



Super-Resolving Coarse-Resolution Weather Forecasts with Flow Matching

Aymeric Delefosse^{1*}, Anastase Charantonis¹, Dominique Béréziat²

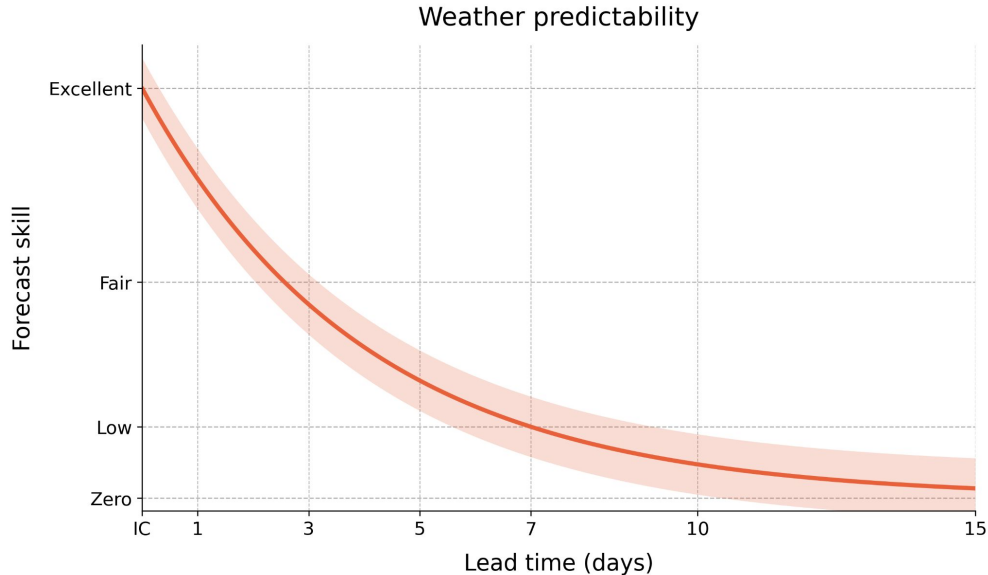
¹Inria, ARCHES, Paris, France

²Sorbonne Université, CNRS, LIP6, Paris, France

*Email: aymeric.delefosse@inria.fr



For medium-range forecasting (~7 days), is spatial resolution a **feature** or a **cost**?




Adapted from J. Lukas, E. Gawthrop and T. Barnston

- Still requires **lots** of resources
- Predictability comes from **initial conditions of the atmosphere**
- At medium-range, **large-scale dynamics dominate**



Coarse-resolution models achieve competitive skill, at a fraction of the cost



Training cost (A100-days)



0.25° (fine resolution)


- GenCast: 480 days
- FourCastNet3: 7147 days




1.5° (coarse resolution)

- ArchesWeatherGen: 23 days

But coarse-resolution is insufficient



Weather warnings

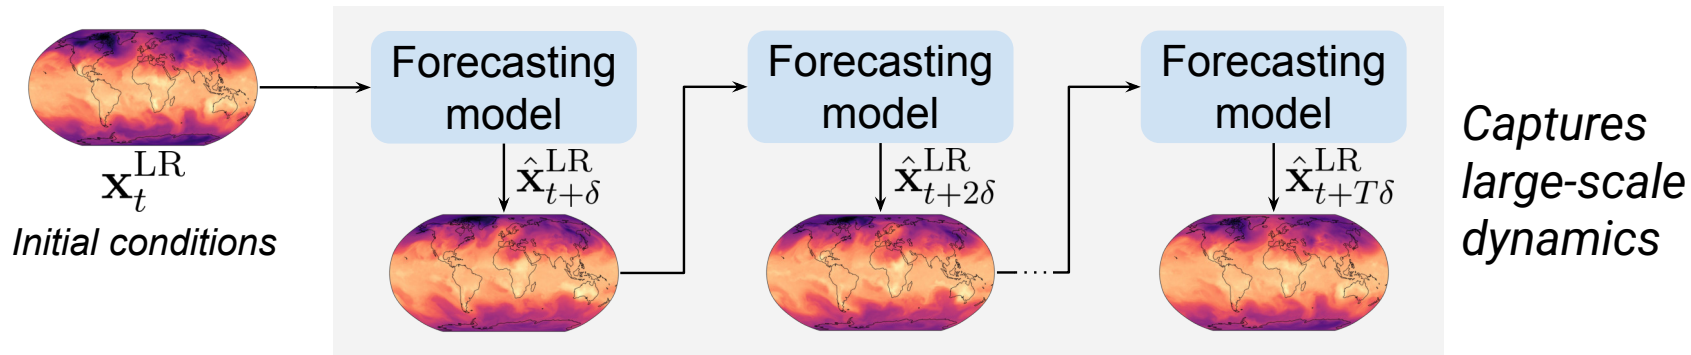


Energy forecasting

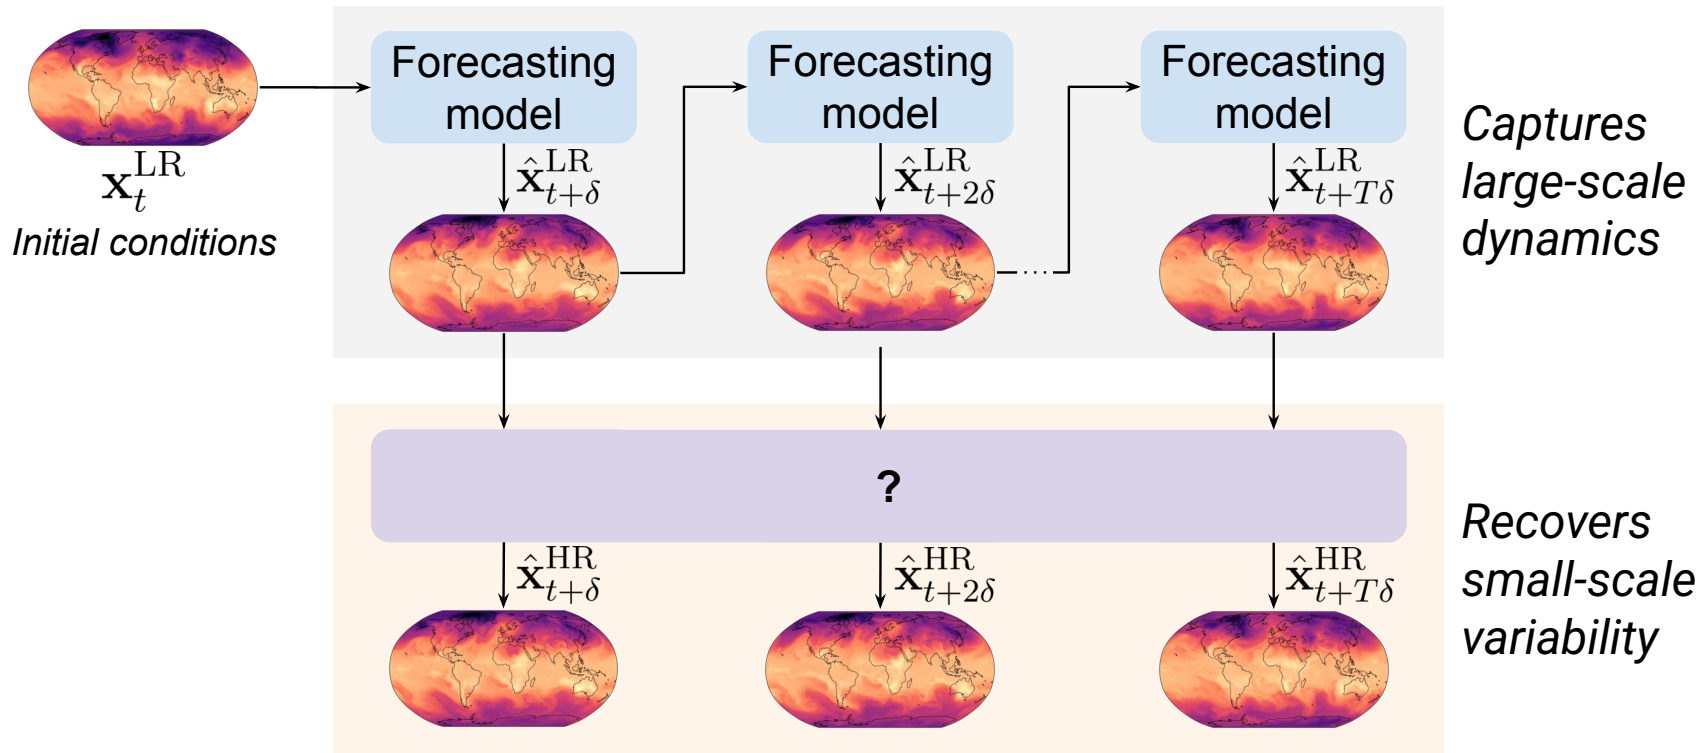


And many more!


What if they were **two separate problems**?




What if they were **two separate problems**?




Downscaling as an affordable approach?



Low resolution (1.5°)




High resolution (0.25°)



Ill-posed

Reframing the problem



$\uparrow \hat{\mathbf{x}}_{t+\delta}^{\text{LR}}$
Bicubic
upsampling

$\mathbf{r}_{t+\delta}$
Learned residual


$\hat{\mathbf{x}}_{t+\delta}^{\text{HR}}$
Reconstructed
high-resolution state

model $p(\mathbf{r}_t \mid \uparrow \mathbf{x}_t^{\text{LR}})$ via flow matching

- Preserves **large-scale dynamics**
- Concentrates on **high-frequency corrections**

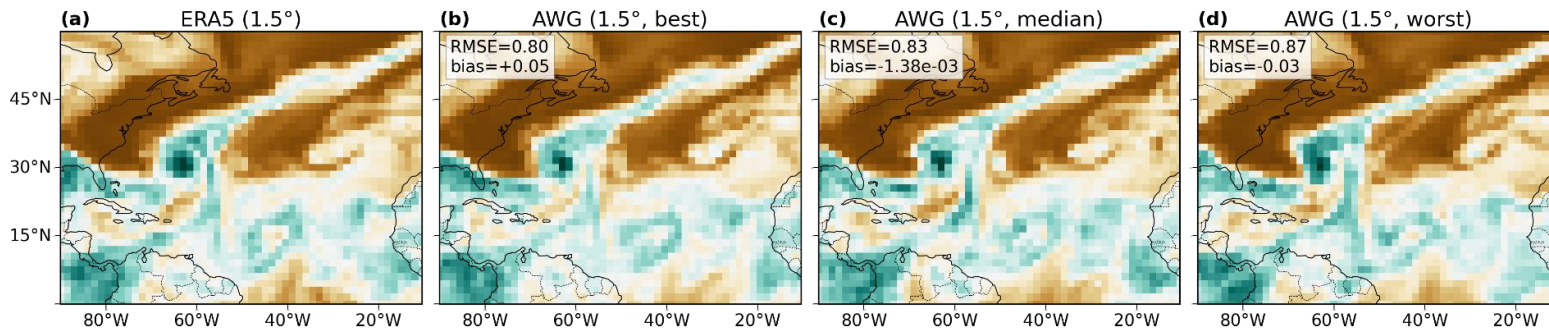
Method

- Trained on **ERA5 pairs on 82 variables**, where $\mathbf{x}_t^{\text{LR}} = \downarrow \mathbf{x}_t^{\text{HR}}$
 - 1979–2018 (train), 2019 (val), 2020 (test)
- $1.5^\circ \Rightarrow 0.25^\circ$ (**6× downscaling**)
- Applied **zero-shot** to ArchesWeatherGen



Case study Hurricane Teddy

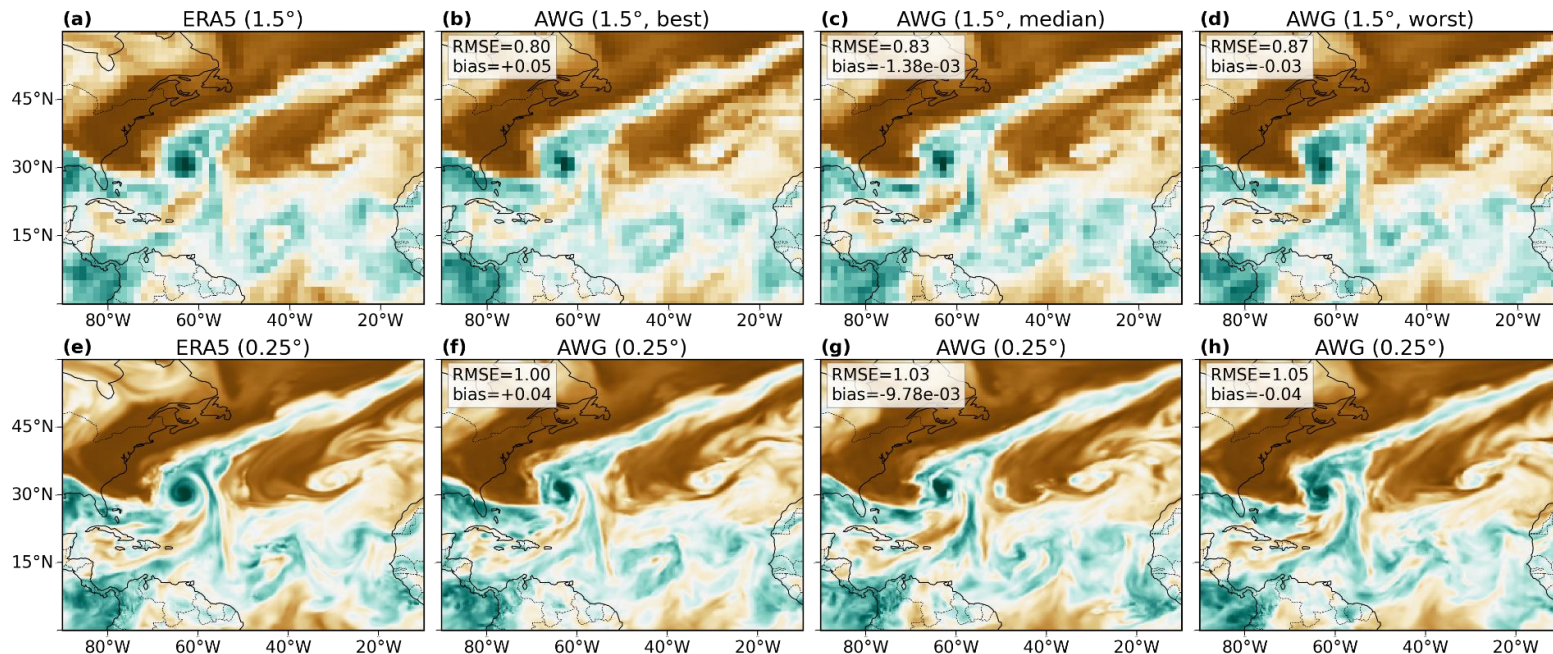
2020-09-21 12:00 UTC, +24h lead time
Specific humidity @ 700 hPa



At 1.5°, ArchesWeatherGen captures the storm's position

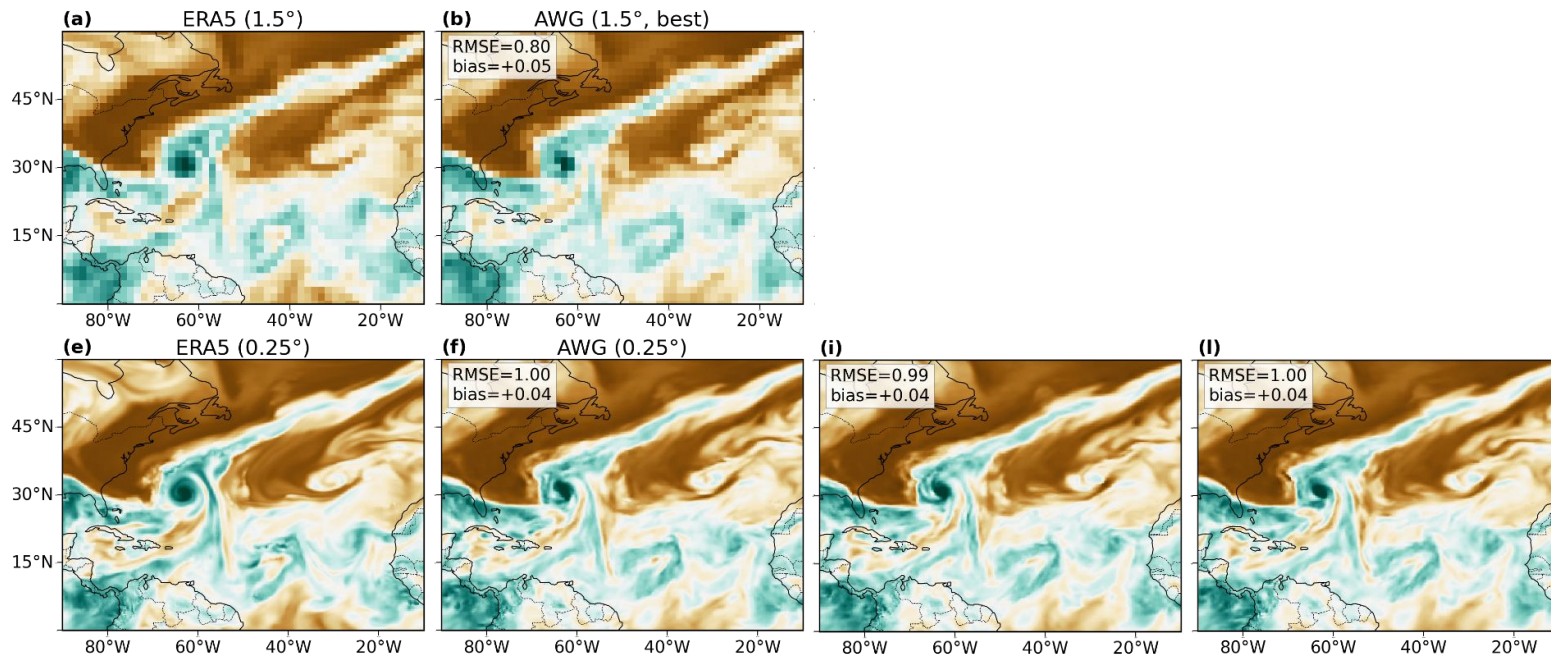
Case study Hurricane Teddy

2020-09-21 12:00 UTC, +24h lead time
Specific humidity @ 700 hPa






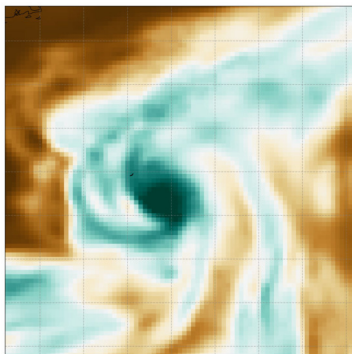
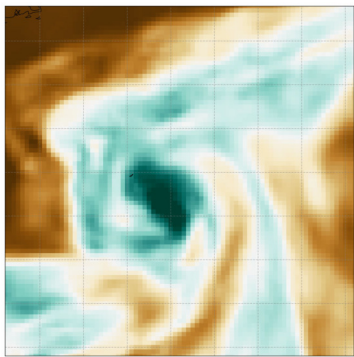


Case study Hurricane Teddy

2020-09-21 12:00 UTC, +24h lead time
Specific humidity @ 700 hPa

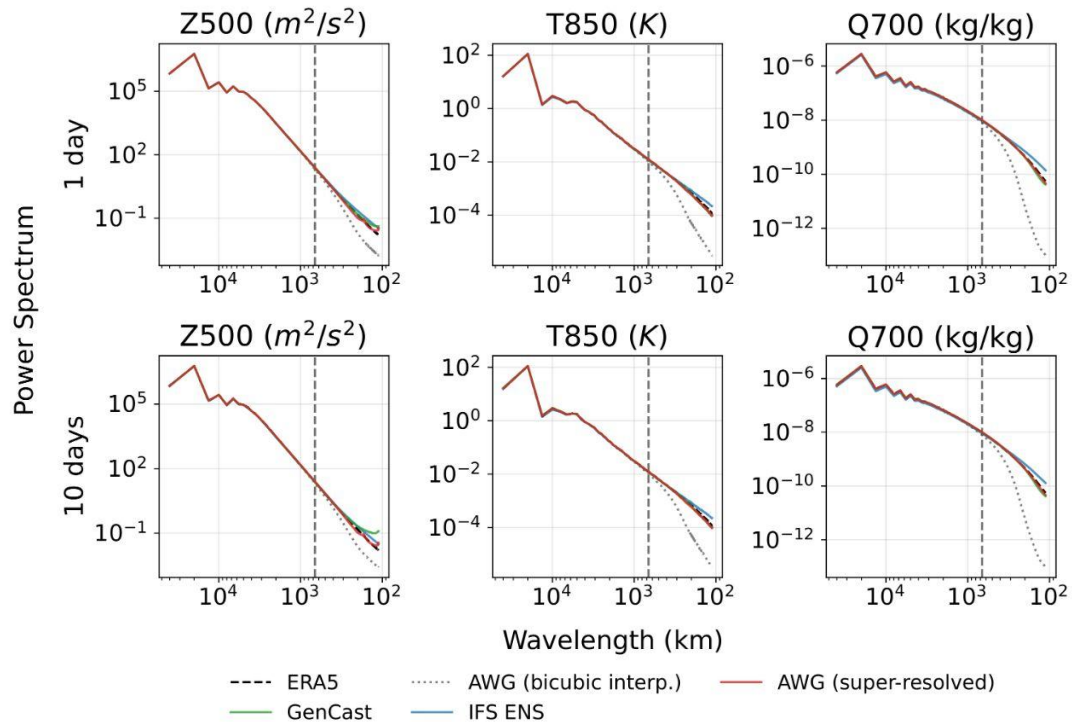
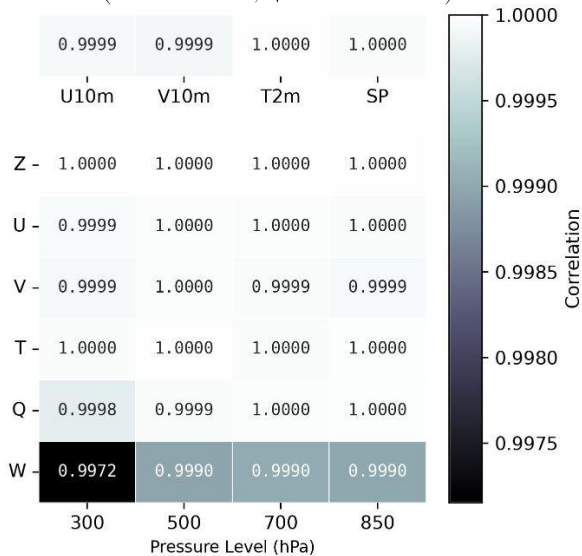


Case study Hurricane Teddy



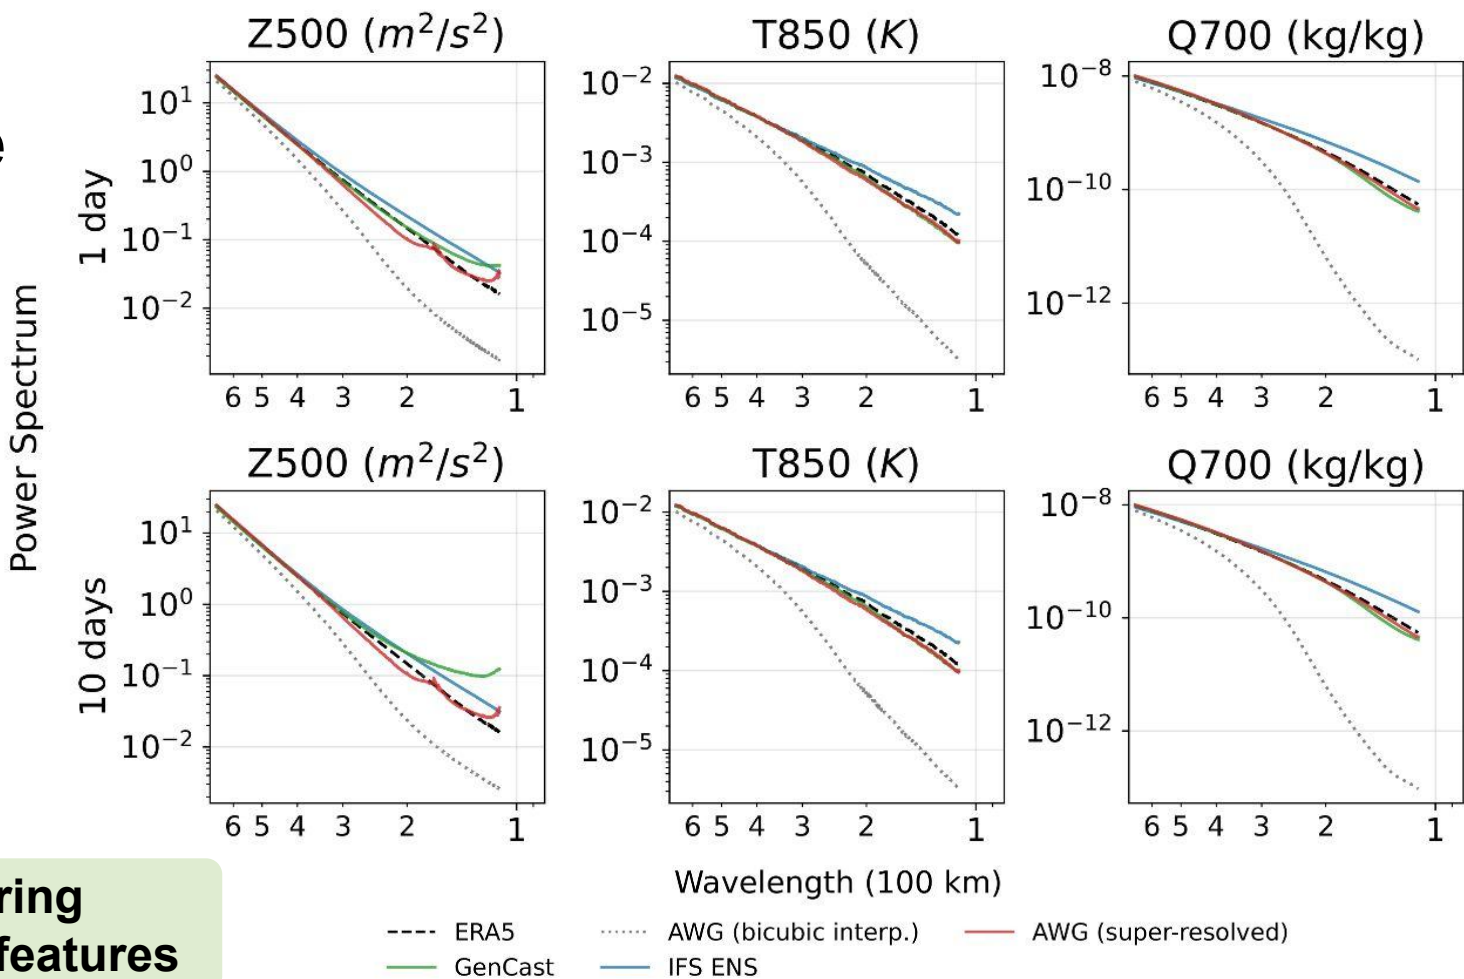
Design validation

$$\text{corr}(\text{AWG}^{1.5^\circ}, \downarrow \text{AWG}^{0.25^\circ})$$




Super-resolution leaves the **resolved-scale dynamics unchanged...**

Fine-scale realism



...while recovering
subgrid-scale features

Weather forecast skill



Averaged over WeatherBench 2 headline variables
(Z500, T850, Q700, U850, V850, T2m, SP, U10m, V10m)

Conclusion and perspectives

Contributions

- **Modular framework**
- Recovers physically realistic fine-scale variability
- Naturally **decomposes large-scale and fine-scale uncertainty**
- **7 A100-days** (30 % of AWG cost)
- \approx 8 min for a 10-day forecast on a **single V100 32GB**

Limitations and future directions

- **Sub-daily** temporal downscaling
- Inference acceleration via **distillation**
- Application to **subseasonal** and climate modelling
- Training on **kilometer-scale models** like ICON as HR reference


Thanks for listening

Any questions?



Application to ArchesWeatherGen

ArchesWeatherGen
(1.5°)



Super-resolution
(0.25°)

