Pangu-Weather

Accurate Medium-Range Global Weather Forecasting with 3D Deep Neural Networks

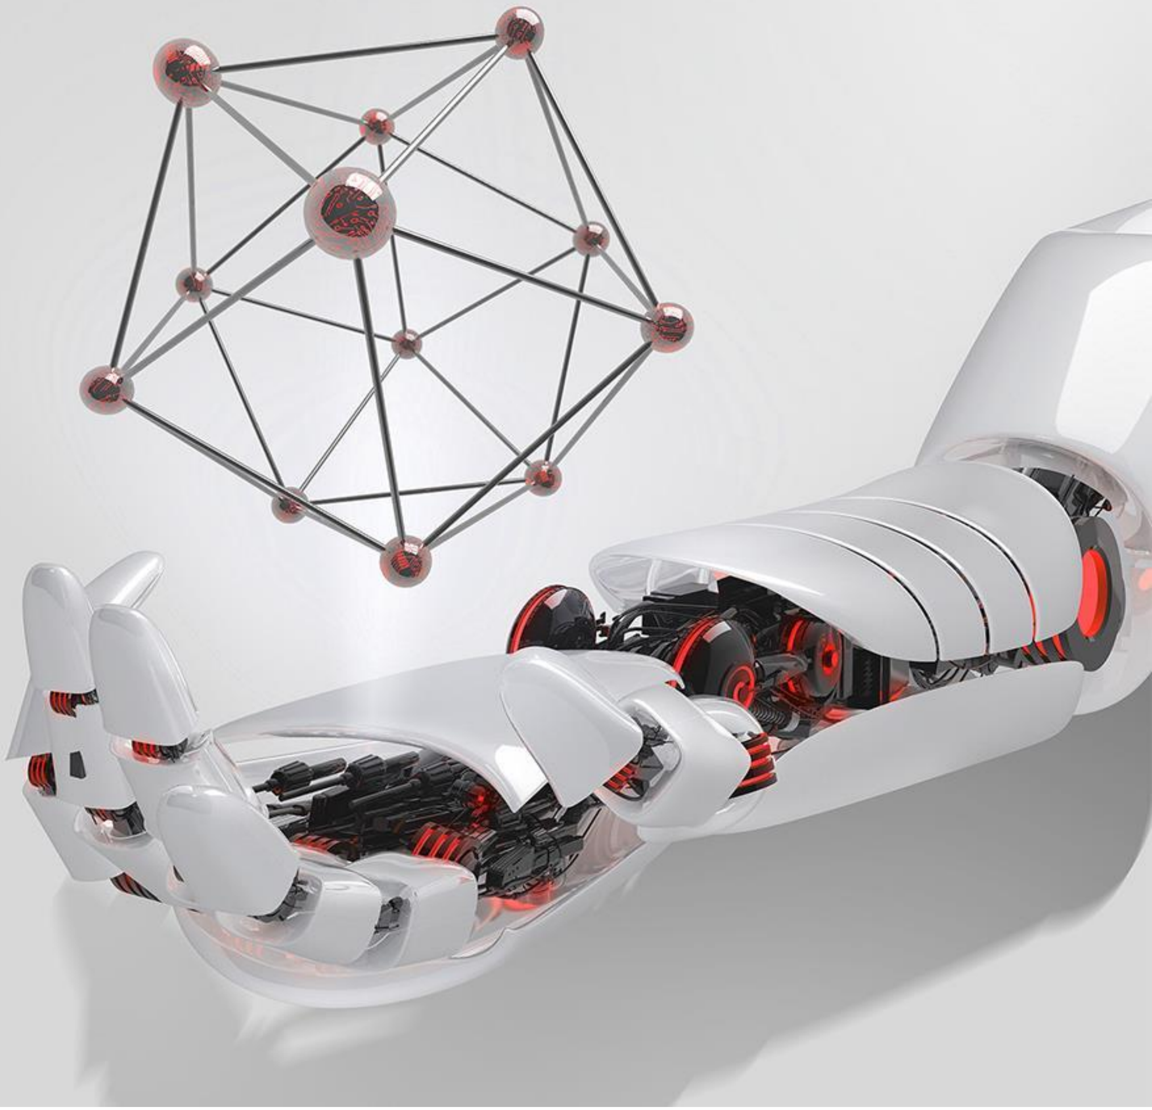

Speaker: Lingxi Xie (谢凌曦)

Sep 4th, 2023

HUAWEI

Outline

- Background: Why Using AI for NWP?
- Pangu-Weather: 3D Deep Networks for Accurate Weather Forecasting
- Results: Determinstic/Ensemble Forecast, Extreme Weather Forecast
- Future Perspectives

Numerical Weather Prediction (NWP)

- **Task description**
  - Given a set of weather variables at the current time, predict the variables at a specified time in the future (*e.g.*, 5 days later)
  - Mathematically, the task is to learn a function $f(\cdot)$ that uses the current weather data as input and produces future weather data as output
- **Two main lines of research for NWP (to be detailed later)**
  - Simulation-based methods: approximating $f(\cdot)$ with a set of partial differential equations (PDEs)[A,B]
  - AI-based methods: estimating $f(\cdot)$ using deep neural networks[C,D]

[A] **E. Kalnay** *et al.*, Atmospheric Modeling, Data Assimilation and Predictability, *Cambridge University Press*, 2003.
[B] **P. Bauer** *et al.*, The Quiet Revolution of Numerical Weather Prediction, in *Nature*, 2015.
[C] **J. Pathak** *et al.*, FourCastNet: A Global Data-driven High-resolution Weather Model Using Adaptive Fourier Neural Operators, in *arXiv preprint:2202.11214*, 2022.
[D] **K. Bi** *et al.*, Pangu-Weather: A 3D High-Resolution Model for Fast and Accurate Global Weather Forecast, in *arXiv preprint:2211.02556*, 2022.

Examples of NWP with Reanalysis Data

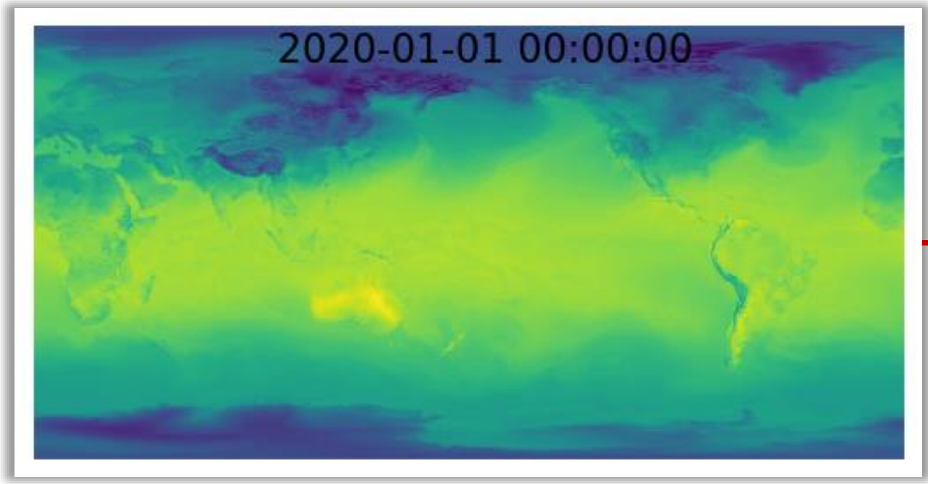

Input: 2m temperature (T2M)

NWP algorithm

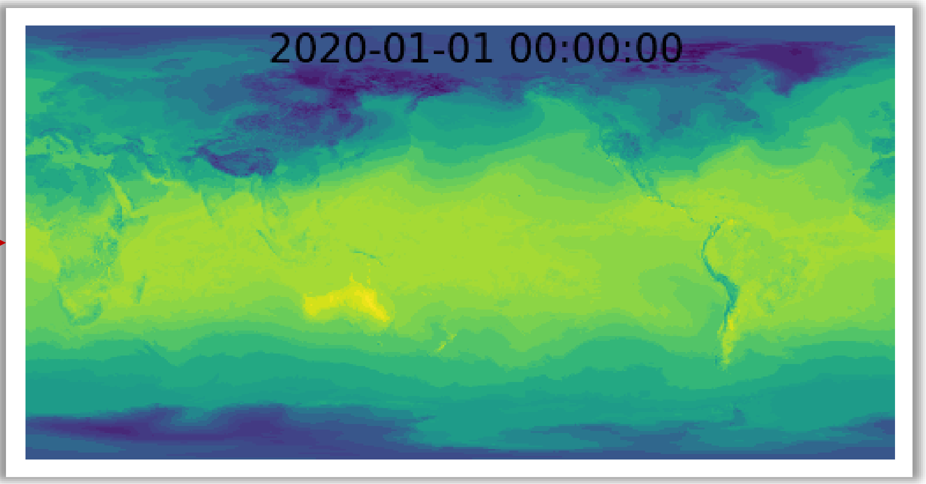

Ground-truth for the next 3 days

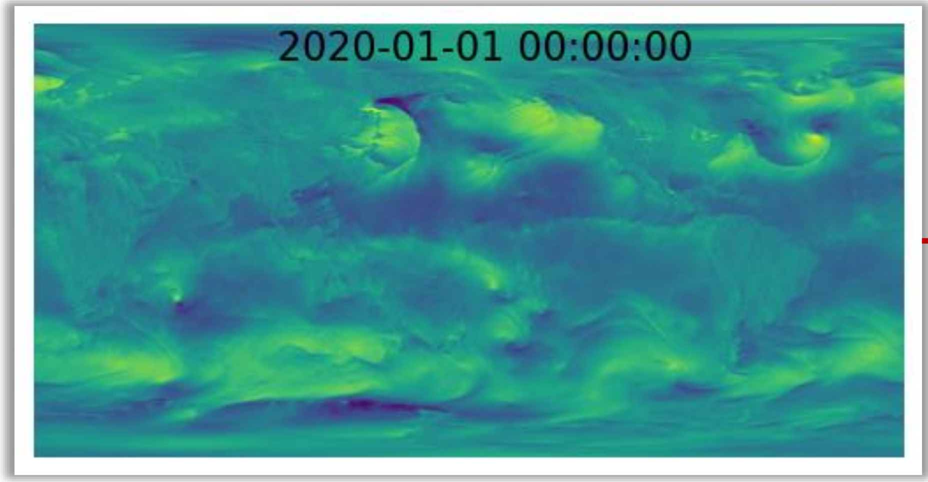

Input: 10m *u*-wind speed (U10)

NWP algorithm

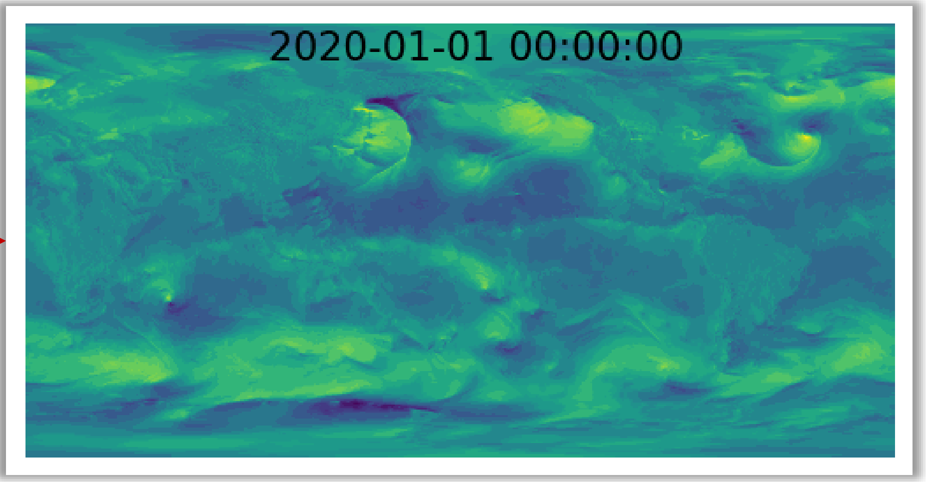

Ground-truth for the next 3 days

Credit: all data are from the 5th generation of the ECMWF reanalysis (ERA5) data[A].

[A] **H. Hersbach** *et al.*, The ERA5 Global Analysis, in *Quarterly Journal of the Royal Meteorological Society*, 2020.

Artificial Intelligence

- **On replicating human intelligence with mathematical methods**
  - AI has largely changed the way how people work
  - Key problem: finding the **relationship** between input and output data

Sortation system

Customer service

Weather forecasting

Playing Go and chess

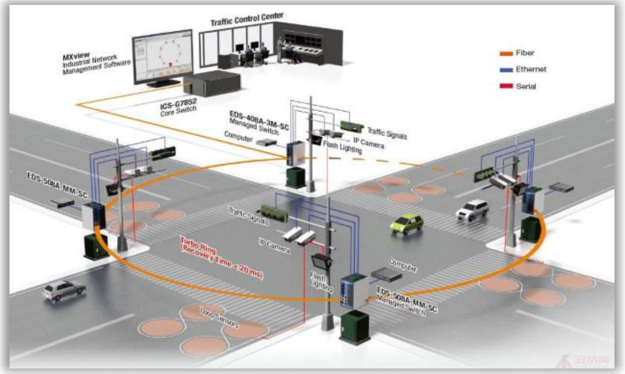

Security monitoring

Medical diagnosis

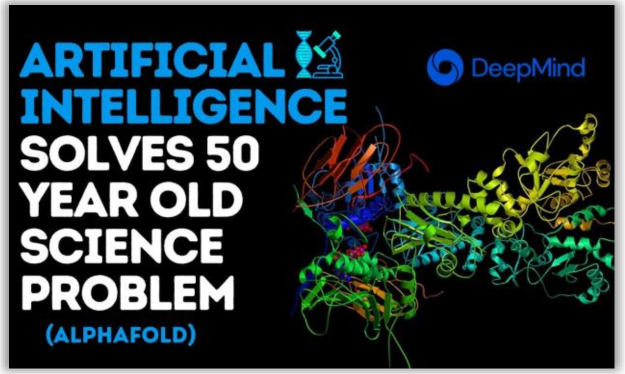

Protein structure prediction

Deep Learning

- **Representing complex functions with deep neural networks**
  - A large number of parameters to **approximate** the complex function
  - End-to-end optimization, making the model hard to interprete
  - Widely applied to many AI problems: CV, NLP, RL, *etc.*

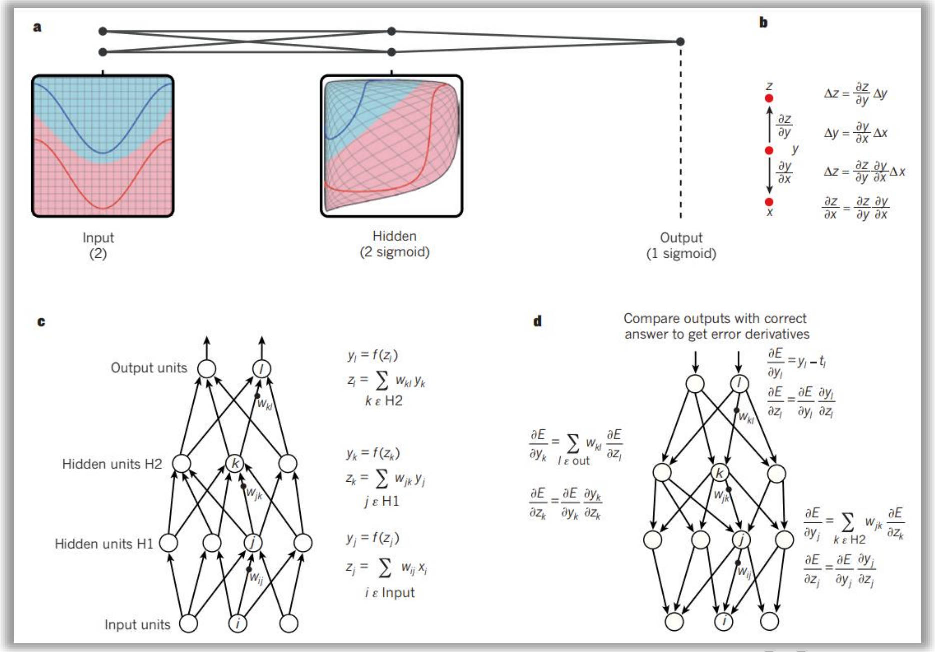

Optimization of deep learning[A]

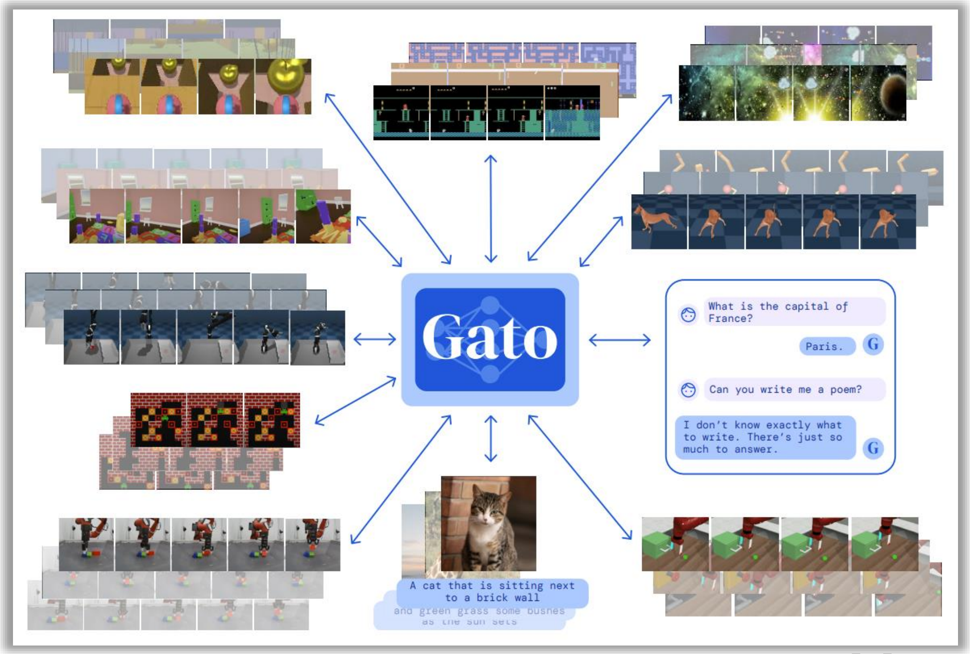

Deep network as a generalized agent[B]

[A] Y. LeCun *et al.*, Deep Learning, in *Nature*, 2015.
[B] S. Reed *et al.*, A Generalized Agent, in *TMLR*, 2022.

AI for Science (AI4S), Part I

- **Solving scientific problems with AI methods**
  - AI brings new solutions and opportunities to old problems

Category	Previous	Now	Examples
Playing chess or video games	Heuristic search with manually designed heuristic functions	Deep learning as heuristics, plus reinforcement learning	DQN[A], AlphaGo[B], AlphaStar[C], *etc.*
Life and biology	Solving thermodynamic or kinetic simulation of protein physics	Deep learning incorporated with the constraints of protein structures	AlphaFold[D], *etc.*
Geophysics	Heuristics for finding patterns in aftershock data	Deep learning for analyzing aftershock patterns	DeVries *et al.*[E], *etc.*
Nuclear fusion	Engineering based on physics to solve the Tokamak magnetic control problem	Deep reinforcement learning for Tokamak magnetic controller design	Degrave *et al.*[F], *etc.*

[A] **V. Mnih** *et al.*, Human-level Control through Deep Reinforcement Learning, in *Nature*, 2015.
[B] **D. Silver** *et al.*, Mastering the Game of Go with Deep Neural Networks and Tree Search, in *Nature*, 2015.
[C] **O. Vinyals** *et al.*, Grandmaster Level in StarCraft II Using Multi-agent Reinforcement Learning, in *Nature*, 2019.
[D] **J. Jumper** *et al.*, Highly Accurate Protein Structure Prediction with AlphaFold, in *Nature*, 2021.
[E] **P. DeVries** *et al.*, Deep Learning of Aftershock Patterns Following Large Earthquakes, in *Nature*, 2018.
[F] **J. Degrave** *et al.*, Magnetic Control of Tokamak Plasmas Through Deep Reinforcement Learning, in *Nature*, 2022.

AI for Science (AI4S), Part II

- **Solving scientific problems with AI methods**
  - AI brings new solutions and opportunities to old problems

Category	Previous	Now	Examples
Mathematics	Relying on mathematicians' intuition to find theorems	AI algorithms assisting mathematicians to find and prove new theorems	Ramanujan Machine[A], Davies *et al.*[B], *etc.*
Computing theory	Designing efficient computing algorithms manually	Reinforcement learning to find efficient computing algorithms	AlphaTensor[C], *etc.*
Weather nowcasting	Simulating weather with partial differential equations	Deep generative model to find patterns from radar data	DGMR[D], *etc.*
Medium-range weather forecasting	Simulating weather with partial differential equations	Deep learning to fit global reanalysis weather data	FourCastNet[E], Pangu-Weather[F], *etc.*

[A] **G. Raayoni** *et al.*, Generating Conjectures on Fundamental Constants with the Ramanujan Machine, in *Nature*, 2021.
[B] **A. Davies** *et al.*, Advancing Mathematics by Guiding Human Intuition with AI, in *Nature*, 2021.
[C] **A. Fawzi** *et al.*, Discovering Faster Matrix Multiplication Algorithms with Reinforcement Learning, in *Nature*, 2022.
[D] **S. Ravuri** *et al.*, Skillful Precipitation Nowcasting using Deep Generative Models of Radar, in *Nature*, 2021.
[E] **J. Pathak** *et al.*, FourCastNet: A Global Data-driven High-resolution Weather Model Using Adaptive Fourier Neural Operators, in *arXiv preprint:2202.11214*, 2022.
[F] **K. Bi** *et al.*, Accurate Medium-range Global Weather Forecasting with 3D Neural Networks, in *Nature*, 2023.

Why Using AI for NWP?

- **Two major principles of using AI (apply to generic scenarios)**
  - Human experience is **insufficient** to formulate the complex system
    - In NWP, the manually-designed PDE systems involve **approximation** (*e.g.*, in dealing with unresolved processes) and **parameterization** (*e.g.*, in formulating convection), and the error can accumulate with forecast time
  - Data is **sufficient** for training models with a large number of parameters
    - In NWP, a large-scale dataset (*i.e.*, ERA5[A]) is available, offering global weather data from 1940s to date
- **The advantages of AI in NWP**
  - Much faster: **10,000 times faster** than simulation-based methods
  - Offering individual forecasts that are complementary to PDEs

[A] H. Hersbach *et al.*, The ERA5 Global Analysis, in *Quarterly Journal of the Royal Meteorological Society*, 2020.

Prior Work: Simulation-based NWP

- **Solving partial differential equations (PDEs) for simulation**
  - Governing a complex PDE system, modeling discretization, *etc.*
  - May be sensitive to the approximation and parameterization of equations and the noise of observation data

Using the variables defined above, the flux-form Euler equations can be written as

$$\partial_t U + (\nabla \cdot \mathbf{V} u) + \mu_d \alpha \partial_x p + (\alpha/\alpha_d) \partial_\eta p \partial_x \phi = F_U \tag{2.8}$$

$$\partial_t V + (\nabla \cdot \mathbf{V} v) + \mu_d \alpha \partial_y p + (\alpha/\alpha_d) \partial_\eta p \partial_y \phi = F_V \tag{2.9}$$

$$\partial_t W + (\nabla \cdot \mathbf{V} w) - g[(\alpha/\alpha_d) \partial_\eta p - \mu_d] = F_W \tag{2.10}$$

$$\partial_t \Theta_m + (\nabla \cdot \mathbf{V} \theta_m) = F_{\Theta_m} \tag{2.11}$$

$$\partial_t \mu_d + (\nabla \cdot \mathbf{V}) = 0 \tag{2.12}$$

$$\partial_t \phi + \mu_d^{-1}[(\mathbf{V} \cdot \nabla \phi) - gW] = 0 \tag{2.13}$$

$$\partial_t Q_m + (\nabla \cdot \mathbf{V} q_m) = F_{Q_m} \tag{2.14}$$

with the diagnostic equation for dry hydrostatic pressure

$$\partial_\eta \phi = -\alpha_d \mu_d \tag{2.15}$$

and the diagnostic relation for the full pressure (dry air plus water vapor)

$$p = p_0 \left(\frac{R_d \theta_m}{p_0 \alpha_d} \right)^\gamma . \tag{2.16}$$

Using these redefined momentum variables, the governing prognostic equations (2.8)-(2.14) including map factors can be written as

$$\begin{aligned} &\partial_t U + m_x[\partial_x(Uu) + \partial_y(Vu)] \\ &\quad + \partial_\eta(\Omega u) + (m_x/m_y)[\mu_d \alpha \partial_x p + (\alpha/\alpha_d)\partial_\eta p \partial_x \phi] = F_U \end{aligned} \tag{2.18}$$

$$\begin{aligned} &\partial_t V + m_y[\partial_x(Uv) + \partial_y(Vv)] \\ &\quad + (m_y/m_x)\partial_\eta(\Omega v) + (m_y/m_x)[\mu_d \alpha \partial_y p + (\alpha/\alpha_d)\partial_\eta p \partial_y \phi] = F_V \end{aligned} \tag{2.19}$$

$$\partial_t W + m_x[\partial_x(Uw) + \partial_y(Vw)] + \partial_\eta(\Omega w) - m_y^{-1} g[(\alpha/\alpha_d)\partial_\eta p - \mu_d] = F_W \tag{2.20}$$

$$\partial_t \Theta_m + m_x m_y[\partial_x(U\theta_m) + \partial_y(V\theta_m)] + m_y \partial_\eta(\Omega \theta_m) = F_{\Theta_m} \tag{2.21}$$

$$\partial_t \mu_d + m_x m_y[U_x + V_y] + m_y \partial_\eta(\Omega) = 0 \tag{2.22}$$

$$\partial_t \phi + \mu_d^{-1}[m_x m_y(U\partial_x \phi + V\partial_y \phi) + m_y \Omega \partial_\eta \phi - m_y g W] = 0 \tag{2.23}$$

$$\partial_t Q_m + m_x m_y \partial_x(U q_m) + \partial_y(V q_m)] + m_y \partial_\eta(\Omega q_m) = F_{Q_m}, \tag{2.24}$$

which are solved together with the diagnostic equations (2.15) and (2.16).

A small part of equations used in the Advanced Research WRF Model Version 4[A].

[A] W. C. Skamarock *et al.*, A Description of the Advanced Research WRF Model Version 4, by *National Center for Atmospheric Research: Boulder, CO, USA*, 2019.

Prior Work: FourCastNet

- **First AI-based method for high-resolution global weather forecasting**
  - Adaptive Fourier Neural Operator with vision transformers
  - Fine-tuning to alleviate iterative errors
  - Able to perform high-resolution forecasting, but still producing inferior results to operational IFS

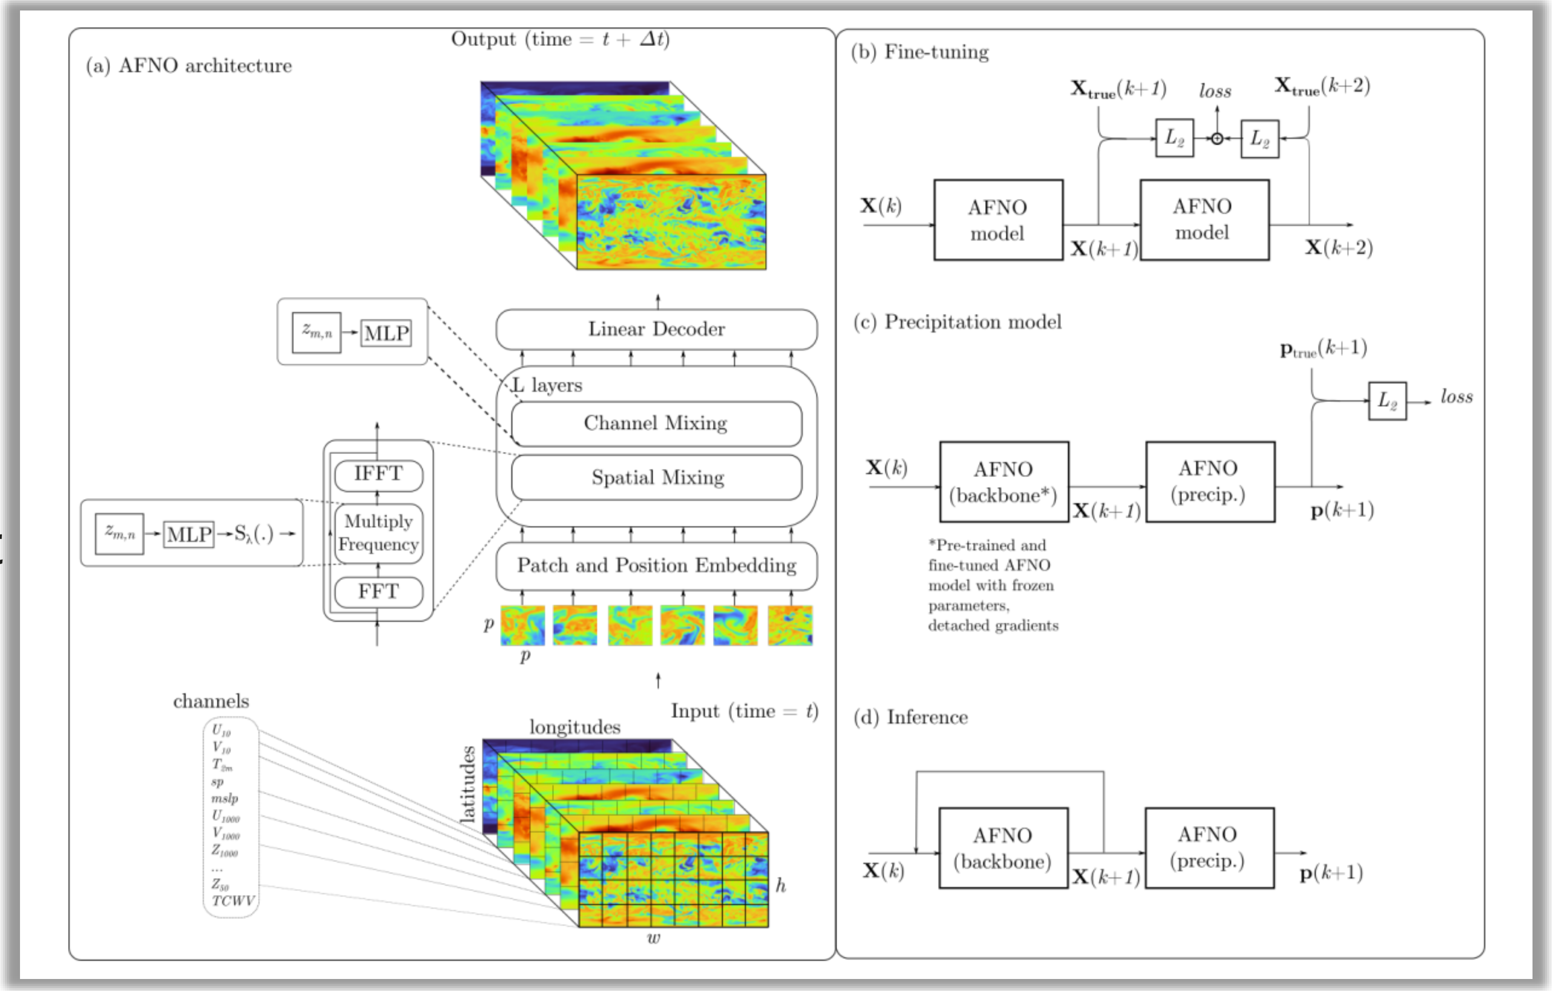

[A] J. Pathak *et al.*, FourCastNet: A Global Data-driven High-resolution Weather Model Using Adaptive Fourier Neural Operators, in *arXiv preprint:2202.11214*, 2022.

Concurrent Work: GraphCast

- **State-of-the-art accuracy on reanalysis data**
  - 2D encoder-decoder with multi-level, graph-based message passing
  - Tested 252 variables, most metrics are SOTA
  - But, it smoothifies the forecasts, weakening its ability in ensemble forecast

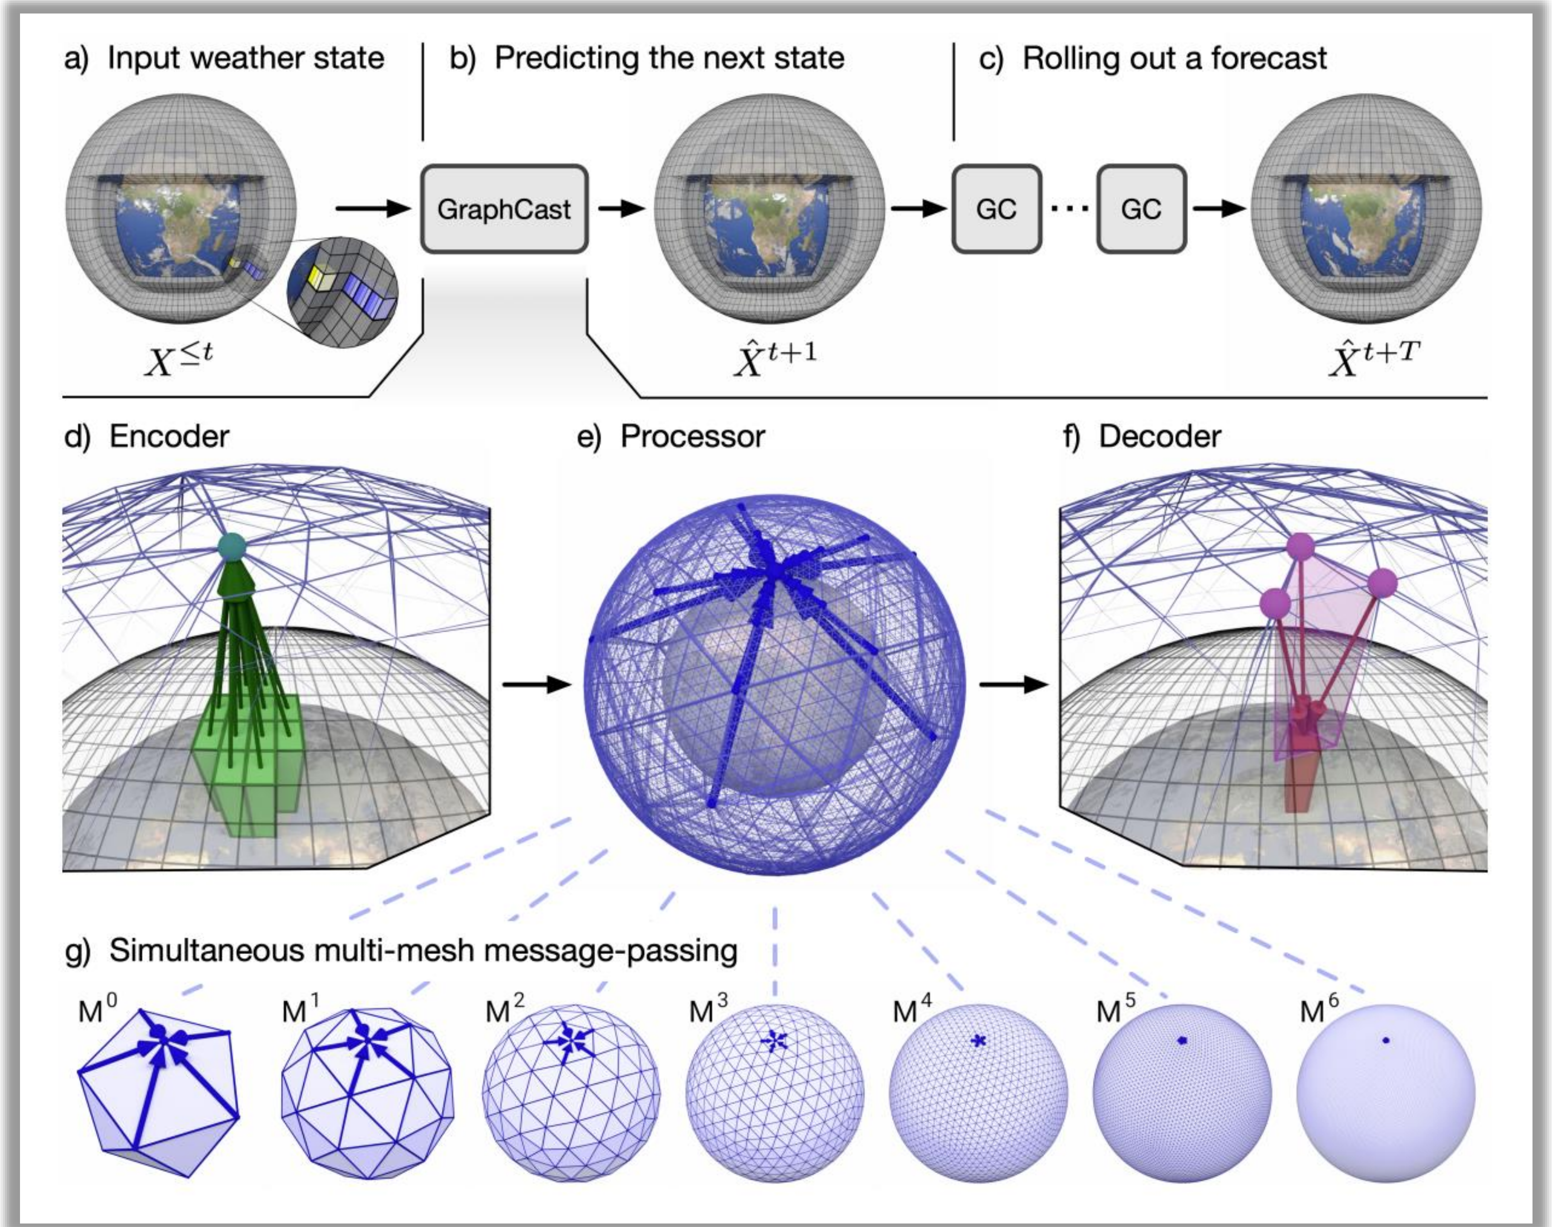

[A] R. Lam *et al.*, GraphCast: Learning Skillful Medium-range Global Weather Forecasting, in *arXiv preprint:2212.12794*, 2015.

Pangu-Weather: Overall Pipeline

- **On the path towards next-generation numerical weather prediction**
  - Given data, AI algorithms can also deal with assimilation and post-processing

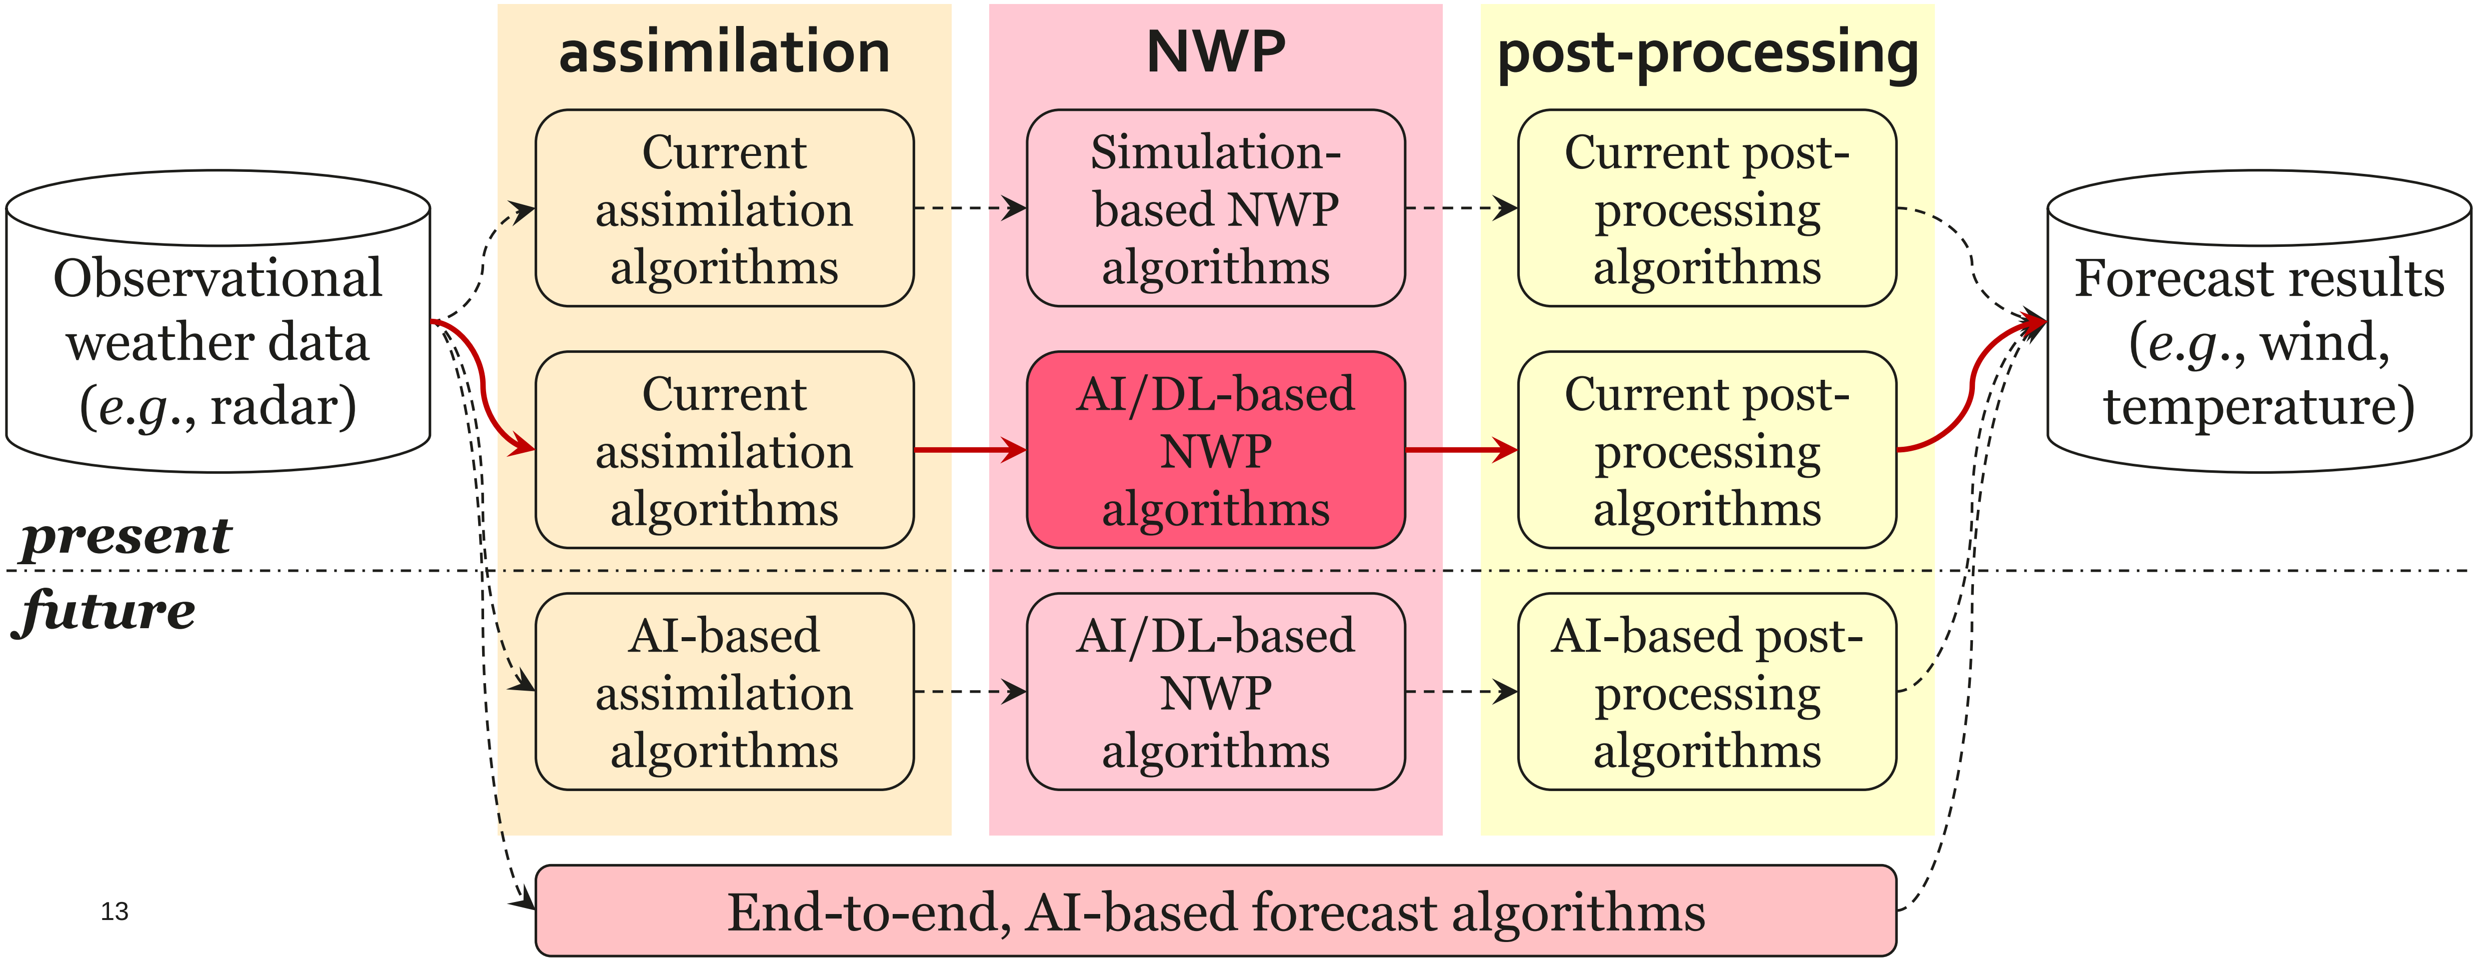

Data and Settings

- **Dataset details**
  - The 5th generation of ECMWF **reanalysis** (ERA5) data – publically available
  - Hourly reanalysis data from 1940 onwards
    - We used: the **1979–2017** data for training, the **2019** data for validation, the **2018, 2020, 2021** data for testing (to be fairly compared to WeatherBench[A])
  - A number of surface variables, plus upper-air variables at 37 pressure levels
    - We used: **four surface variables** (2m temperature, u- and v-components of 10m wind speed, mean sea-level pressure) with and **five upper-air variables** (geopotential, specific humidity, temperature, u- and v-components of wind speed) at **13 pressure levels** (50hPa, 100hPa, 150hPa, 200hPa, 250hPa, 300hPa, 350hPa, 400hPa, 500hPa, 600hPa, 700hPa, 850hPa, 925hPa, 1000hPa)
  - The full dataset is over **2000TB**, we used ~60TB of data

[A] S. Rasp *et al.*, WeatherBench: A Benchmark Data Set for Data-driven Weather Forecasting, in *JAMES*, 2020.

Input and Output for Deep Networks

- **Input and output data have the same shape**
  - More than 70 million variables: $1440 \times 720 \times (13 \times 5 + 4)$
  - Seems large, but the resolution ($0.25°$, or about 25km around the Tropics) is **barely enough** for some important events (*e.g.*, tropical cyclones)
  - Larger data in the future
    - ECMWF annouced the 6th generation of reanalysis data with **8 × resolution**
    - More pressure levels and/or weather factors to be studied
- **Difficulties for machine learning**
  - Complex relationship between input and output
  - Solution: given **a fixed lead time** Δt, train a deep network $f(\cdot)$ that receives weather data at time t and predicts weather data at time $t + \Delta t$

Architecture: 3D Earth-Specific Transformer

- **A 3D vision transformer to process volumetric data**
  - Swin transformer[A] to accelerate computation (standard window attentions)
  - Reduced network depth and width **(larger models can be better!)**

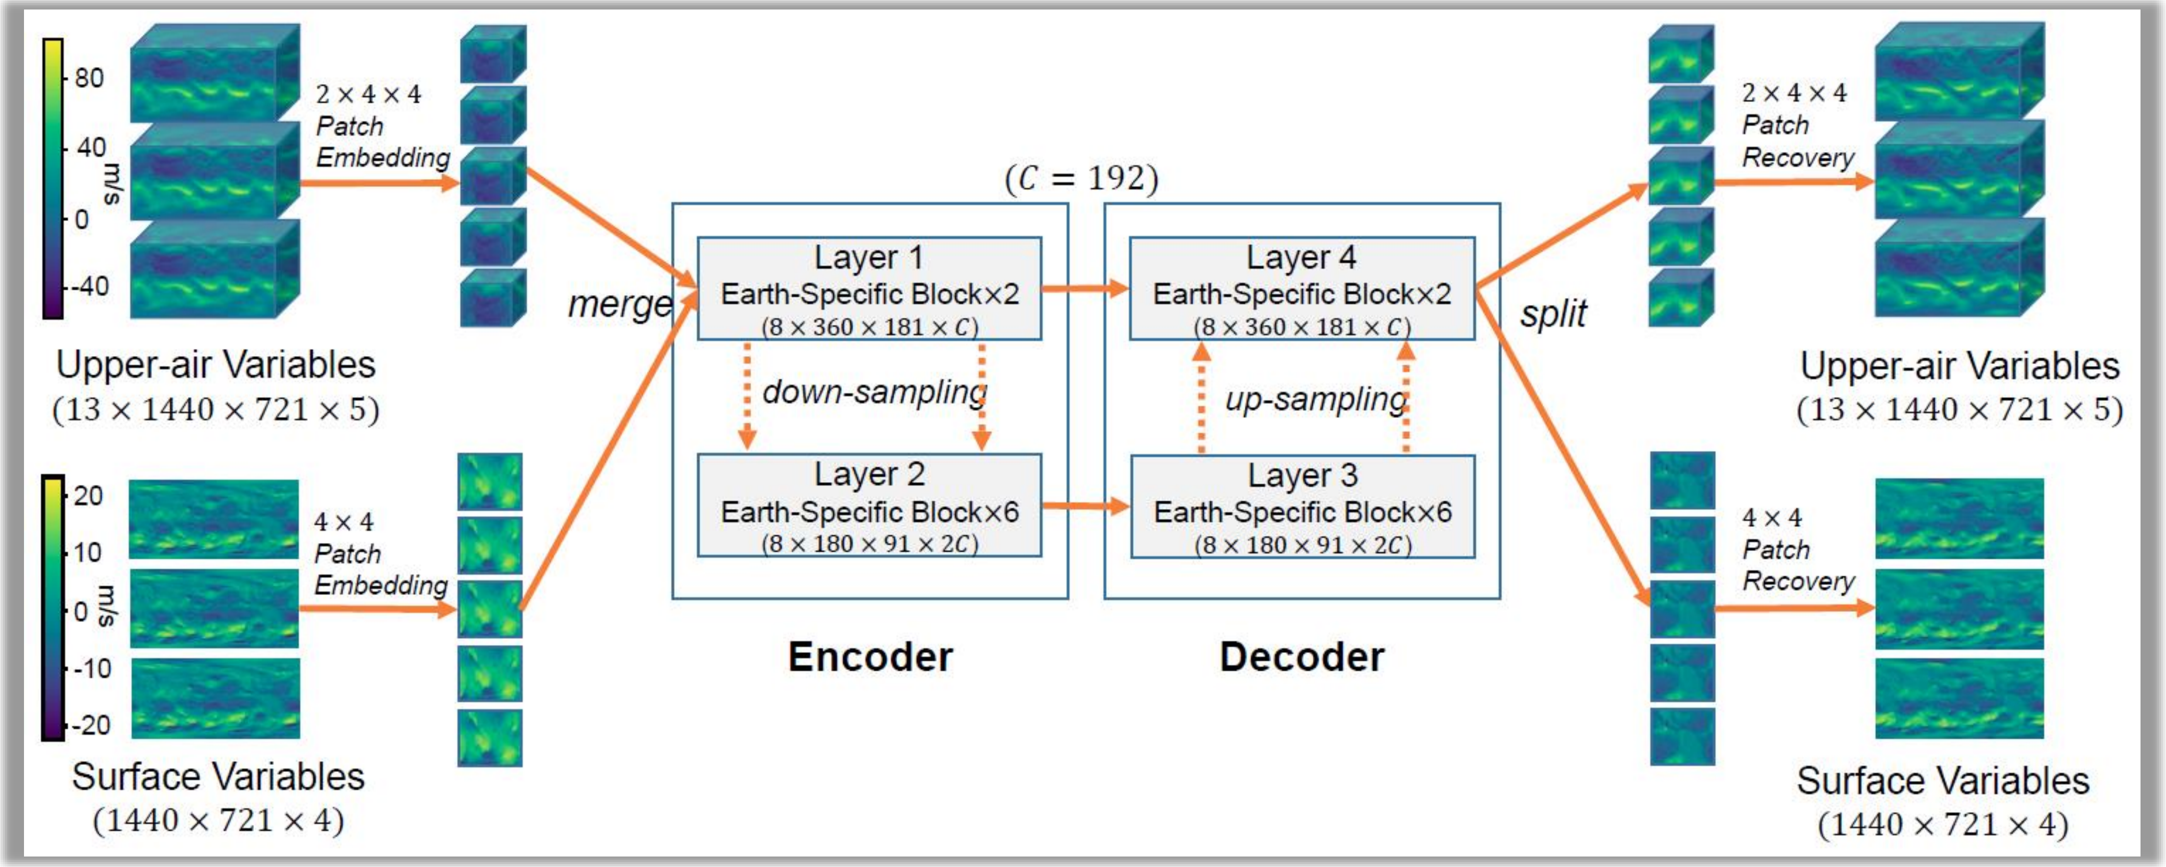

[A] Z. Liu *et al.*, Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, in *ICCV*, 2021.

Why 3D Deep Neural Networks?

- **3D networks can integrate richer information**
  - Each neuron in a 3D network is aware of the height information
  - The spacing and distribution of atmospheric states and the relationship between atmospheric patches change rapidly across pressure levels
  - Many weather processes (*e.g.*, radiation, convection, *etc.*) can only be completely formulated in the 3D space
- **3D networks are faster in inference**
  - No need to process each pressure level individually and perform combination

Injecting Earth-Specific Priors

- **Motivation: variables are closely related to the absolute coordinate**
- **Working mechanism: modifying the positional bias**
  - Replacing relative positional bias with absolute positional bias

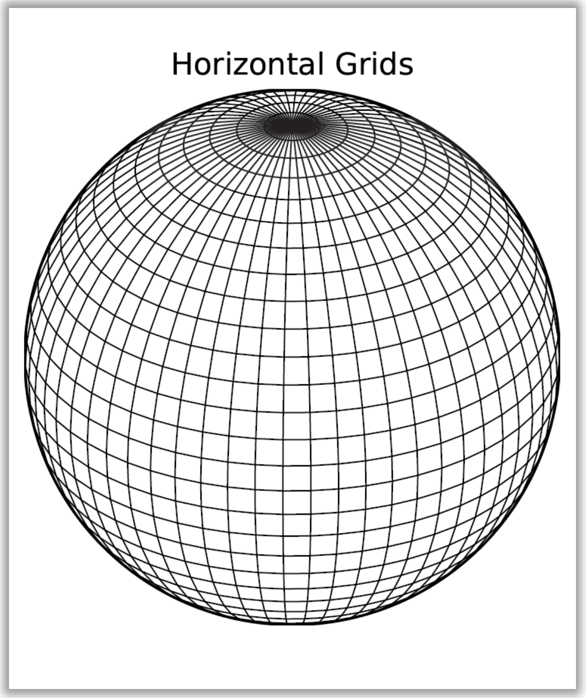

Uneven grid distribution on Earth's sphere

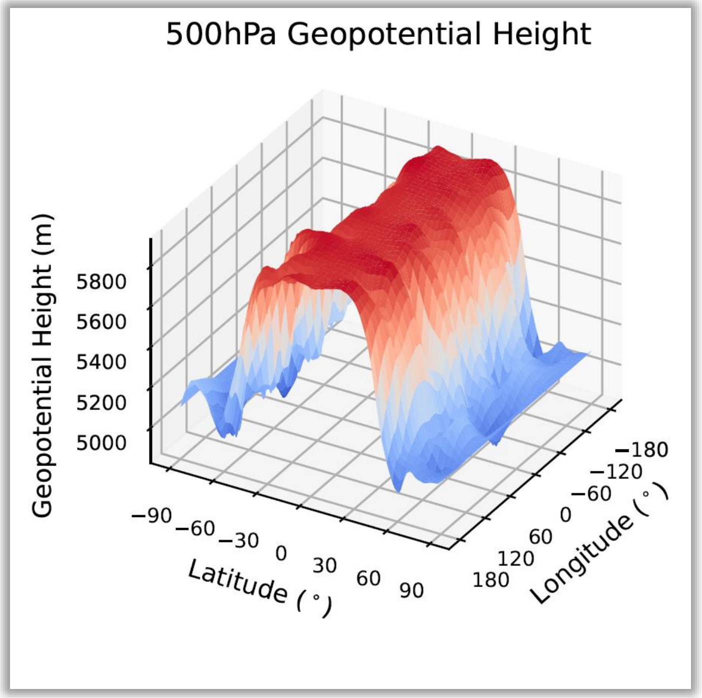

Geopotential height *w.r.t* latitude

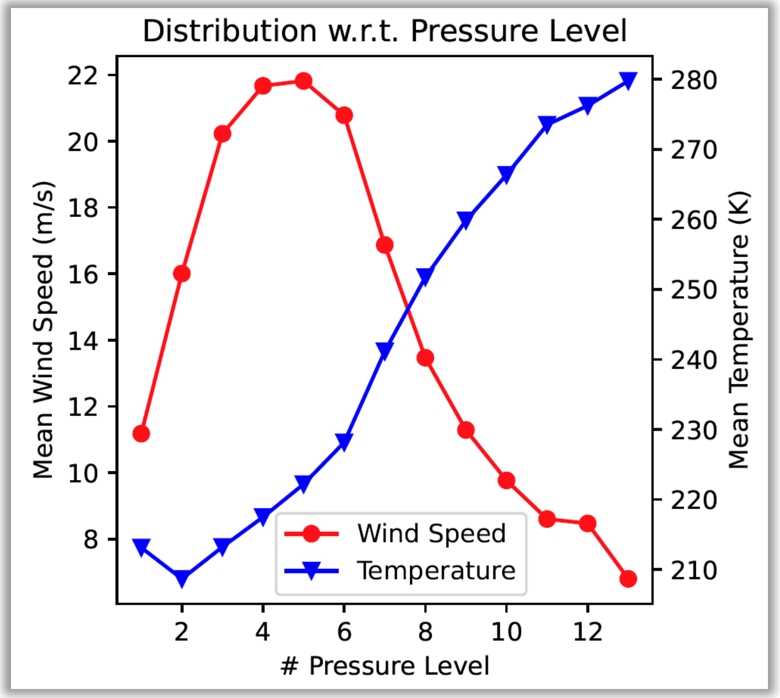

Wind speed and temperature *w.r.t* height (pressure level)

Hierarchical Temporal Aggregation

- Forecast error grows fast with iteration
  - The key is to reduce the number of iterations!
- Hierarchical temporal aggregation
  - Training **4 models** with lead times being 1 hour, 3 hours, 6 hours, 24 hours, respectively
  - A **greedy algorithm** to choose the model with max allowed lead time for the next iteration
  - Example: a 7-day forecast needs 7 iterations

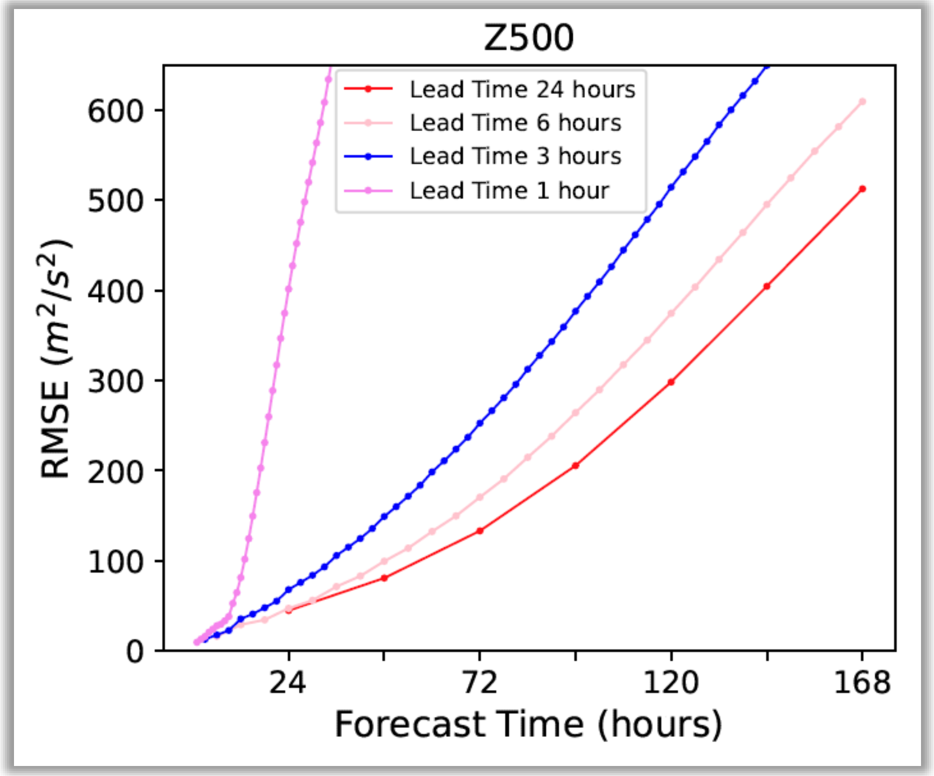

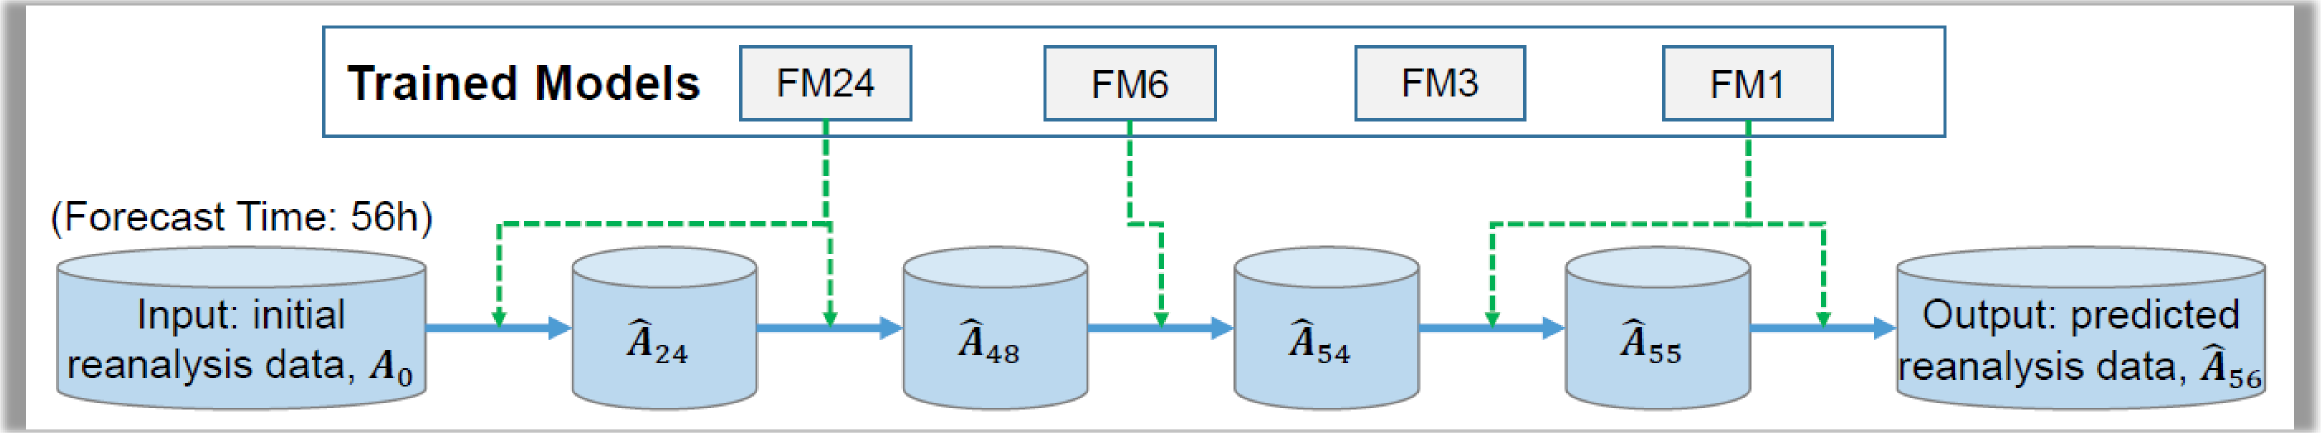

Computational Costs

- **Training**
  - Each forecast model has around **64 million** parameters
  - Each forecast model is trained for **100 epochs** (not converged yet), taking about **16 days** on **192 NVIDIA Tesla-V100 GPUs**
- **Inference**
  - Each forecast takes around **1.4 seconds** on **a single V100 GPU**
    - The inference can also be executed on CPU, taking longer time
  - Example: performing 7-day global forecasting requires executing the 24-hour model 7 times, requiring **less than 10 seconds** in total
  - With faster inference, ensemble forecast is made easier (see later slides)
- **Trained models were released[A] for research use**

[A] https://github.com/198808xc/Pangu-Weather

Deterministic Forecast: by Variables

- **First AI algorithm to surpass operational IFS**
  - For Z500, 3-day and 5-day RMSEs (unit: m^2/s^2) are 152.8 and 333.7 for operational IFS, and 134.5 (12% smaller) and 296.7 (11% smaller) for Pangu
  - For T850, 3-day and 5-day RMSEs (unit: K) are 1.37 and 2.06 for operational IFS, and 1.14 (17% smaller) and 1.79 (13% smaller) for Pangu

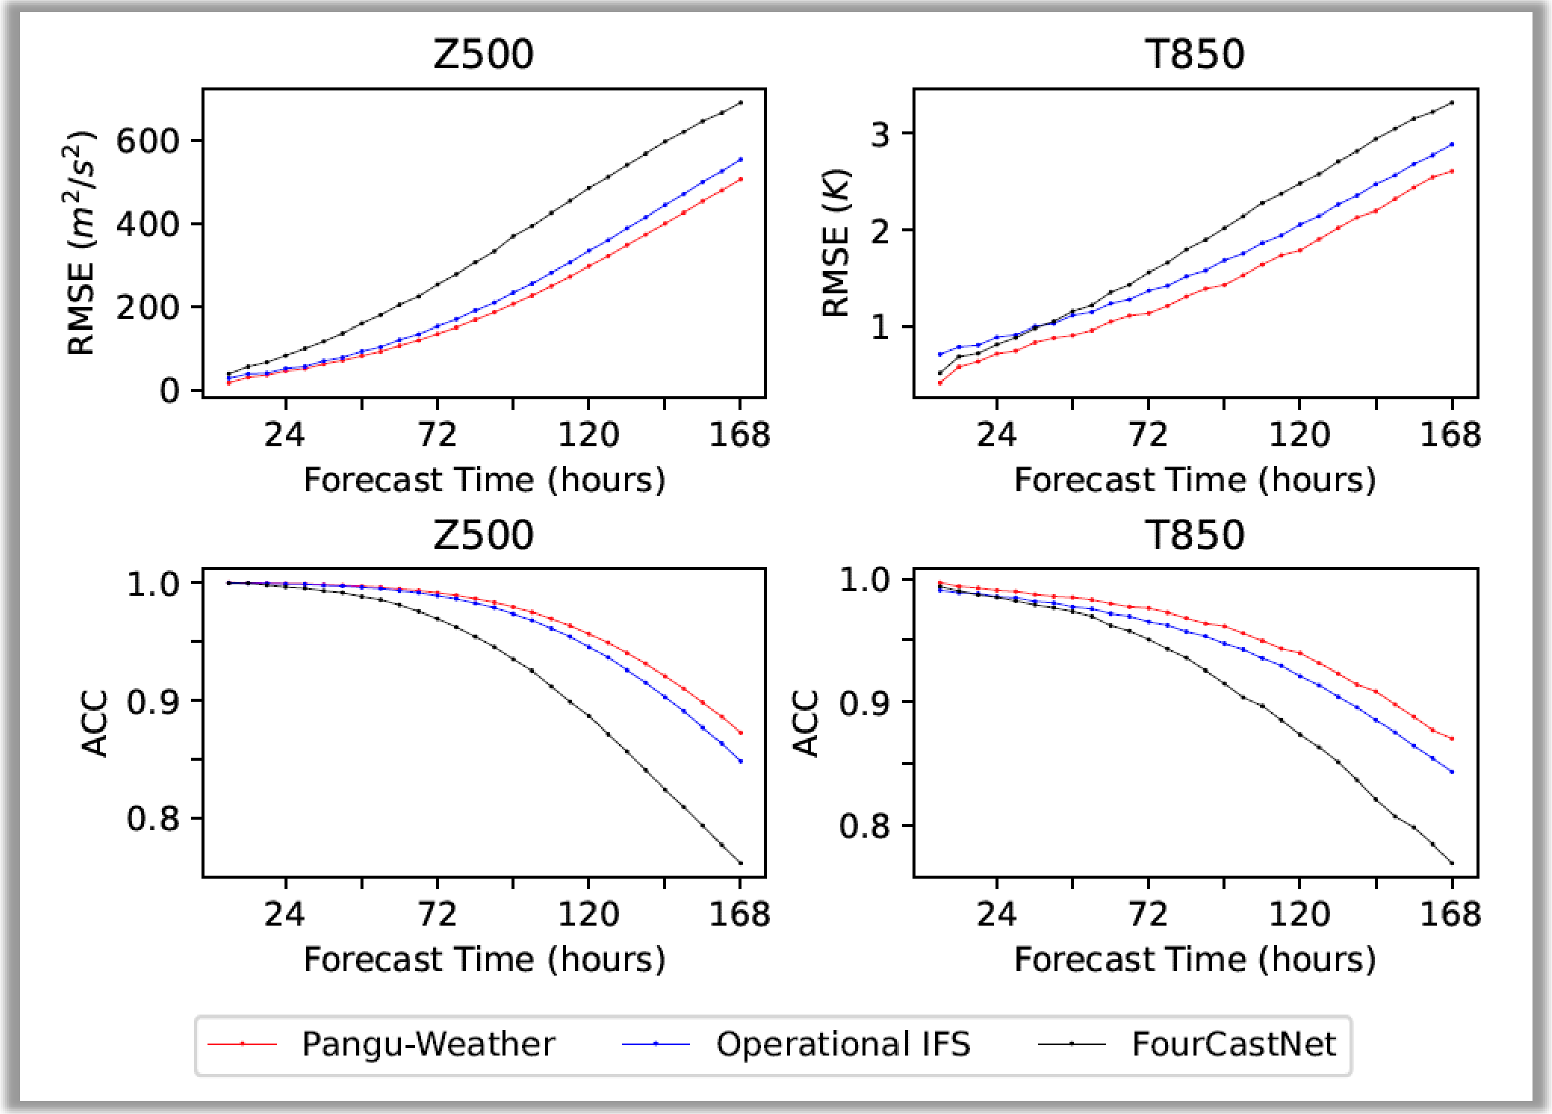

Deterministic Forecast: by Variables (cont.)

- **First AI algorithm to surpass operational IFS**
  - For T2M, 3-day and 5-day RMSEs (unit: K) are 1.34 and 1.75 for operational IFS, and 1.05 (22% smaller) and 1.53 (13% smaller) for Pangu
  - For U10, 3-day and 5-day RMSEs (unit: m/s) are 1.94 and 2.90 for operational IFS, and 1.61 (17% smaller) and 2.53 (13% smaller) for Pangu

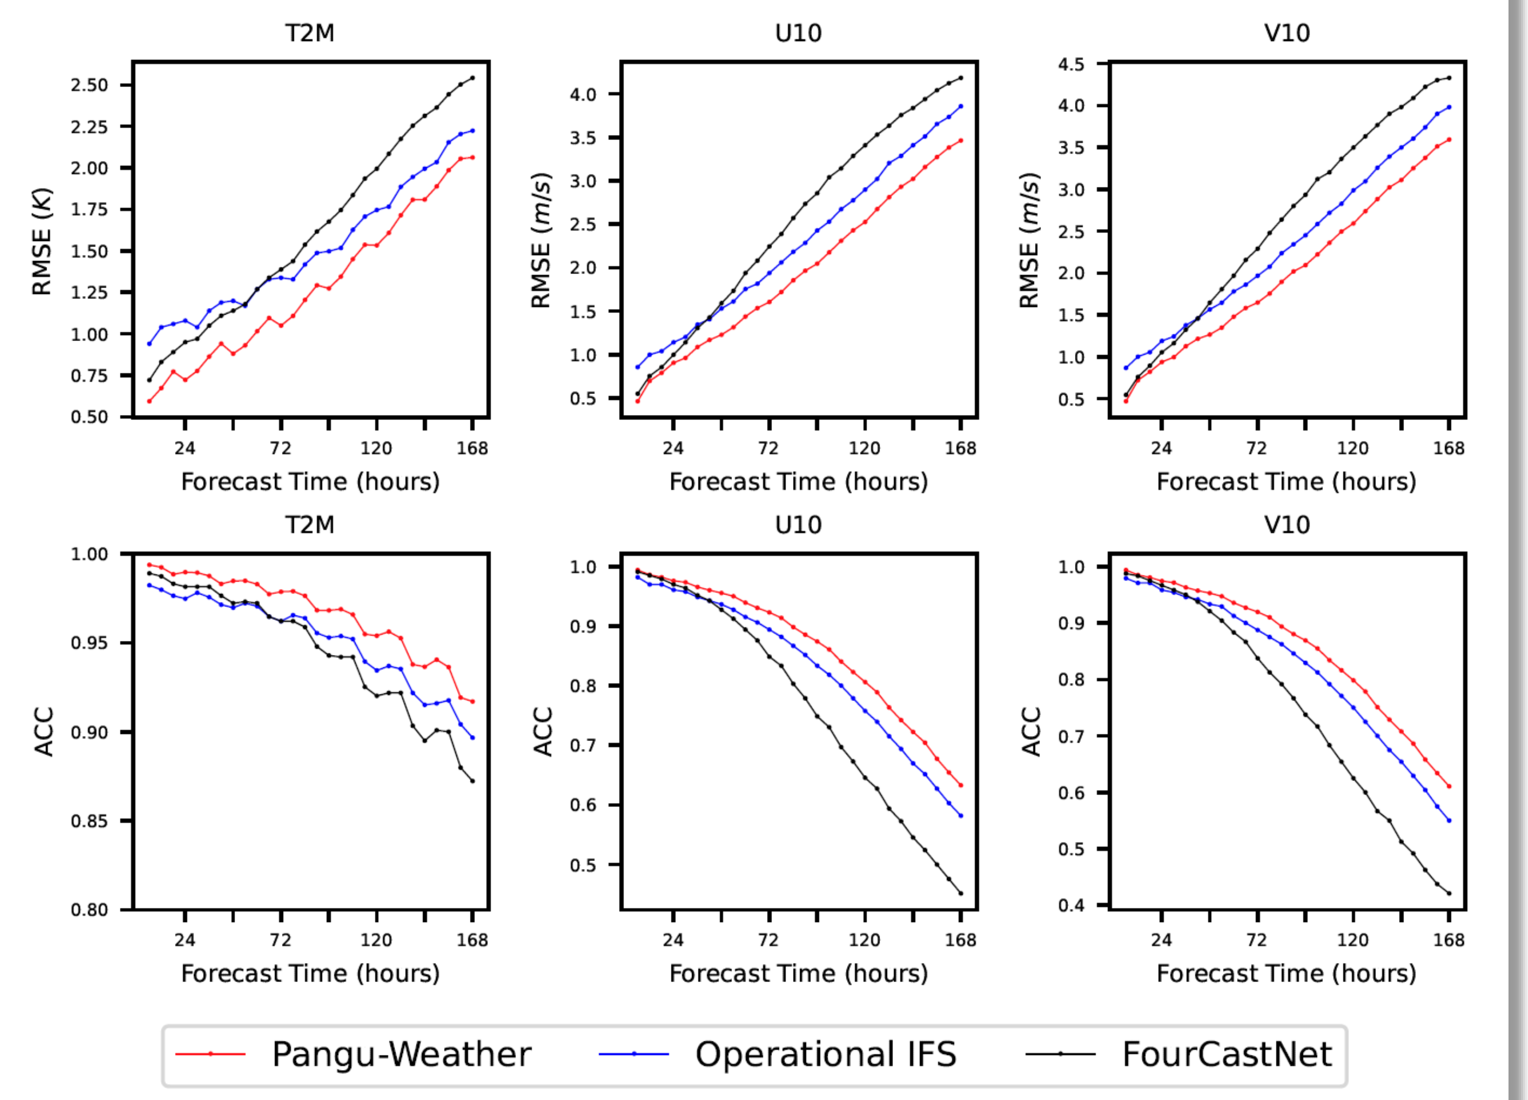

Deterministic Forecast: by Variables (cont.)

- **Significant advantage: "forecast time gain"**
  - If Pangu's forecast error at 7 days (168 hours) is equivalent to another method's forecast error at $168 - \Delta t$ hours, then Δt is called the "forecast time gain" of Pangu over the specified method

Variable	Gain over operational IFS (h)	Gain over FourCastNet (h)
Z500	10.45	43.23
T850	15.37	41.05
T2M	18.19	43.11
U10	19.68	43.81
V10	19.10	42.78
Z850	10.62	N/A
T850	13.66	N/A
Q500	31.00	N/A
U500	17.52	N/A
V500	16.16	N/A

Deterministic Forecast: by Regions

- **In the Northern Hemisphere**
  - Latitude between +20° (exclusive) and +90° (inclusive)

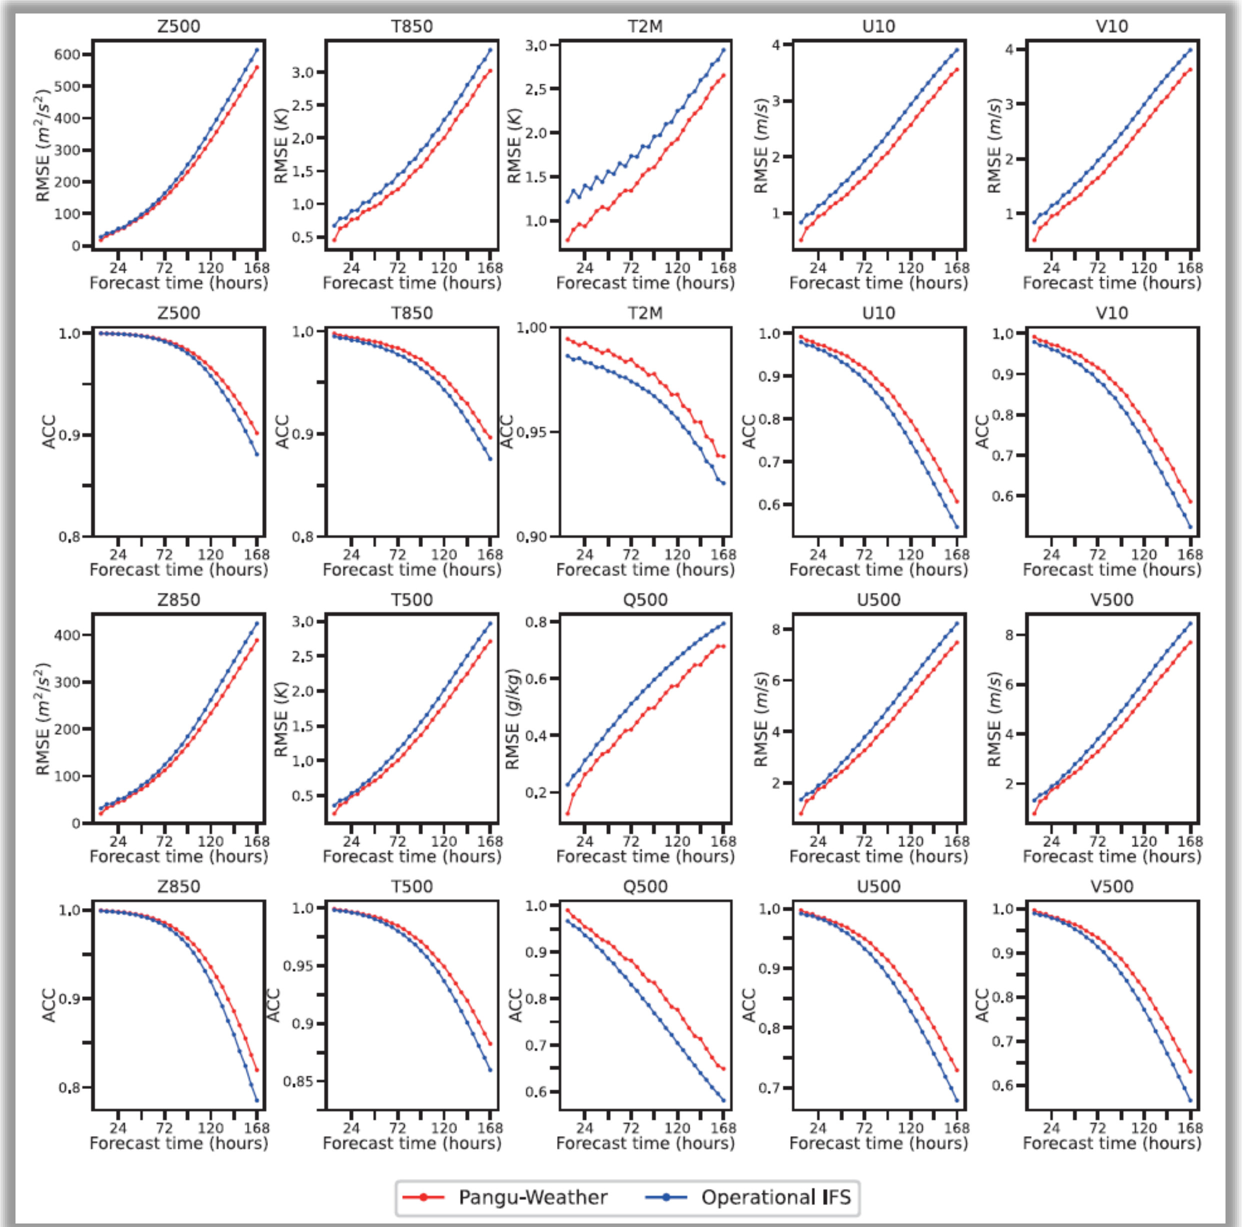

Deterministic Forecast: by Regions (cont.)

- **In the Southern Hemisphere**
  - Latitude between −20° (exclusive) and −90° (inclusive)

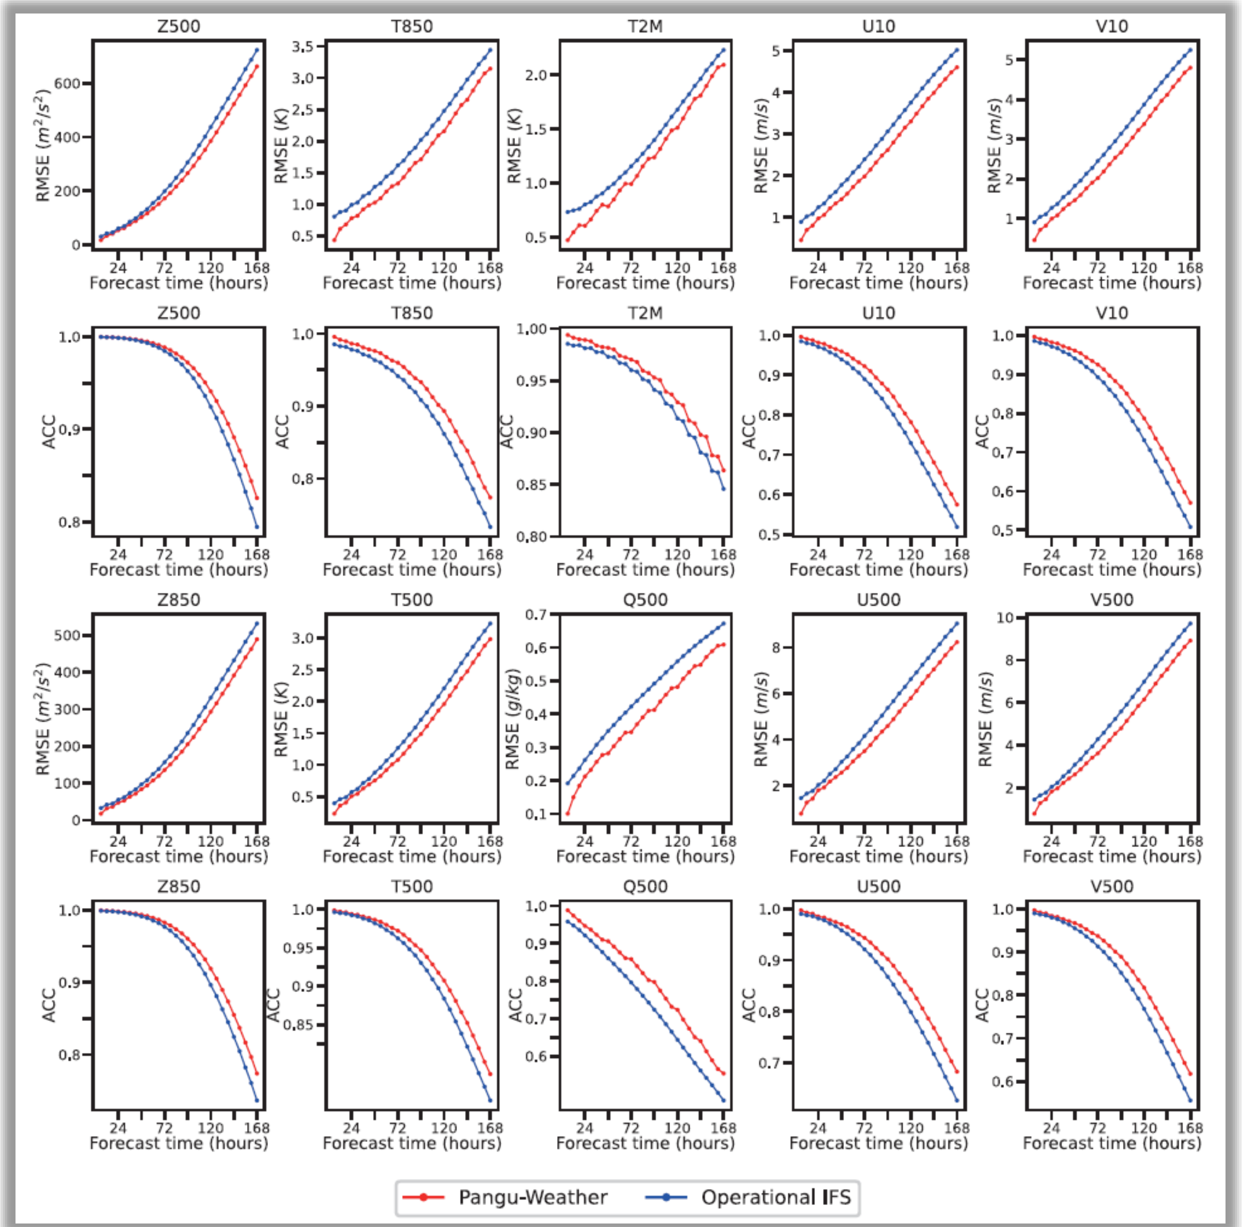

Deterministic Forecast: by Regions (cont.)

- **In the Tropics**
  - Latitude between +20° (inclusive) and −20° (inclusive)
  - Larger accuracy gain compared to the results in the Northern/Southern Hemispheres, corresponding to the significant advantages in tracking tropical cyclones (see later)

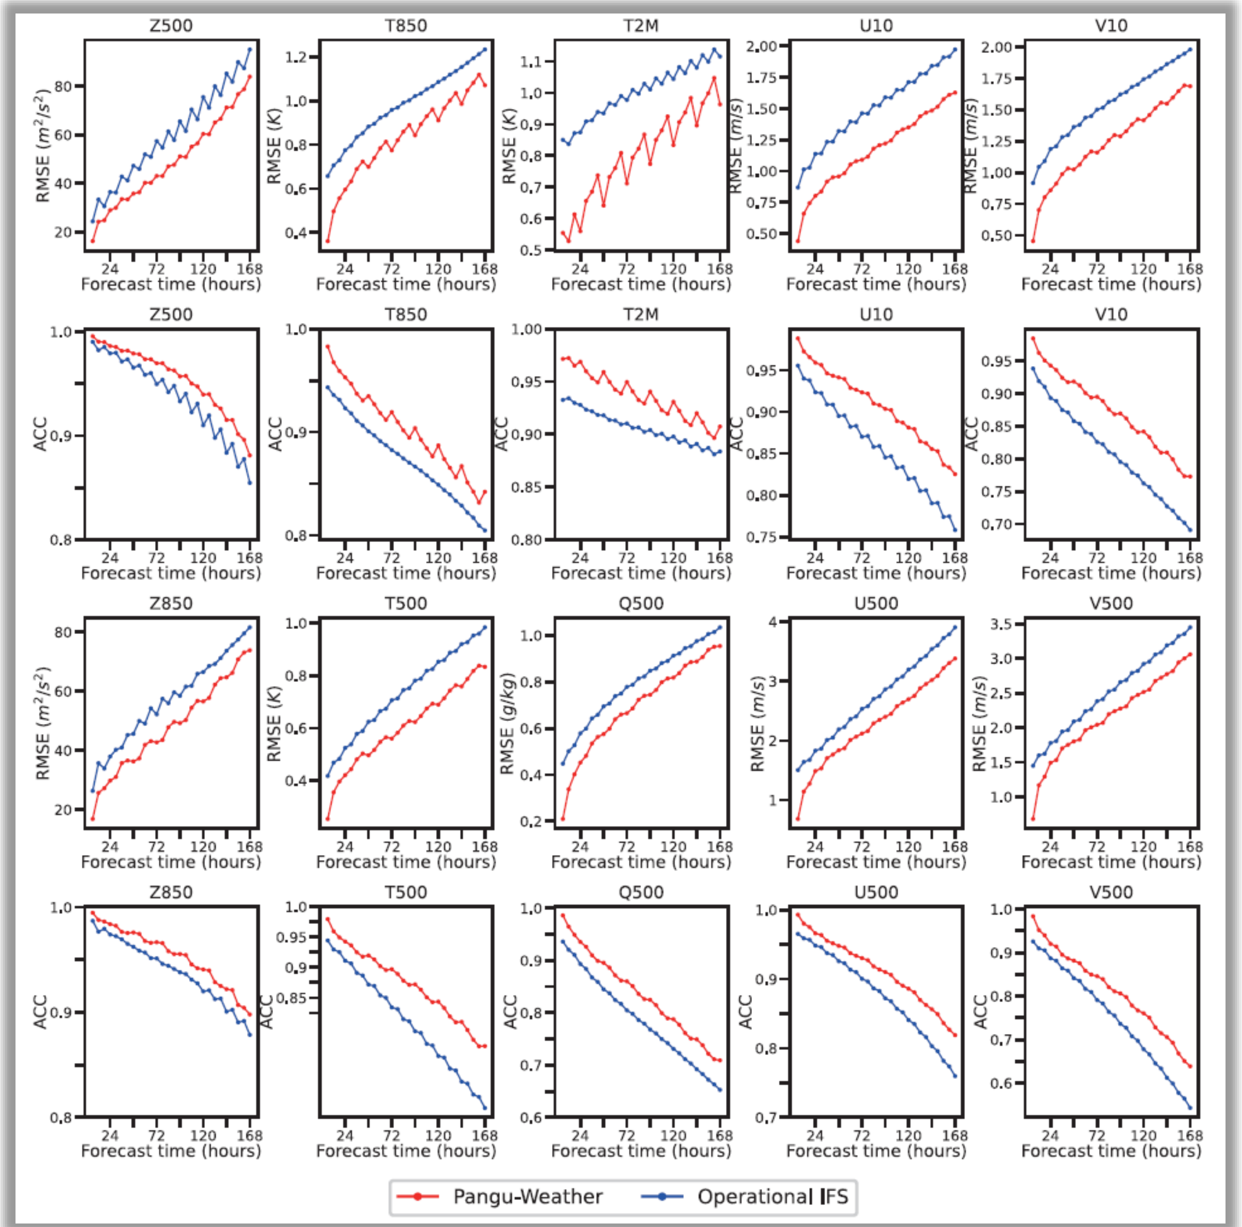

Deterministic Forecast: by Year

- **Comparison: 2018, 2020, 2021**
  - Consistent trends of RMSE and ACC in different years, indicating Pangu's stable forecast skill

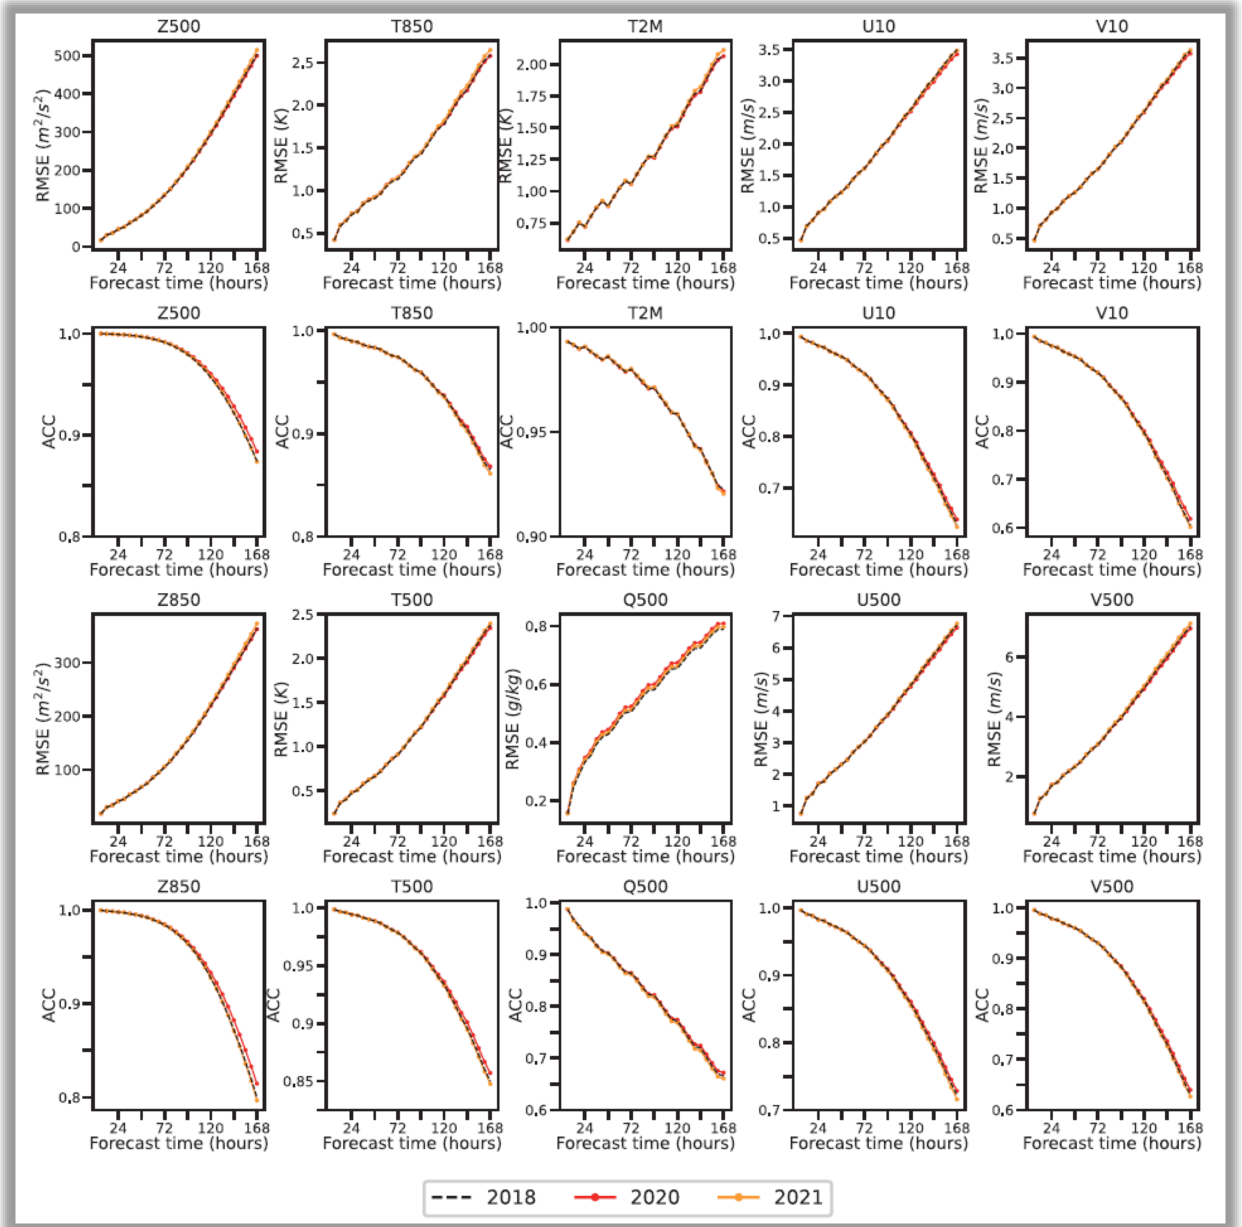

Deterministic Forecast: Visualization

- **Sufficiently close to ground-truth, with visible differences**
  - Pangu tends to produce smoother results (a typical behavior of AI algorithms)

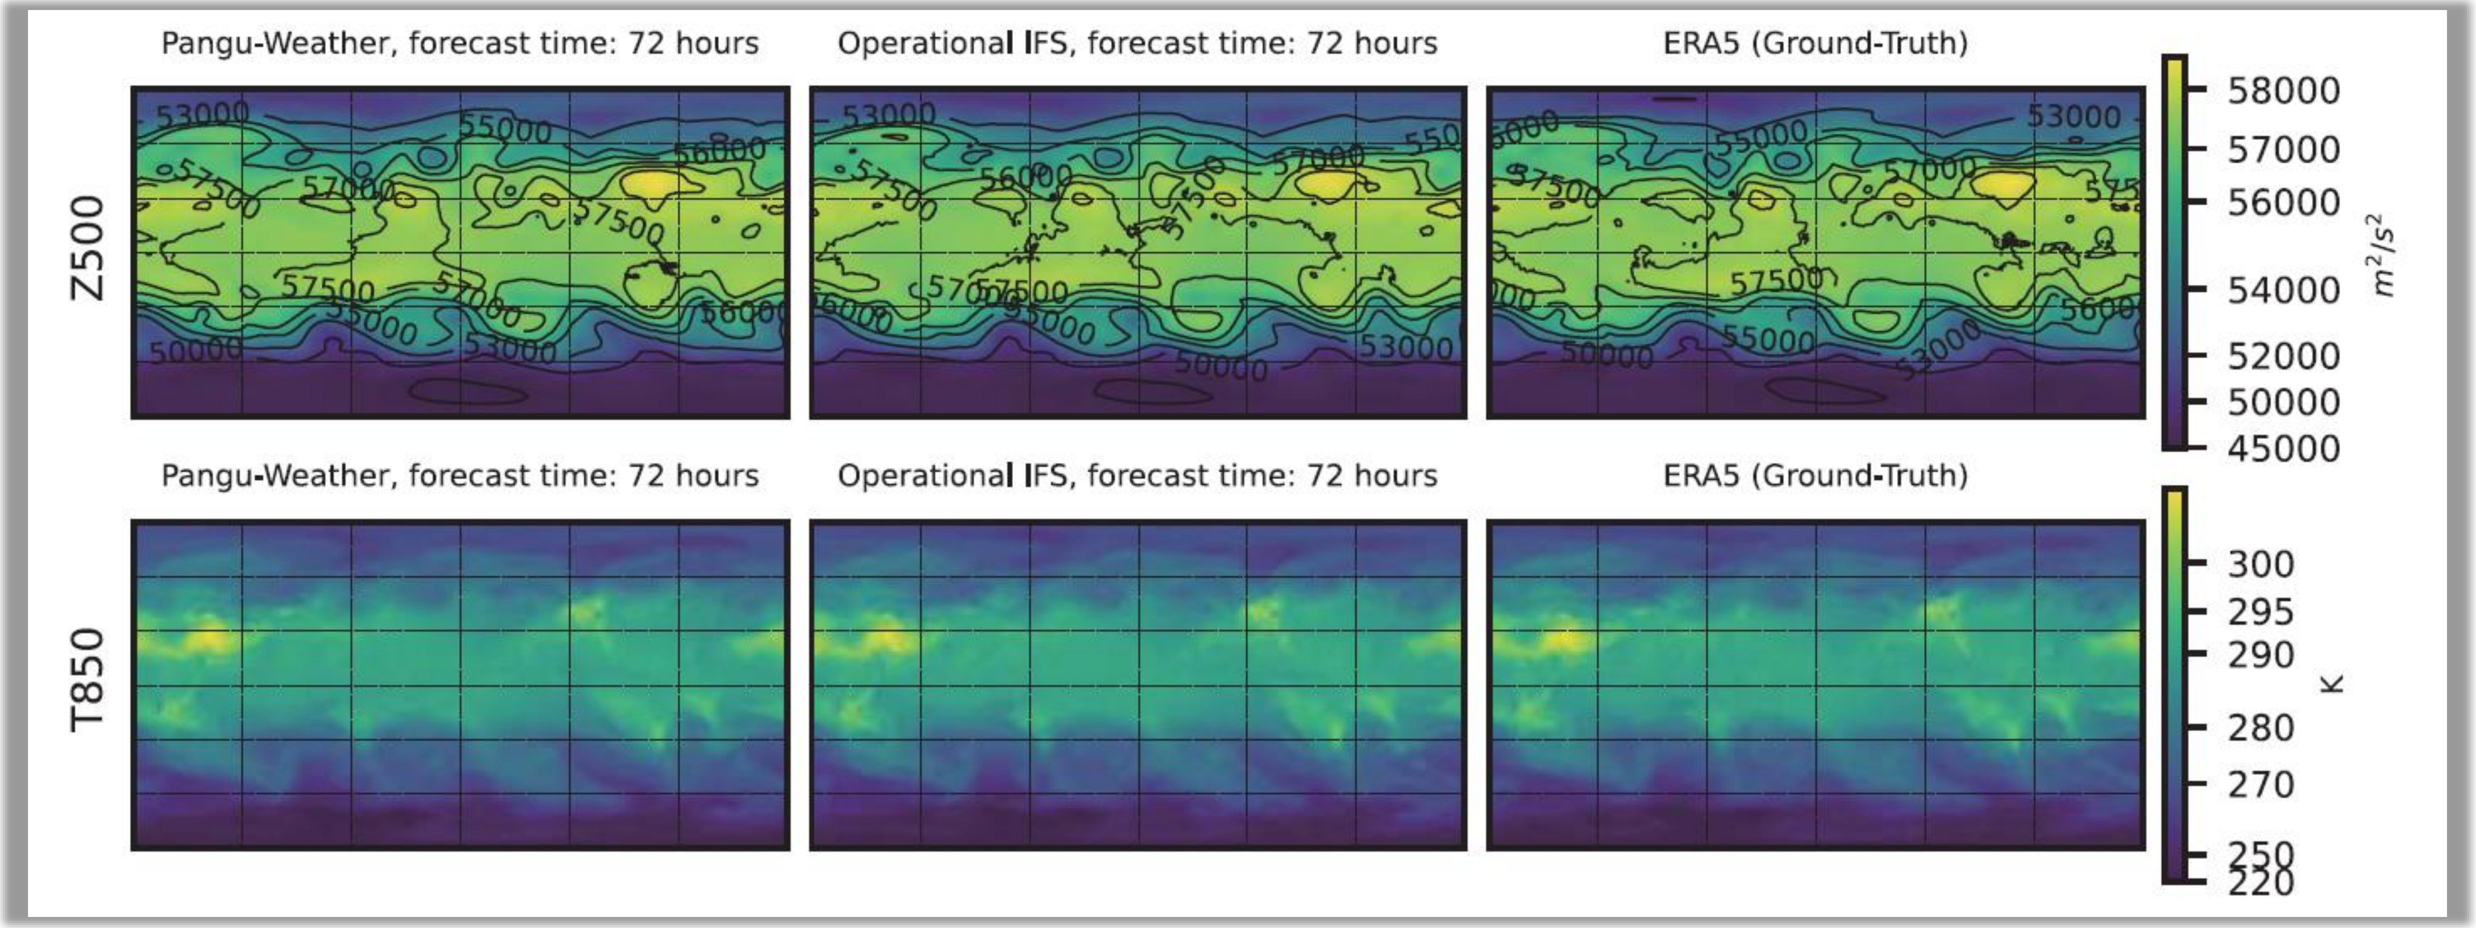

Visualization of 3-day forecast of two upper-air variables (Z500 and T850)

Deterministic Forecast: Visualization

- **Sufficiently close to ground-truth, with visible differences**
  - Pangu tends to produce smoother results (a typical behavior of AI algorithms)

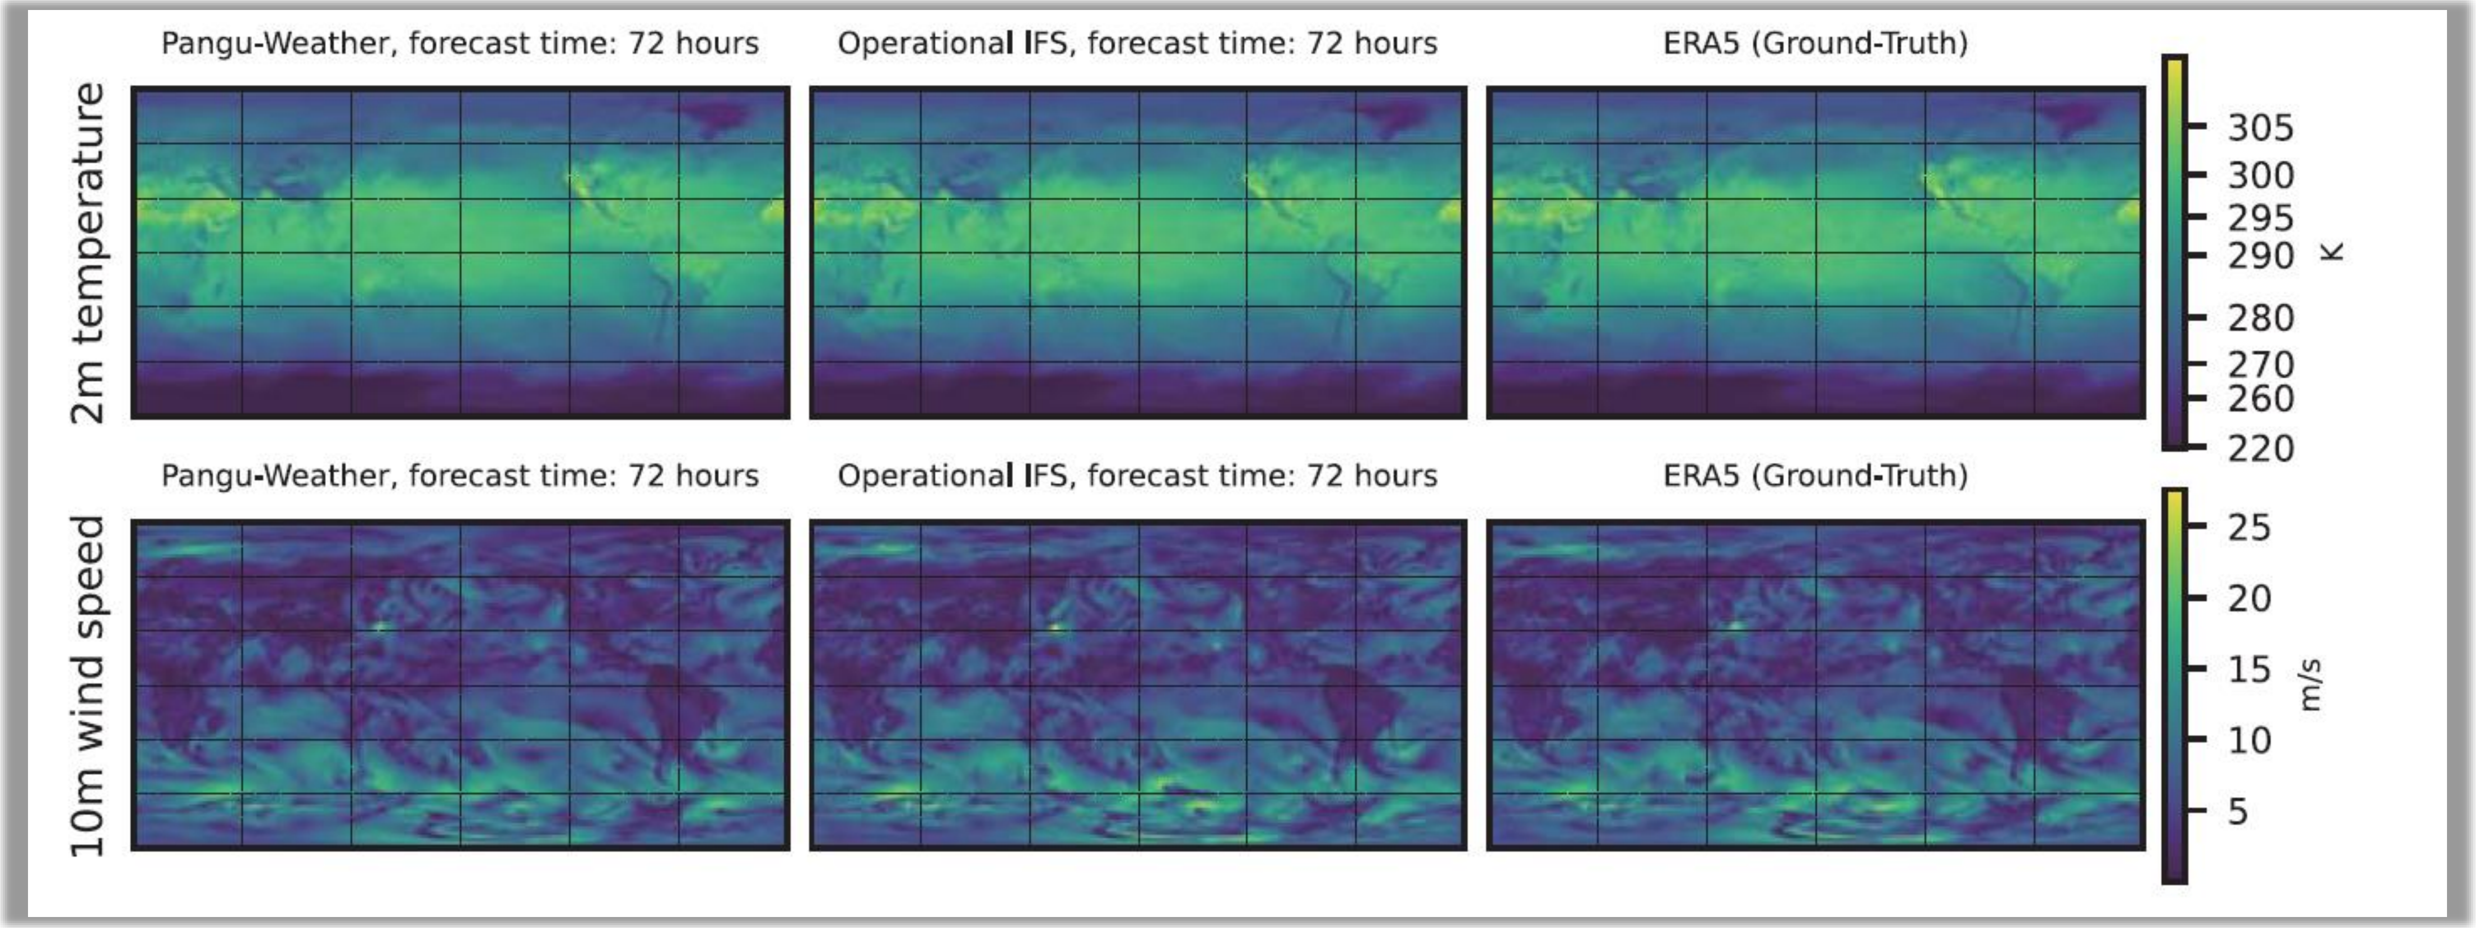

Visualization of 3-day forecast of two surface variables (T2m and U10)

Ensemble Forecast Results

- **Improved medium-range forecast results**
  - In 7-day forecast, the RMSE of Z500 is reduced from 500.3 to 450.6 (10% smaller), and that of U10 reduced from 3.48 to 2.96 (15% smaller)
  - Short-range (*e.g.*, 2-day) forecast results are not improved or even deteriorated

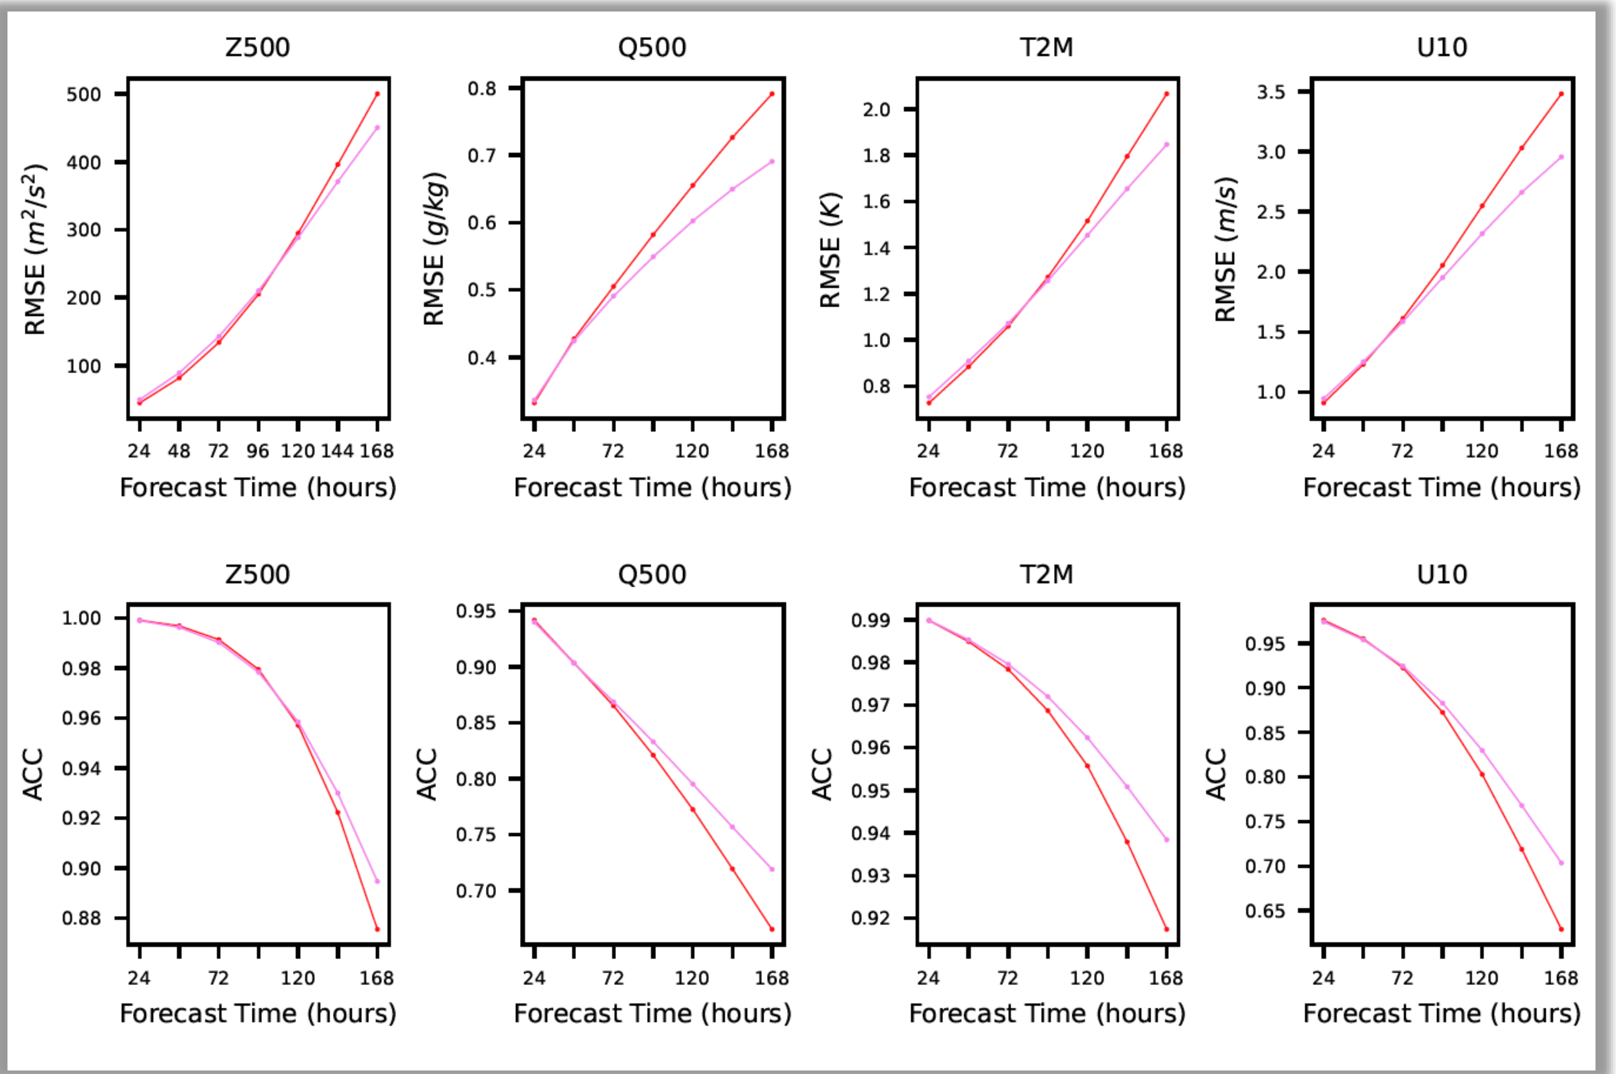

Results of 100-member ensemble forecast

Extreme Weather Forecast

- **Trend of relative quantiles with respect to lead time**
  - Pangu often reports lower quantile values because AI algorithms tend to produce smooth forecasts

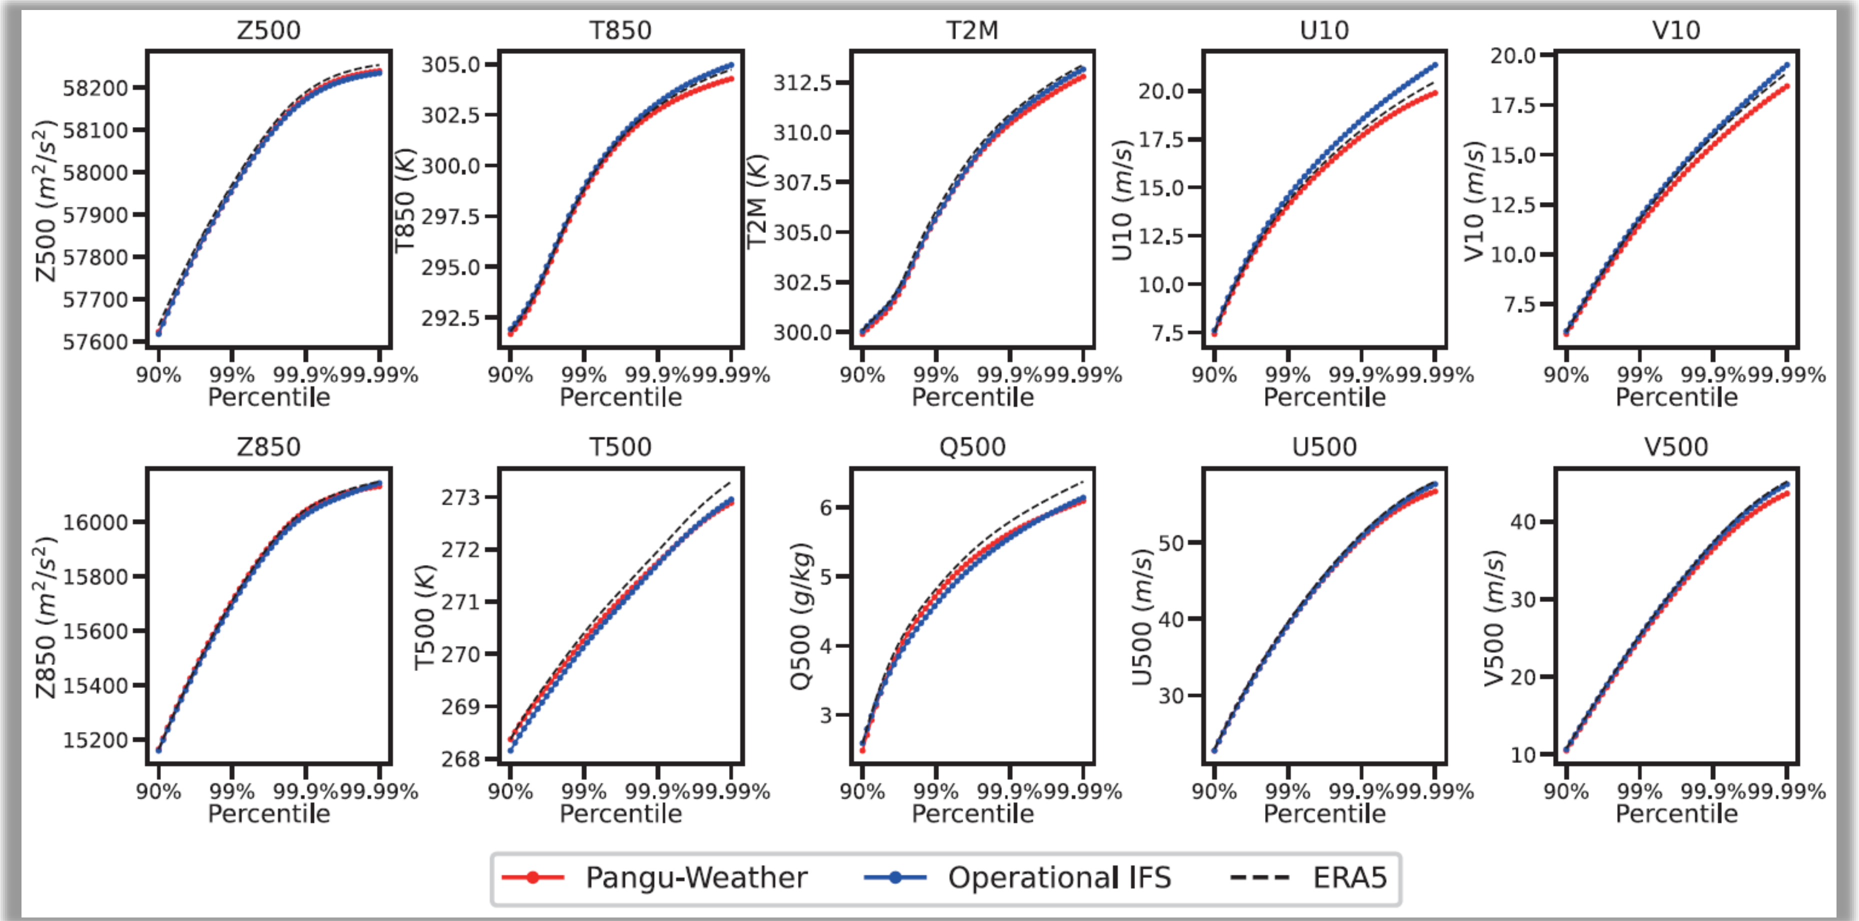

Extreme Weather Forecast (cont.)

- **Comparison between Pangu, FourCastNet, and operational IFS**
  - AI algorithms (Pangu and FourCastNet) tend to underestimate extremes
  - For Q500, Pangu is better than IFS due to the much better deterministic results

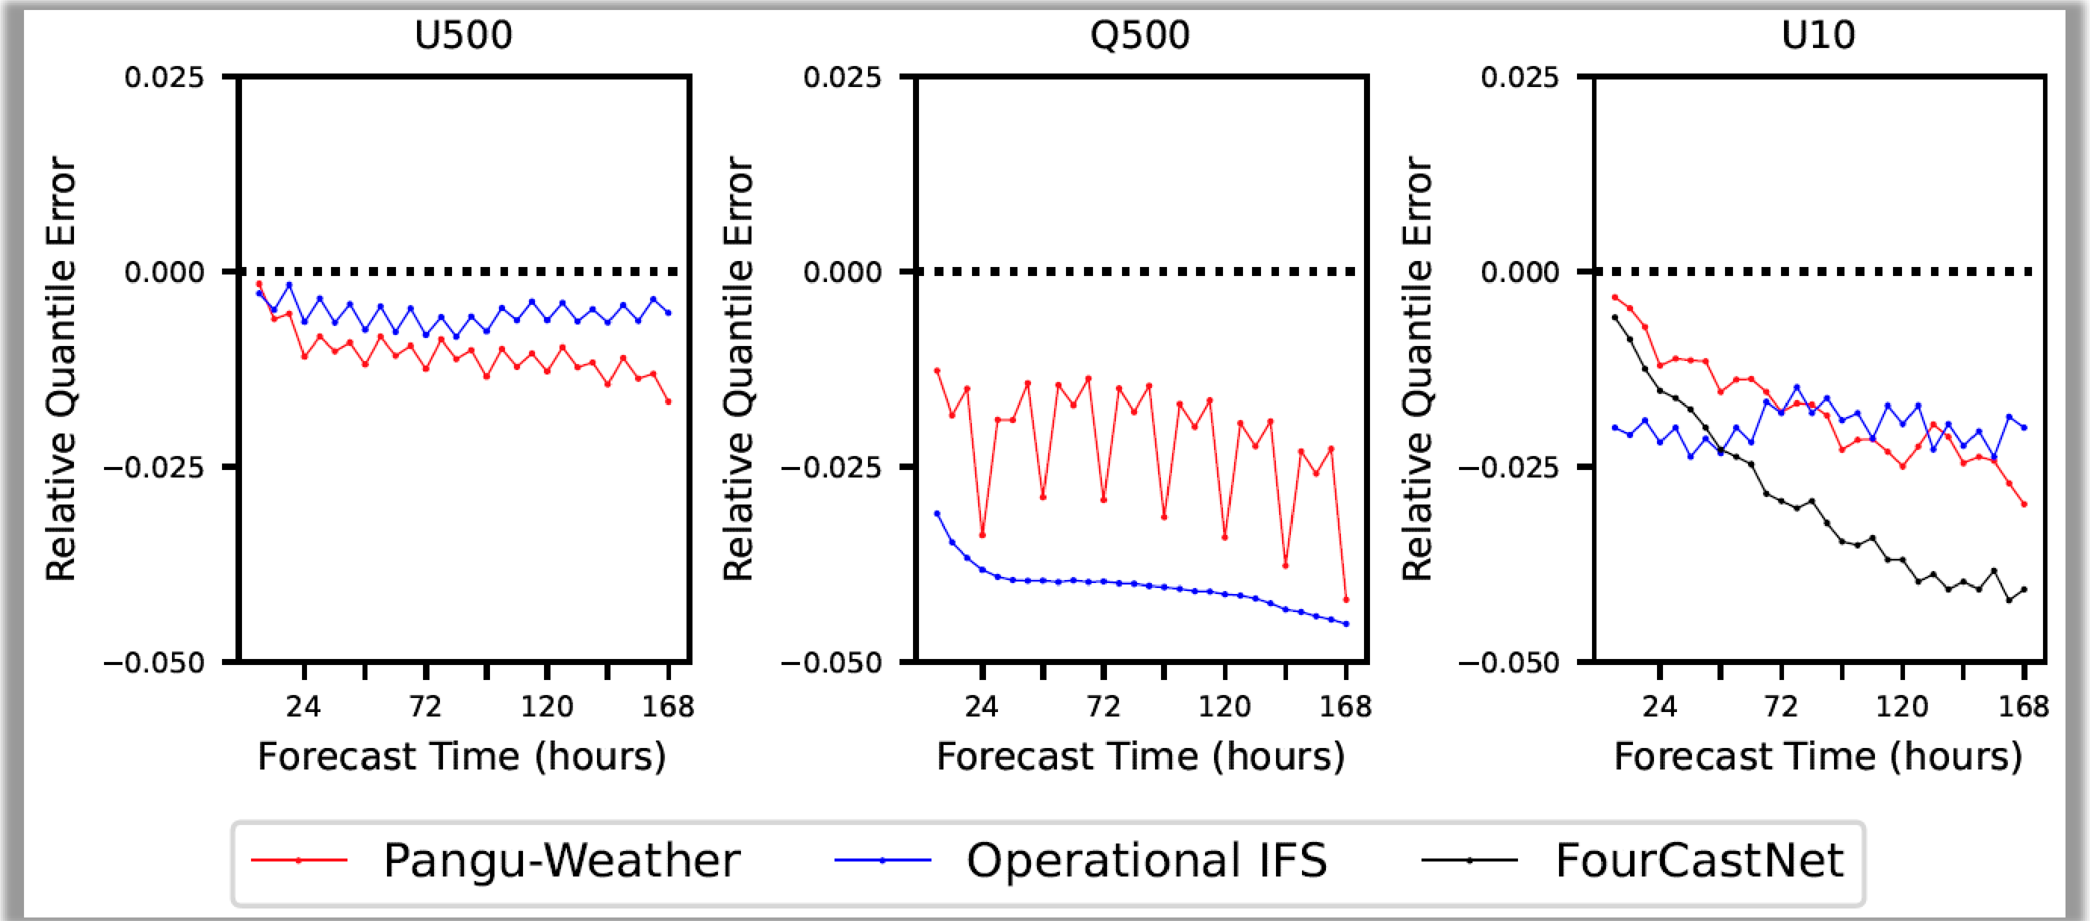

Tracking Tropical Cyclones

- **Tracking algorithm[A] and illustration**
  - Initial time and position are given
  - Given lead time (6 hours), finding a local minimum of mean sea-level pressure (MSLP) within 445km, satisfying:
    - There is a maximum of 850hPa vorticity within a radius of 278km, with absolute value $> 5 \times 10^{-5}$
    - There is a maximum of thickness between 850hPa and 200hPa within a radius of 278km
    - The maximum 10m wind speed is larger than 8m/s within a radius of 278km when the cyclone is on land

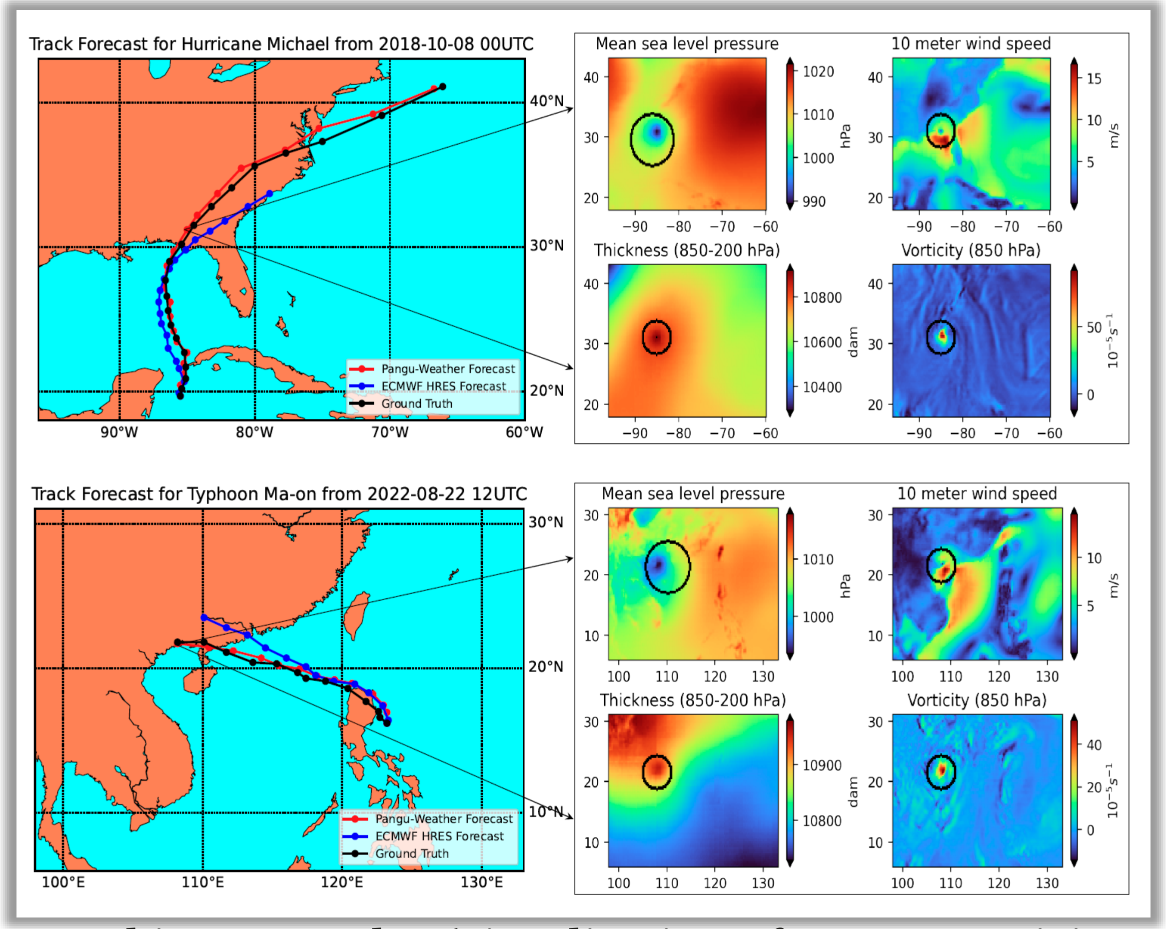

Tracking examples (visualization of MSLP, vorticity, thickness, 10m wind speed) for Hurricane Michael (2018-13) and Typhoon Ma-on (2022-09)

[A] P. White, Newsletter No. 102 – Winter 2004/05, by *ECMWF*, 2005.

Tracking Tropical Cyclones (cont.)

- **Overall better results over 88 cyclones, 2018**
  - Advantage gets larger with lead time
  - More accurate by regions and by intensity levels

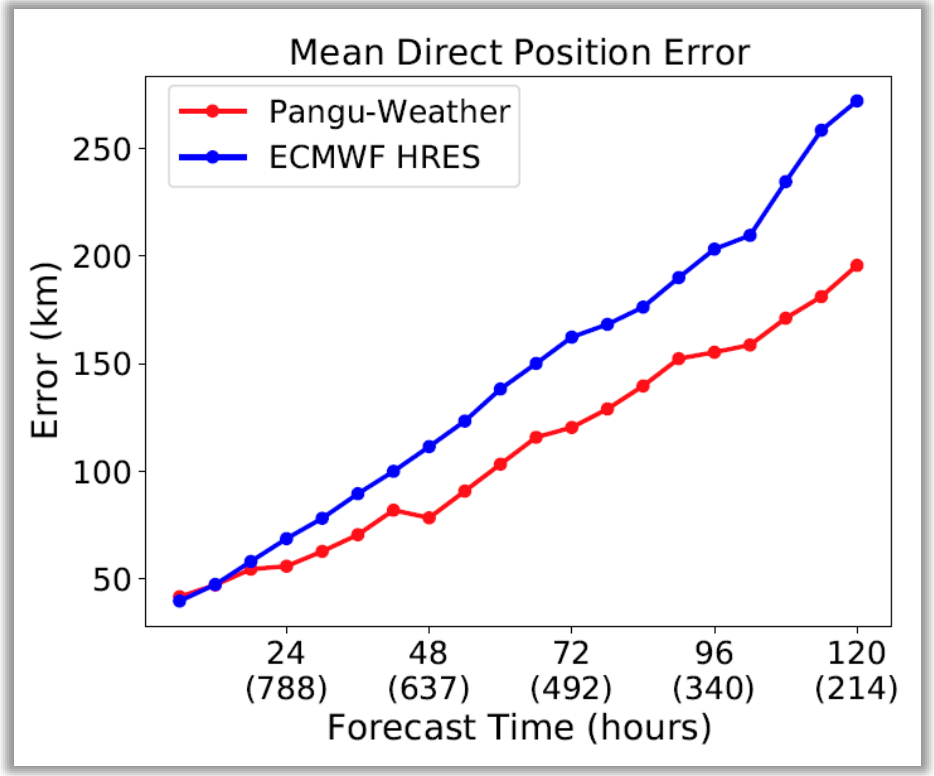

Overall errors

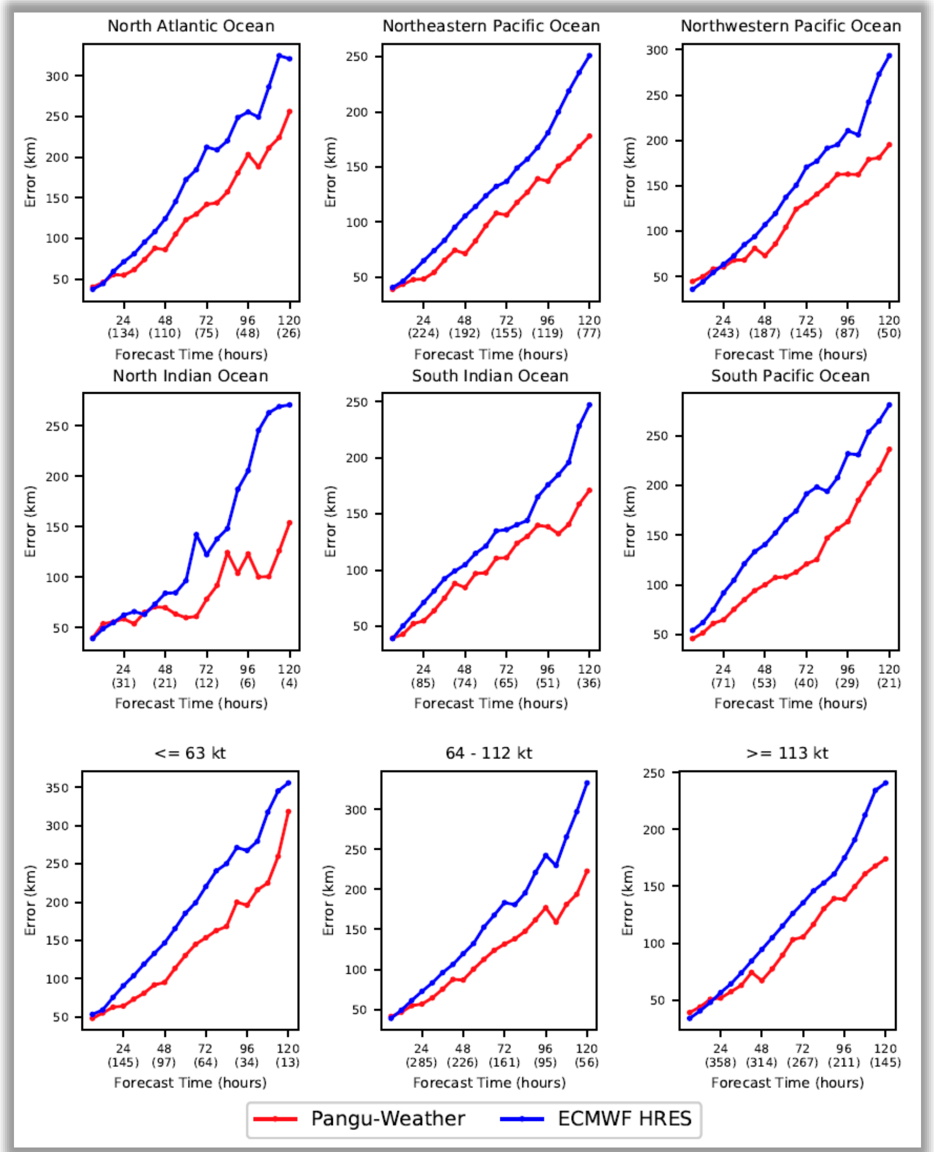

Errors in different subsets

Tracking Tropical Cyclones (cont.)

- **Example: Typhoon Kong-rey (2018-25) and Yutu (2018-26)**
  - Strongest typhoons in Western Pacific Ocean, 2018
  - Much more accurate forecasts compared to ECMWF-HRES
    - Potential reason: large advantage of deterministic forecasts in the Tropics

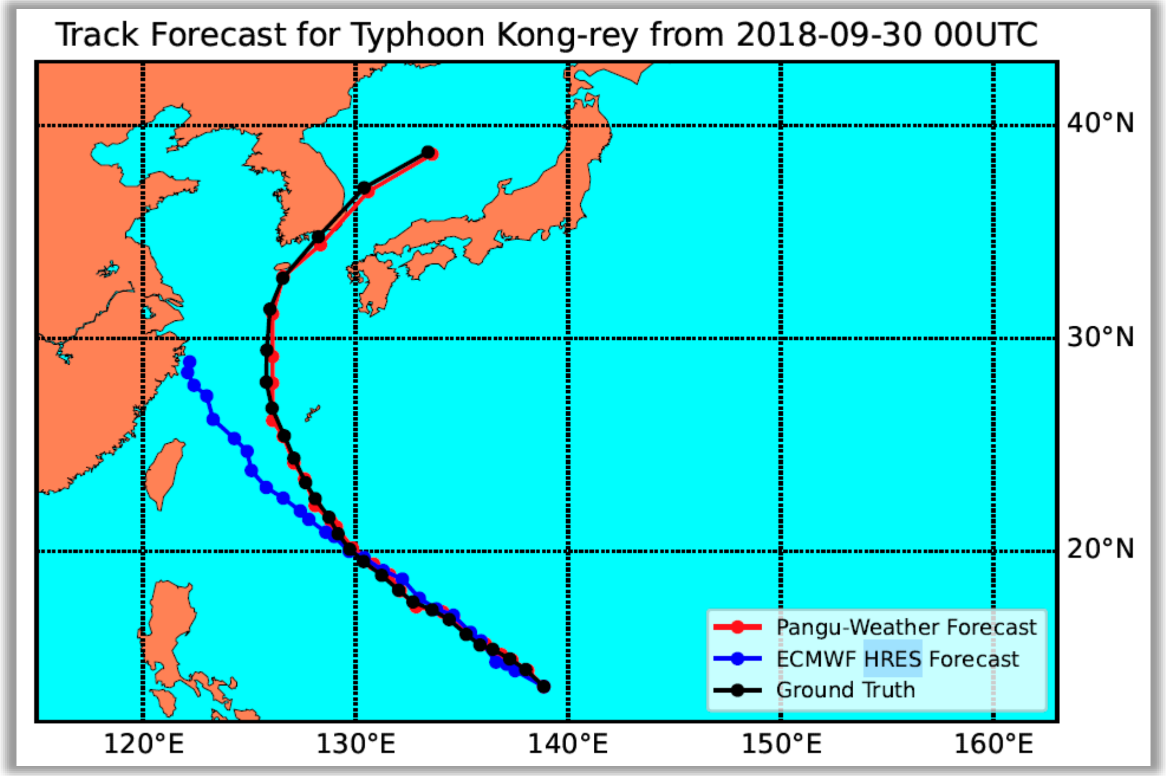

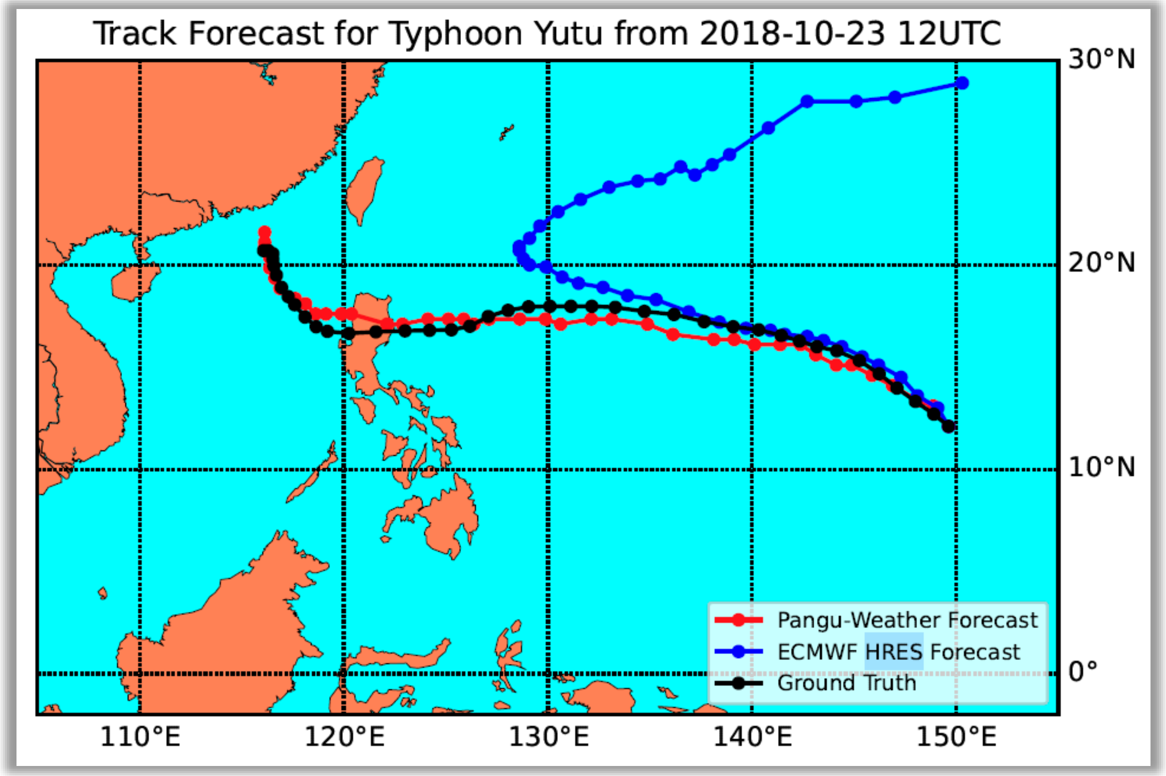

Summary and Takeaway

- **AI algorithms show great potentials in NWP**
  - For the first time, AI algorithms surpass operational IFS in reanalysis data
    - There is still **a long way** before AI can "replace" conventional NWP
  - Method: 3D deep neural networks with a bit **meteorological priors**
    - More meteorological expertise can be helpful
- **Experimental results**
  - Stronger results on reanalysis data
    - Preliminary studies show good results using IFS initial data as input
  - Orders of magnitude faster for ensemble forecast
  - Competitive in extreme weather forecasts, *e.g.*, tracking tropical cyclones
    - New paradigm in AI: **pre-trained models** for downstream tasks

Future of AI for NWP

- **More data**
  - Reanalysis data in **a finer spatial resolution** (*e.g.*, ERA6)
  - Complete observational data, allowing for **an end-to-end system**
- **Stronger models (upon more powerful hardware)**
  - **Deeper** and **wider** neural networks, trained for **more epochs**
  - From 3D to 4D: incorporating the time dimension can help
- **Better metrics**
  - RMSE and ACC are not so meaningful for long-range weather forecasts
  - Calling for **a new metric** (clues in the AI community?)

Discussion: Metrics Can Lie!

- **Consider a small-scale weather event, *e.g.*, a tropical cyclone**
  - In medium-range (*e.g.*, 5-day) weather forecasting, the chance of accurately predicting its path is very low
  - Fact: the contribution to RMSE by a prediction with a shifted cyclone is almost $2\times$ of the contribution by a prediction that directly eliminates the cyclone
  - Hence, a "smart" AI can learn to eliminate (or weaken) uncertain events: this strategy can improve quantitative metrics, but it is meaningless!
- **The bias gets larger with lead time**
  - In long-range (*e.g.*, 10-day) forecasting, we desire a new metric

Pangu-Weather Was Accepted by *Nature*[A]

- **Referees agreed with the contribution and potential of our work**
  - Referee #1: "*I am convinced that the paper makes* ***an important contribution to the field*** *and that* ***it is scientifically and technically sound.***"
  - Referee #2: "*The results themselves are* ***a significant step beyond previous results.*** *This work will, in my opinion,* ***make people reevaluate what forecasting models might look like in the future.***"
  - Referee #3: "*Our group picked one trained model, the 24h forecast model, and I can confirm that* ***it is very easy to download and run it****. It just took us one afternoon to get this to work, and it executed quickly on even a desktop computer. This means that anyone in the meteorological community can now run and test these models to their heart's desire.* ***What a great opportunity*** *for the community to explore how well the model predicts specific meteorological phenomena. Now THAT's going to* ***help with progress in the field.***"

[A] K. Bi *et al.*, Accurate Medium-range Global Weather Forecasting with 3D Neural Networks, in *Nature*, 2023.

Pangu-Weather Models Were Released[A]

- **ECMWF tested Pangu-Weather models and developed a real-time toolkit based on it**
  - Results are promising; believed to be part of "***the rise of data-driven weather forecasting***"[B]
  - *"We have **very much enjoyed running Pangu-Weather at ECMWF**, to better understand the possible power of ML forecasts. To make it easier for us to run, we have created a tool[C] which gathers the necessary data for running Pangu-Weather (and other similar models) from our repositories (e.g. CDS), and write the output data to the GRIB file format."*

[A] https://github.com/198808xc/Pangu-Weather
[B] **Z. Ben-Bouallegue** *et al.*, The Rise of Data-driven Weather Forecasting, in *arXiv preprint*:2307.10128, 2023.
[C] https://github.com/ecmwf-lab/ai-models-panguweather

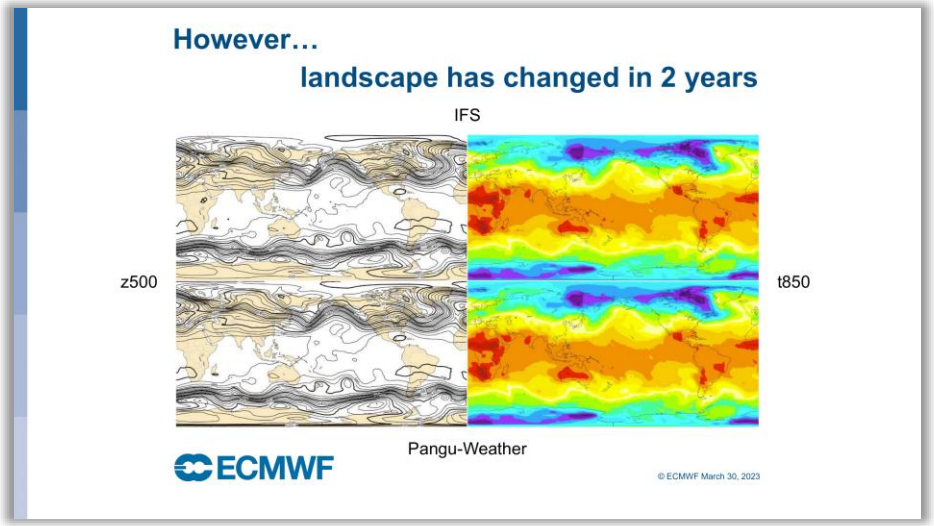

On forecasting common weather variables

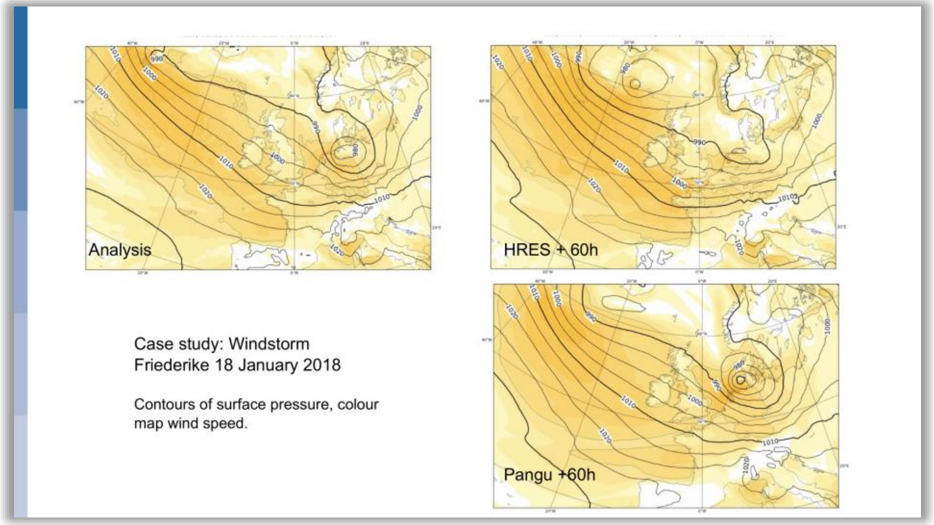

On forecasting Windstorm Friederike

Pangu-Weather at ECMWF Charts[A]

- **ECMWF released Pangu-Weather as part of their operational suite**
  - Search "PANGU" on ECMWF Charts[B]
  - Choose a set of variables you are interested in (*e.g.*, MSLP for cyclone tracking)
  - Choose the region and time you are interested in and view the results

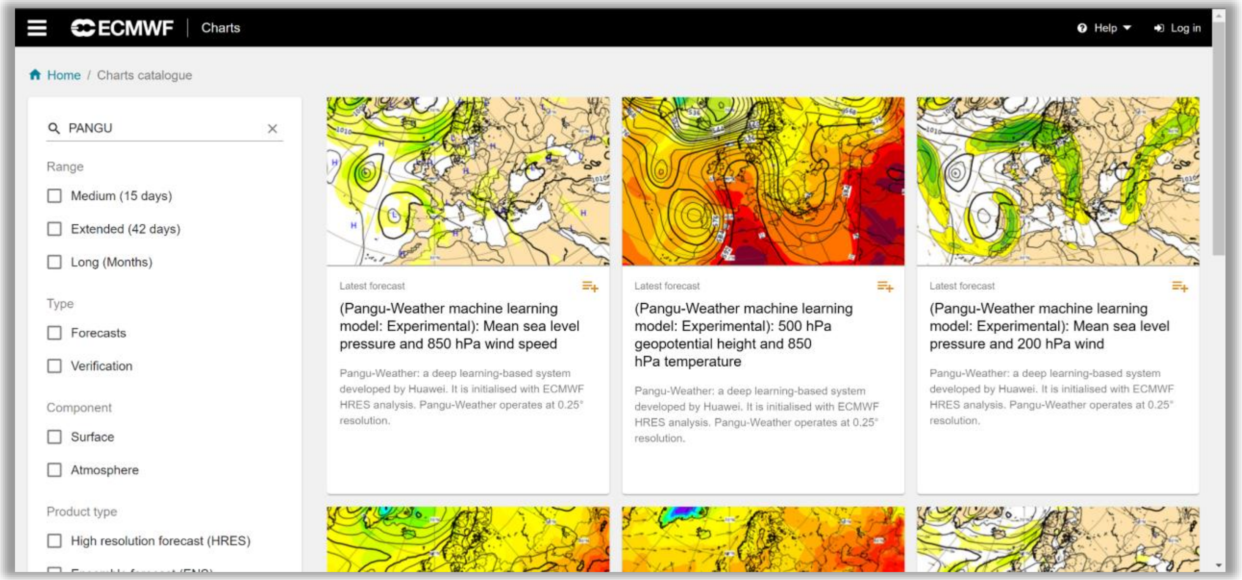

The homepage of ECMWF Charts

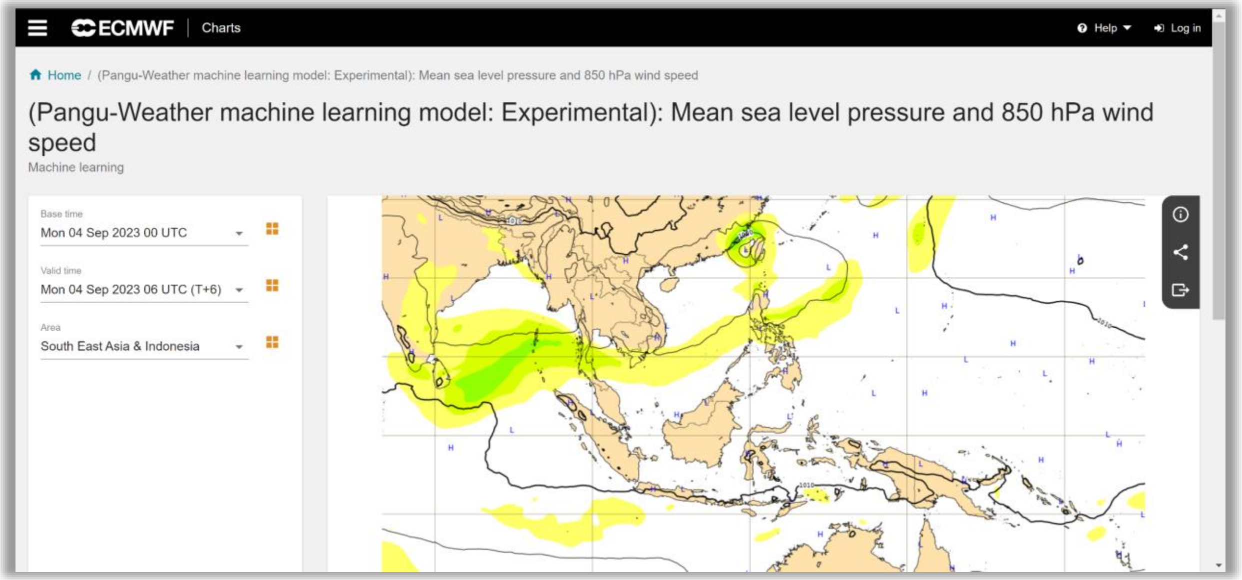

Visualization of MSLP and wind speed at 850hPa

[A] https://charts.ecmwf.int/
[B] https://charts.ecmwf.int/?query=PANGU

A Lite Version of Pangu-Weather

- **Approaching operational IFS using 1% computational costs!**
  - Strategy: using fewer data (11 rather than 39 years, daily sampling, less than 1TB data), training for 50 or 100 epochs
  - Adjusting the network architecture (heavier down-sampling)
  - We will release the training details in the GitHub repository[A]

Model	RMSE, Z500		RMSE, T850		RMSE, T2M		RMSE, U10		#Years	Patch Size	#Epochs	GPU-days (#G x days)
	3-day	5-day	3-day	5-day	3-day	5-day	3-day	5-day				
Operational IFS	152.8	333.7	1.37	2.06	1.34	1.75	1.94	2.90	/	/	/	/
Pangu-Weather	134.5	296.7	1.14	1.79	1.05	1.53	1.61	2.53	39	$2 \times 4 \times 4$	100	192×16*
Pangu-Weather-L1	163.1	338.2	1.29	1.96	1.16	1.64	1.80	2.74	11	$2 \times 8 \times 8$	100	8×6
Pangu-Weather-L2	177.9	357.5	1.36	2.05	1.24	1.71	1.90	2.84	11	$2 \times 8 \times 8$	50	8×3

[A] https://github.com/198808xc/Pangu-Weather

* We trained four models in the original release, but only one (24-hour) model for lite versions.

Thanks!

- Questions, please? (Welcome to contact me: 198808xc@gmail.com)

Check out our *Nature* paper!

Check out the code and models!