Physical Perturbation of Atmospheric States at 1° Resolution

Siddharth Rout[1,2], Eldad Haber[1,2], Stéphane Gaudreault[3]

[1]Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia
[2]Institute of Applied Mathematics, University of British Columbia
[3]Atmospheric Numerical Weather Prediction Research, Environment and Climate Change Canada

UBC Earth, Ocean and Atmospheric Sciences

INSTITUTE OF APPLIED MATHEMATICS

Environment and Climate Change Canada

Environnement et Changement climatique Canada

Conventional Numerical Weather Prediction Models are Expensive!!

"The time-to-solution for training FourCastNet measured on JUWELS Booster on 3,072 GPUs is 67.4 minutes, resulting in an 80,000 times faster time-to-solution relative to state-of-the-art NWP, in inference."

Latency and Energy consumption for a 24-hour 100-member ensemble forecast				
	IFS	FCN - 30km (actual)	FCN - 18km (extrapolated)	IFS / FCN(18km) Ratio
Nodes required	3060	1	2	**1530**
Latency (Node-seconds)	984000	7	22	**44727**
Energy Consumed (kJ)	271000	7	22	**12318**

Table 2: The FourCastNet model can compute a 100-member ensemble forecast on a single 4GPU A100 node. In comparison, the IFS model needs 3060 nodes for such a forecast. In this table, we provide information about latency and energy consumption for the FourCastNet model in comparison with the IFS model. The FourCastNet model at a 30km resolution is about 145,000 times faster on a single-node basis than the IFS model. We can also estimate the cost of generating an 18km resolution forecast using FourCastNet. Such a hypothetical 18km model would be about 45,000 times faster than the IFS on a single-node basis. The FourCastNet model at 30km resolution uses 24,000 times less energy to compute the ensemble forecast than the IFS model, while a hypothetical FourCastNet model at 18km resolution would use 12000 times less energy.

*FourCastNet: https://doi.org/10.1145/3592979.3593412, https://doi.org/10.48550/arXiv.2202.11214

Semi-Lagrangian Advection Operator for Deep Learning

Large advection in grid scale is a problem to model.

Convolution Neural Networks (CNN) are not sufficient!

CNNs' depth needs to scale linearly with the scale of advection.

NSL: Semi-Lagrangian push operator to learn large advections.

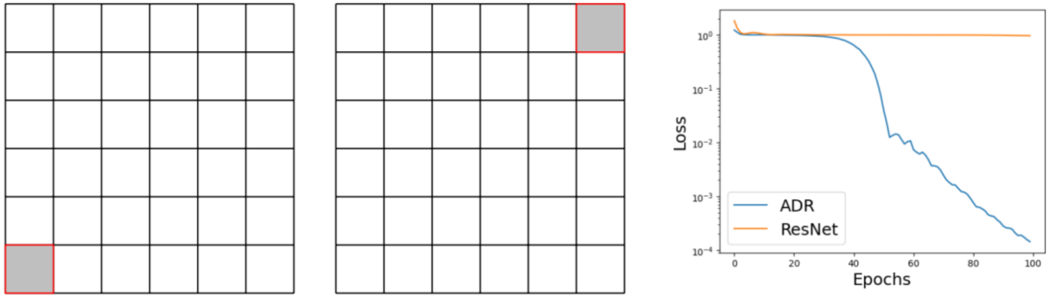

(A) Source image (B) Target image (C) Convergence

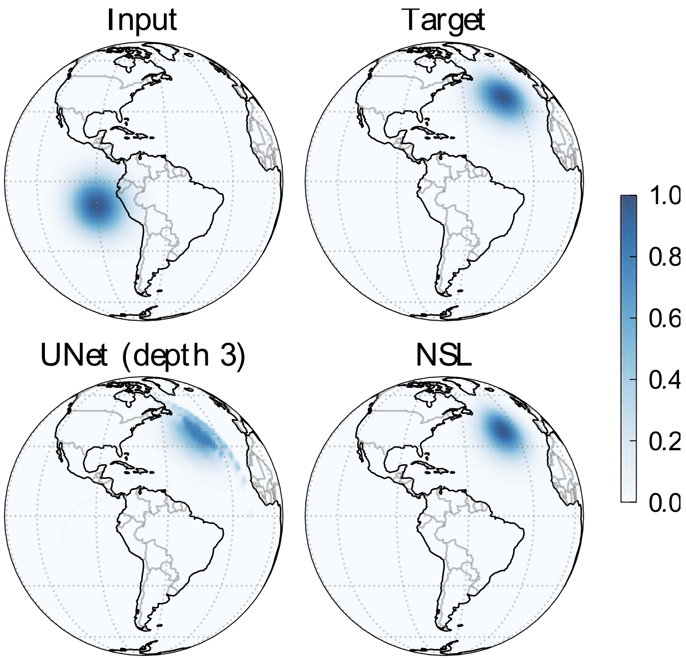

NSL: Neural Semi Lagrangian
ADR: NSL Based CNN doi.org/10.52202/079017-1110
ResNet: doi.org/10.1109/CVPR.2016.90
UNet: doi.org/10.1007/978-3-319-24574-4_28

PARADIS: Global Weather AI Model at 1° Resolution

Table 5. Input and output variable inventory for the PARADIS model. Atmospheric variables are evaluated at 13 pressure levels (50, 100, 150, 200, 250, 300, 400, 500, 600, 700, 850, 925, 1000 hPa).

Category	Variable Name	Input	Output
Atmospheric	Geopotential	✓	✓
	Cartesian horizontal wind vector	✓	✓
	Specific humidity	✓	✓
	Temperature	✓	✓
	Vertical velocity	✓	✓
Surface	10m Cartesian horizontal wind vector	✓	✓
	2m Temperature	✓	✓
	Mean sea-level pressure	✓	✓
	Surface pressure	✓	✓
	Total column water	✓	✓
Forcings	TOA incident solar radiation	✓	✗
	Time of day (sin, cos)	✓	✗
	Year progress (sin, cos)	✓	✗
Constants	Geopotential at surface	✓	✗
	Land-sea mask	✓	✗
	Topographic slope	✓	✗
	Sub-grid topographic standard deviation	✓	✗
	Inverse longitude spacing	✓	✗
	Latitude / longitude	✓	✗
	$\cos(\phi)$	✓	✗
	$\sin(\lambda), \cos(\lambda)$	✓	✗

PARADIS: doi.org/10.48550/arXiv.2601.21151

PARADIS vs ECMWF HRES vs Google GraphCast:

Z500

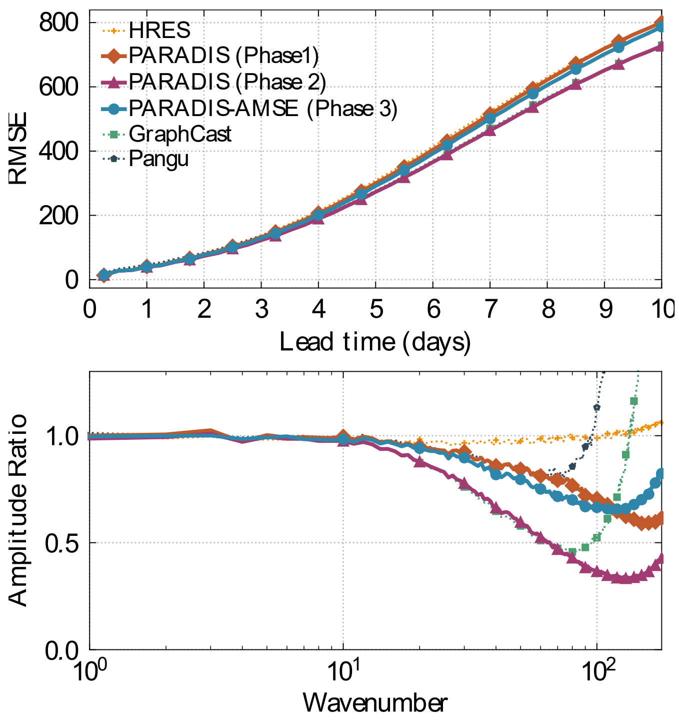

Cyclones Tracked in 2020

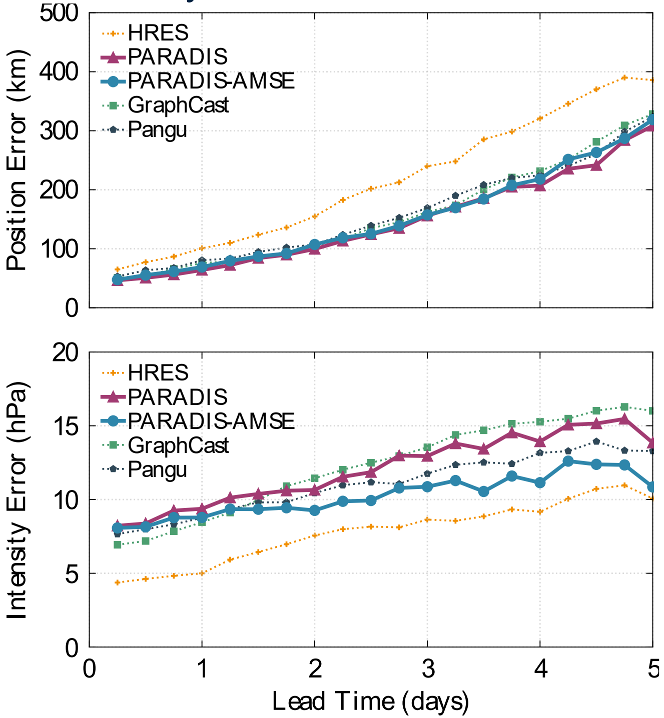

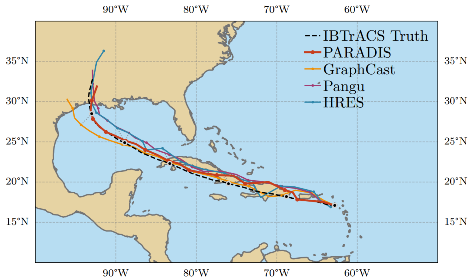

Figure 1. Tracking of cyclone eye for Hurricane Laura (August 2020) using different models and the observed trajectory (IBTrACS dataset (Knapp et al., 2010; Gahtan et al., 2024)).

PARADIS: doi.org/10.48550/arXiv.2601.21151
HRES: doi.org/10.5065/D68050ZV
GraphCast: doi.org/10.1126/science.adi2336
Pangu: doi.org/10.1038/s41586-023-06185-3

Predicting Hurricane Laura (August 2020)

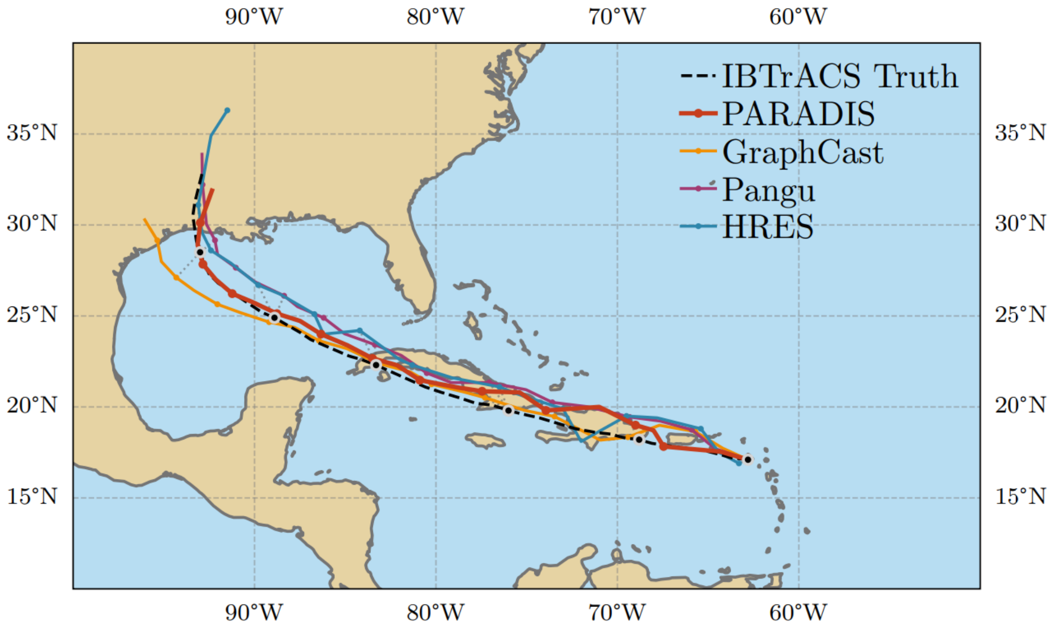

Figure 1. Tracking of cyclone eye for Hurricane Laura (August 2020) using different models and the observed trajectory (IBTrACS dataset (Knapp et al., 2010; Gahtan et al., 2024)).

PARADIS: doi.org/10.48550/arXiv.2601.21151
HRES: doi.org/10.5065/D68050ZV
GraphCast: doi.org/10.1126/science.adi2336
Pangu: doi.org/10.1038/s41586-023-06185-3

"All models are wrong, but some are useful"

George E. P. Box

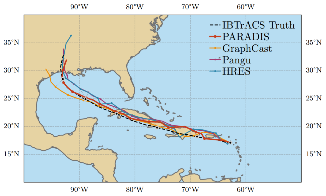

Figure 1. Tracking of cyclone eye for Hurricane Laura (August 2020) using different models and the observed trajectory (IBTrACS dataset (Knapp et al., 2010; Gahtan et al., 2024)).

Need for Larger Set of Ensemble Forecasts

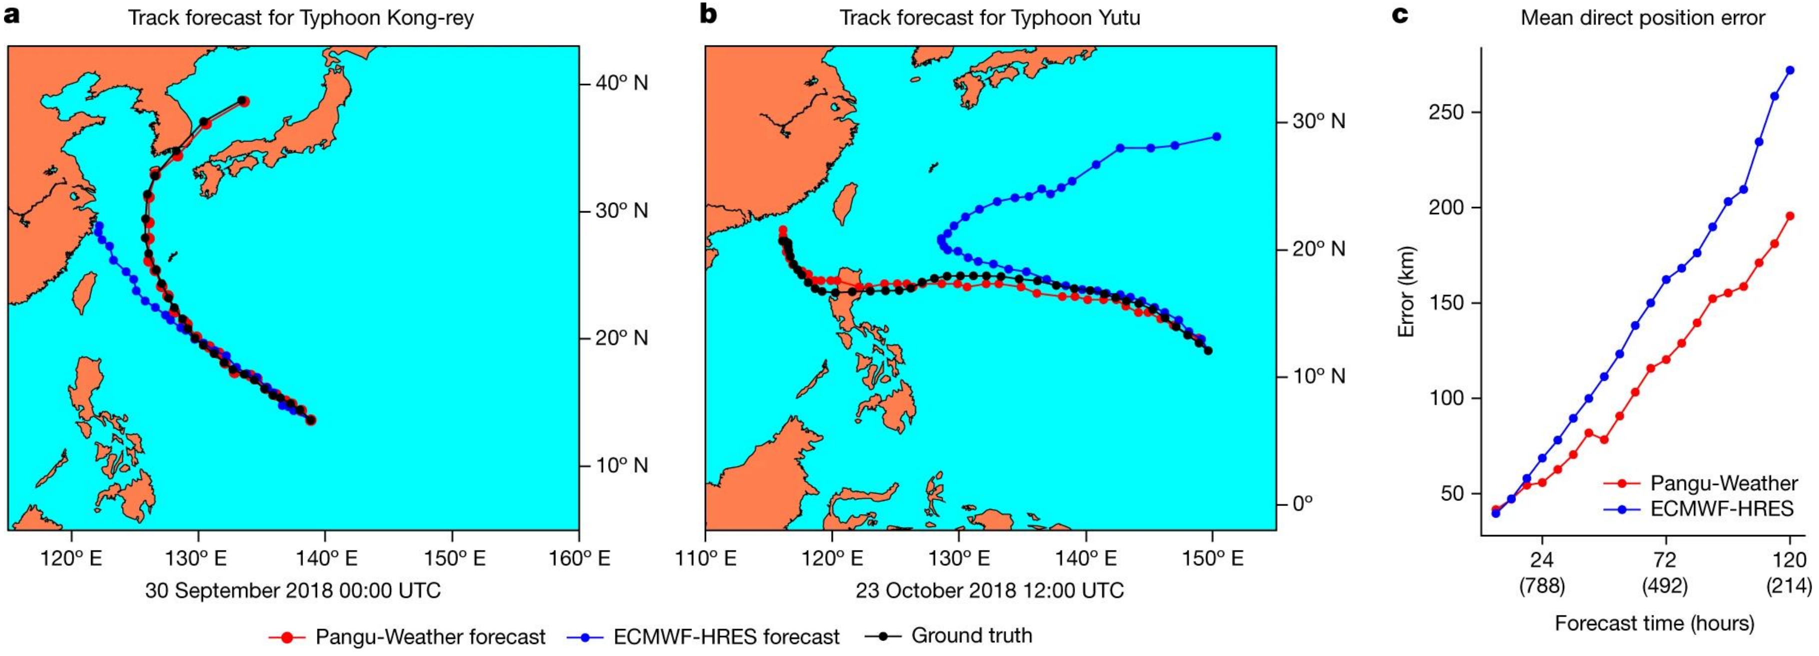

HRES: doi.org/10.5065/D68050ZV

Pangu: doi.org/10.1038/s41586-023-06185-3

Easy to Generate Random Noise
But Not Easy to Generate Images :

The backbone of generative AI

Distribution of Complex Data

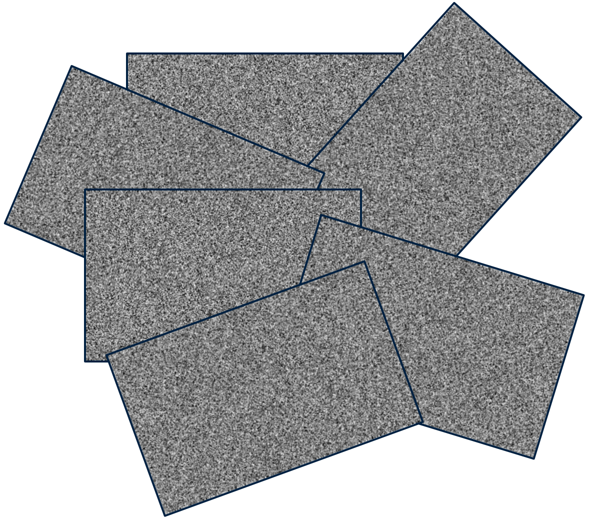

Distribution of Noises

Ideally,

Generative Model (Noise) = Real Data

Leveraging Flow Based Generative Models to Perturb Data Realistically

Invertible

Reconstruction

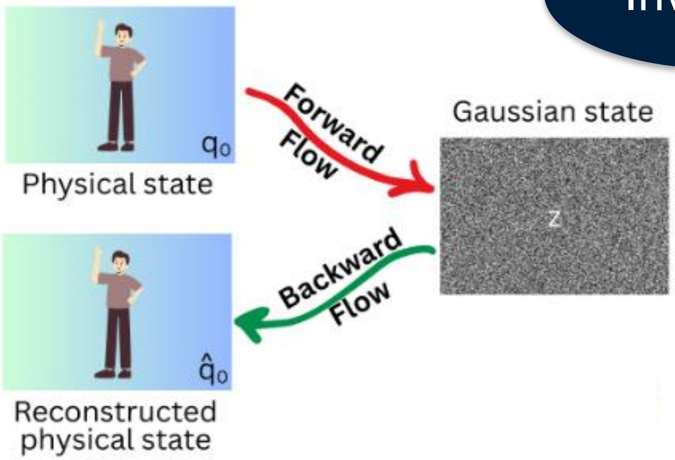

Perturbation

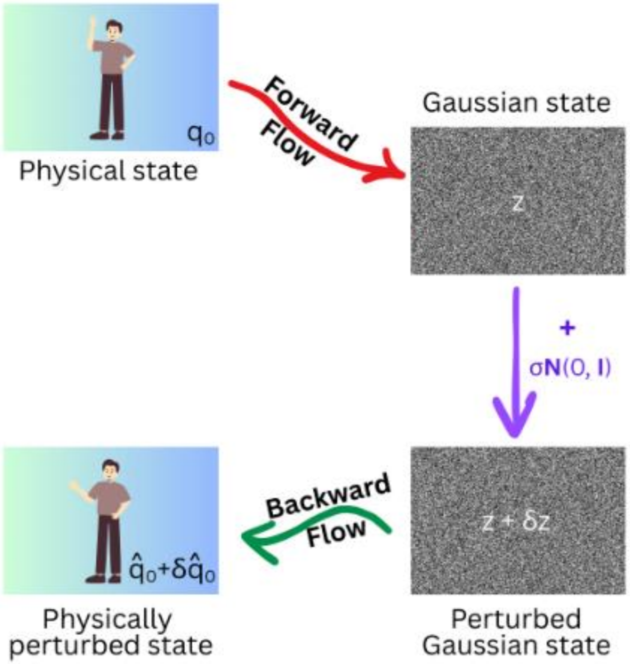

Forward Flow:
Generative Model (Noise) = Realistic Data
Backward Flow:
Generative Model^{-1} (Realistic Data) = Noise

*https://doi.org/10.48550/arXiv.2508.01101

Leveraging Flow Based Generative Models to Perturb Data Realistically

Forward Flow:
Generative Model (Noise) = Realistic Data
Backward Flow:
Generative Model^{-1} (Realistic Data) = Noise

Logic:

- F is a trained generative flow model.
- Let X is a sample from physical data.

$$\mathbf{F(F^{-1}(X)) = X}$$

- We don't know the physical perturbation dX, but since $F^{-1}(X)$ is a Gaussian Noise, I can perturb $F^{-1}(X)$.
- $$\mathbf{X' = F(F^{-1}(X) + \sigma \mathcal{N}(O,I))}$$
- **physical perturbation** (dX) = X' - X

Perturbation

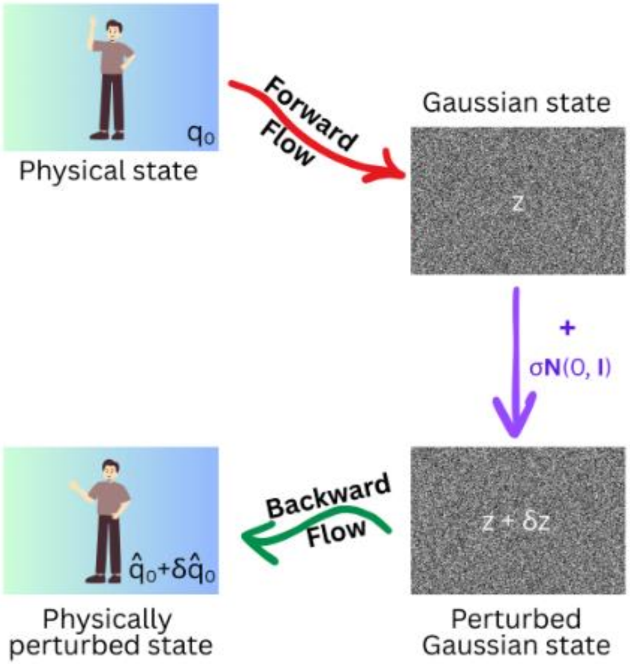

*https://doi.org/10.48550/arXiv.2508.01101

Leveraging Flow Based Generative Models to Perturb Data Realistically

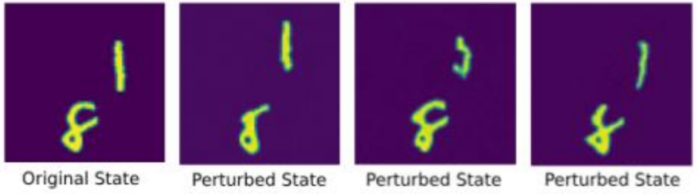

Figure 9: Three random perturbed states of a Moving MNIST sample state.

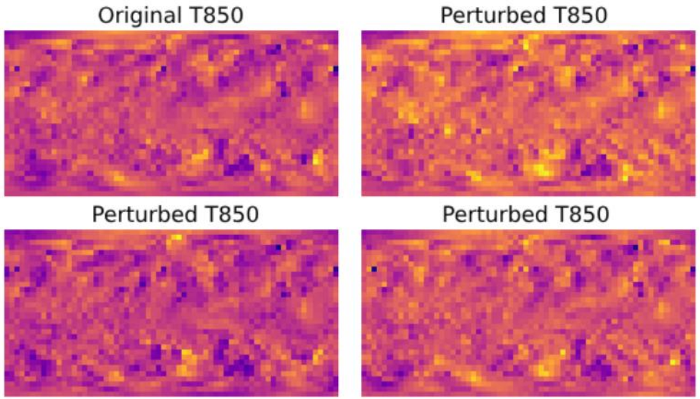

Figure 11: Perturbed states of a WeatherBench T850 state.

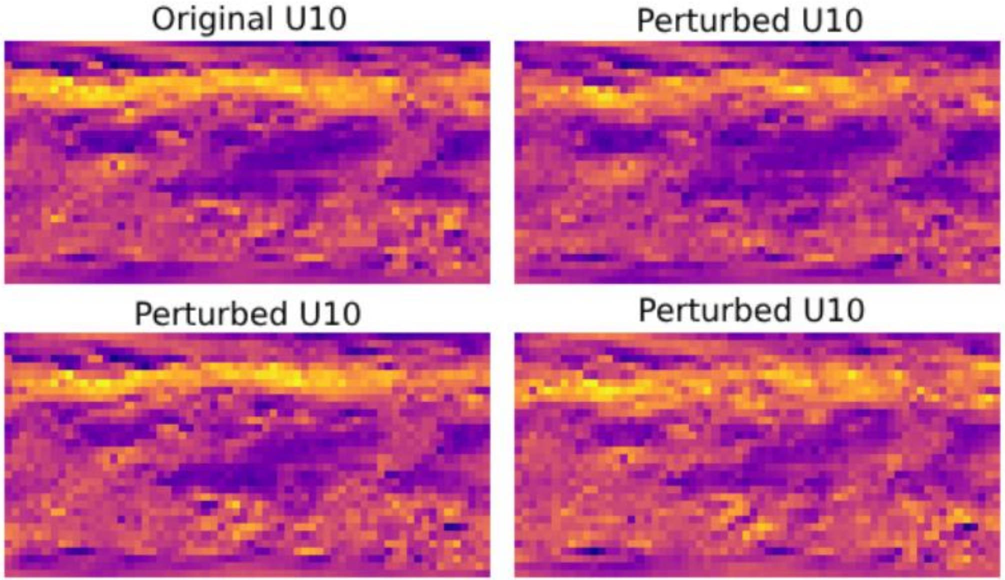

Figure 10: Three random perturbed states of a WeatherBench U10 state.

An Example from WeatherBench 1.40625deg

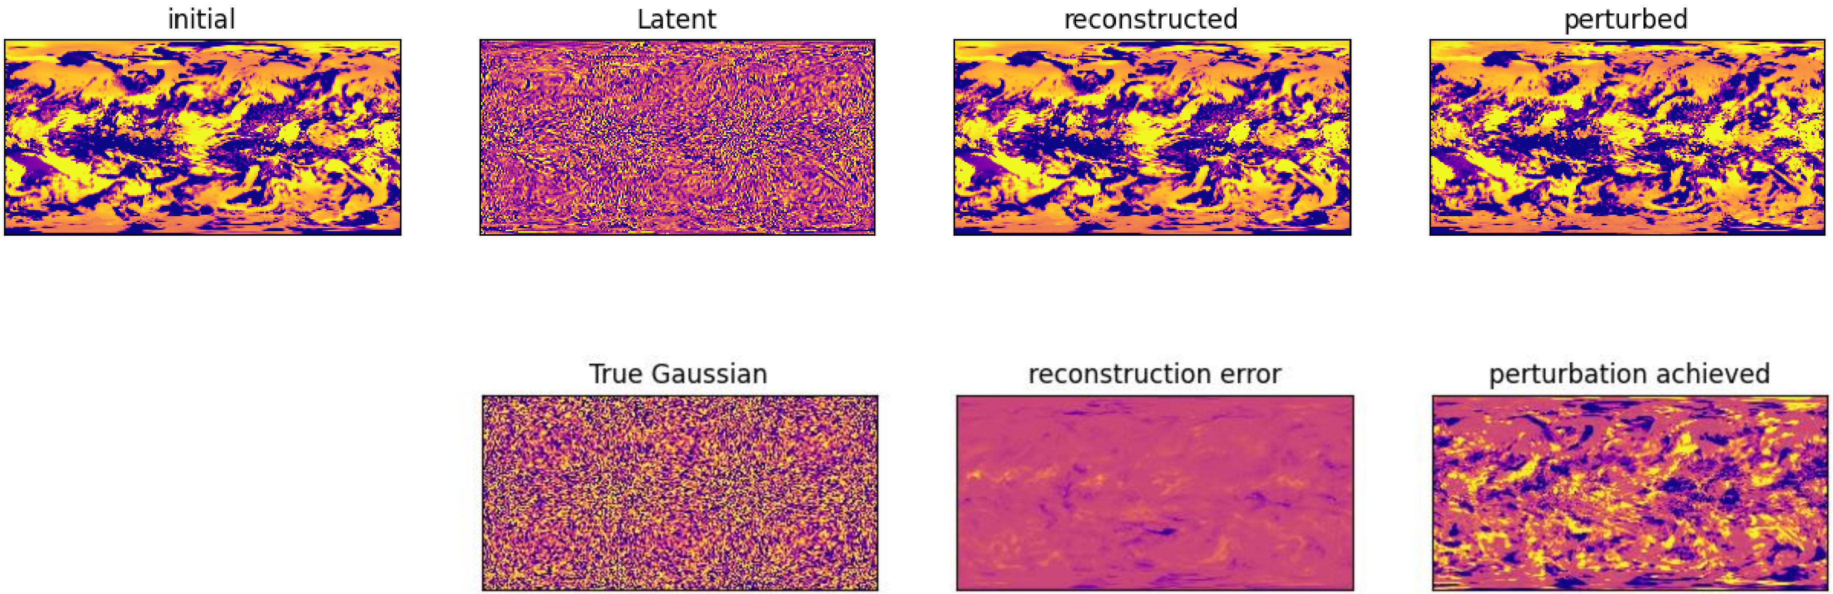

Conclusion

- We got a technique to generate multiple perturbed initial states for AI and NWP models.
- With Deep Learning, we can easily generate 1000 ensemble members in minutes.
- So, easily produce 1000 ensemble predictions.

- We engineering PARADIS 2.0 with ensemble forecasting system, to be released by the Summer 2026.

Special Thanks to

Primary References

- Zakariaei, N., Rout, S., Haber, E. and Eliasof, M., 2024.
  Advection augmented convolutional neural networks. Advances in Neural Information Processing Systems, 37, pp.35225-35256.

- Rout, S., Haber, E. and Gaudreault, S., 2025.
  Flow Matching for Probabilistic Learning of Dynamical Systems from Missing or Noisy Data. arXiv preprint arXiv:2508.01101.

- Pereira, C.A., Gaudreault, S., Dallerit, V., Subich, C., Panday, S., Wei, S., Zhang, S.,Rout, S., Haber, E., Spiteri, R.J. and Millard, D., 2026.
  Learning to Advect: A NeuralSemi-Lagrangian Architecture for Weather Forecasting. arXiv preprintarXiv:2601.21151.

I Invite collaborators in Applications and Research


UBC
THE UNIVERSITY OF BRITISH COLUMBIA
THE UNIVERSITY OF BRITISH