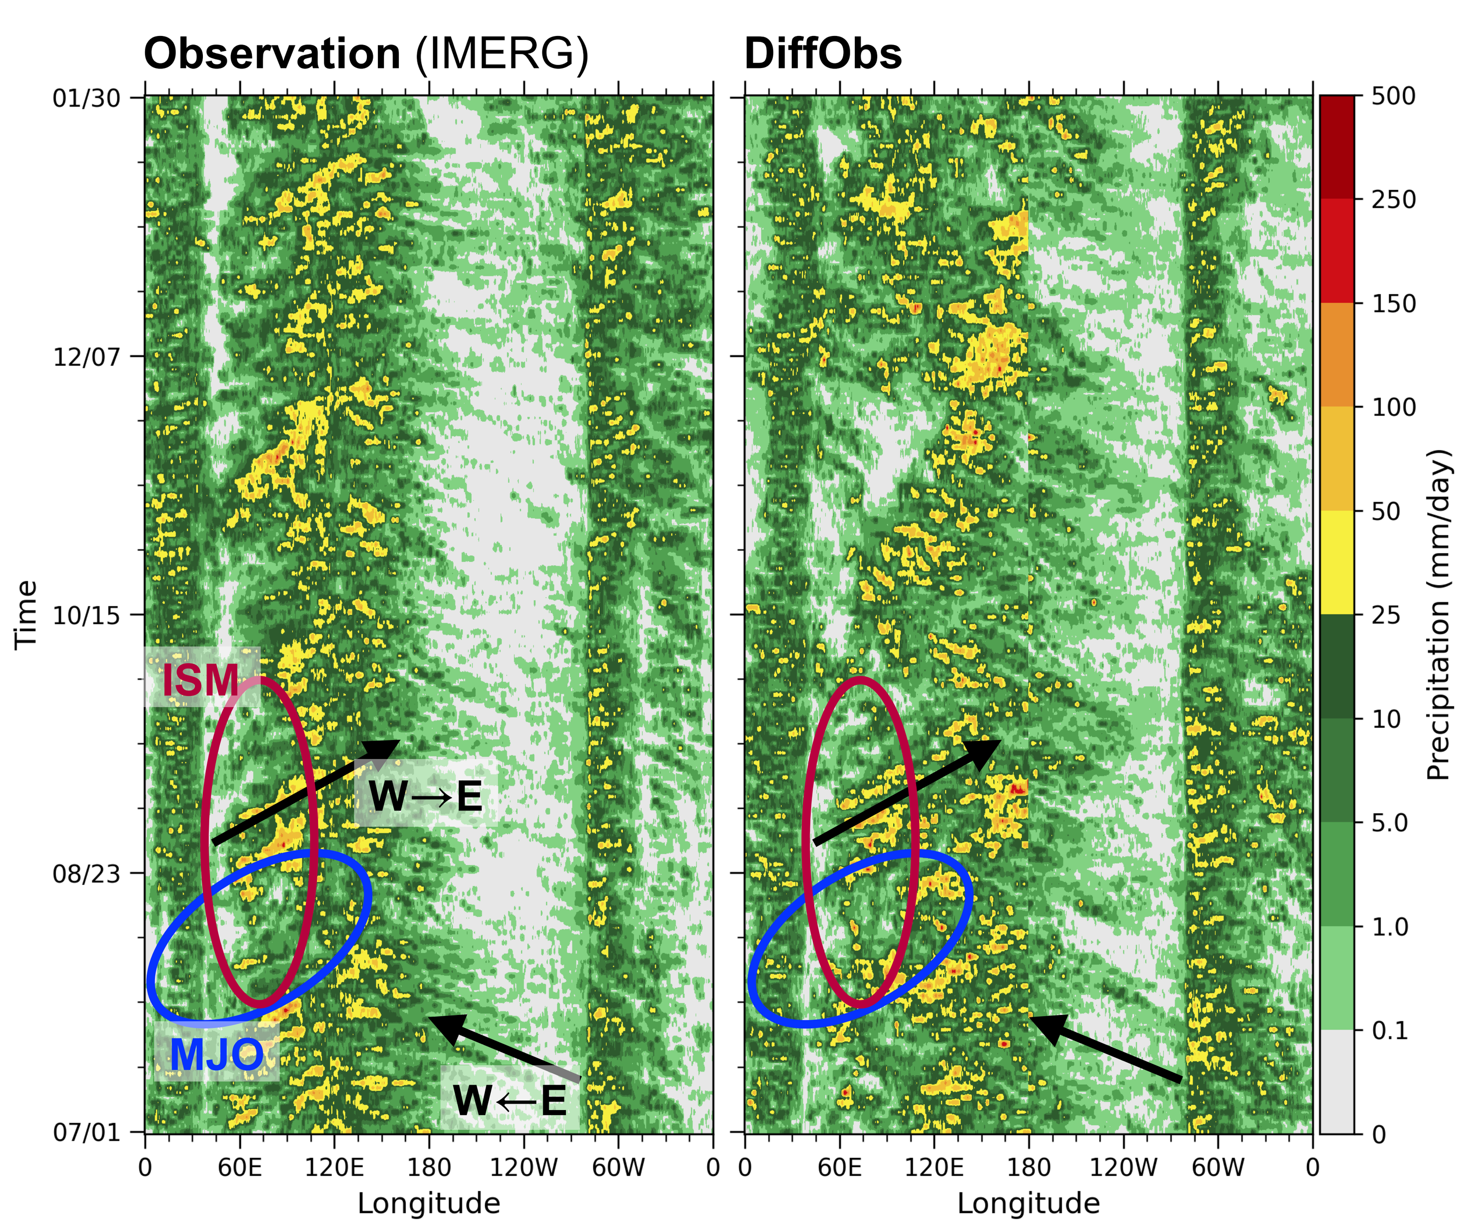

Figure 3: Hovmöller diagrams (between 5°S and 5°N) initialized on July 1, 2020.

Hovmöller Diagram: we generate long rollouts from an initial state, averaging the outputs around the equator and stacking temporally to identify patterns.
Results $\rightarrow$ superposition of eastward- and westward-propagating tropical disturbances, modulated by a large-scale envelope of slow, eastward moving variability characteristic of the Madden–Julian oscillation (MJO).

Additional Features (?): secondary model with **(a)** *temporal conditioning* by zonal averaging the cosine of solar zenith angle as a function of the input condition's date and latitudes, and **(b)** *coordinate conditioning* with static spatially bound $\sin/\cos(\text{lon})$ to counter the network's rotational equivariance.
Results $\rightarrow$ suboptimal performance with inconsistent atmospheric dynamics, including missing landmass influence, tropical wave modes, and directionality.

[2] work done during an internship at NVIDIA
stock@rams.colostate.edu | cs.colostate.edu/~stock

Methodology

Goal: probabilistically forecast day-ahead precipitation, estimating $p(\mathbf{x}_t \mid \mathbf{x}_{t-1})$ without incorporating any additional priors.

Soln: train a 13.6M param conditional EDM diffusion model (adapted UNet architecture) on a cluster of $256 \times 80\text{GB}$ H100 NVIDIA GPUs (32 nodes) using a global batch size of 1,024 for 12.5M total steps.

Diffusion Details: consider the following **forward** SDE

$$\mathrm{d}\mathbf{x} = \sqrt{2\dot{\sigma}(t)\sigma(t)}\,\mathrm{d}\omega_t \quad t \in [0,1],$$

then the **reverse**-time SDE (Figure 2) is given by

$$\mathrm{d}\mathbf{x} = -2\dot{\sigma}(t)\sigma(t)\,\nabla_{\mathbf{x}}\log p\left(\mathbf{x};\sigma(t)\right)\mathrm{d}t + \sqrt{2\dot{\sigma}(t)\sigma(t)}\,\mathrm{d}\bar{\omega}_t.$$

We estimate $\nabla_{\mathbf{x}}\log p(\mathbf{x};\sigma) = \left(D_\theta(\mathbf{x},\sigma) - \mathbf{x}\right)/\sigma^2$ with a denoising neural network D_θ conditioned on $\mathbf{x}_{t-1}$ (via channel-wise concatenated) by minimizing

$$\min_\theta \mathbb{E}_{\mathbf{x}_{t,t-1}\sim p_{\text{data}}}\mathbb{E}_{\sigma\sim p_\sigma}\mathbb{E}_{\mathbf{n}\sim\mathcal{N}(\mathbf{0},\sigma^2\mathbf{I})}\left[\lambda(\sigma)\|D_\theta\left(\mathbf{x}_t+\mathbf{n},\mathbf{x}_{t-1};\sigma\right) - \mathbf{x}_t\|_2^2\right]$$

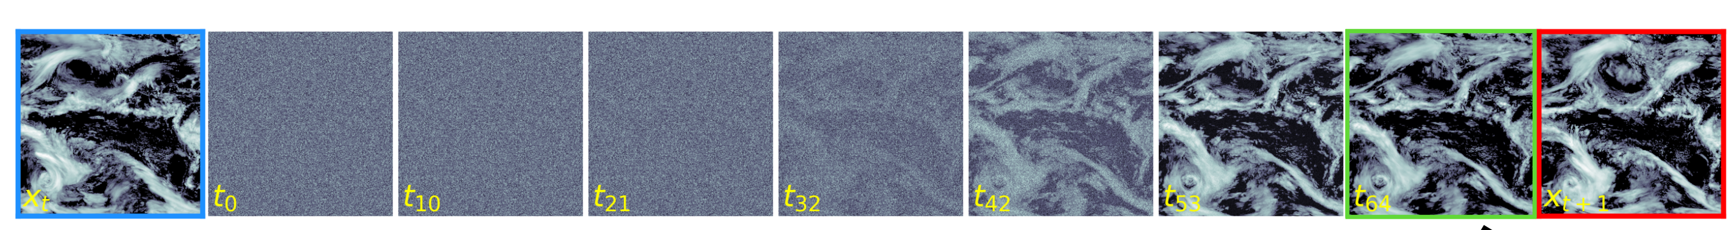

Figure 2: Reverse diffusion with the condition, sampling steps, next step estimate, and target output

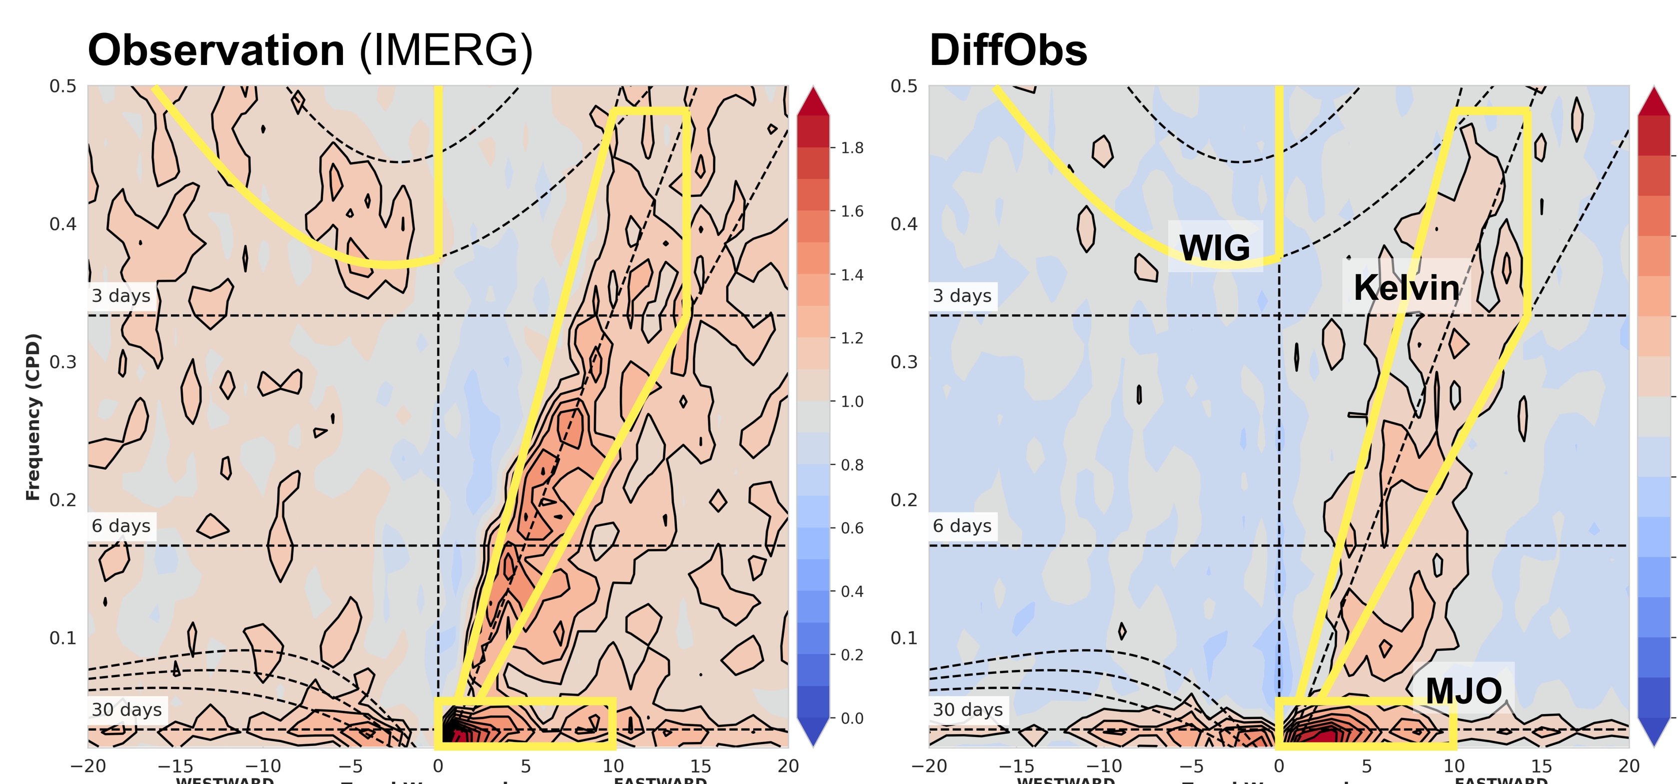

Figure 4: Symmetric / Background Wheeler-Kiladis space-time spectra between $15°$S and $15°$N

Spectral Analysis: we generate 80 yrs of data on one-day intervals, initially conditioning 1 yr rollouts on Jan 1 for years 2017–2021 and sample with perturbed noise, concatenating the results temporally prior to analysis.
Results $\rightarrow$ discovery of Kelvin wave and strong MJO spectral signals within the signal-to-noise ratio of the equatorially-symmetric component.

Conclusions

Overview: we autoregressively generate multi-month, high-resolution rollouts of univariate observations, showing with domain-specific diagnostics (Hovmöller / Wheeler-Kiladis) stable rollouts and a realistic spectrum of tropical wave modes.

Takeaways: global diffusion models trained on sparse observations of the world show promise for applications in subseasonal and climate prediction.

Predictability (?): *not our main focus*, yet we find the best skill (RMSE / FSS) out to 3- to 5-day lead times when compared to persistence and climatology.

Future Work: leverage lower-level products, explicitly model the temporal distributions, and compare to numerical weather prediction and reanalysis.