T. HOEFLER

High-performance Climate Simulations and AI - from Earth Virtualization to Data Compression

with contributions by B. Stevens, O. Fuhrer, L. Huang, P. Groenquist, P. Dueben, F. Scheidl, and the whole SPCL team, CSCS, ECMWF, EVE Summit Attendees, industry friends, and others

Keynote at Supercomputing Asia 2024, Sydney, Australia

Singapore Airlines SQ 319 10/25/22, 280 pax

Source: MENTOUR PILOT (youtube)

T. HOEFLER

High-performance Climate Simulations and AI - from Earth Virtualization to Data Compression

with contributions by B. Stevens, O. Fuhrer, L. Huang, P. Groenquist, P. Dueben, F. Scheidl, and the whole SPCL team, CSCS, ECMWF, the EVE Summit Attendees, industry friends, and many others

Keynote at Supercomputing Asia 2024, Sydney, Australia

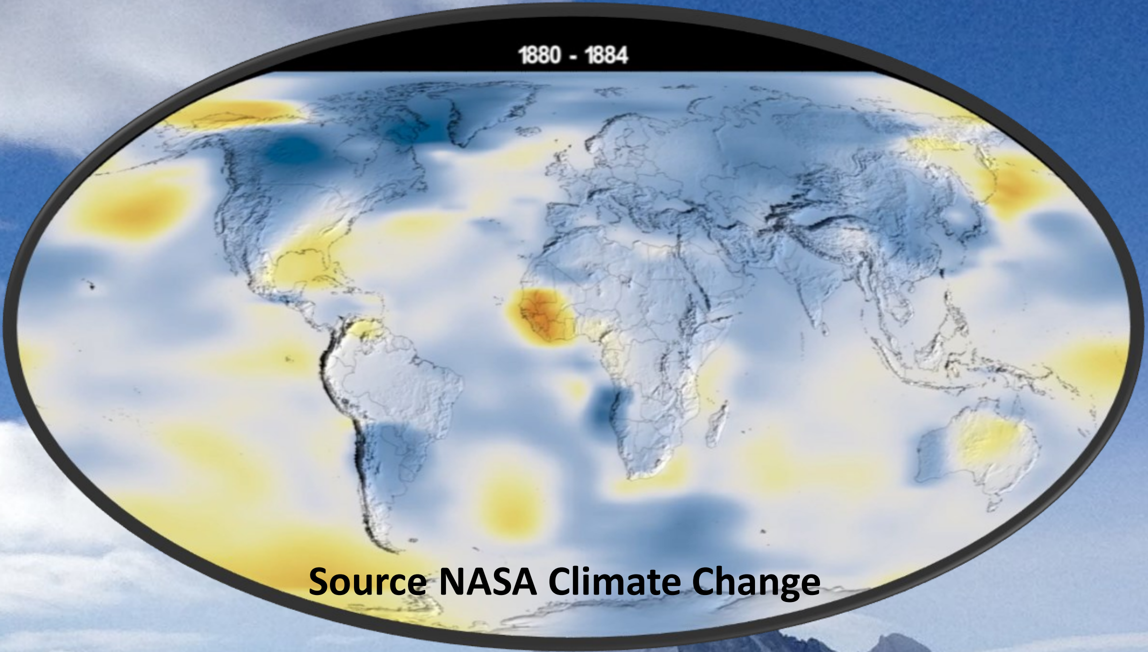

Source NASA Climate Change


Jensen Huang, Cofounder + CEO NVIDIA

"Climate simulation is basically impossible today."

"Predicting the average temperature is possible. However, the world doesn't care about average. You care about your own region." (Huang, Nov. 2023)

Earth Virtualization Engines Summit in Berlin

140 scientists from 93 institutions

Stevens et al.: "Earth Virtualization Engines (EVE)" (https://essd.copernicus.org/preprints/essd-2023-376/)

Climate prediction is extremely demanding (“impossible” – decades long)

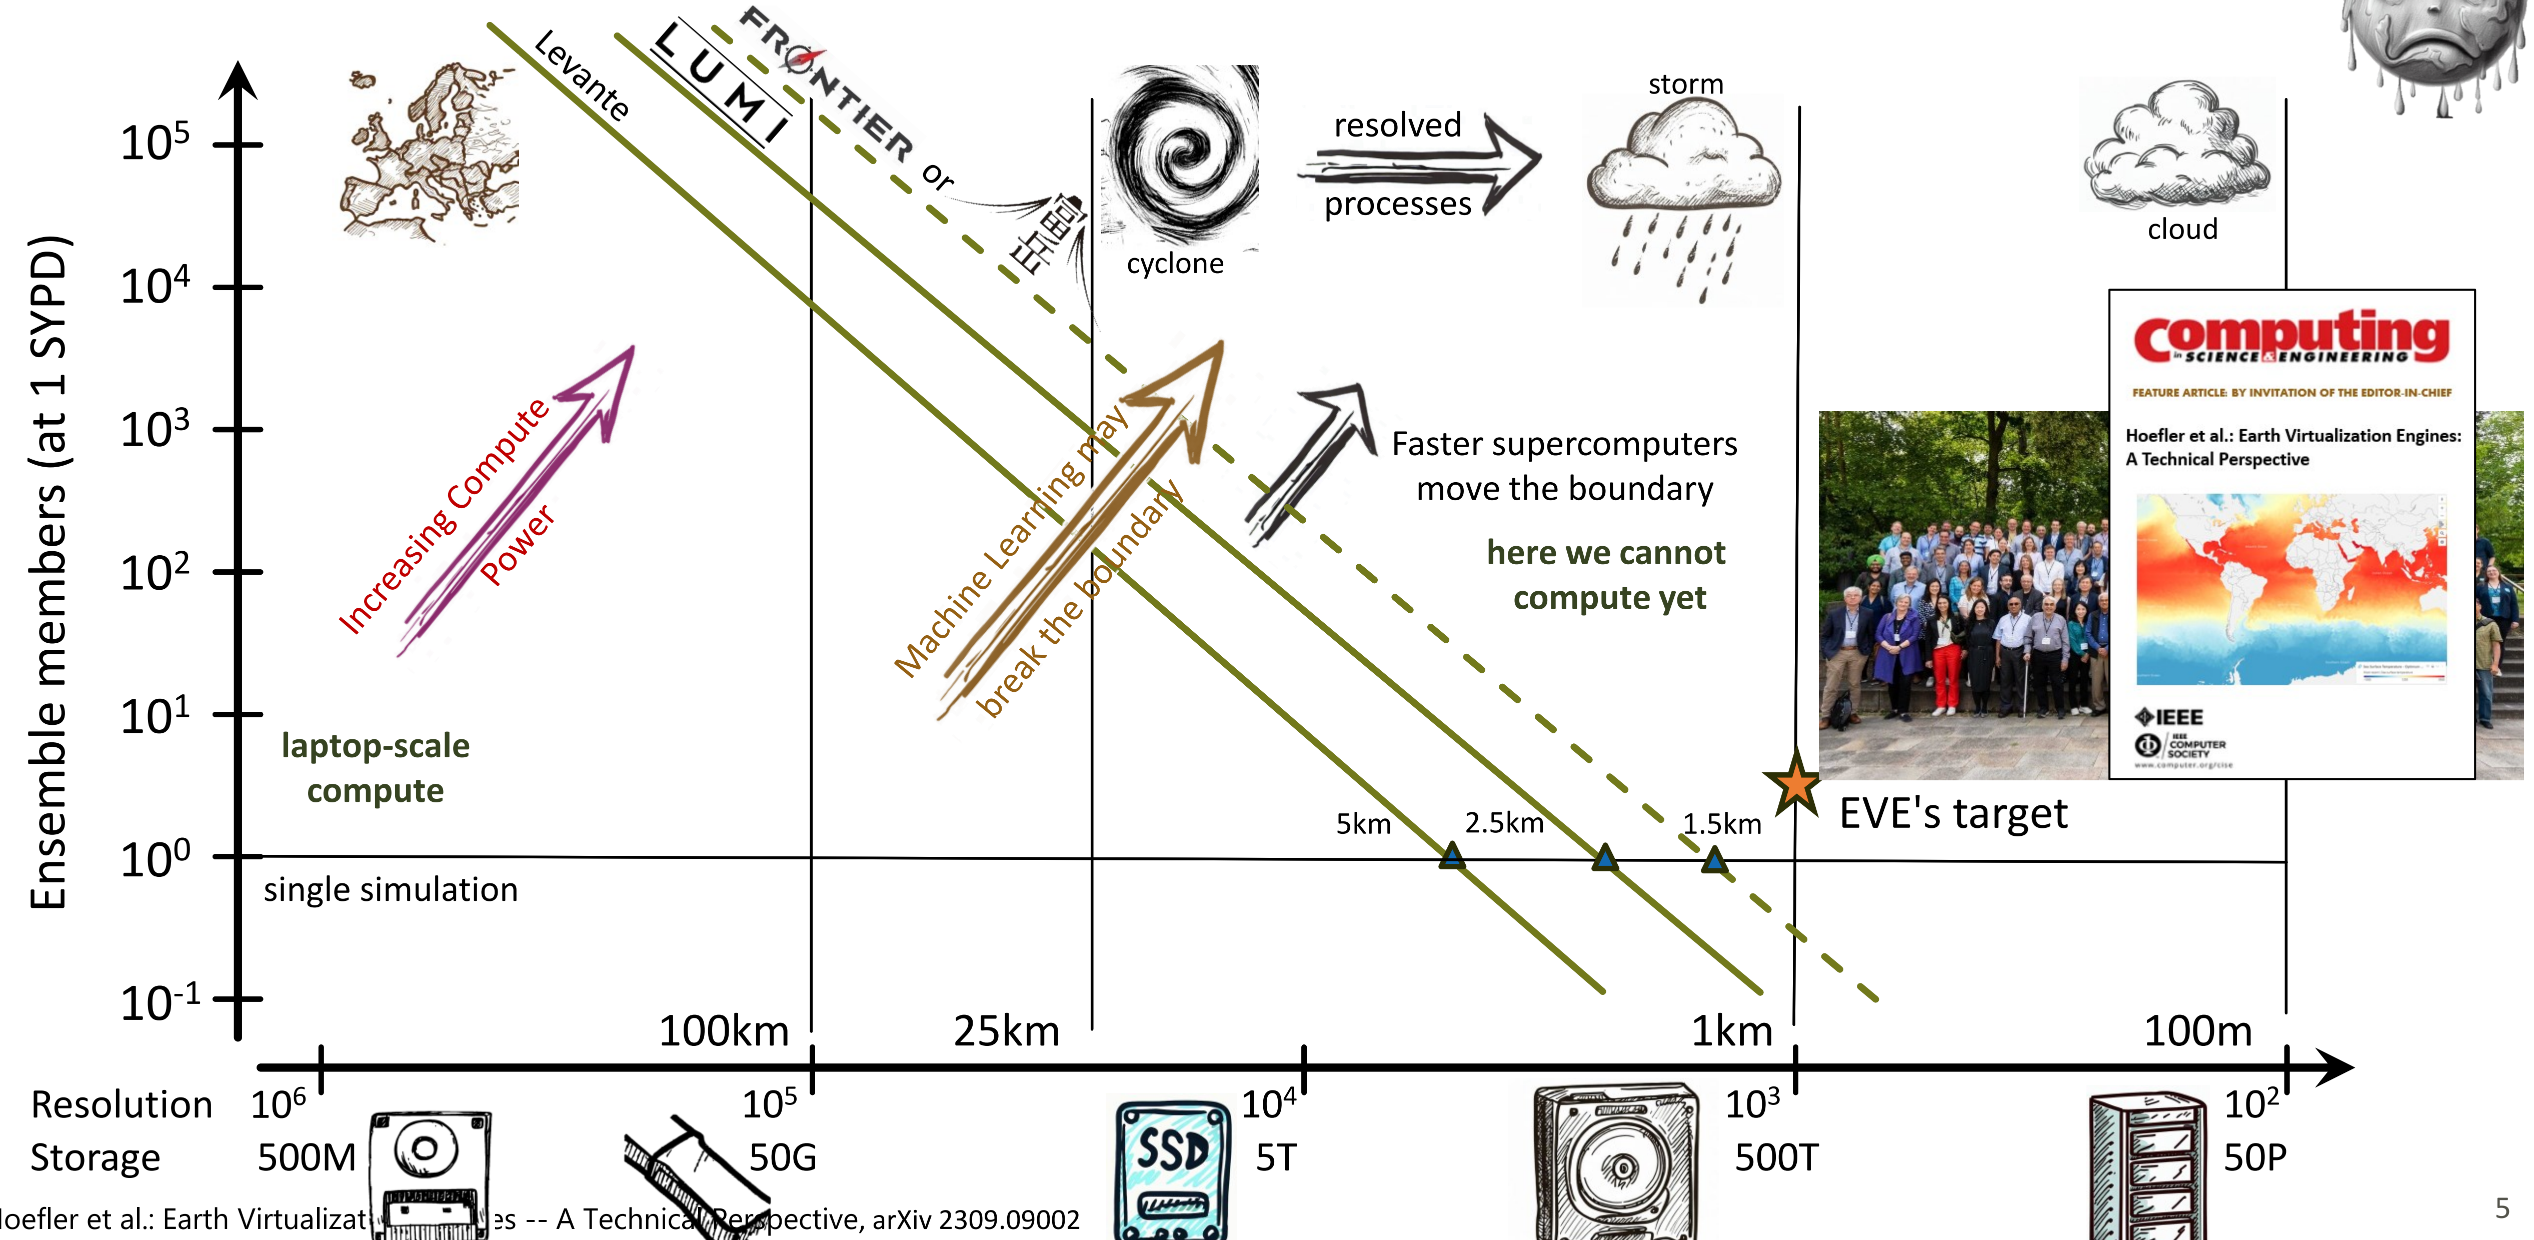

Earth Virtualization Engines Summit in Berlin

140 scientists from 93 institutions

Icosahedral Nonhydrostatic Weather and Climate Model (ICON): >630k SLOC Fortran

30-year simulation at 1.23km (1 trajectory): **11 Zops, 13 days on Exascale** (1% efficiency, 2.4 SYPD)

16k MPI processes, 3.25 GiB/s comm, **344 PiB of storage (7 vars, 15 min), 306 GiB/s average file I/O**

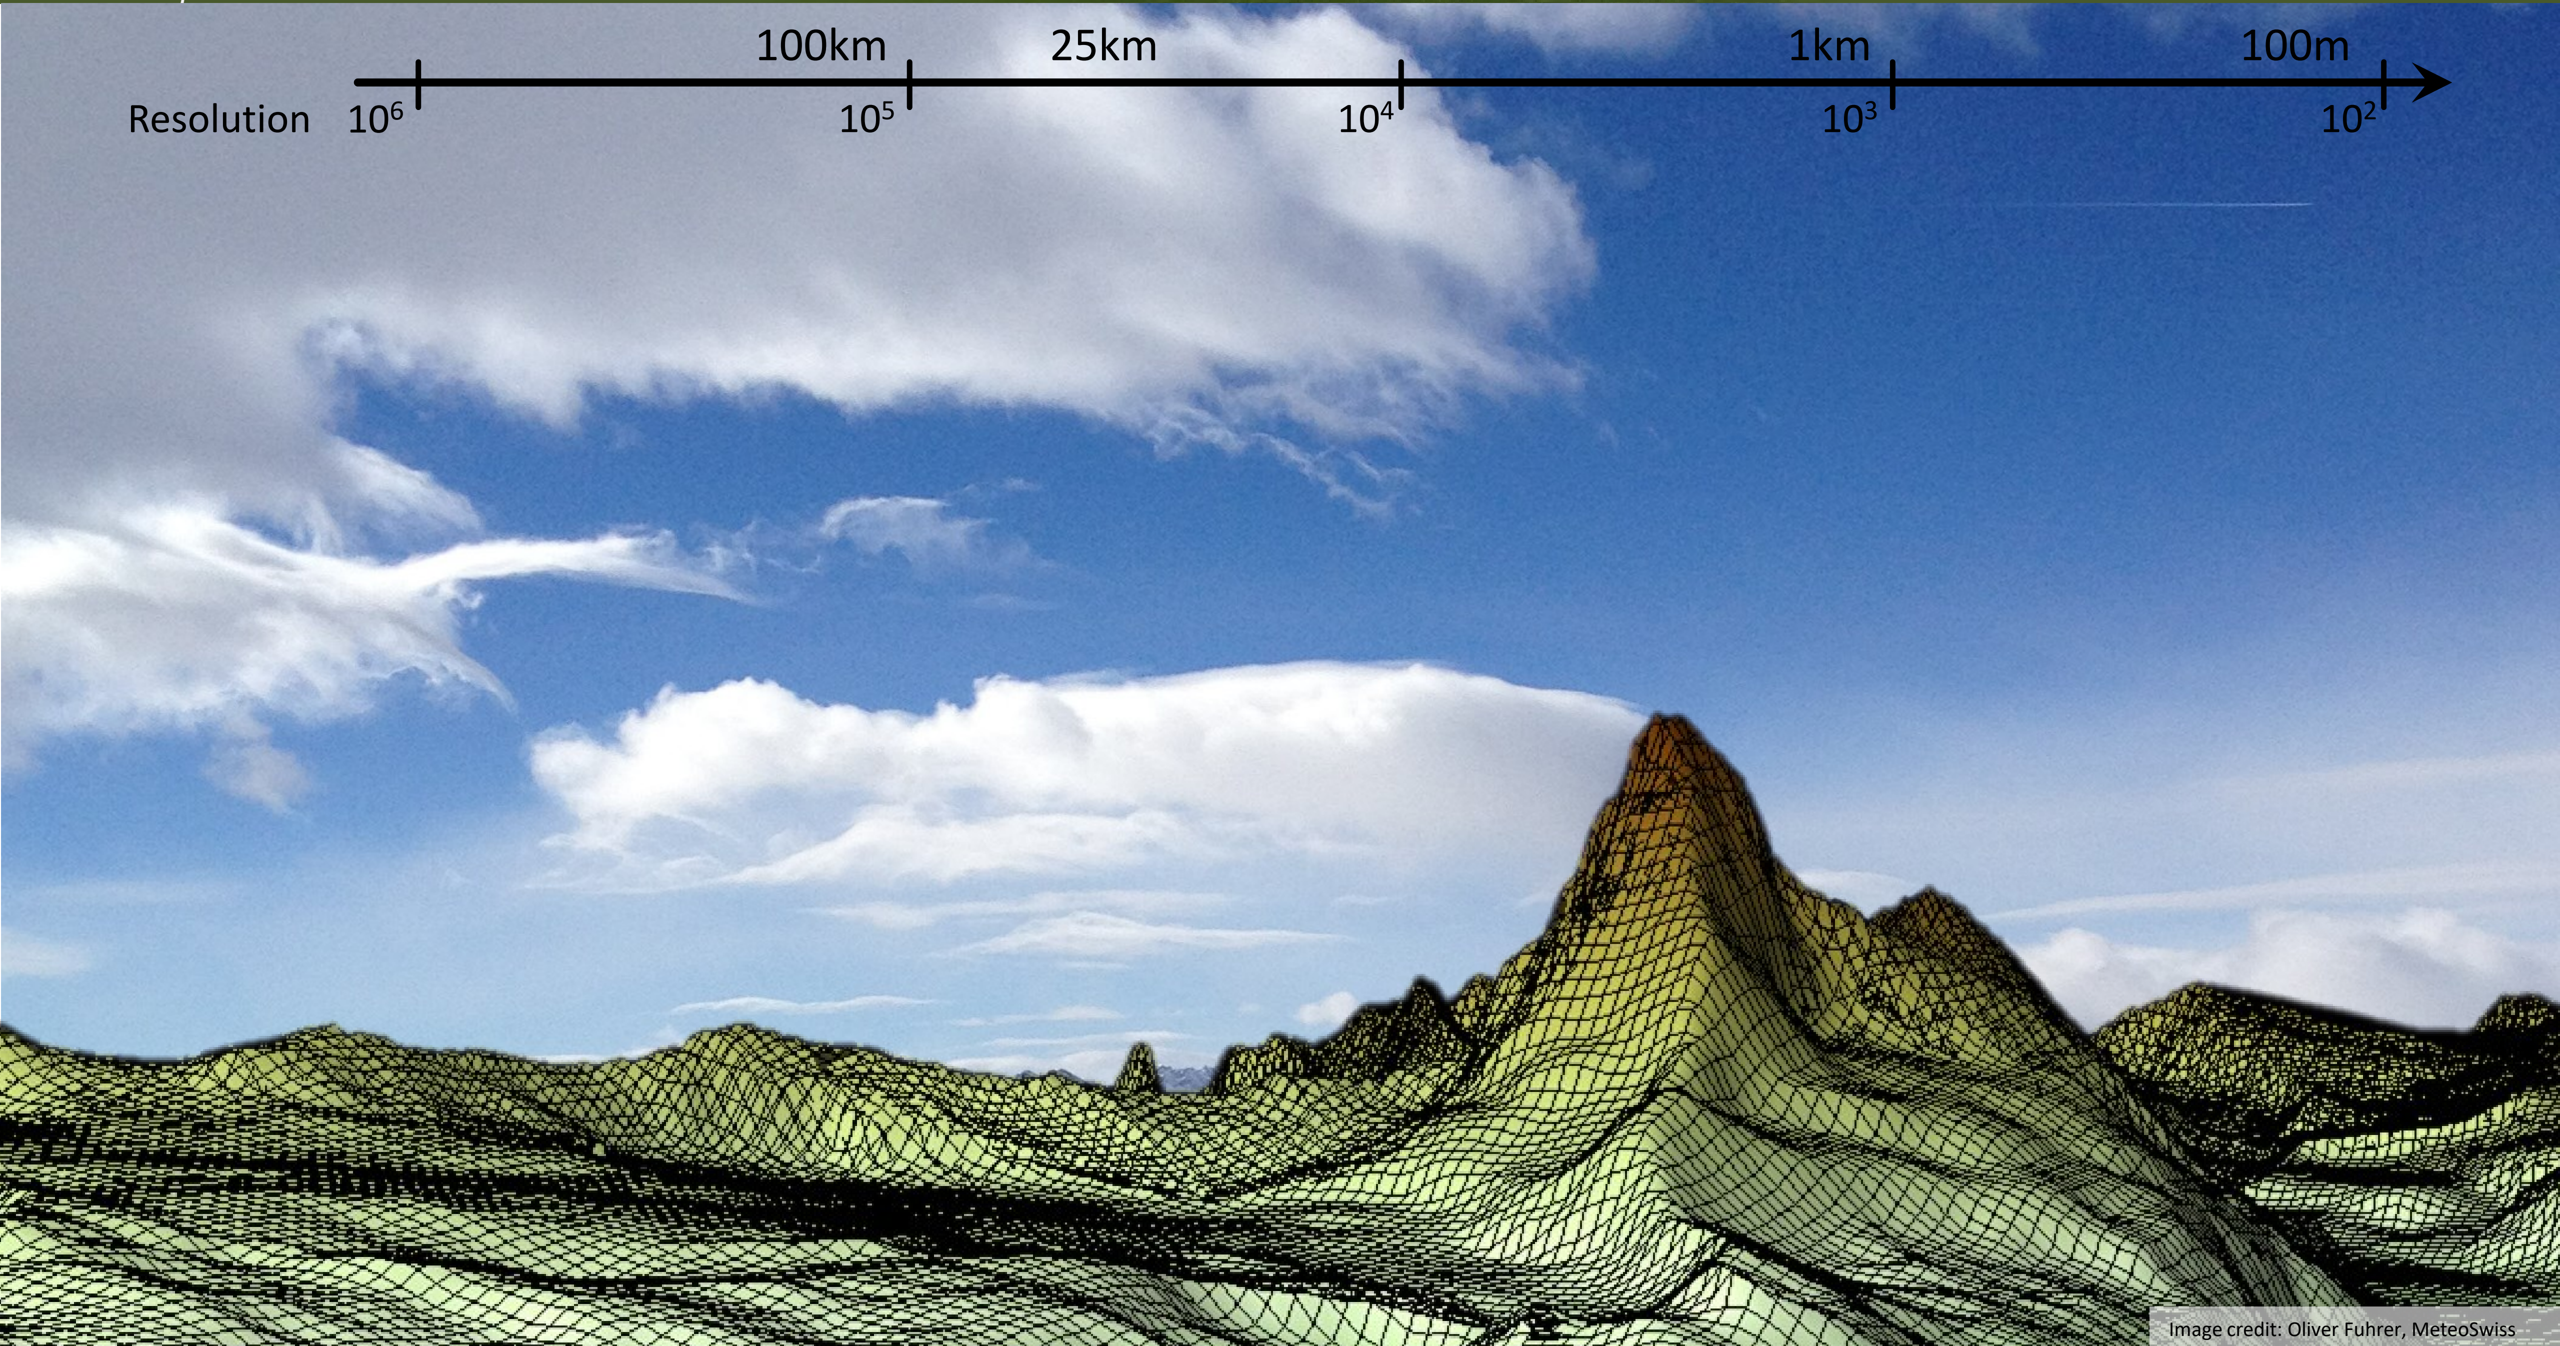

Image credit: Oliver Fuhrer, MeteoSwiss

	100km	25km		1km	100m
Resolution	10^6	10^5	10^4	10^3	10^2
Storage	500M	50G	5T	500T	50


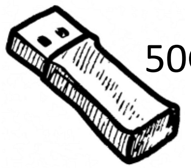

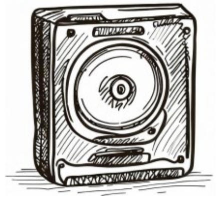
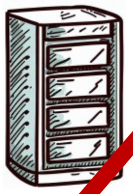

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

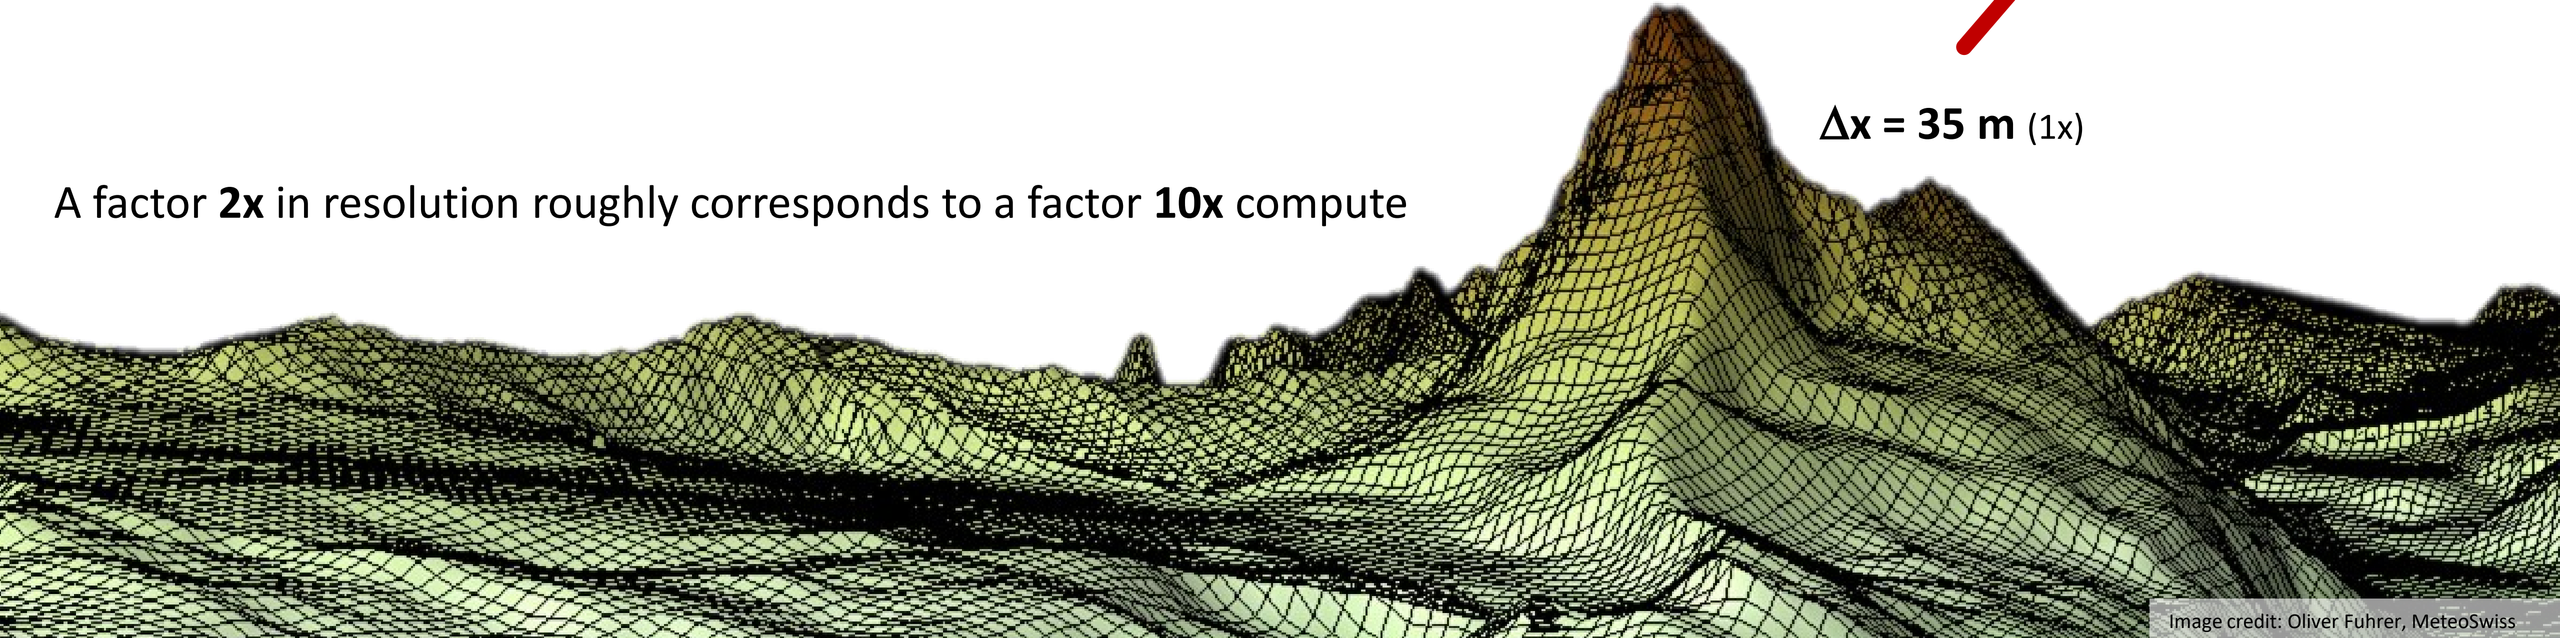

Image credit: Oliver Fuhrer, MeteoSwiss

	100km	25km		1km	100m
Resolution	10^6	10^5	10^4	10^3	10^2

Storage

500M  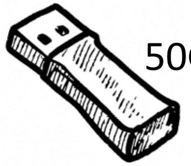 50G 

5T 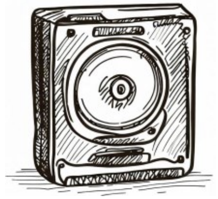 500T 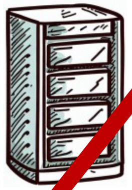 50P

Δ**x = 70 m** (10x)

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

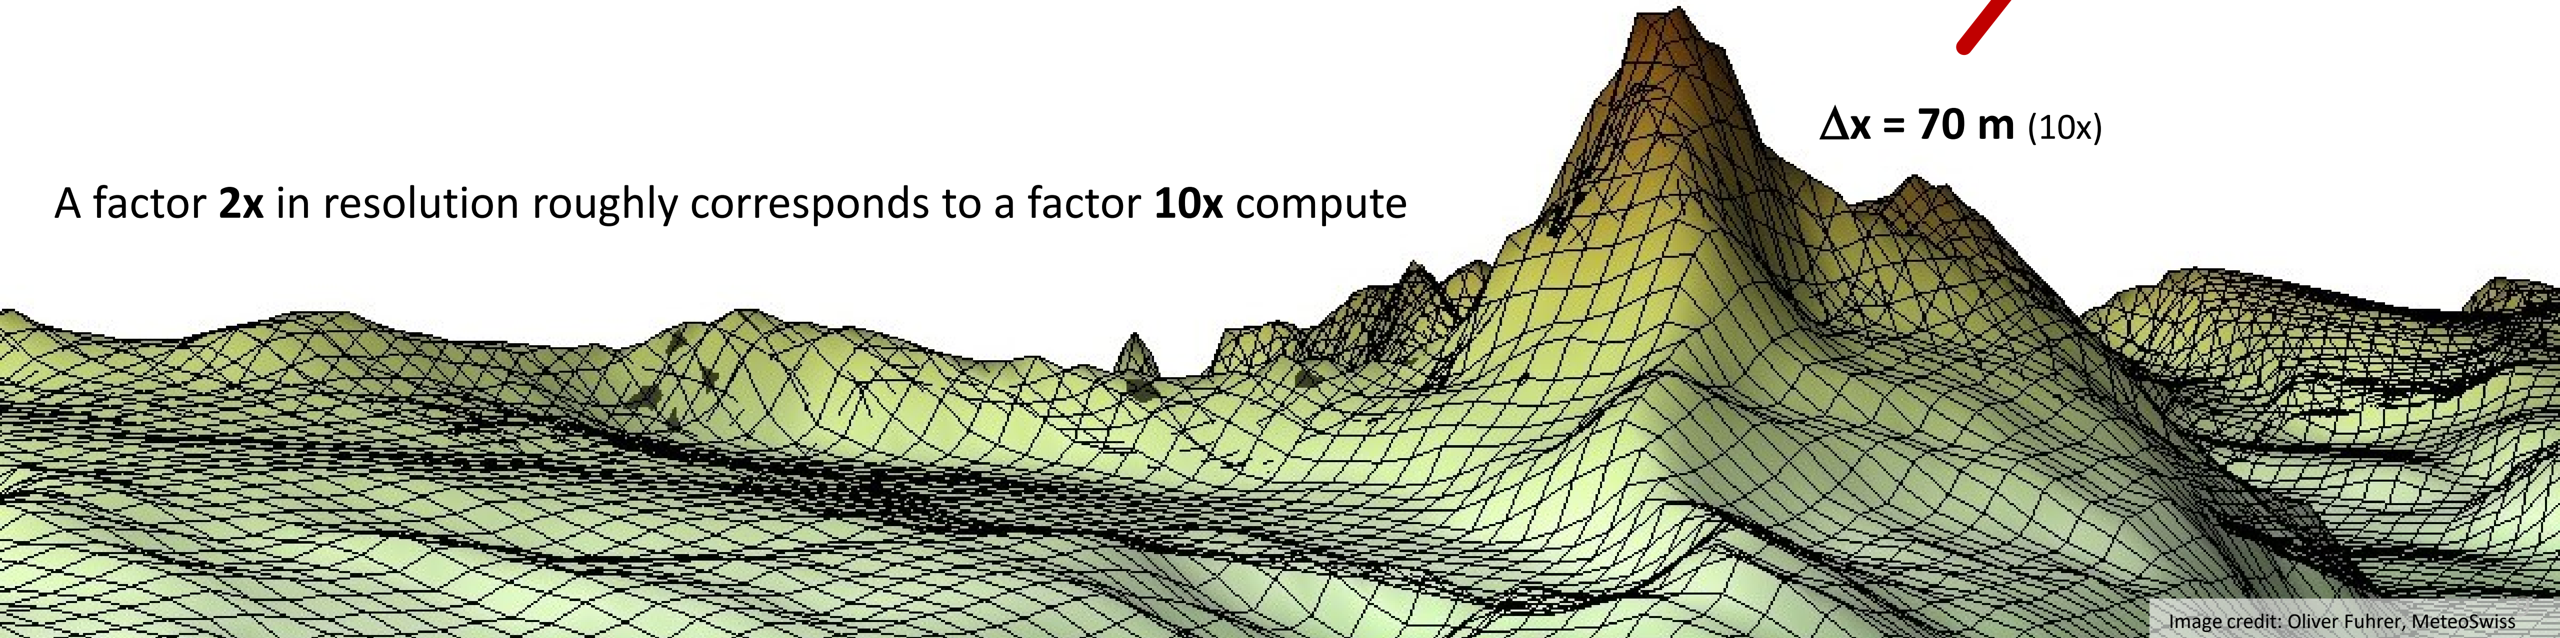

Image credit: Oliver Fuhrer, MeteoSwiss

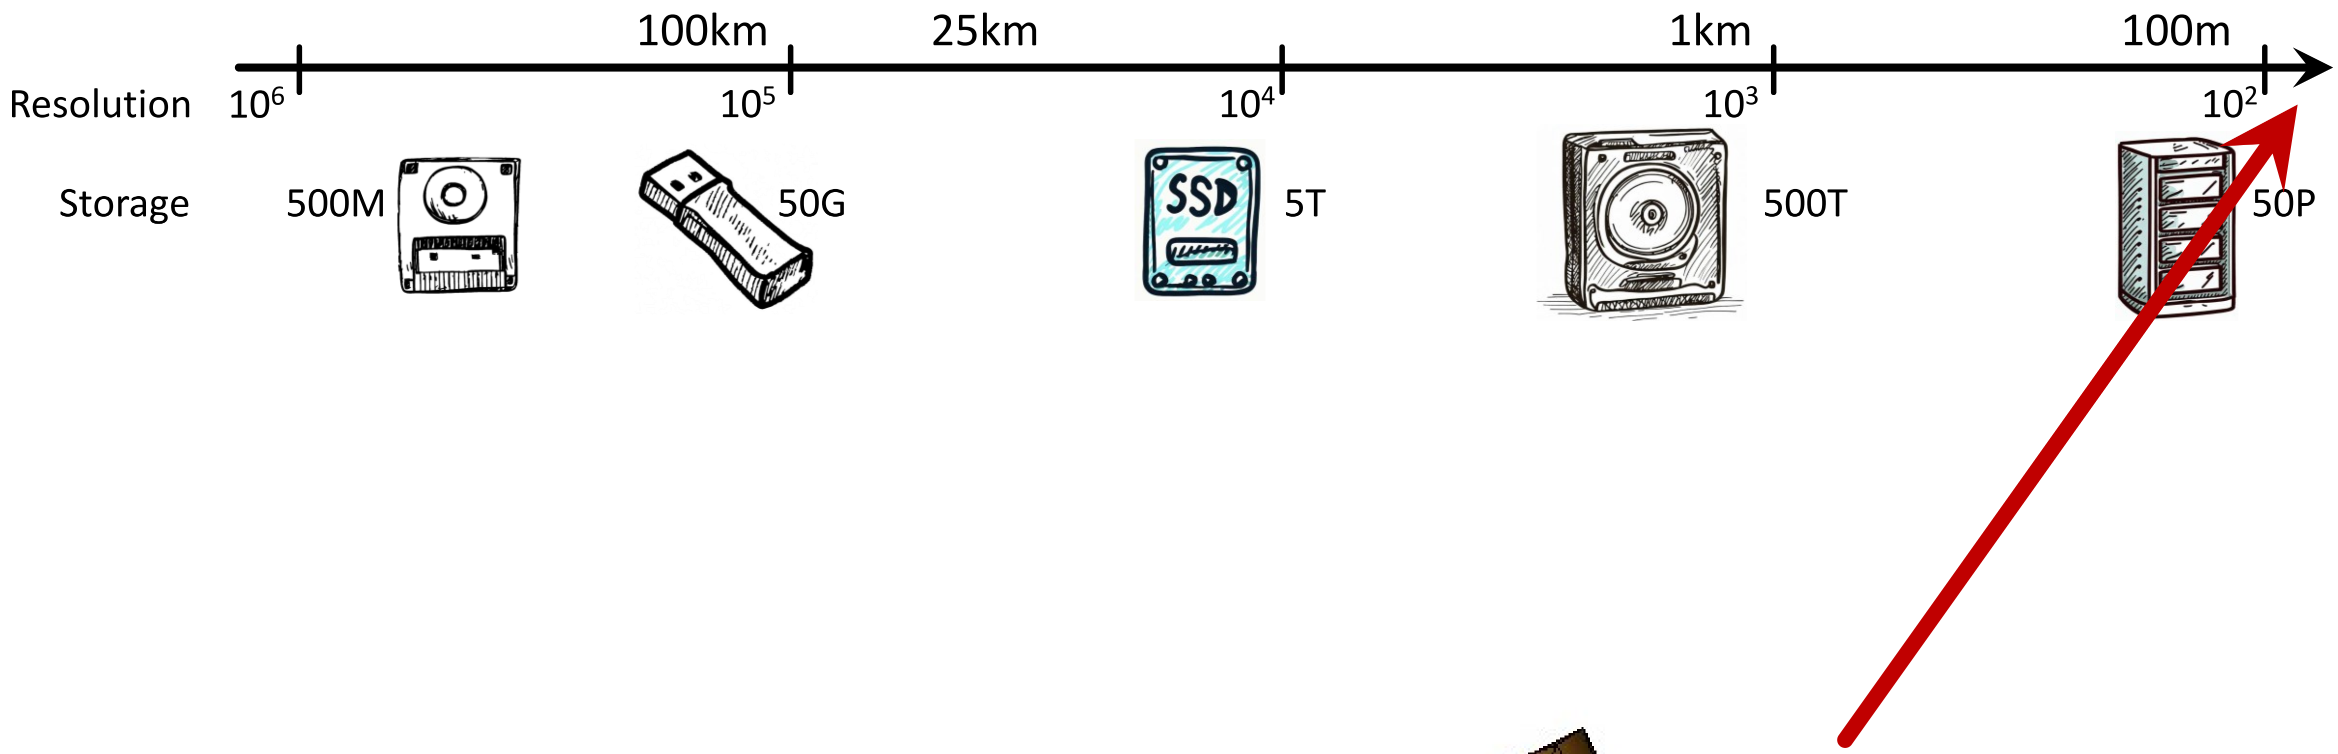

Δx = 140 m (100x)

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

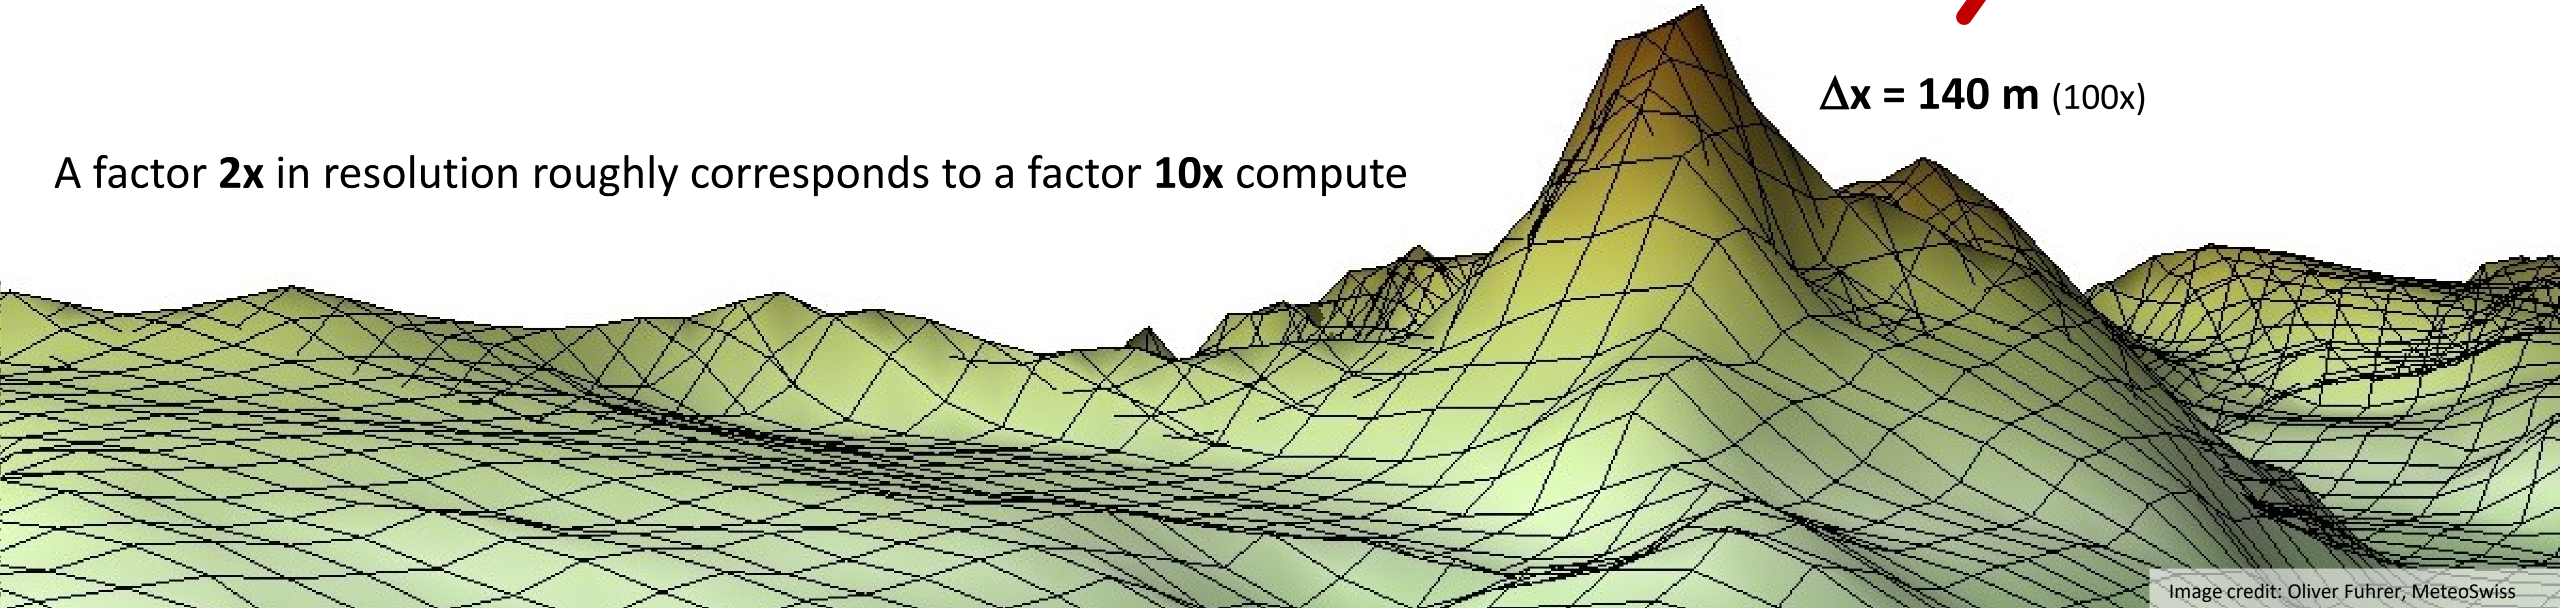

Image credit: Oliver Fuhrer, MeteoSwiss

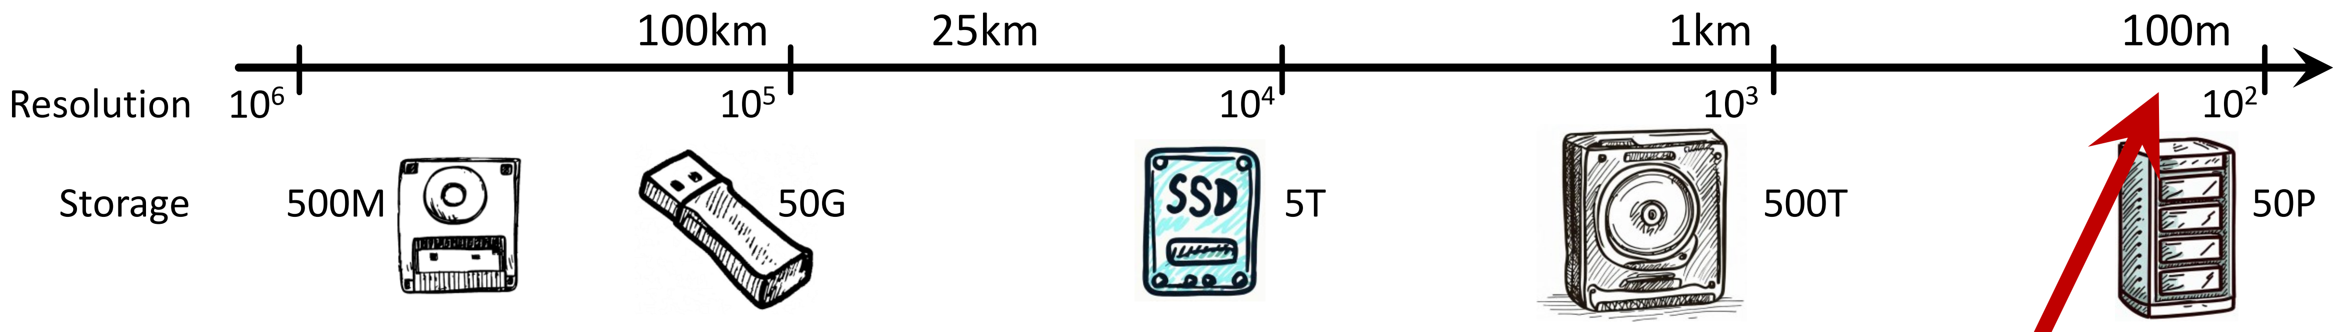

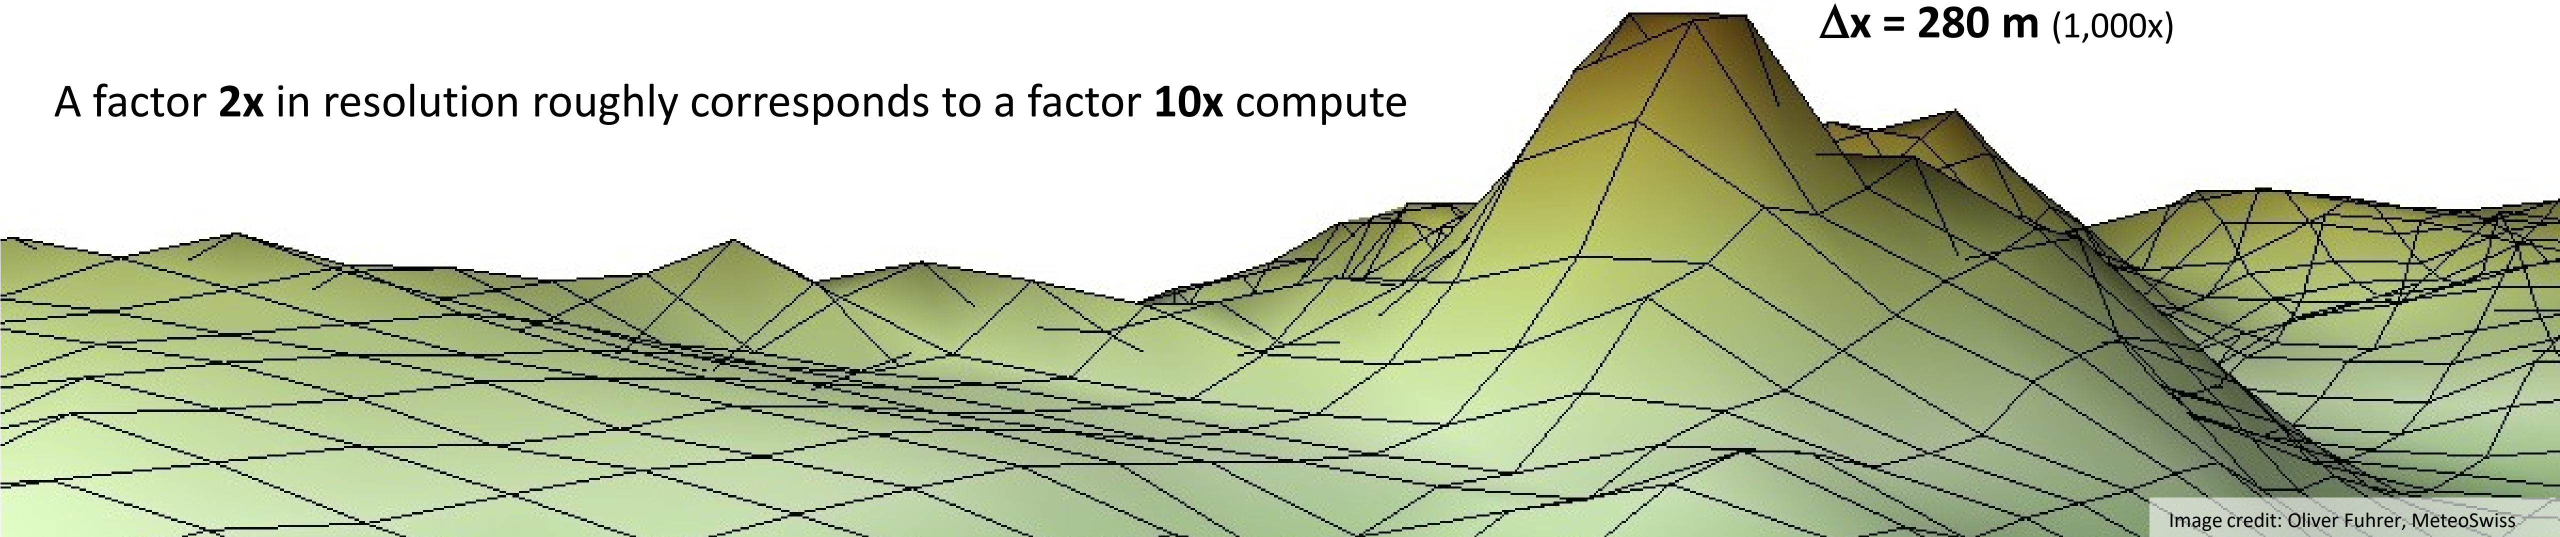

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

Image credit: Oliver Fuhrer, MeteoSwiss

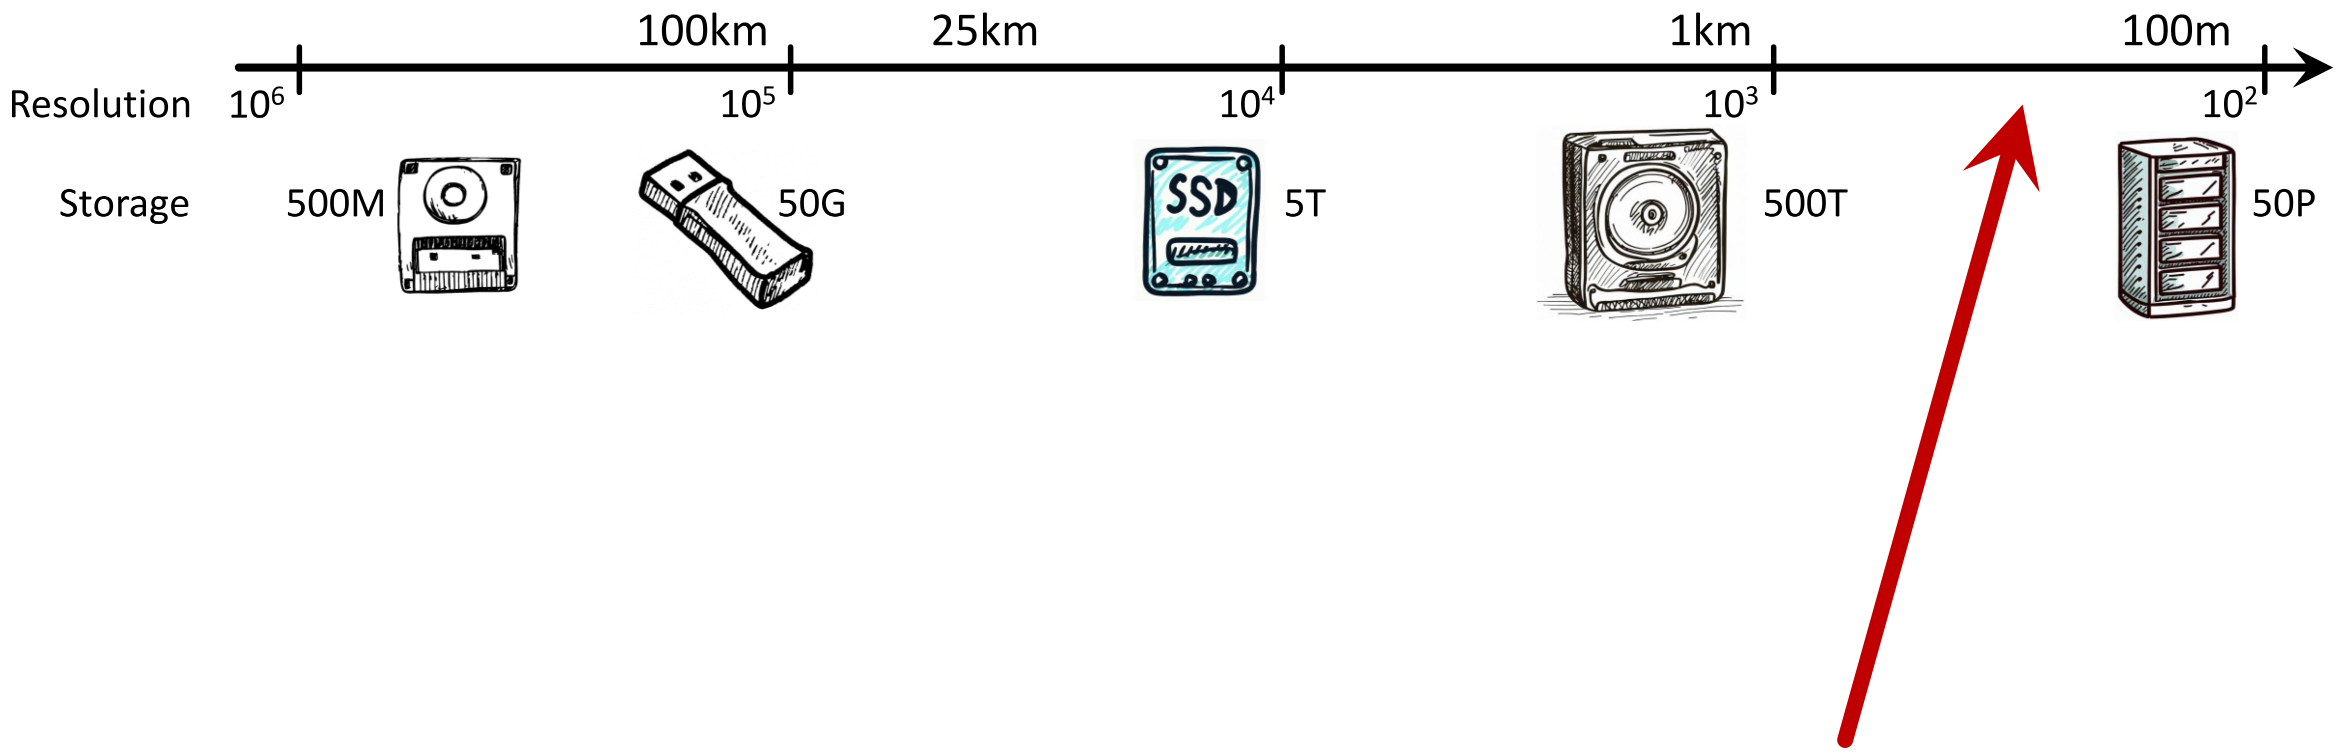

Δx = 550 m (10,000x)

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

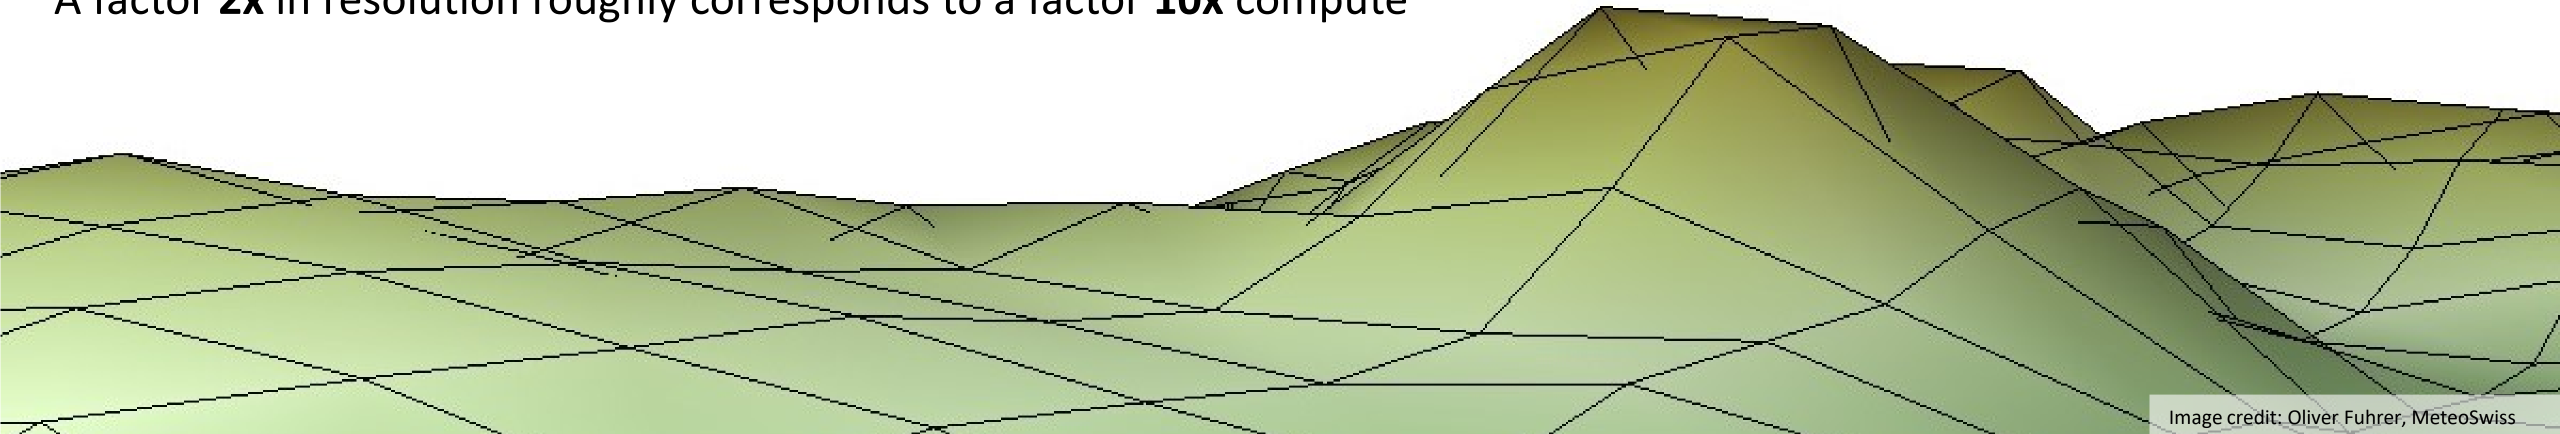

Image credit: Oliver Fuhrer, MeteoSwiss

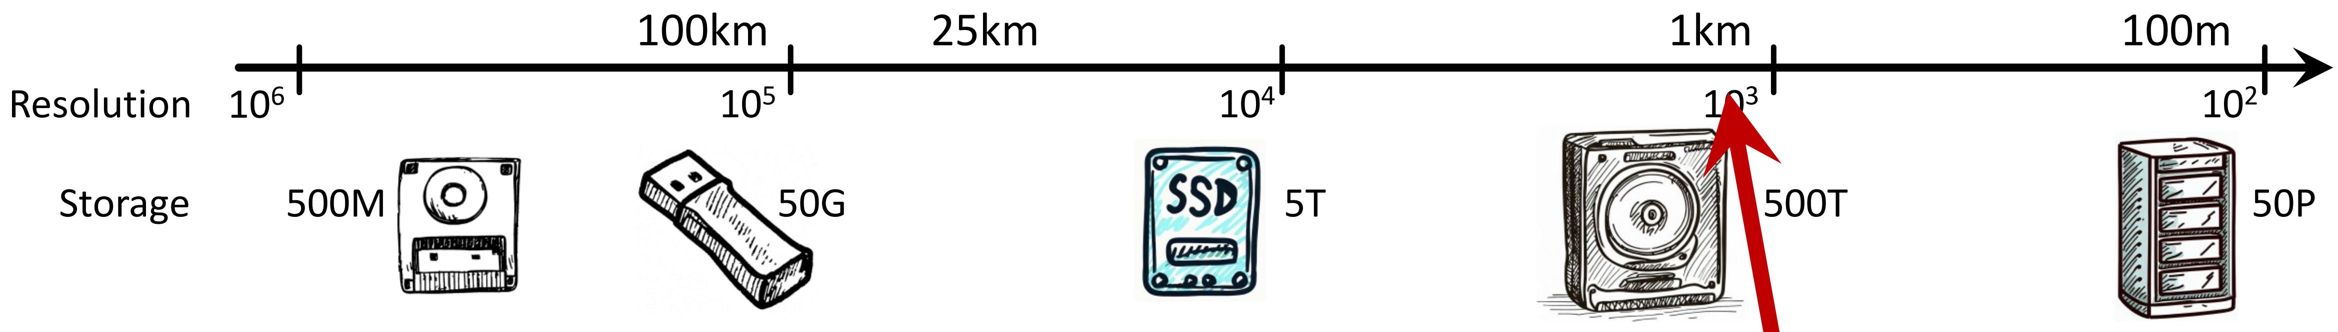

Operational weather model of MeteoSwiss today!

Δx = 1100 m (100,000x)

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

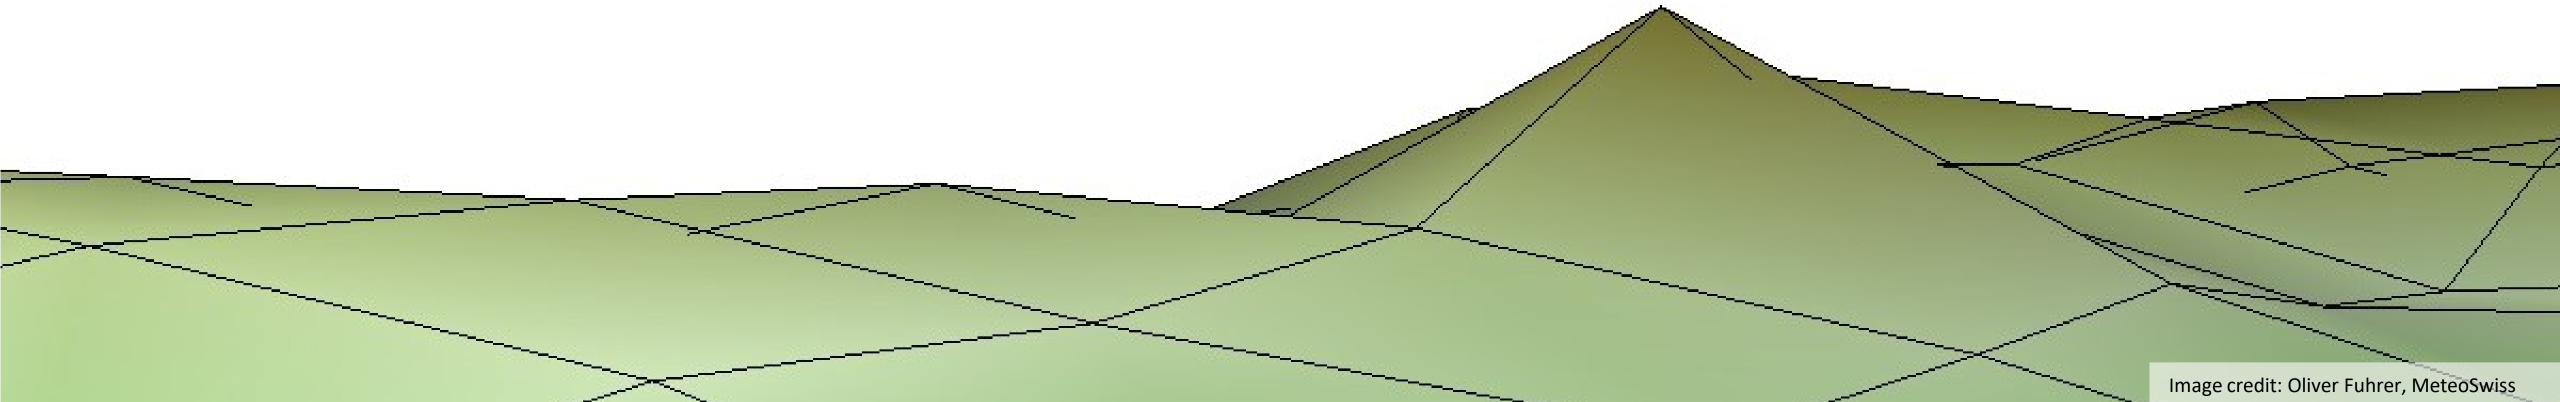

Image credit: Oliver Fuhrer, MeteoSwiss

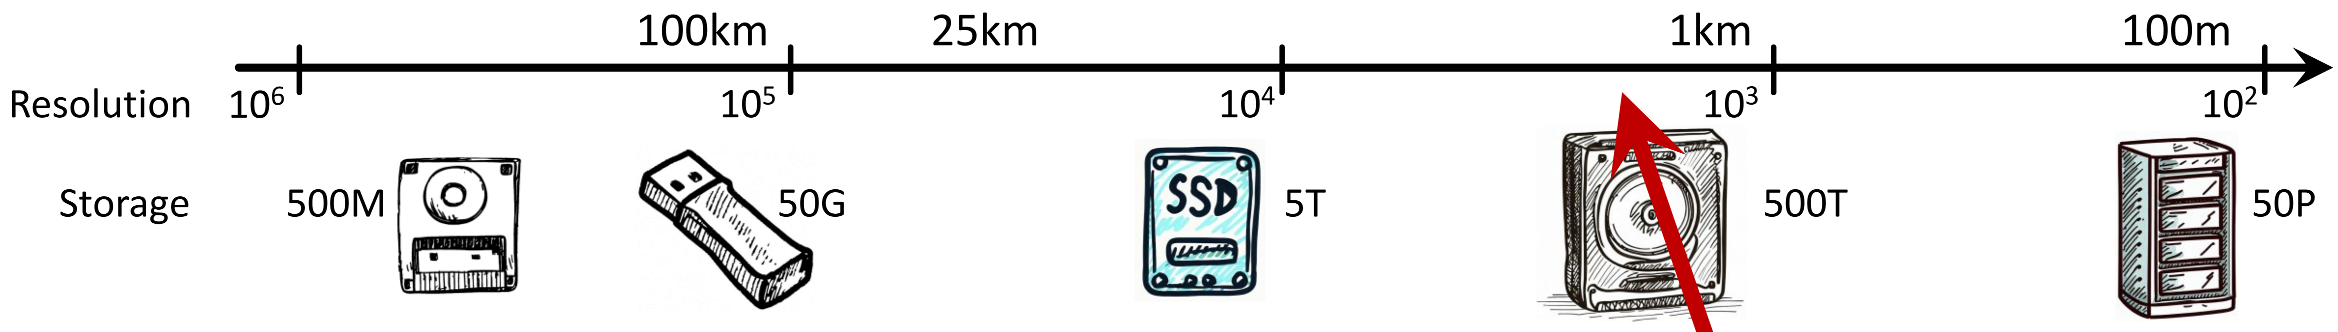

Swiss to simulate weather using GPUs

SCIENTIFIC COMPUTING WORLD

2015

The Swiss Federal Office of Meteorology and Climatology (MeteoSwiss) has announced that it has taken delivery of the first GPU-accelerated supercomputer used to power the numerical weather forecasts.

Tech > Computing

CES 2019: Moore's Law is dead, says Nvidia's CEO

The long-held notion that the processing power of computers increases exponentially every couple of years has hit its limit, according to Jensen Huang.

Shara Tibken
Jan. 9, 2019 11:46 a.m. PT
3 min read

Δx = **2200 m** (1,000,000x)

A factor **2x** in resolution roughly corresponds to a factor **10x** compute

Operational weather model of MeteoSwiss before 2016!

TECHNOLOGY

Nvidia CEO Says AI Can Overcome the Death of Moore's Law

By Tae Kim Follow
March 22, 2023, 1:35 pm EDT

Image credit: Oliver Fuhrer, MeteoSwiss

A Changing High-P…

Brian Lamacraft · Jan 5 · 2 min read · Member-only

Chat GTP-4 Could Pass the Bar Exam

How Our Technology Evolves FAST

Source: https://medium.com/

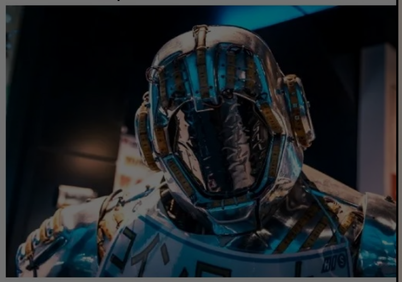

Photo by Maximalfocus on Unsplash

AI chatbot's MBA ex… business schools

ChatGPT earned a solid grade and out…

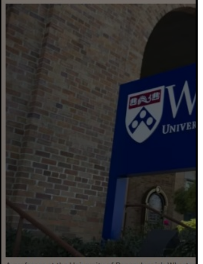

A professor at the University of Pennsylvania's Wharton … © Bloomberg

Source…

Andrew Jack in New Yo… JANUARY 21 2023

Bloomberg | Subscribe

Technology | AI

ChatGPT to Fuel $1.3 Trillion AI Market by 2032, New Report Says

- Bloomberg Intelligence expects generative AI market to soar
- Amazon, Microsoft, Google and Nvidia seen as biggest winners

By Jake Rudnitsky
June 1, 2023 at 3:00 PM GMT+2
Updated on June 1, 2023 at 5:50 PM GMT+2

Microsoft invests $1 billion in OpenAI to pursue holy grail of artificial intelligence

Building artificial general intelligence is OpenAI's ambitious goal

By James Vincent | Jul 22, 2019, 10:08am EDT

…world's new

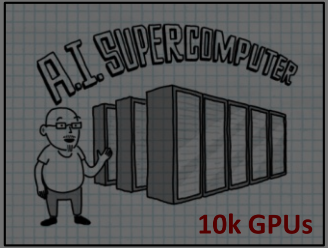

reddit

A robot may __ injure a … to harm.

	0.74	not	1.00
	0.28	sometimes	0.00
always	0.07	always	0.00

layer-wise weight update

Deep Learning Drives Future Computing Architectures!

Small datatypes
(int + fp – 4, 8, 16 bits)
Example: fp8 is 33x faster than fp64

Matrix and vector ops
(tensor cores and vector units)
Example: NVIDIA TCs: 8.3x over CUDA cores

(Structured) Sparsity
(in tensor cores and vector units)
Example: NVIDIA TCs: 2:4 sparsity

T. Ben-Nun, TH: Demystifying parallel and distributed deep learning: An in-depth concurrency analysis, ACM Computing Surveys (CSUR), 2019

Embracing the future of accelerated computation

Accelerate Simulations on ML/AI Hardware

- Systems challenges:
  - Enable efficient GPU support
  - Optimize for data movement
- Algorithm challenges:
  - Make simulations look like low-precision matrix multiplication!

Use ML/AI Techniques (“ML inside/on top”)

- Use ML models to replace simulations completely
  - E.g., GraphCast, FourCastNet, PanGu, FuXi, …
- Use ML to replace parts of the workflow
  - E.g., physics parametrization in simulations, data post processing, analyses, …

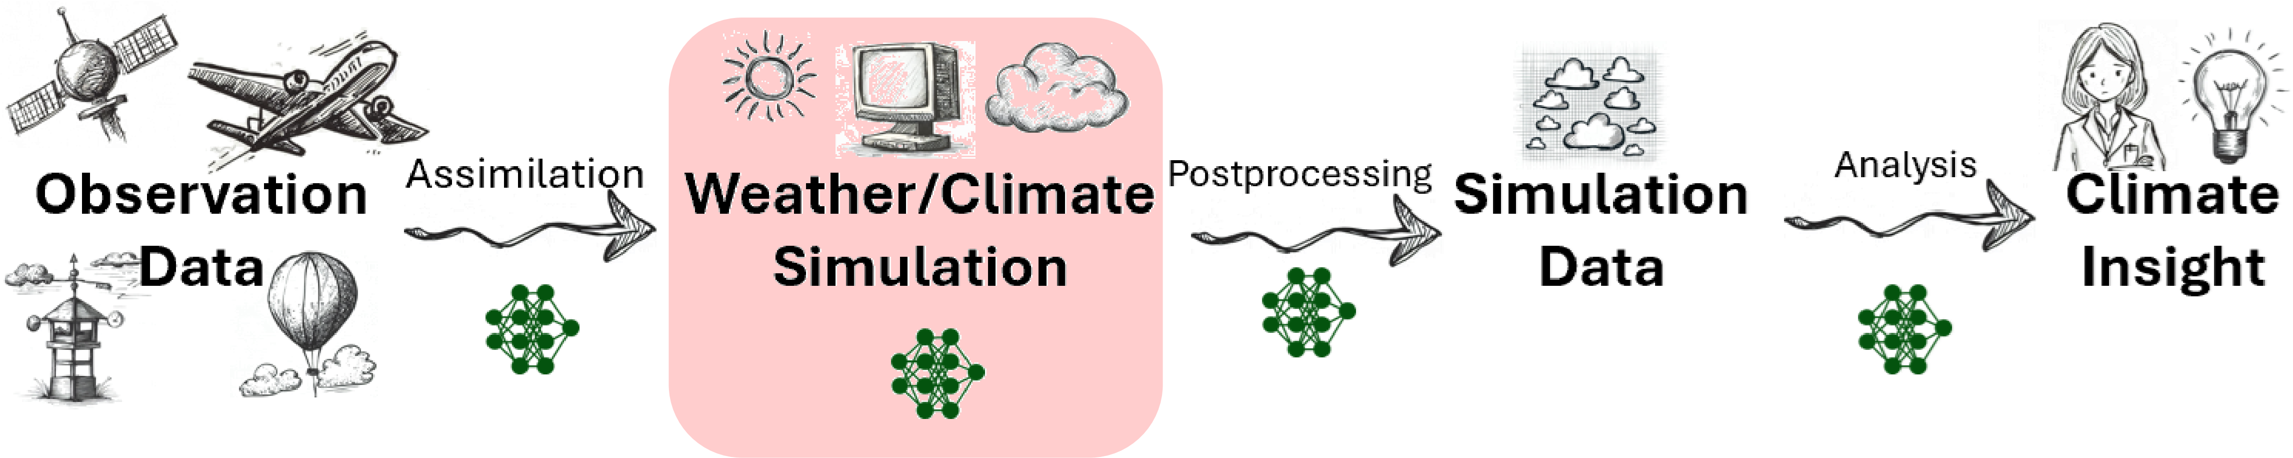

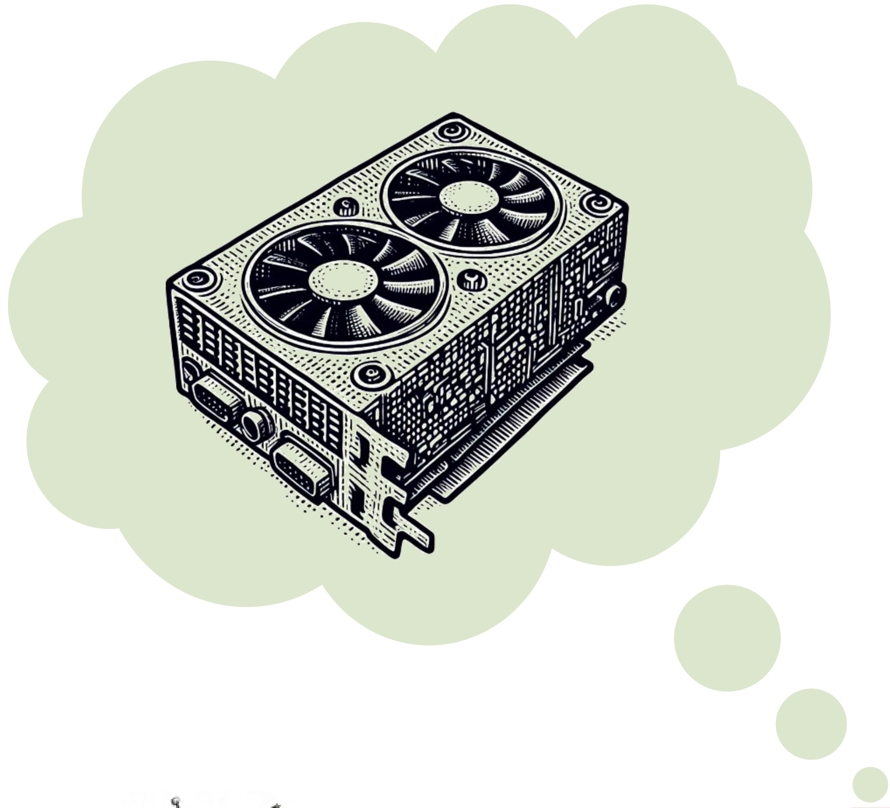

```
!$ACC DATA &
!$ACC PRESENT(density1,energy1) &
!$ACC PRESENT(vol_flux_x,vol_flux_y,volume,mass_flux_x,mass_flux_y,vertexdx,vertexdy) &
!$ACC PRESENT(pre_vol,post_vol,ener_flux)

!$ACC KERNELS

  IF(dir.EQ.g_xdir) THEN

   IF(sweep_number.EQ.1)THEN

!$ACC LOOP INDEPENDENT
     DO k=y_min-2,y_max+2
!$ACC LOOP INDEPENDENT
       DO j=x_min-2,x_max+2
         pre_vol(j,k)=volume(j,k)+(vol_flux_x(j+1,k  )-vol_flux_x(j,k)+vol_flux_y(j  ,k+1)-vol_flux_y(j,k))
         post_vol(j,k)=pre_vol(j,k)-(vol_flux_x(j+1,k  )-vol_flux_x(j,k))
       ENDDO
     ENDDO
   ELSE
!$ACC LOOP INDEPENDENT
     DO k=y_min-2,y_max+2
!$ACC LOOP INDEPENDENT
       DO j=x_min-2,x_max+2
         pre_vol(j,k)=volume(j,k)+vol_flux_x(j+1,k)-vol_flux_x(j,k)
         post_vol(j,k)=volume(j,k)
       ENDDO
     ENDDO

   ENDIF
```

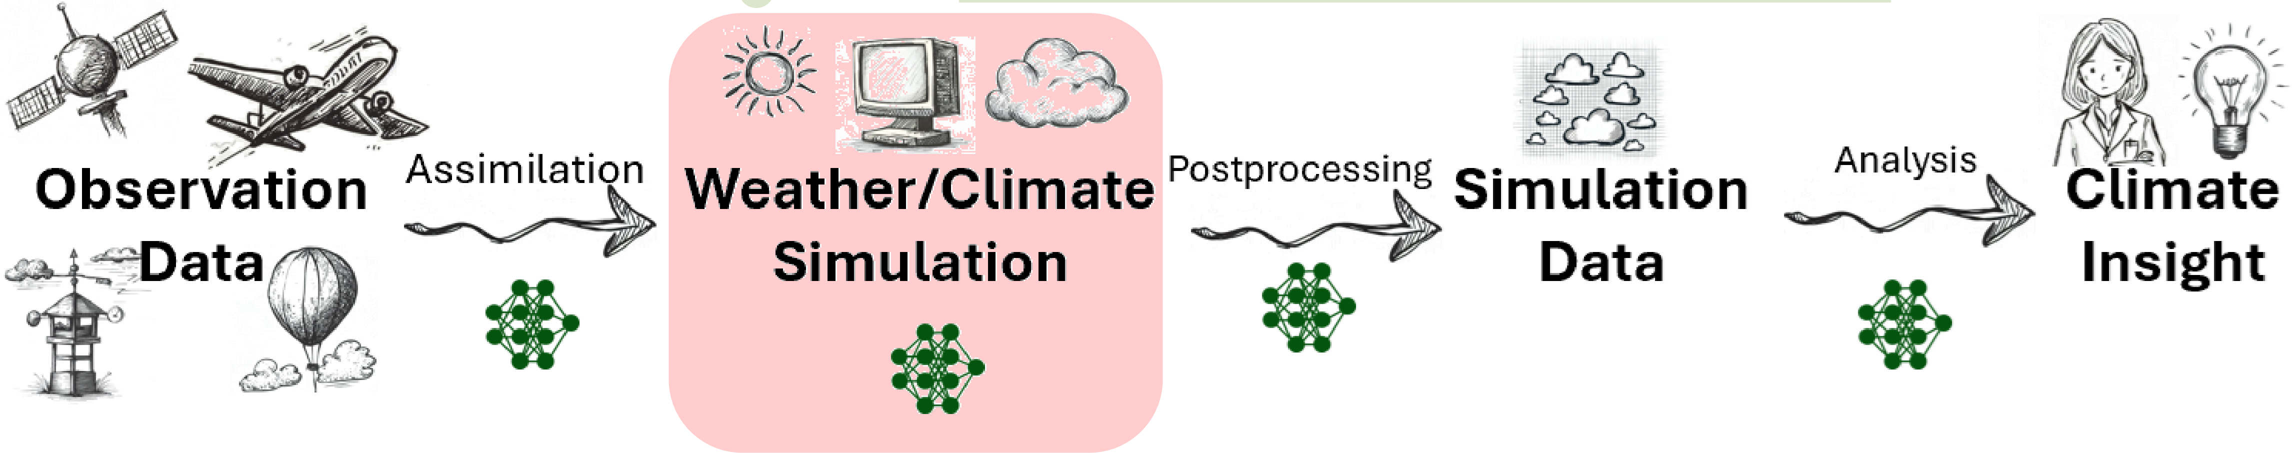

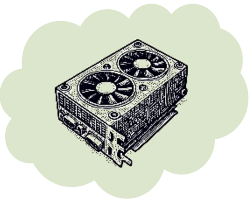

```
!$ACC DATA &
   !$ACC COPY(chunk%tiles(1)%field%density0)   &
   !$ACC COPY(chunk%tiles(1)%field%density1)   &
   !$ACC COPY(chunk%tiles(1)%field%energy0)    &
   !$ACC COPY(chunk%tiles(1)%field%energy1)    &
   !$ACC COPY(chunk%tiles(1)%field%pressure)   &
   !$ACC COPY(chunk%tiles(1)%field%soundspeed) &
   !$ACC COPY(chunk%tiles(1)%field%viscosity)  &
   !$ACC COPY(chunk%tiles(1)%field%xvel0)      &
   !$ACC COPY(chunk%tiles(1)%field%yvel0)      &
   !$ACC COPY(chunk%tiles(1)%field%xvel1)      &
   !$ACC COPY(chunk%tiles(1)%field%yvel1)      &
   !$ACC COPY(chunk%tiles(1)%field%vol_flux_x) &
   !$ACC COPY(chunk%tiles(1)%field%vol_flux_y) &
   !$ACC COPY(chunk%tiles(1)%field%mass_flux_x)&
   !$ACC COPY(chunk%tiles(1)%field%mass_flux_y)&
   !$ACC COPY(chunk%tiles(1)%field%volume)     &
   !$ACC COPY(chunk%tiles(1)%field%work_array1)&
   !$ACC COPY(chunk%tiles(1)%field%work_array2)&
   !$ACC COPY(chunk%tiles(1)%field%work_array3)&
   !$ACC COPY(chunk%tiles(1)%field%work_array4)&
   !$ACC COPY(chunk%tiles(1)%field%work_array5)&
   !$ACC COPY(chunk%tiles(1)%field%work_array6)&
   !$ACC COPY(chunk%tiles(1)%field%work_array7)&
   !$ACC COPY(chunk%tiles(1)%field%cellx)      &
   !$ACC COPY(chunk%tiles(1)%field%celly)      &
   !$ACC COPY(chunk%tiles(1)%field%celldx)     &
   !$ACC COPY(chunk%tiles(1)%field%celldy)     &
   !$ACC COPY(chunk%tiles(1)%field%vertexx)    &
   !$ACC COPY(chunk%tiles(1)%field%vertexdx)   &
   !$ACC COPY(chunk%tiles(1)%field%vertexy)    &
   !$ACC COPY(chunk%tiles(1)%field%vertexdy)   &
   !$ACC COPY(chunk%tiles(1)%field%xarea)      &
   !$ACC COPY(chunk%tiles(1)%field%yarea)      &
   !$ACC COPY(chunk%left_snd_buffer)   &
   !$ACC COPY(chunk%left_rcv_buffer)   &
   !$ACC COPY(chunk%right_snd_buffer)  &
   !$ACC COPY(chunk%right_rcv_buffer)  &
   !$ACC COPY(chunk%bottom_snd_buffer) &
   !$ACC COPY(chunk%bottom_rcv_buffer) &
   !$ACC COPY(chunk%top_snd_buffer)    &
   !$ACC COPY(chunk%top_rcv_buffer)
```

```
,mass_flux_x,mass_flux_y,vertexdx,vertexdy) &
1,k
+1,l
l,k)
```

Sloccount *f90: 6,440

!$ACC: 833 (13%)

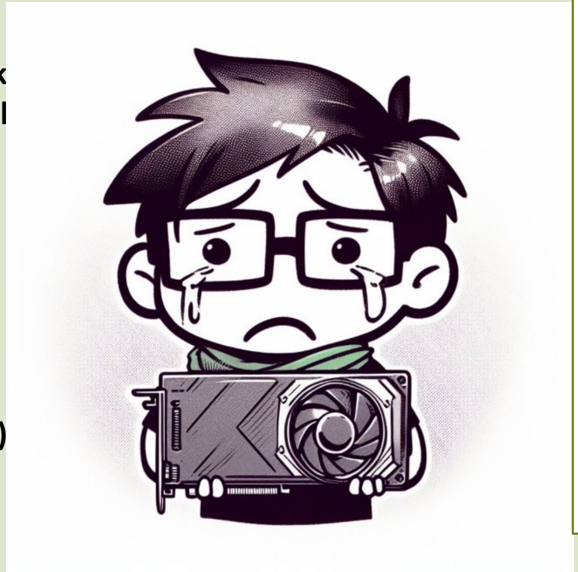

Heitlager et al.: A Practical Model for Measuring Maintainability

ISO 9126 maintainability	source code properties: volume	complexity per unit	duplication	unit size	unit testing
analysability	x		x	x	x
changeability		x	x		
stability					x
testability		x		x	x

ICON: about 1% OpenMP + 3% OpenACC

Performance Metaprogramming for Optimization and Performance Portability

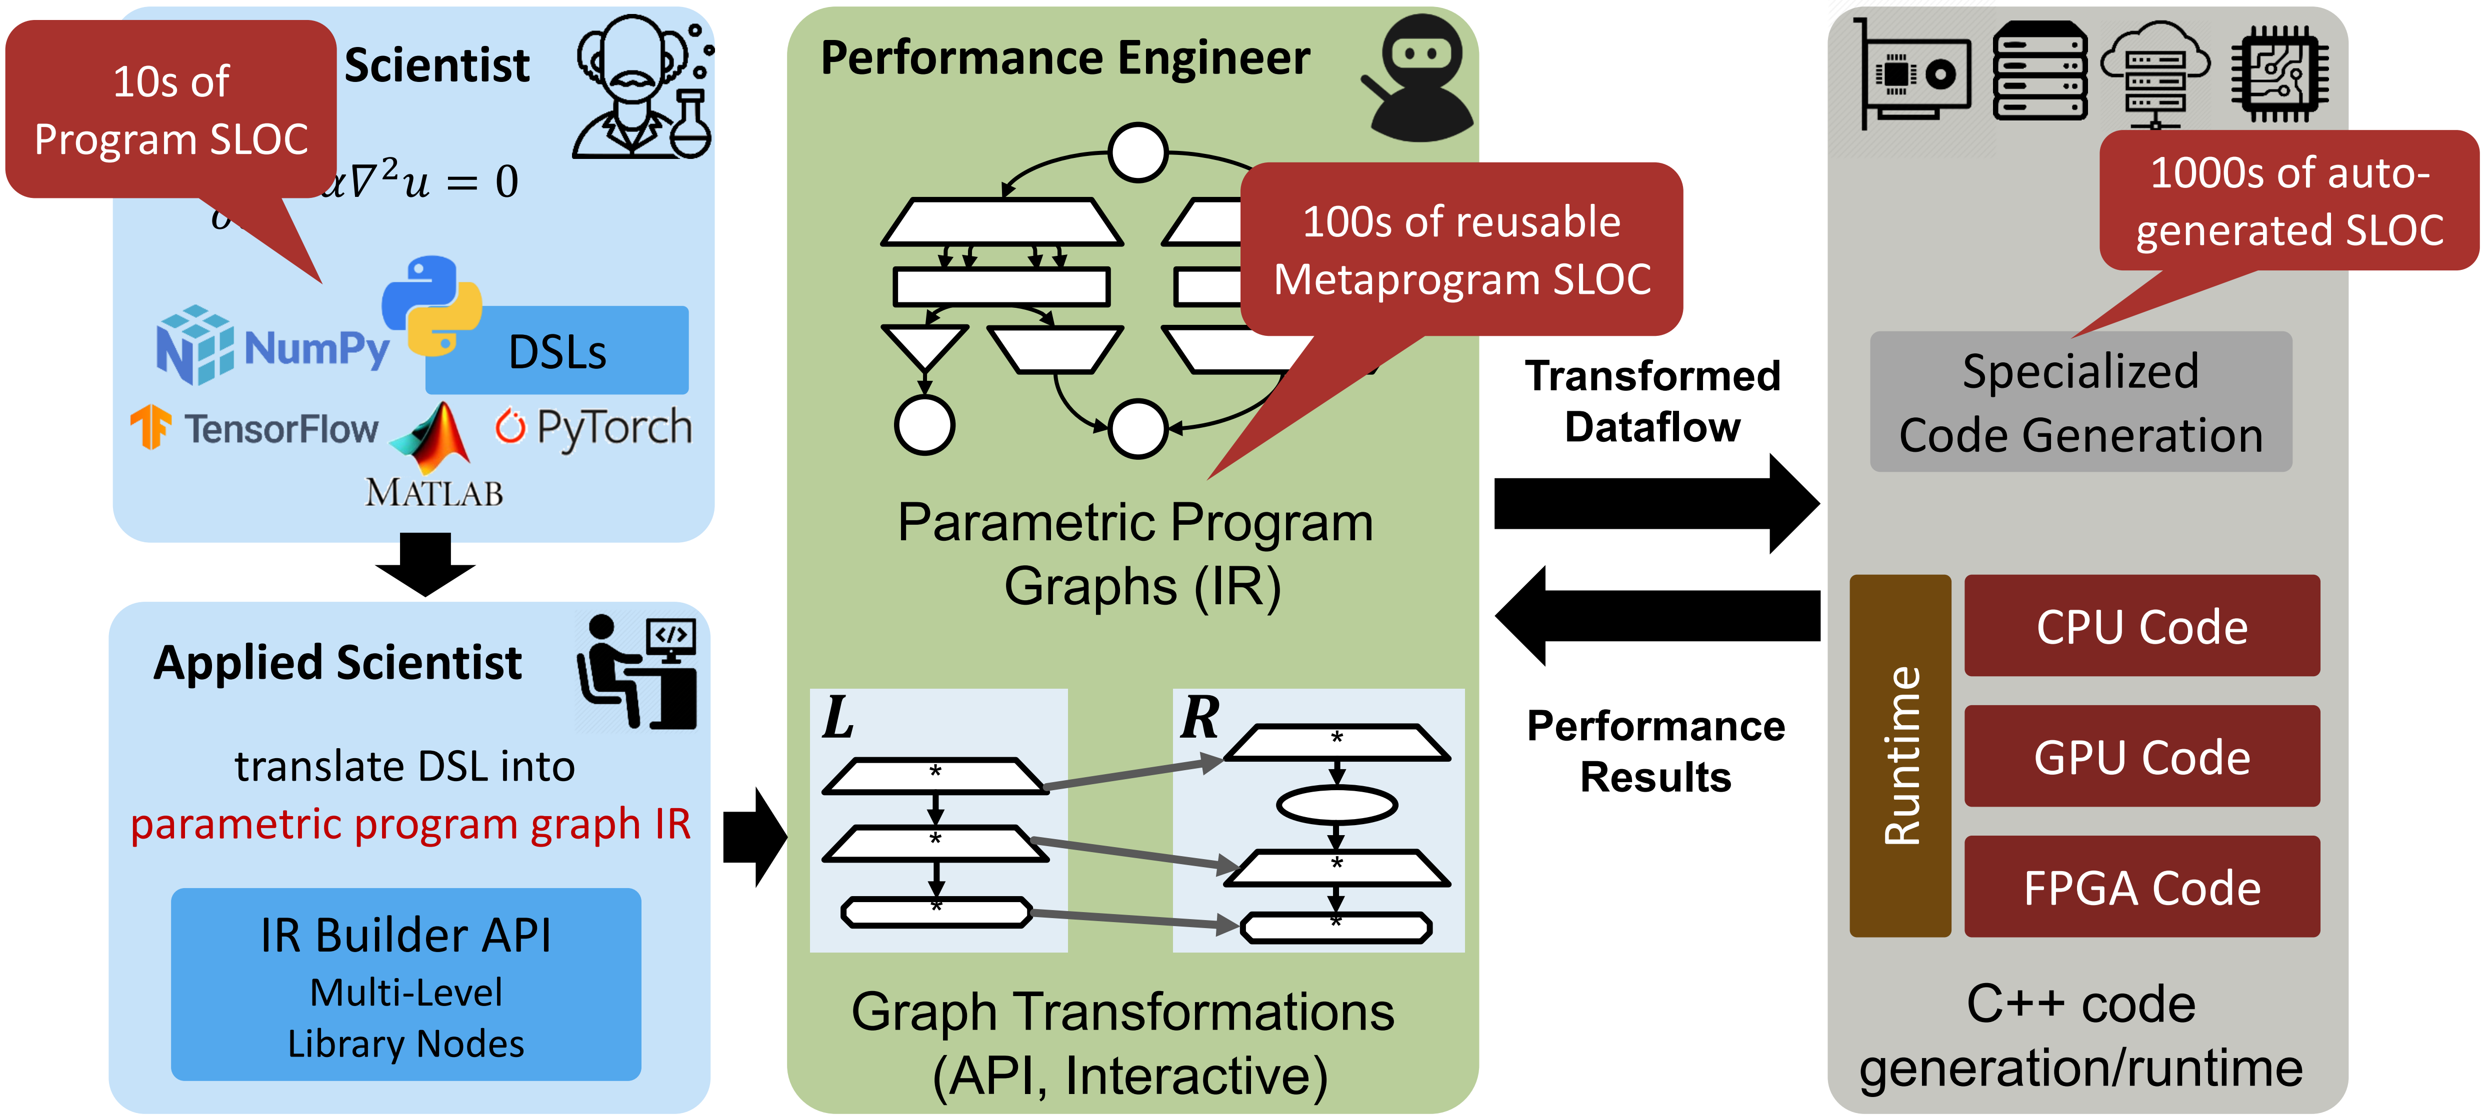

Ben-Nun, de Fine Licht, Ziogas, TH: Stateful Dataflow Multigraphs: A Data-Centric Model for High-Performance Parallel Programs, SC19

The Pace Project

https://github.com/ai2cm/pace

```
Usage: python -m pace.driver.run [OPTIONS] CONFIG_PATH

  Run the driver.

  CONFIG_PATH is the path to a DriverConfig yaml file.

Options:
  ...
```

- **NOOA's FV3 reimagined in Python**
  - Goal: Atmospheric model that can run at scale on modern supercomputers
  - No FORTRAN involved – move to 21st century programming + devops + package management (with similar syntax!)
- **Full dynamical core: 12,450 Python LoC across 36 modules**

  vs. 29,458 in the baseline implementation

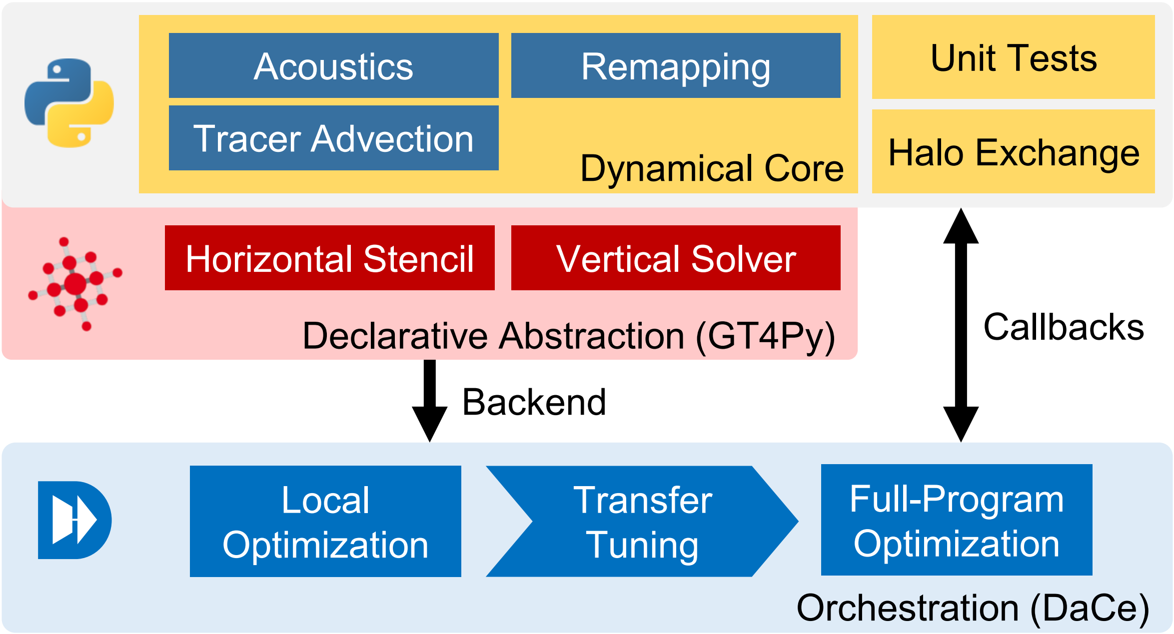

Building around GridTools/Gt4Py

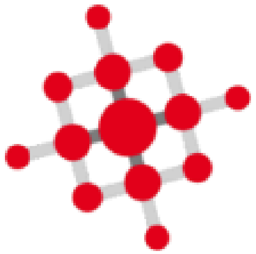

J. Dahm et al., "Pace v0.1: A Python-based Performance-Portable Implementation of the FV3 Dynamical Core". EGUSphere'22

Pace in DaCe for Performance Metaprogramming – 12k SLOC Python

AI-based Transfer Tuning to the Rescue!

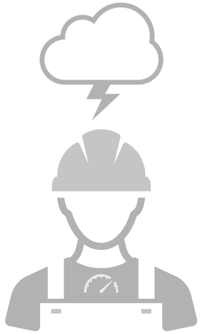

Ben-Nun et al.: Productive Performance Engineering for Weather and Climate Modeling with Python, SC22

Evaluated Systems

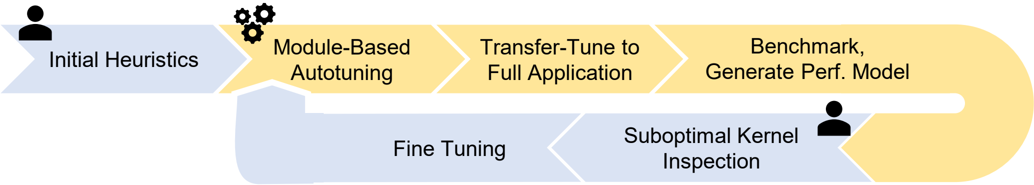

Photo courtesy of the Swiss National Supercomputing Centre

Piz Daint:

- GPU: 1 x NVIDIA Tesla P100 / Node
- CPU: Intel Xeon E5-2690 v3 (12 cores)

Per-node domain size: 192x192x80

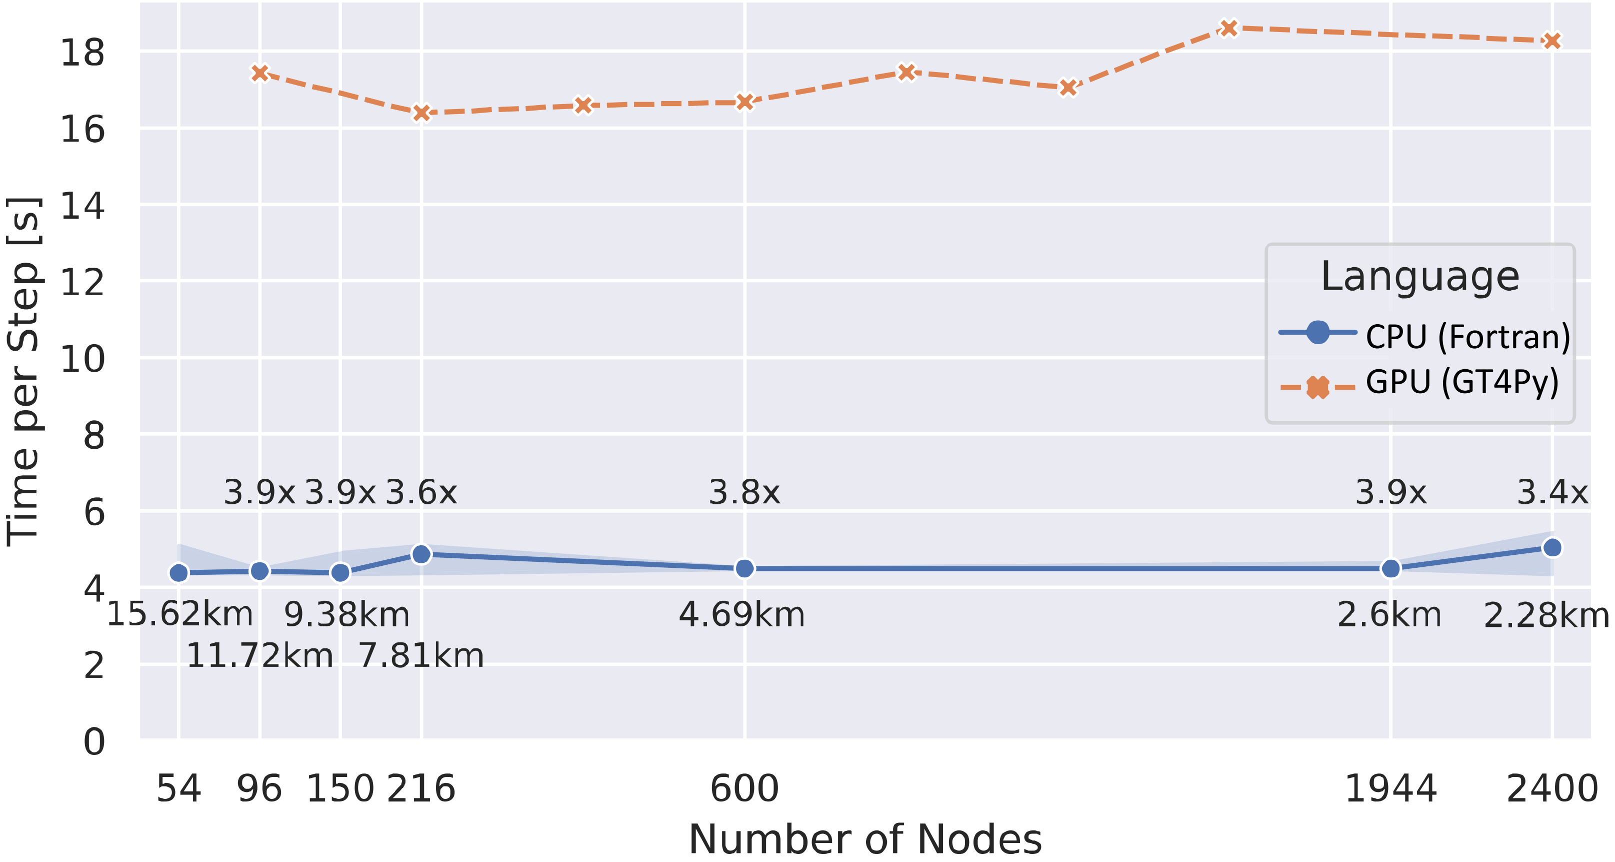

Simulation throughput of **0.12 SYPD** at 2.6 km grid spacing

https://github.com/ai2cm/pace
https://github.com/GridTools/gt4py
https://github.com/spcl/dace

That’s all nice but do we really want to rewrite all codes?

6 weeks of work	10 optimization revisions	4 performance engineers	3.92 – 8.48x speedup vs. production FORTRAN	0 model changes

youtube.com/@spcl

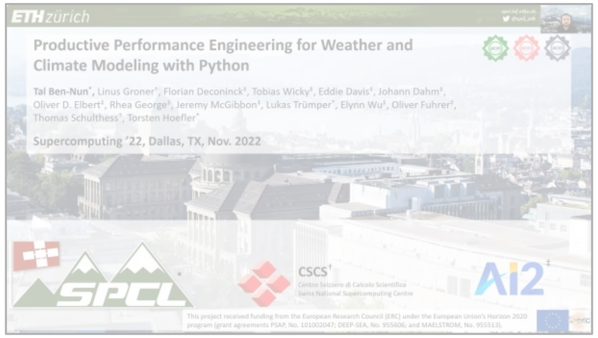

Another real production code … ECMWF's CLOUDSC

```fortran
9
10 SUBROUTINE CLOUDSC &
11  !---input
12  & (KIDIA,    KFDIA,    KLON,    KLEV, &
13  & PTSPHY,&
14  & PT, PQ, tendency_cml,tendency_tmp,tendency_loc, &
15  & PVFA, PVFL, PVFI, PDYNA, PDYNL, PDYNI, &
16  & PHRSW,    PHRLW,&
17  & PVERVEL,  PAP,      PAPH,&
18  & PLSM,     LDCUM,    KTYPE, &
19  & PLU,      PLUDE,    PSNDE,    PMFU,     PMFD,&
20  !---prognostic fields
21  & PA,&
22  & PCLV,  &
23  & PSUPSAT,&
24 !-- arrays for aerosol-cloud interactions
25 !!! & PQAER,    KAER, &
26  & PLCRIT_AER,PICRIT_AER,&
27  & PRE_ICE,&
28  & PCCN,     PNICE,&
29  !---diagnostic output
30  & PCOVPTOT, PRAINFRAC_TOPRFZ,&
31  !---resulting fluxes
32  & PFSQLF,   PFSQIF ,  PFCQNNG,  PFCQLNG,&
33  & PFSQRF,   PFSQSF ,  PFCQRNG,  PFCQSNG,&
34  & PFSQLTUR, PFSQITUR , &
35  & PFPLSL,   PFPLSN,   PFHPSL,   PFHPSN, KFLDX, &
36  & YDCST, YDTHF, YDECLDP)
```

… variable setup/initialization until line 500 ;-)

- **Cloud Microphysics of IFS**
  - Resolve sub-grid features
  - Original 2,525 SLOC of **Fortran 95**
- **Rewritten for performance portability benchmarking (optimization took months!)**
  - 2,635 SLOC C
  - 2,610 SLOC C++/CUDA

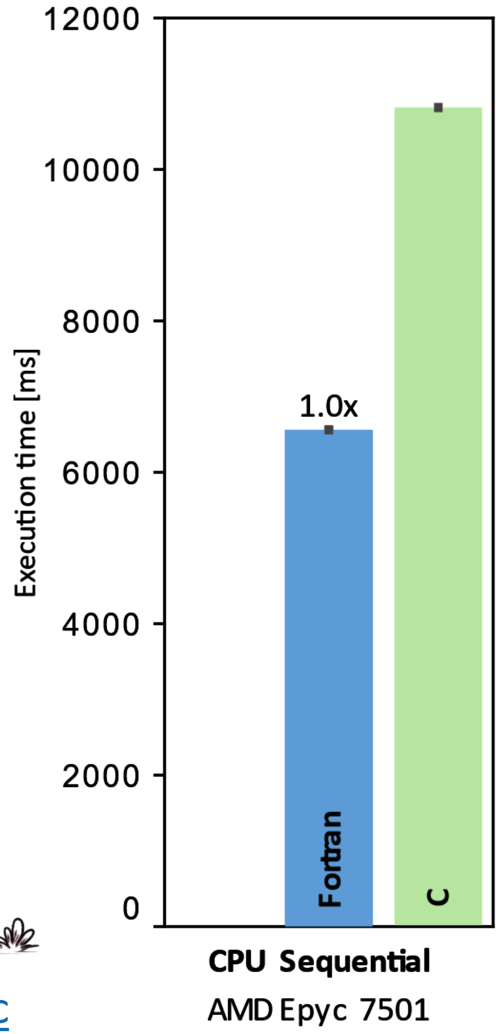

https://github.com/ecmwf-ifs/dwarf-p-cloudsc

Performance Metaprogramming – from the unchanged CLOUDSC Fortran code!

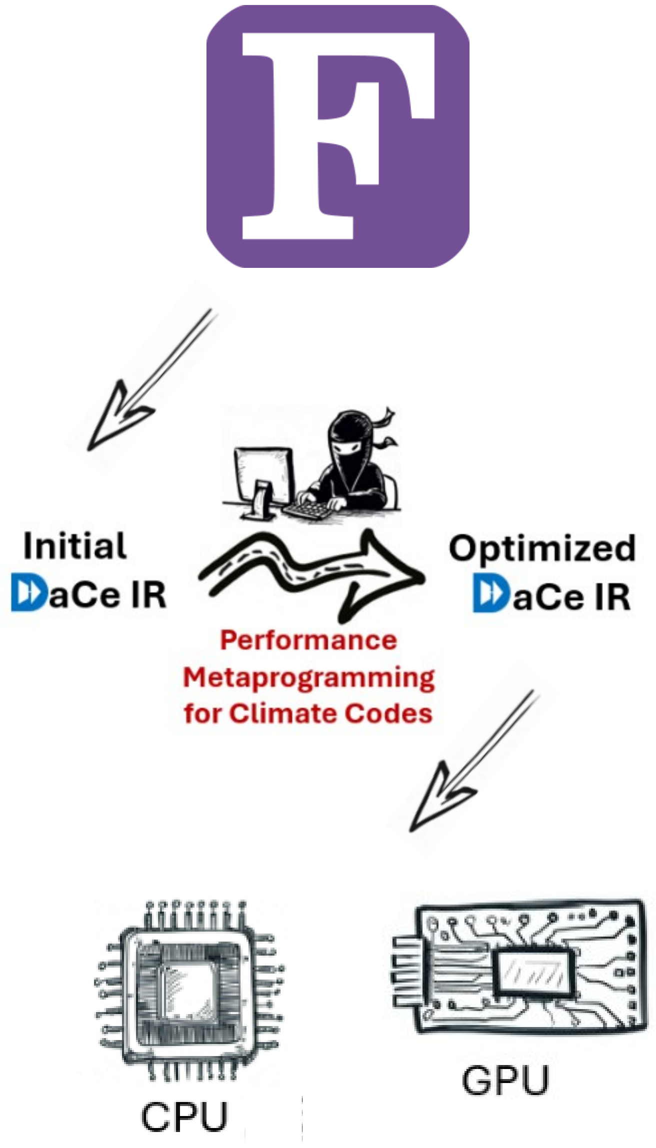

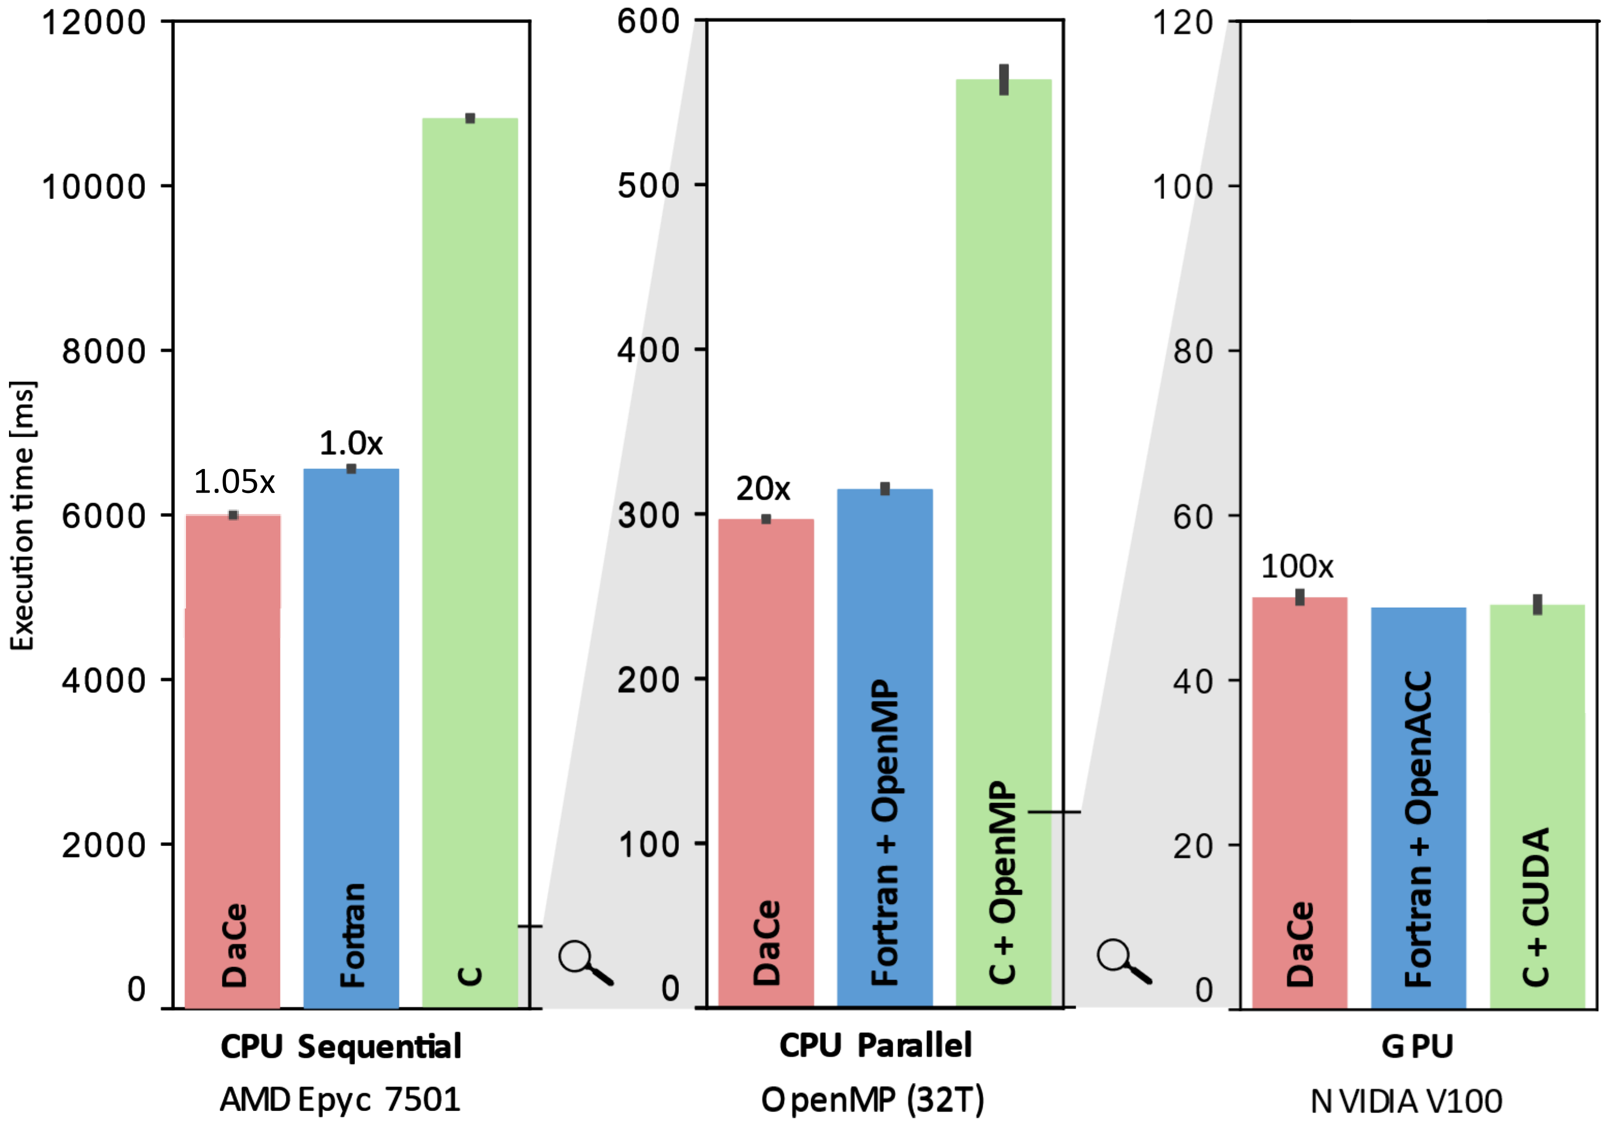

Embracing the future of accelerated computation

Accelerate Simulations on ML/AI Hardware

recisi

Use ML

- Use ML models to ... ely
  - E.g., GraphCas...
- Use ML to replace
  - E.g., physics pa... data post processing

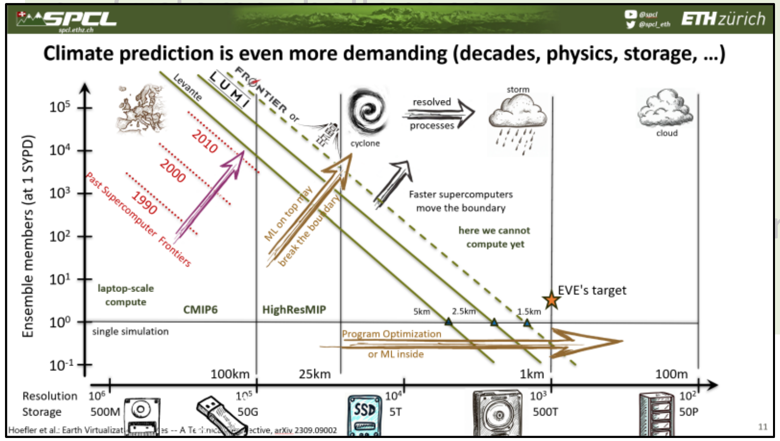

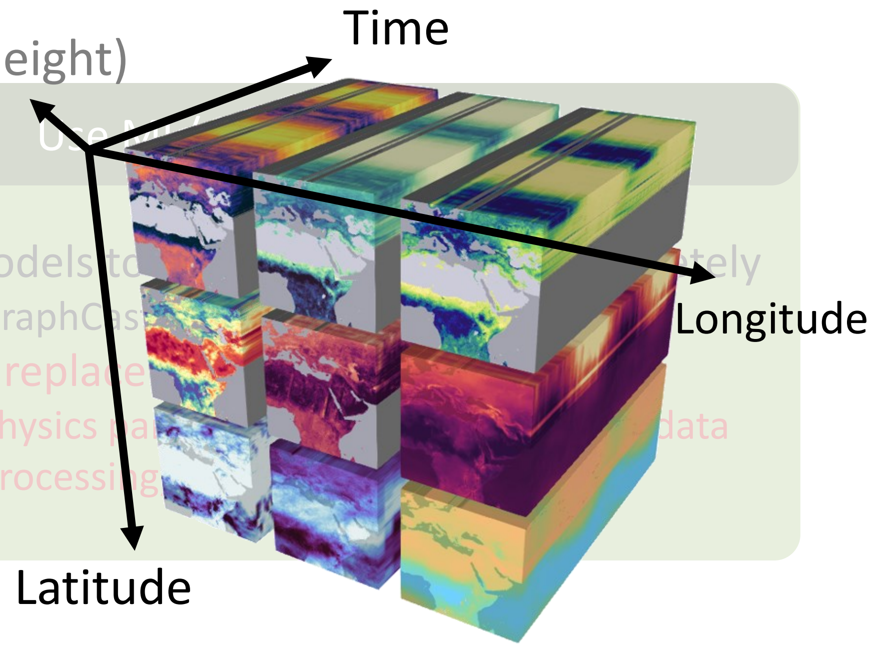

"ML on top"

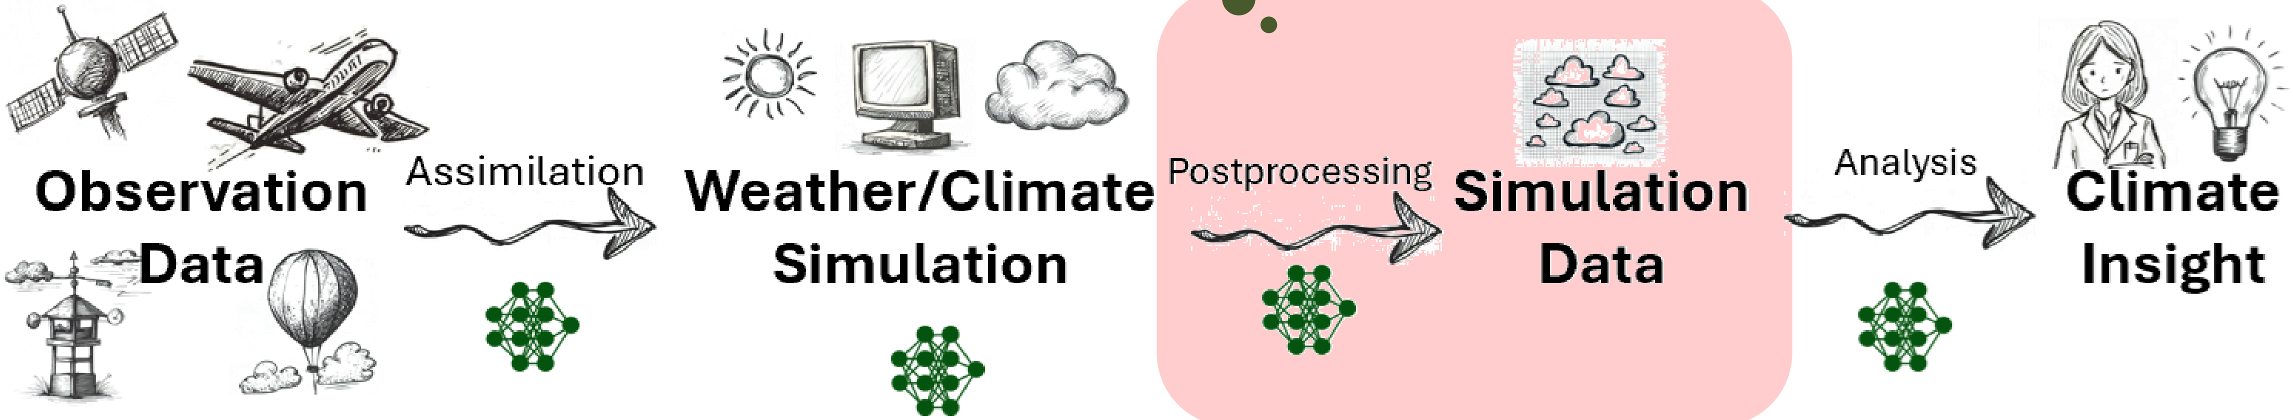

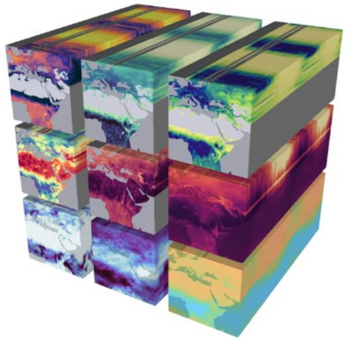

Simulation runs time-stepping forward

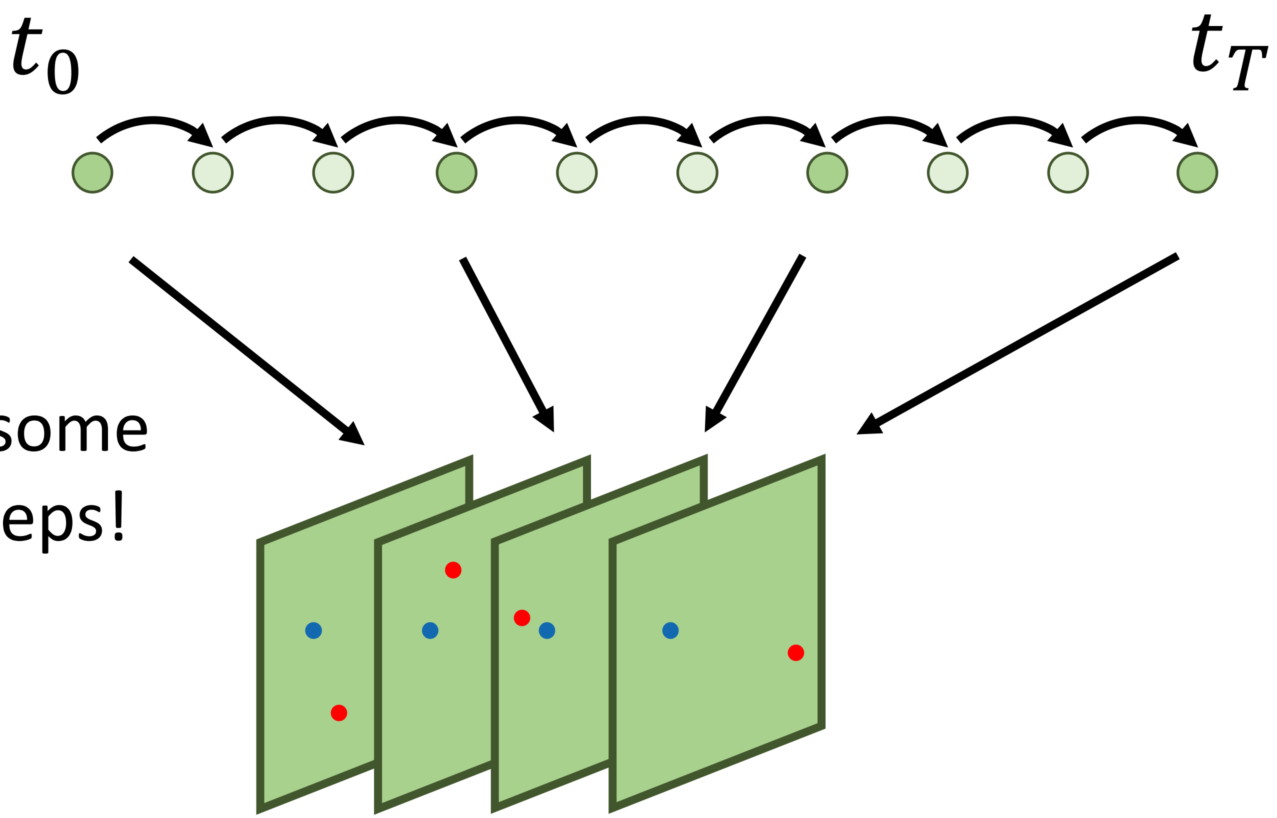

Store some timesteps!

Analysis access pattern is often strided or even random

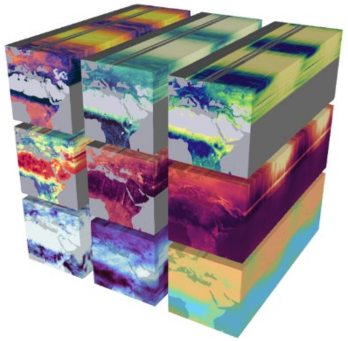

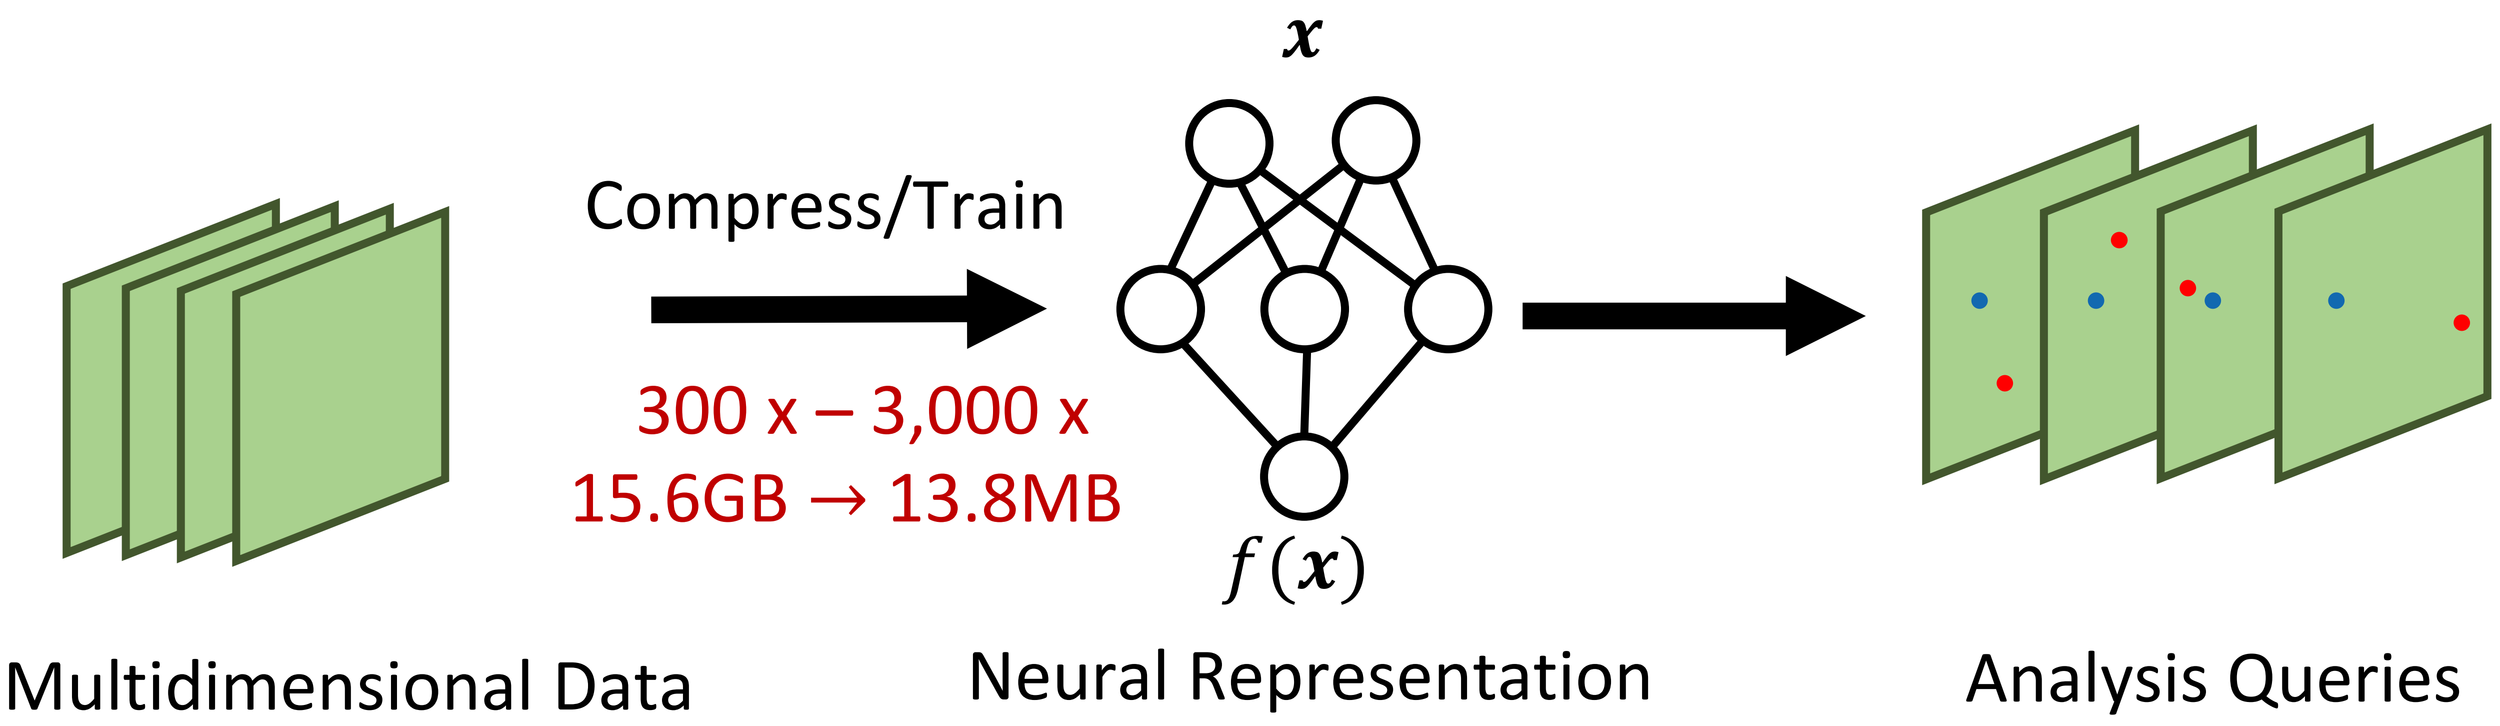
$\boldsymbol{x}$
Compress/Train
300 x – 3,000 x
15.6GB → 13.8MB
$f(\boldsymbol{x})$
Multidimensional Data
Neural Representation
Analysis Queries

Neural Network Structure

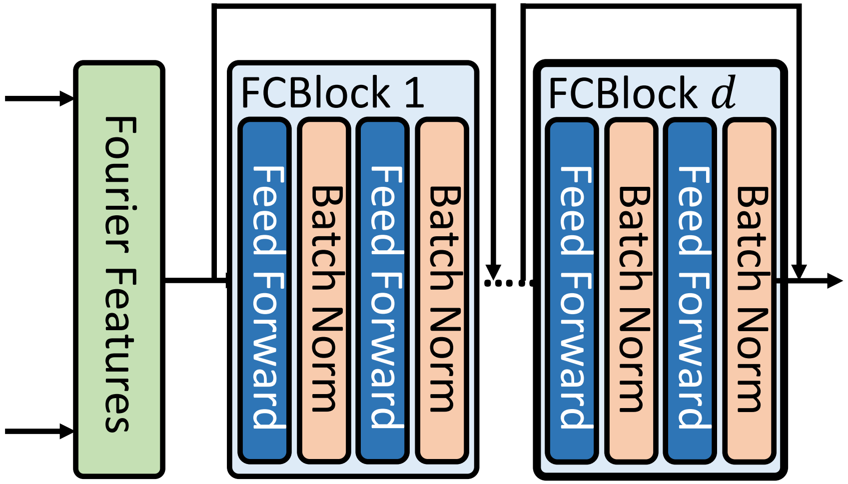

Neural Network Structure

Decompression / Inference

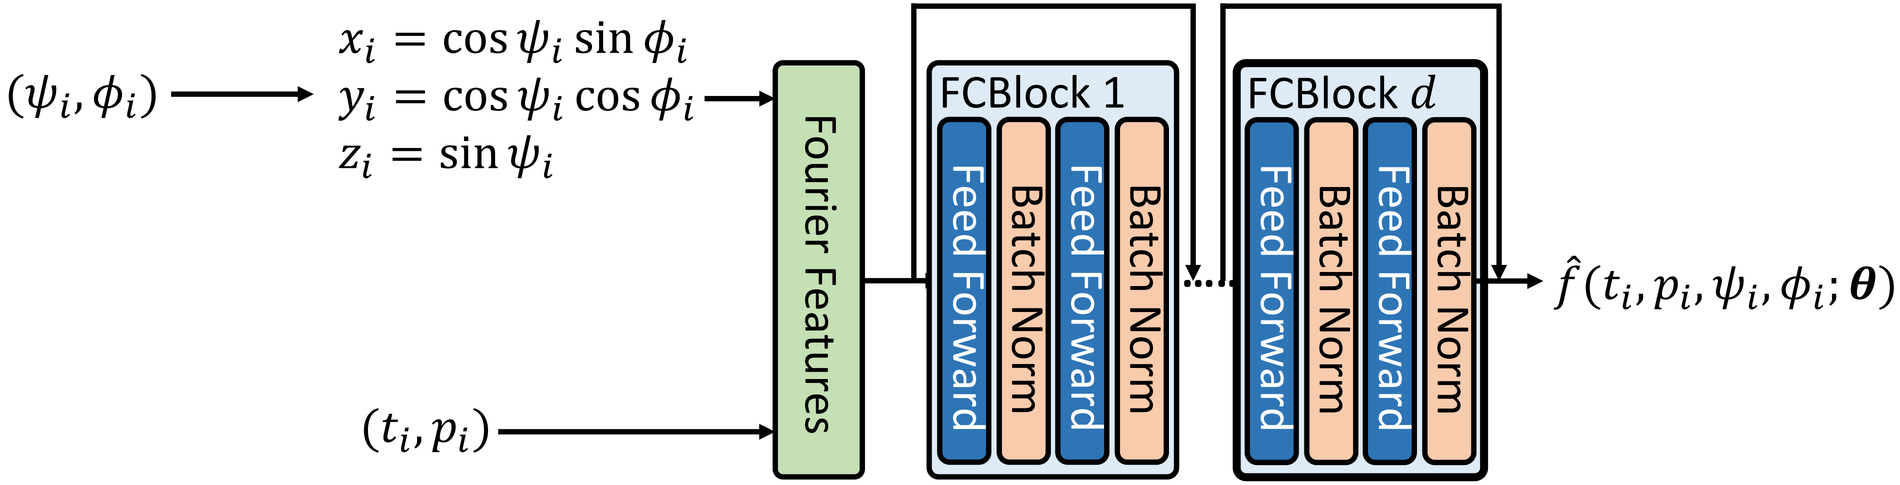

- On-demand decompression
- Fully utilize GPUs

Neural Network Structure

Compression / Training

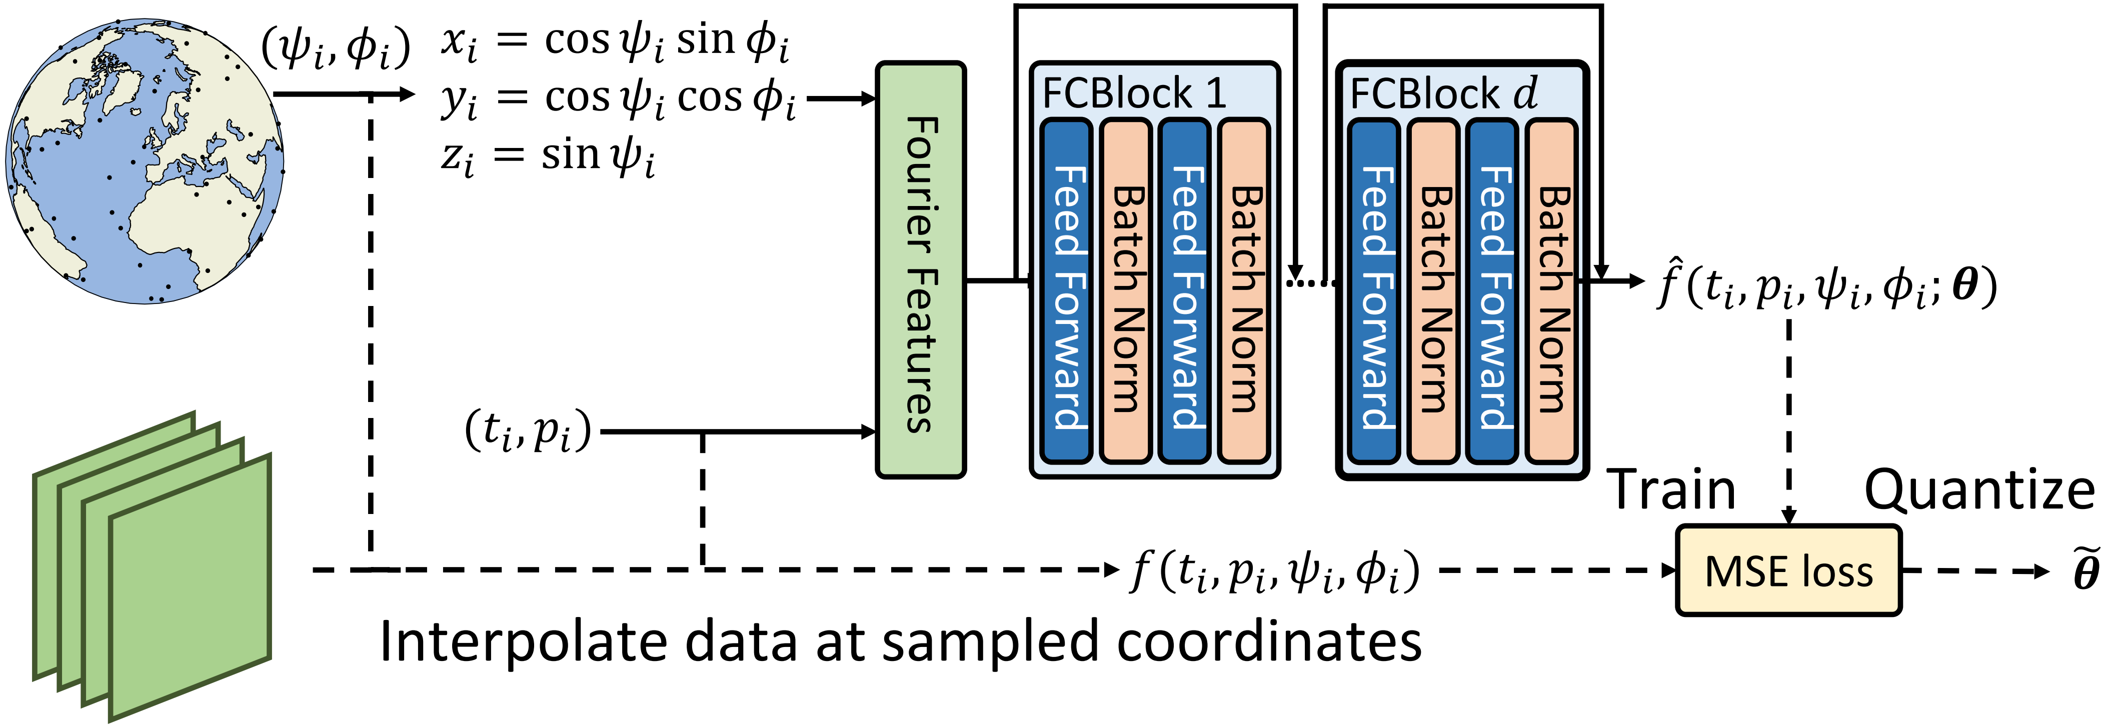

Evaluation: Case Study

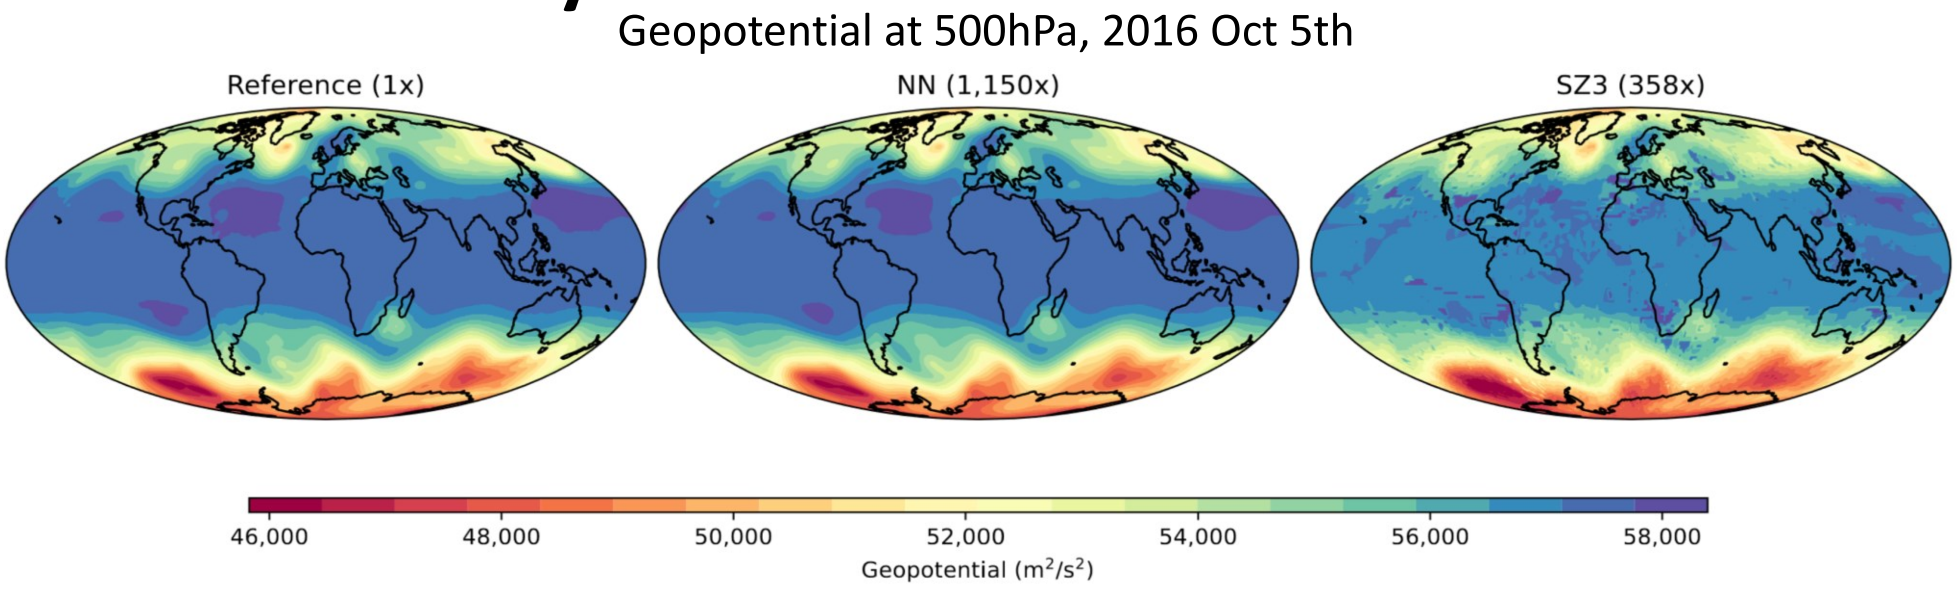

Evaluation: Case Study

Geopotential at 500hPa, 2016 Oct 5th

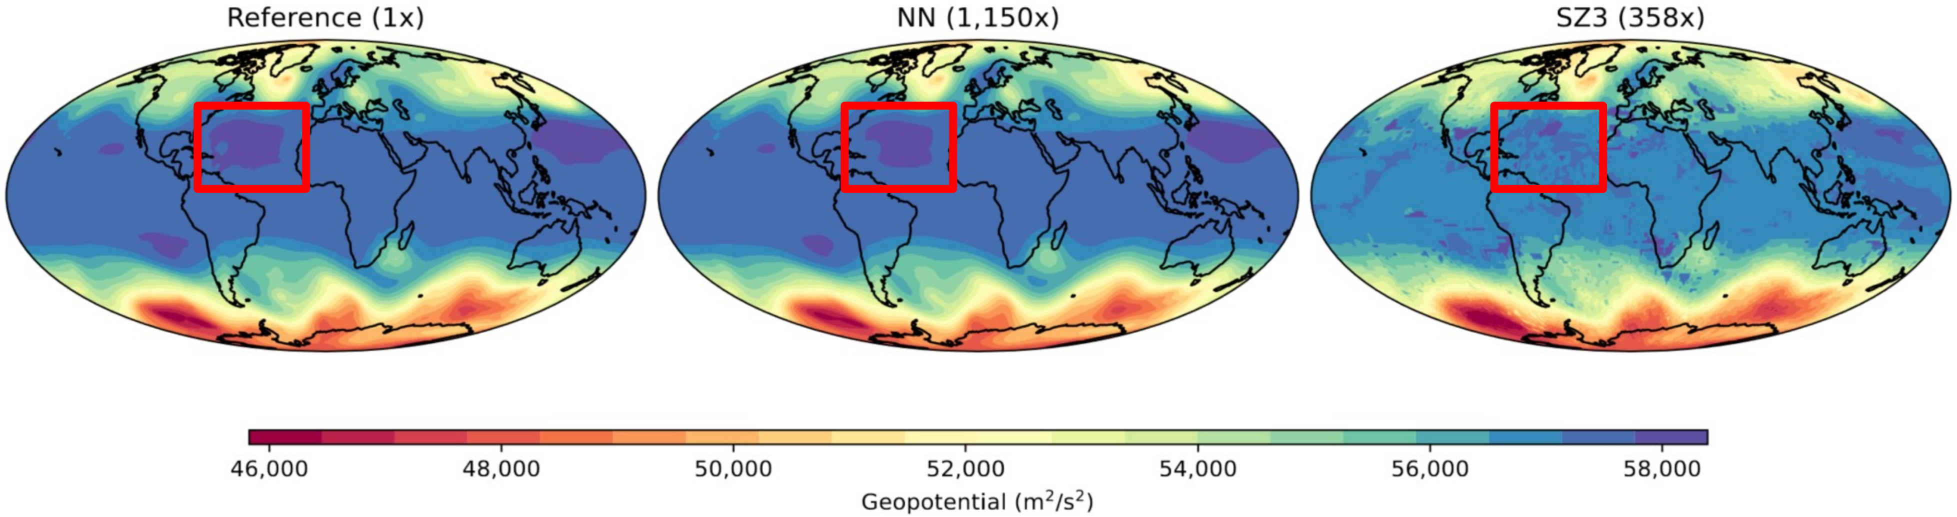

Preserves general shapes of important events and average values without introducing significant artifacts

Evaluation: Case Study

Geopotential at 500hPa, 2016 Oct 5th

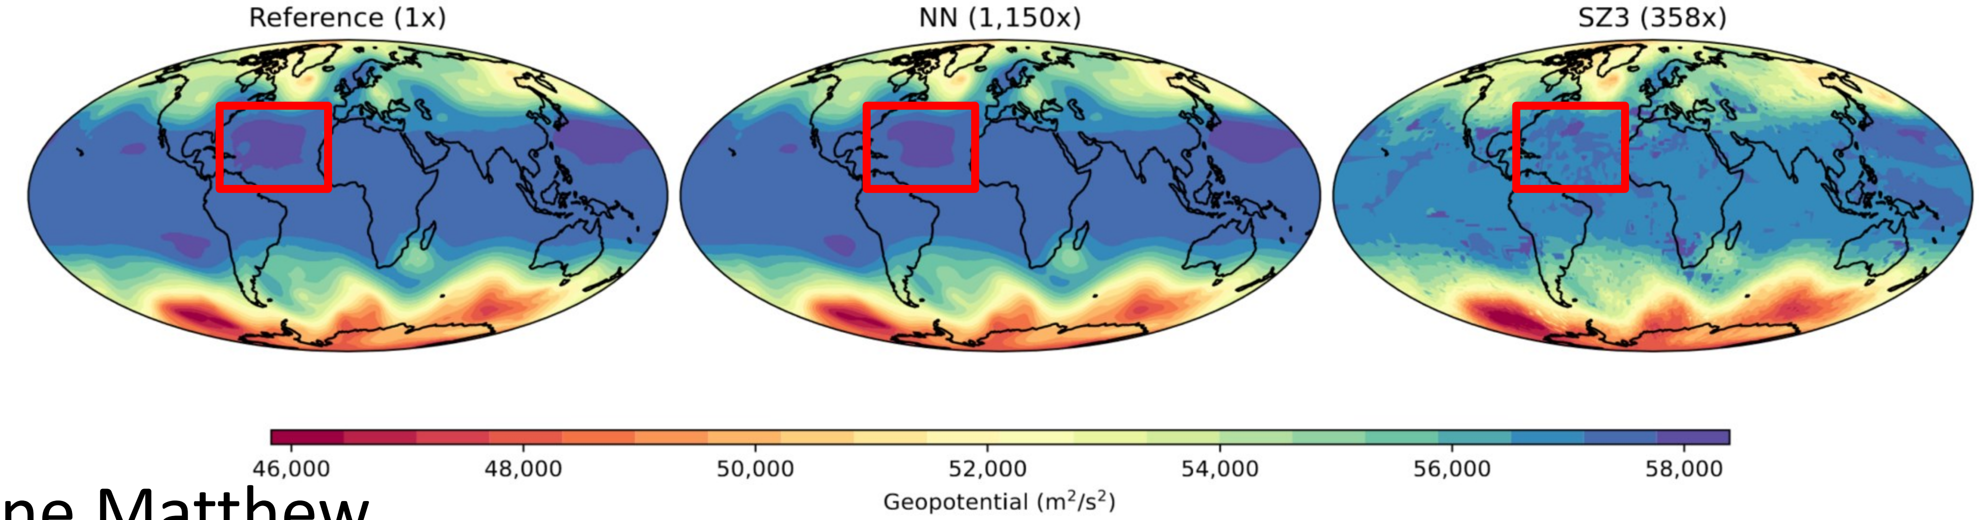

Hurricane Matthew

16.5bn damage
603 fatalities

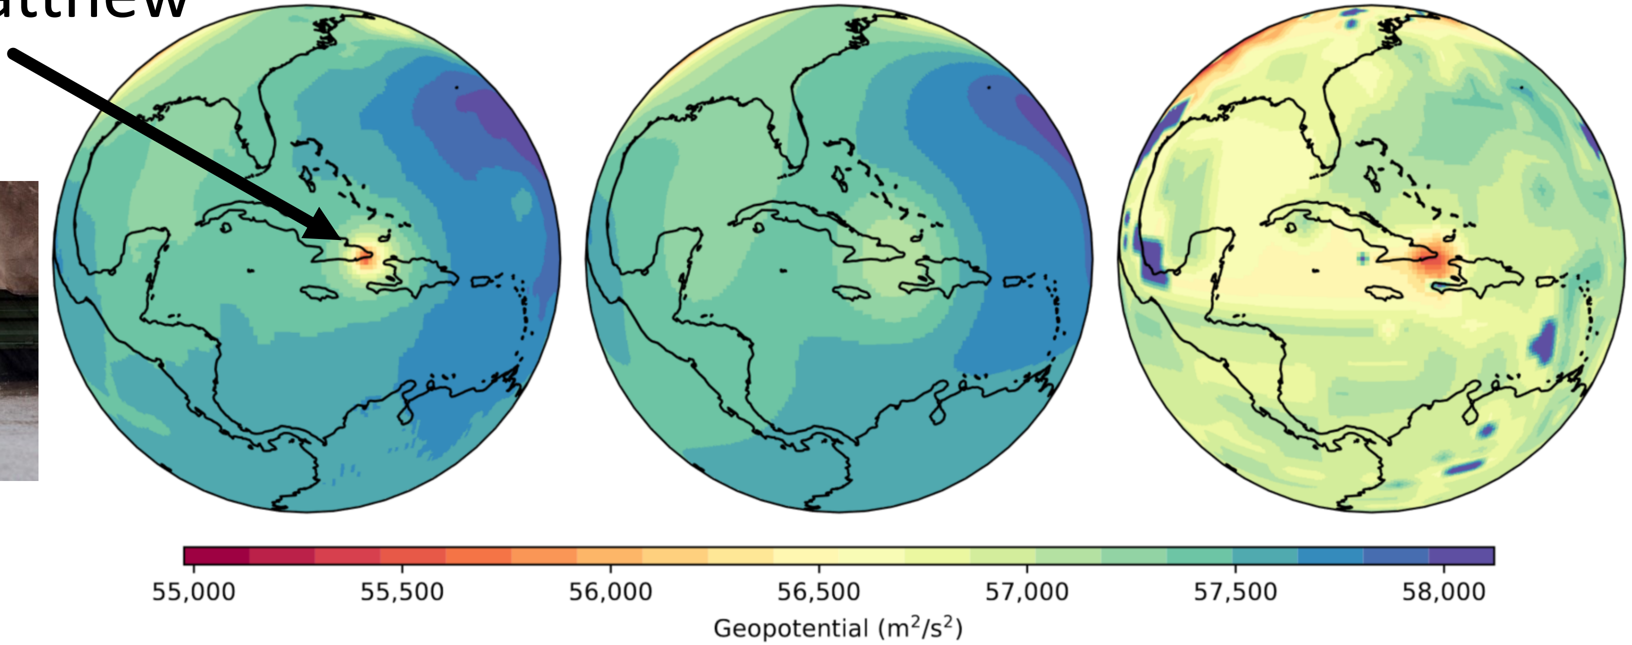

Evaluation: Case Study

Geopotential at 500hPa, 2016 Oct 5th

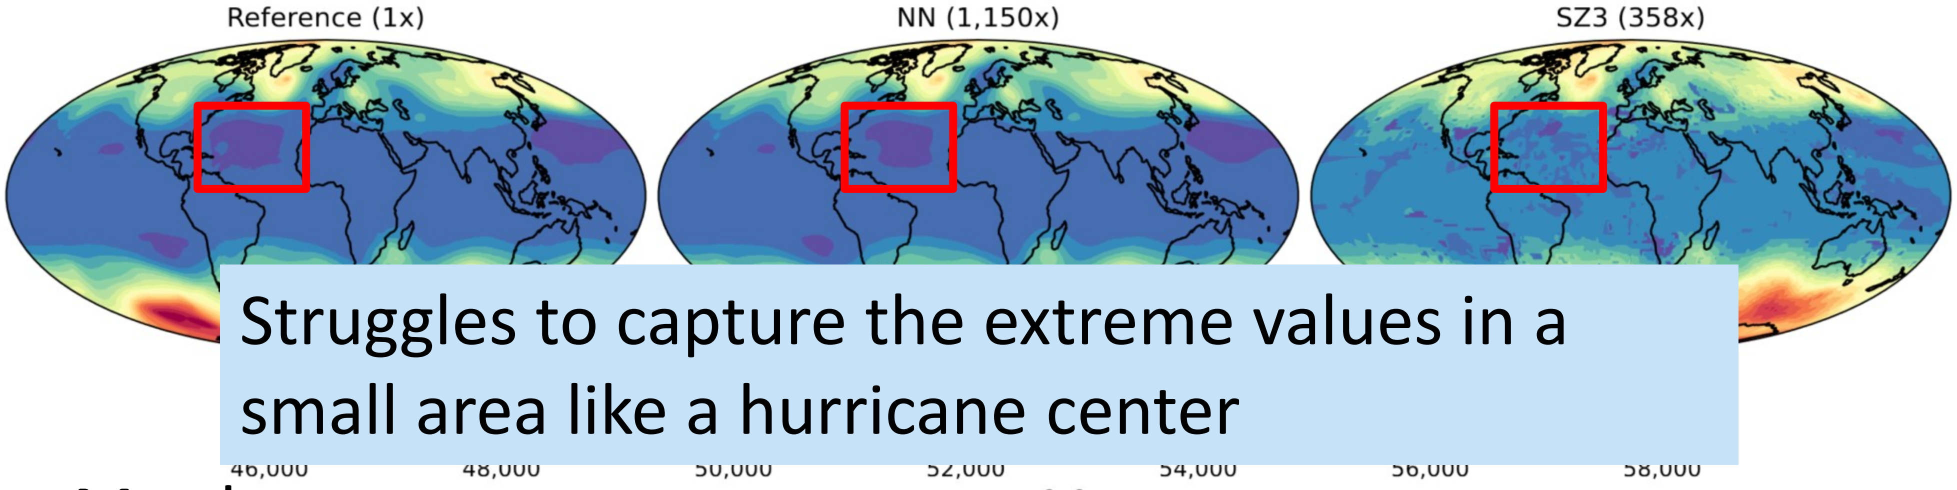

Struggles to capture the extreme values in a small area like a hurricane center

Hurricane Matthew

16.5bn damage
603 fatalities

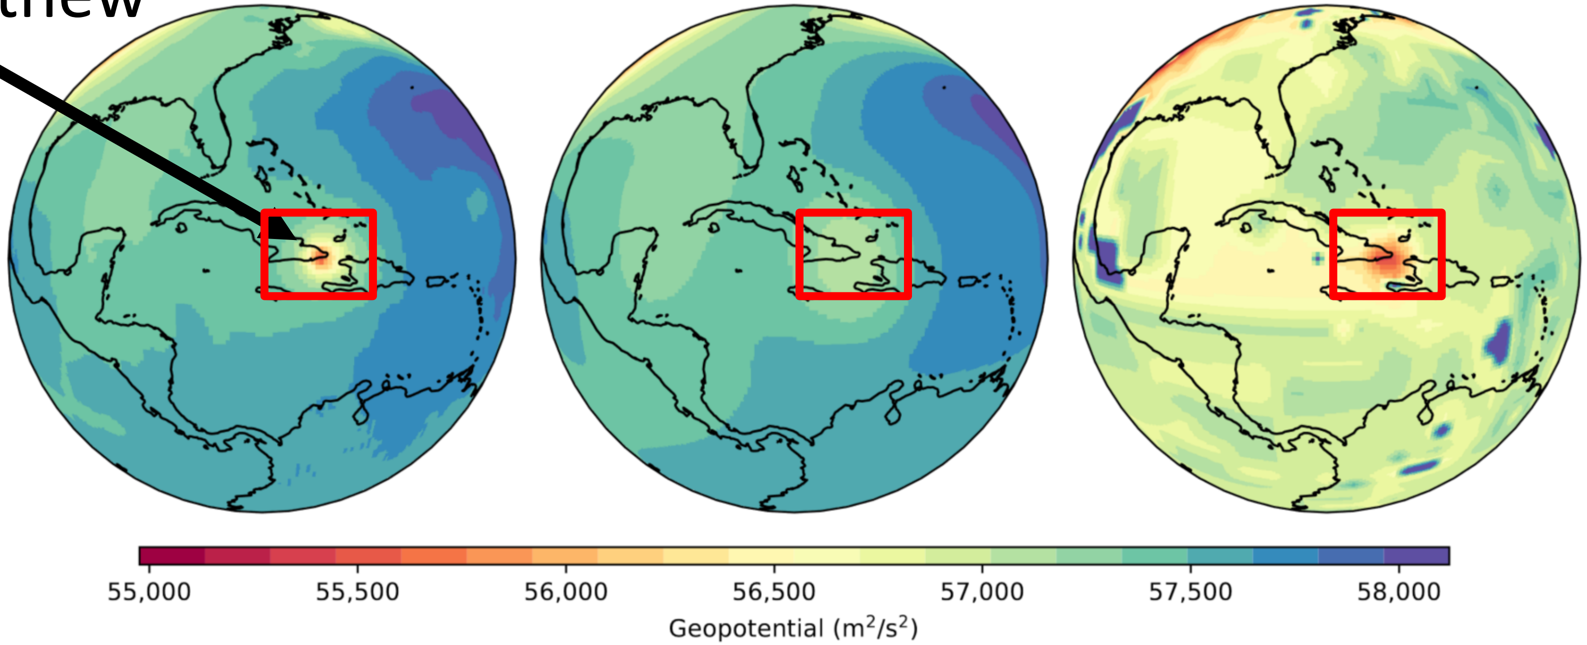

Evaluation: Case Study

Geopotential at 500hPa, 2016

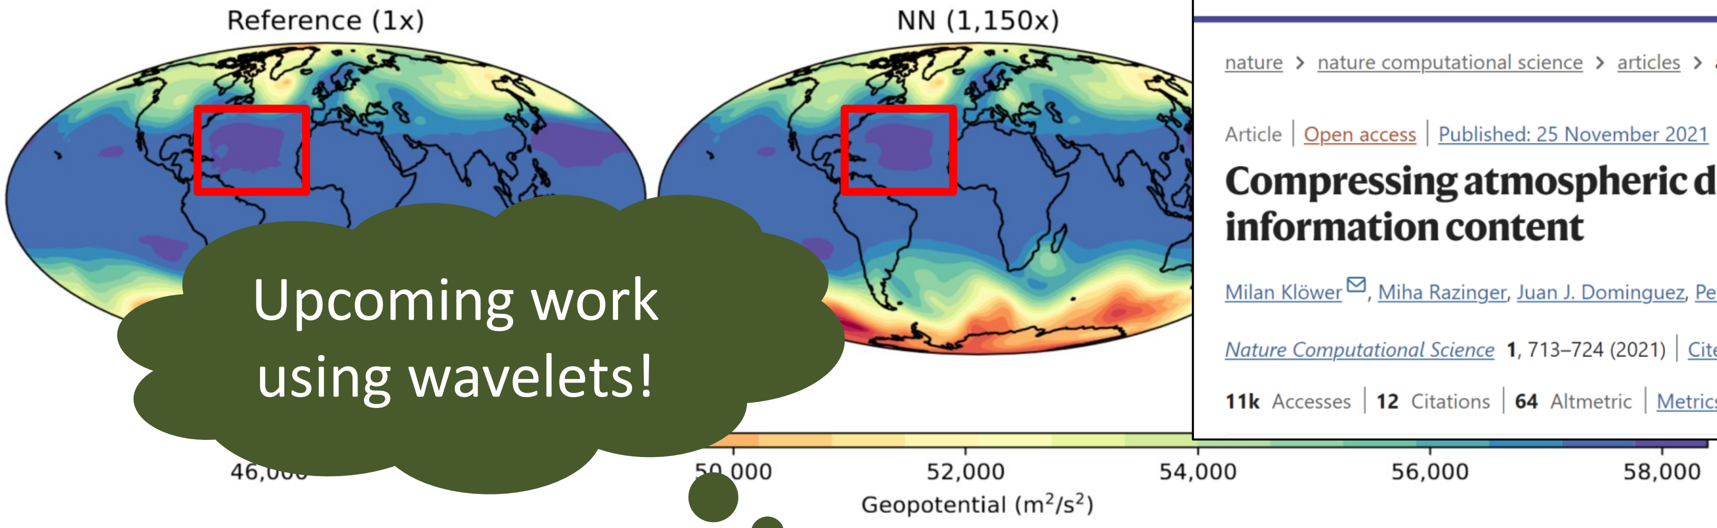

Upcoming work using wavelets!

nature computational science

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › nature computational science › articles › article

Article | Open access | Published: 25 November 2021

Compressing atmospheric data into its real information content

Milan Klöwer, Miha Razinger, Juan J. Dominguez, Peter D. Düben & Tim N. Palmer

Nature Computational Science **1**, 713–724 (2021) | Cite this article

11k Accesses | **12** Citations | **64** Altmetric | Metrics

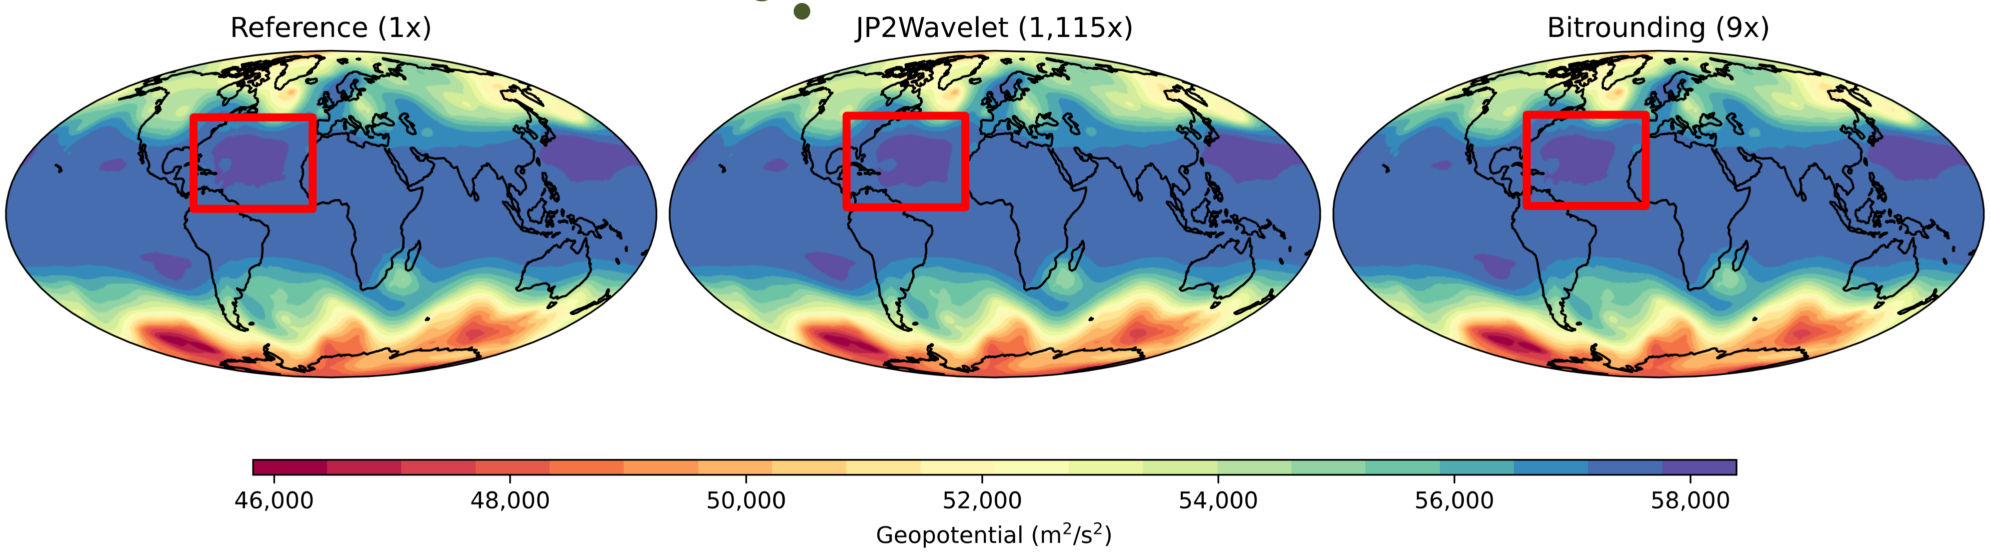

Evaluation: Case Study

Hurricane Matthew

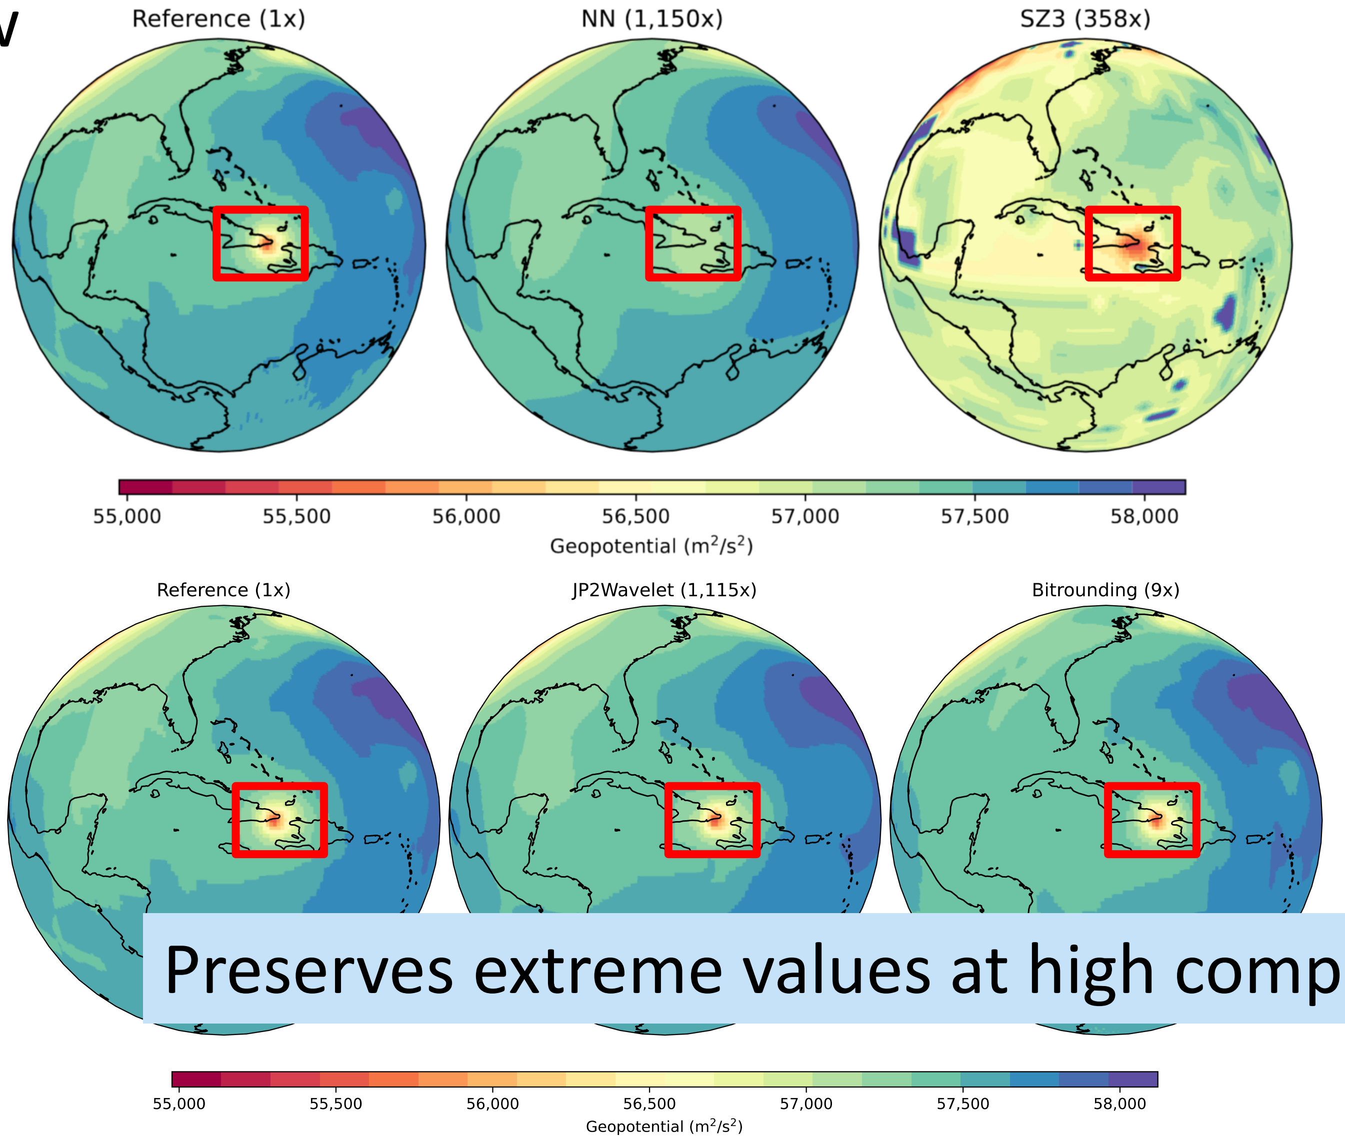

Preserves extreme values at high compression ratio

Summary and Key Points

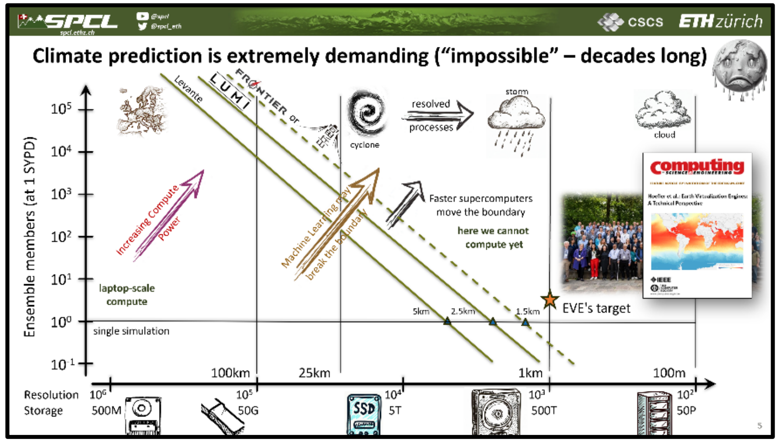

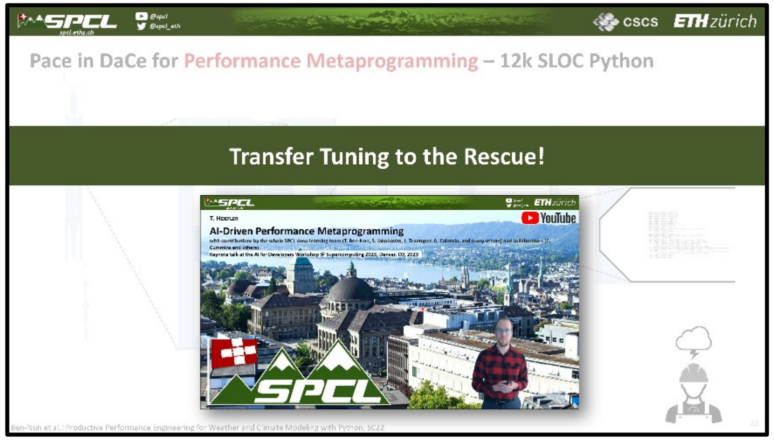

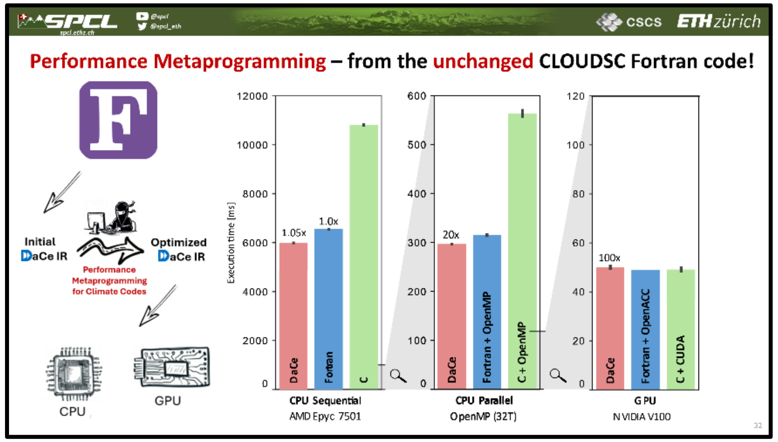

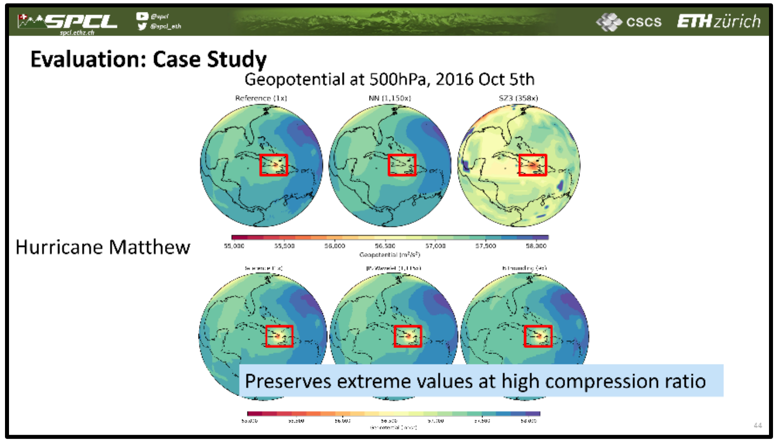

All of **ERA-5 on a USB-drive**! Run your own analyses on your laptop!

More of SPCL's research:

youtube.com/@spcl	150+ Talks
twitter.com/spcl_eth	1.2K+ Followers
github.com/spcl	2K+ Stars

… or spcl.ethz.ch

Join us! We're looking for PhD students, postdocs, and academic visitors in Zurich!

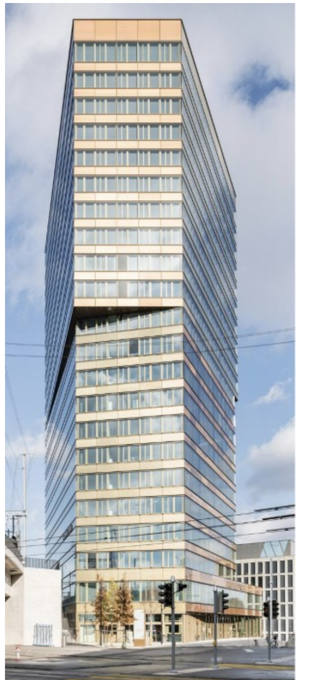

http://spcl.ethz.ch/Jobs/
http://spcl.ethz.ch/Visit/

Uncertainty in forecasting

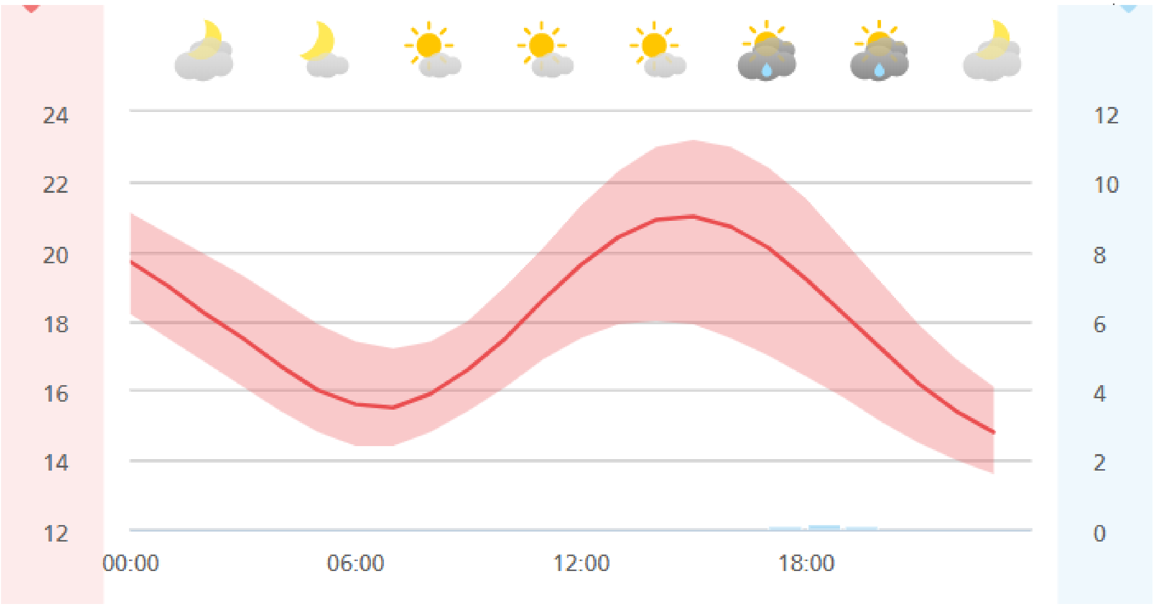

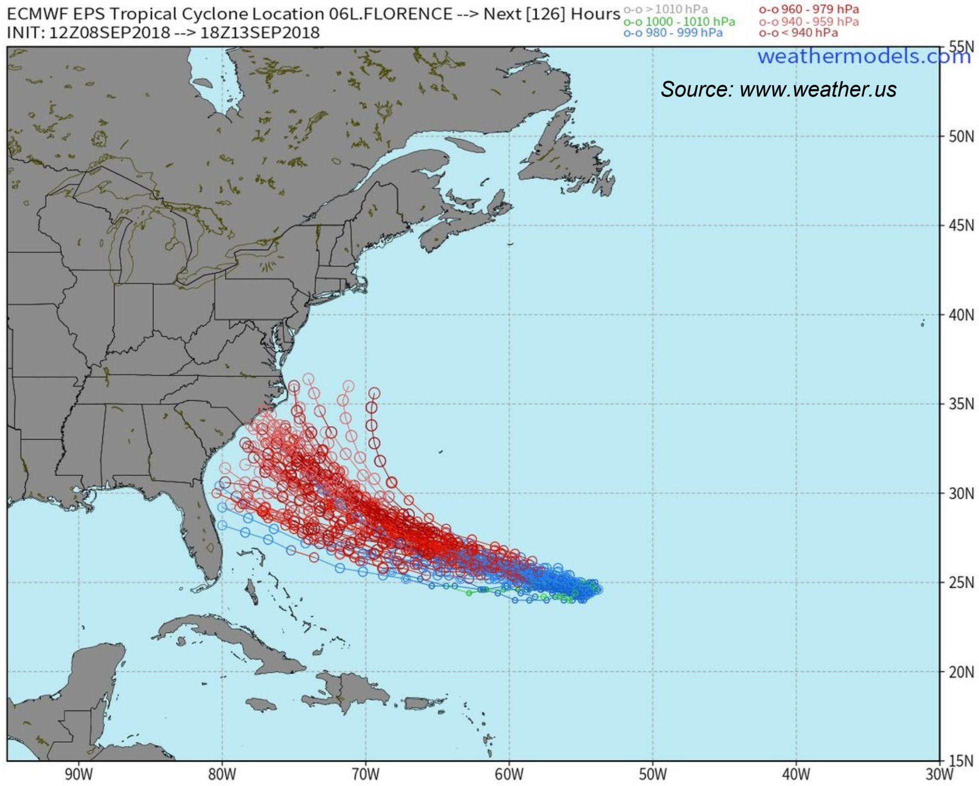

- **Weather is a chaotic system**
  - Minor perturbations affect the outcome the further into the future we predict

- **Solution: Ensemble Prediction Systems – predict weather as a probability distribution**
  - Approximated by partial differential equations with perturbed inputs

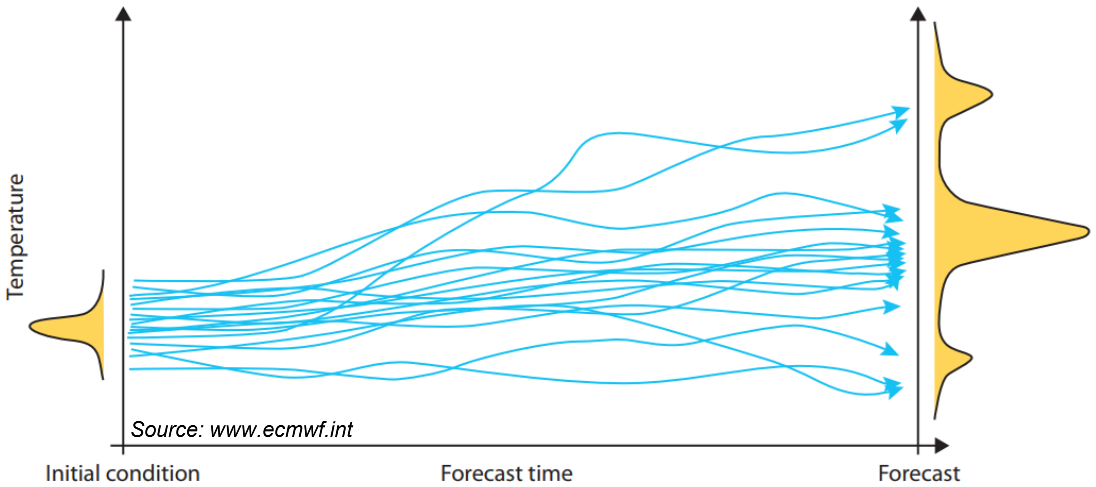

Ensemble Prediction System at ECMWF

- **Initial condition uncertainties result from data assimilation of sensor data**

- **51 ensemble members, 1 control (deterministic), 50 perturbed (stochastic)**
  - Approximate the highest likely trajectory from output distribution D
  - Lower resolution (9km vs. 18km) in order to fit compute budget
    mostly an economic argument

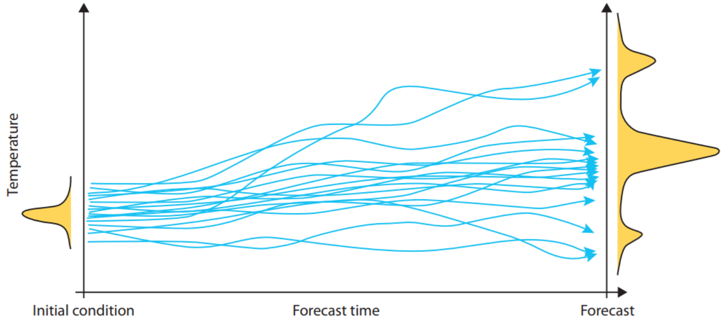

- **Next step in the economic argument:**
  - Could the number of ensemble members be reduced without sacrificing accuracy?
  - Idea I: predict mean and standard deviation (StdDev) of D from a smaller ensemble
    This may allow us to increase resolution at equal cost – better predictions
  - Can we improve prediction quality by learning from ground truth observations?
  - Idea II: learn (local) model bias from observations
    This may allow us to increase accuracy – better predictions

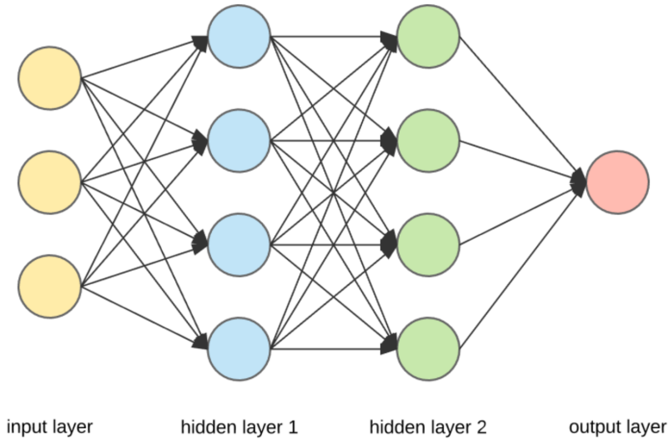

ENS-10: A Dataset For Post-Processing Ensemble Weather Forecasts

Ashkboos et al., NeurIPS'22

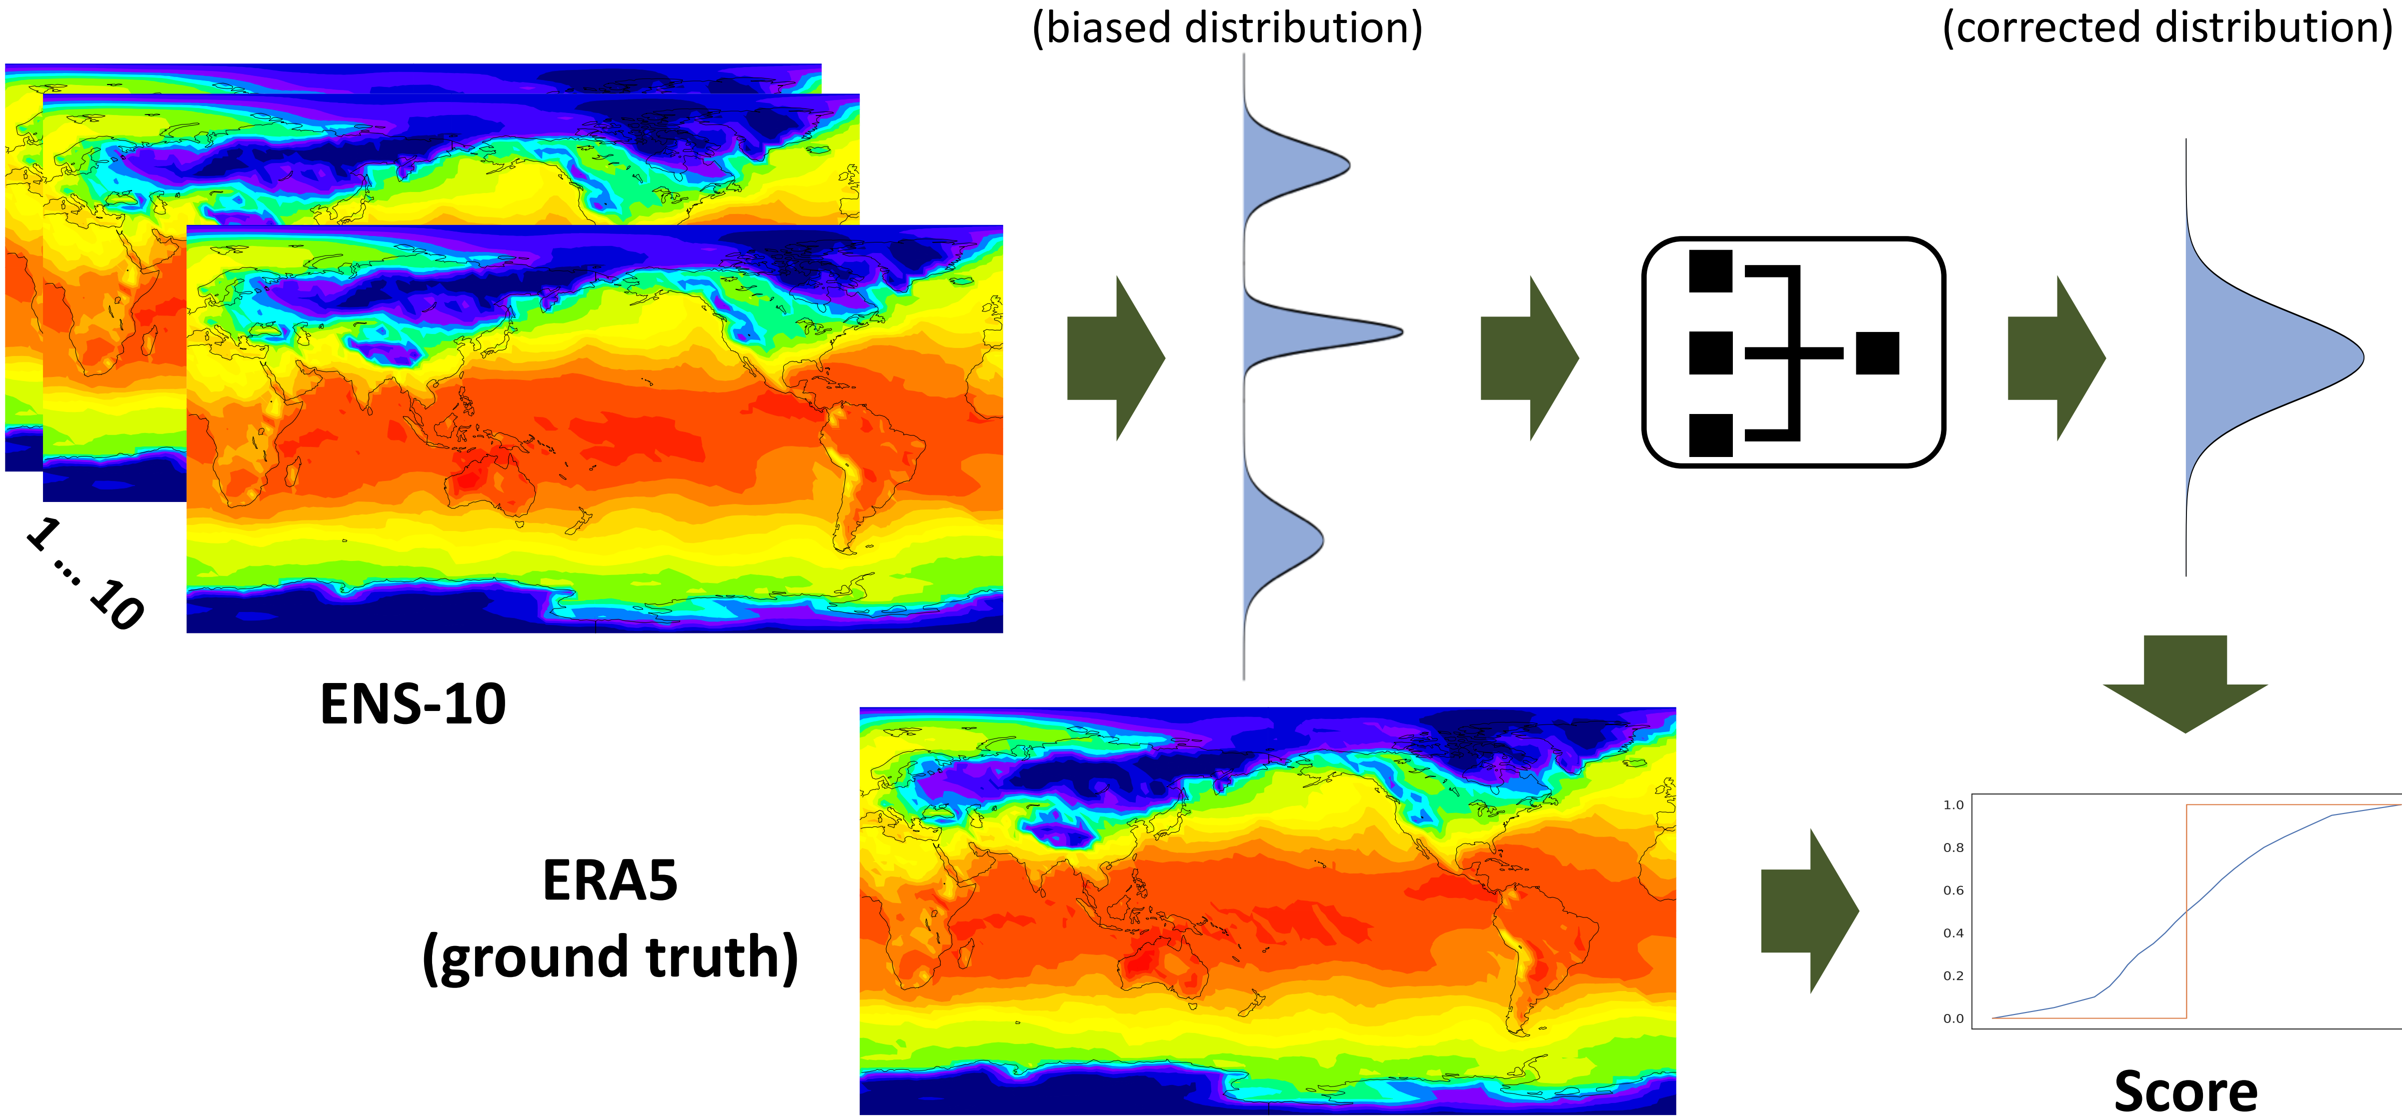

S. Ashkboos et al.: ENS-10: A Dataset For Post-Processing Ensemble Weather Forecasts, NeurIPS'22, arXiv

Uncertainty Quantification Network (based on ResNet)

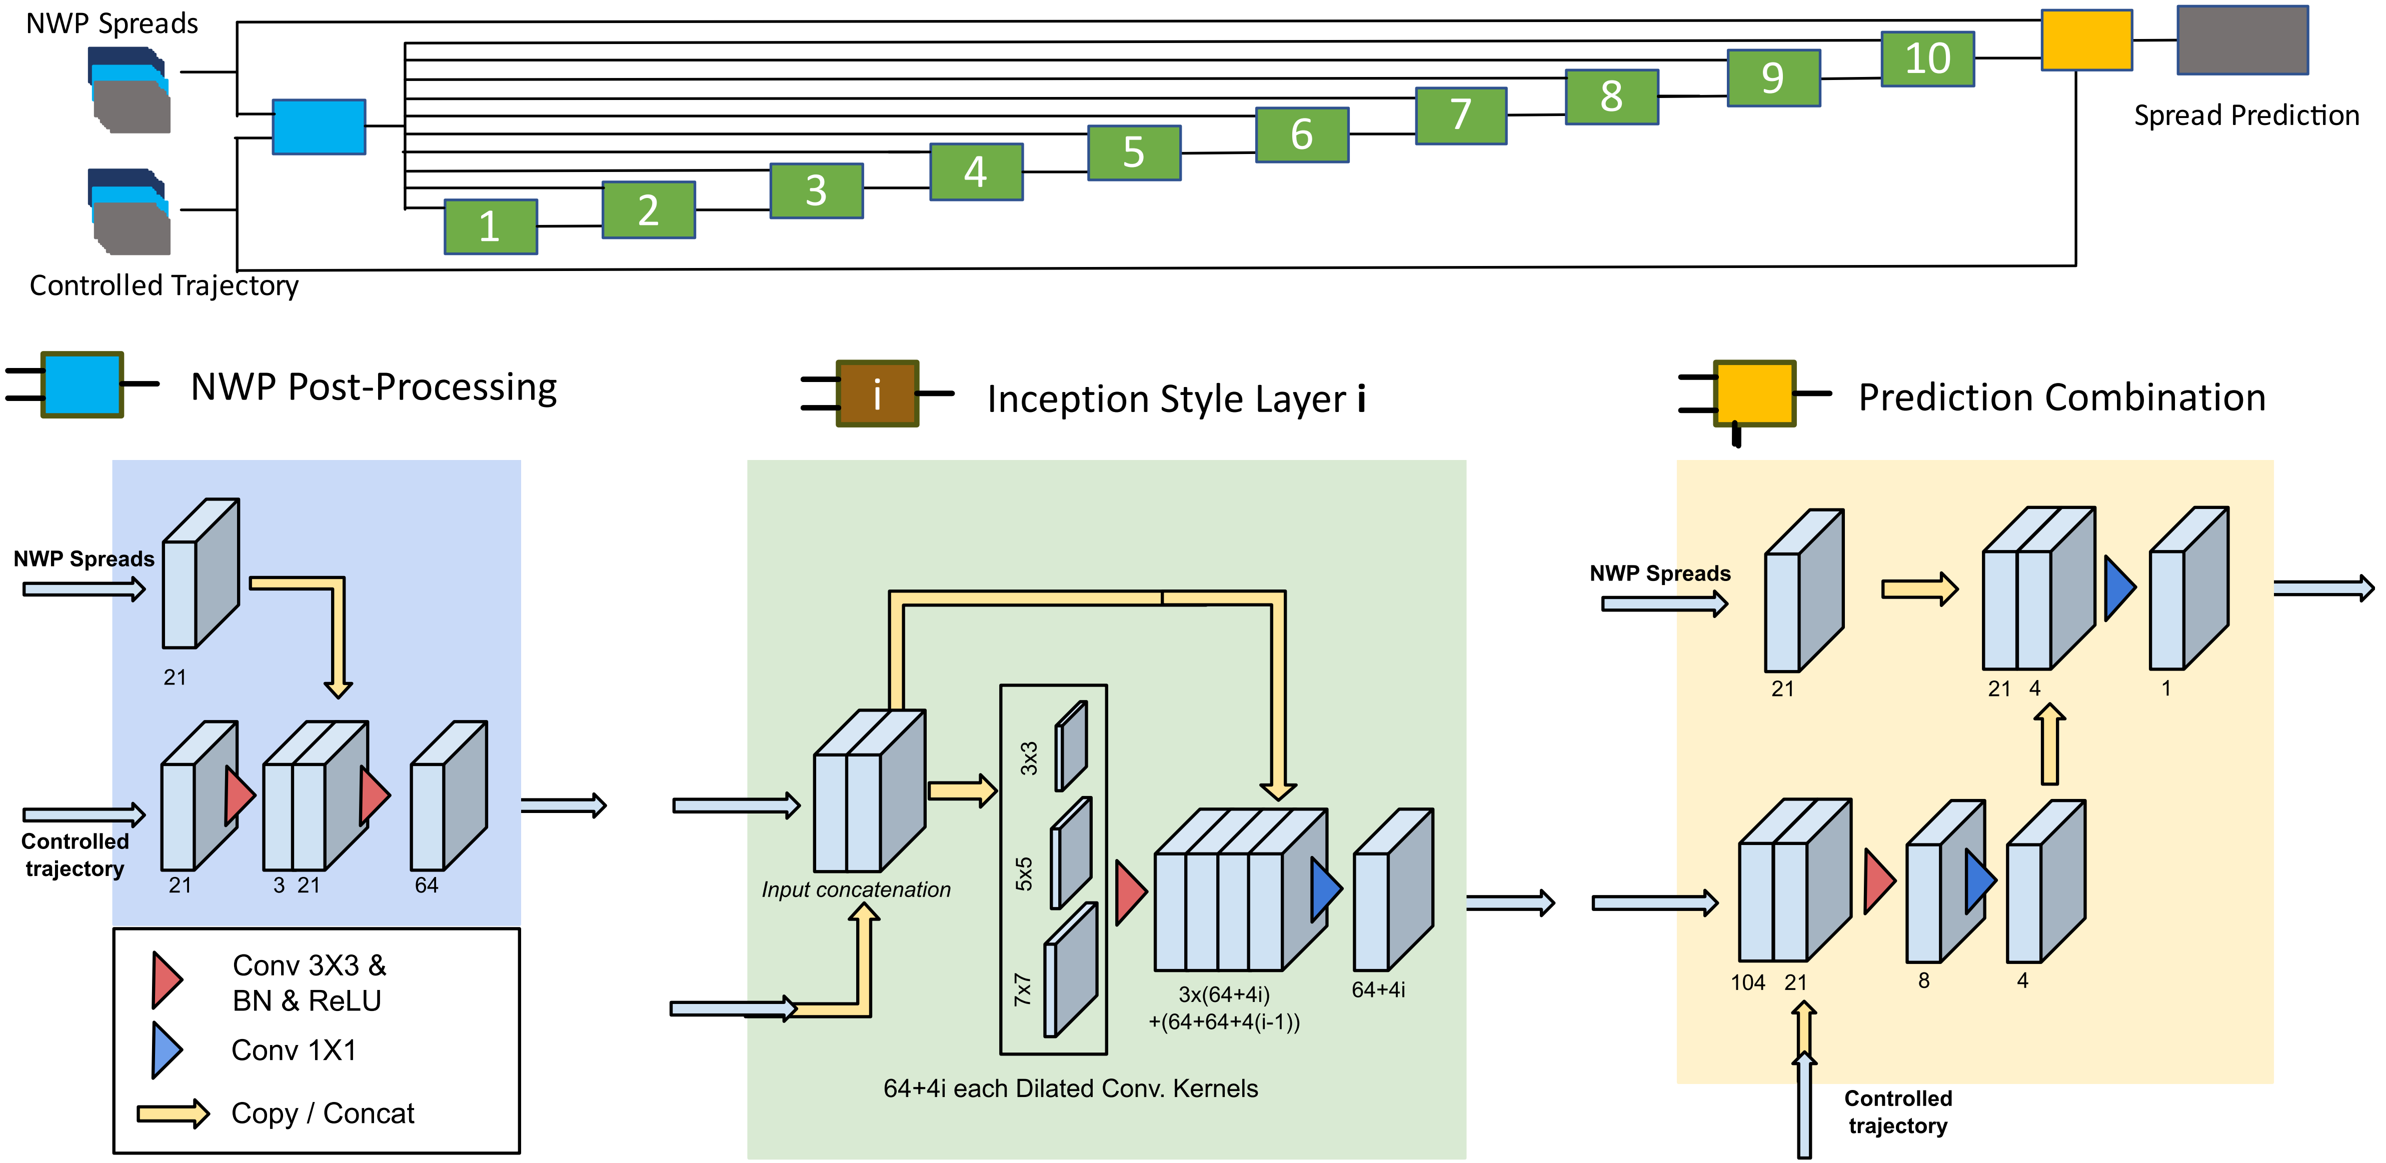

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Bias Correction Network (based on 3D-Unet + LCN)

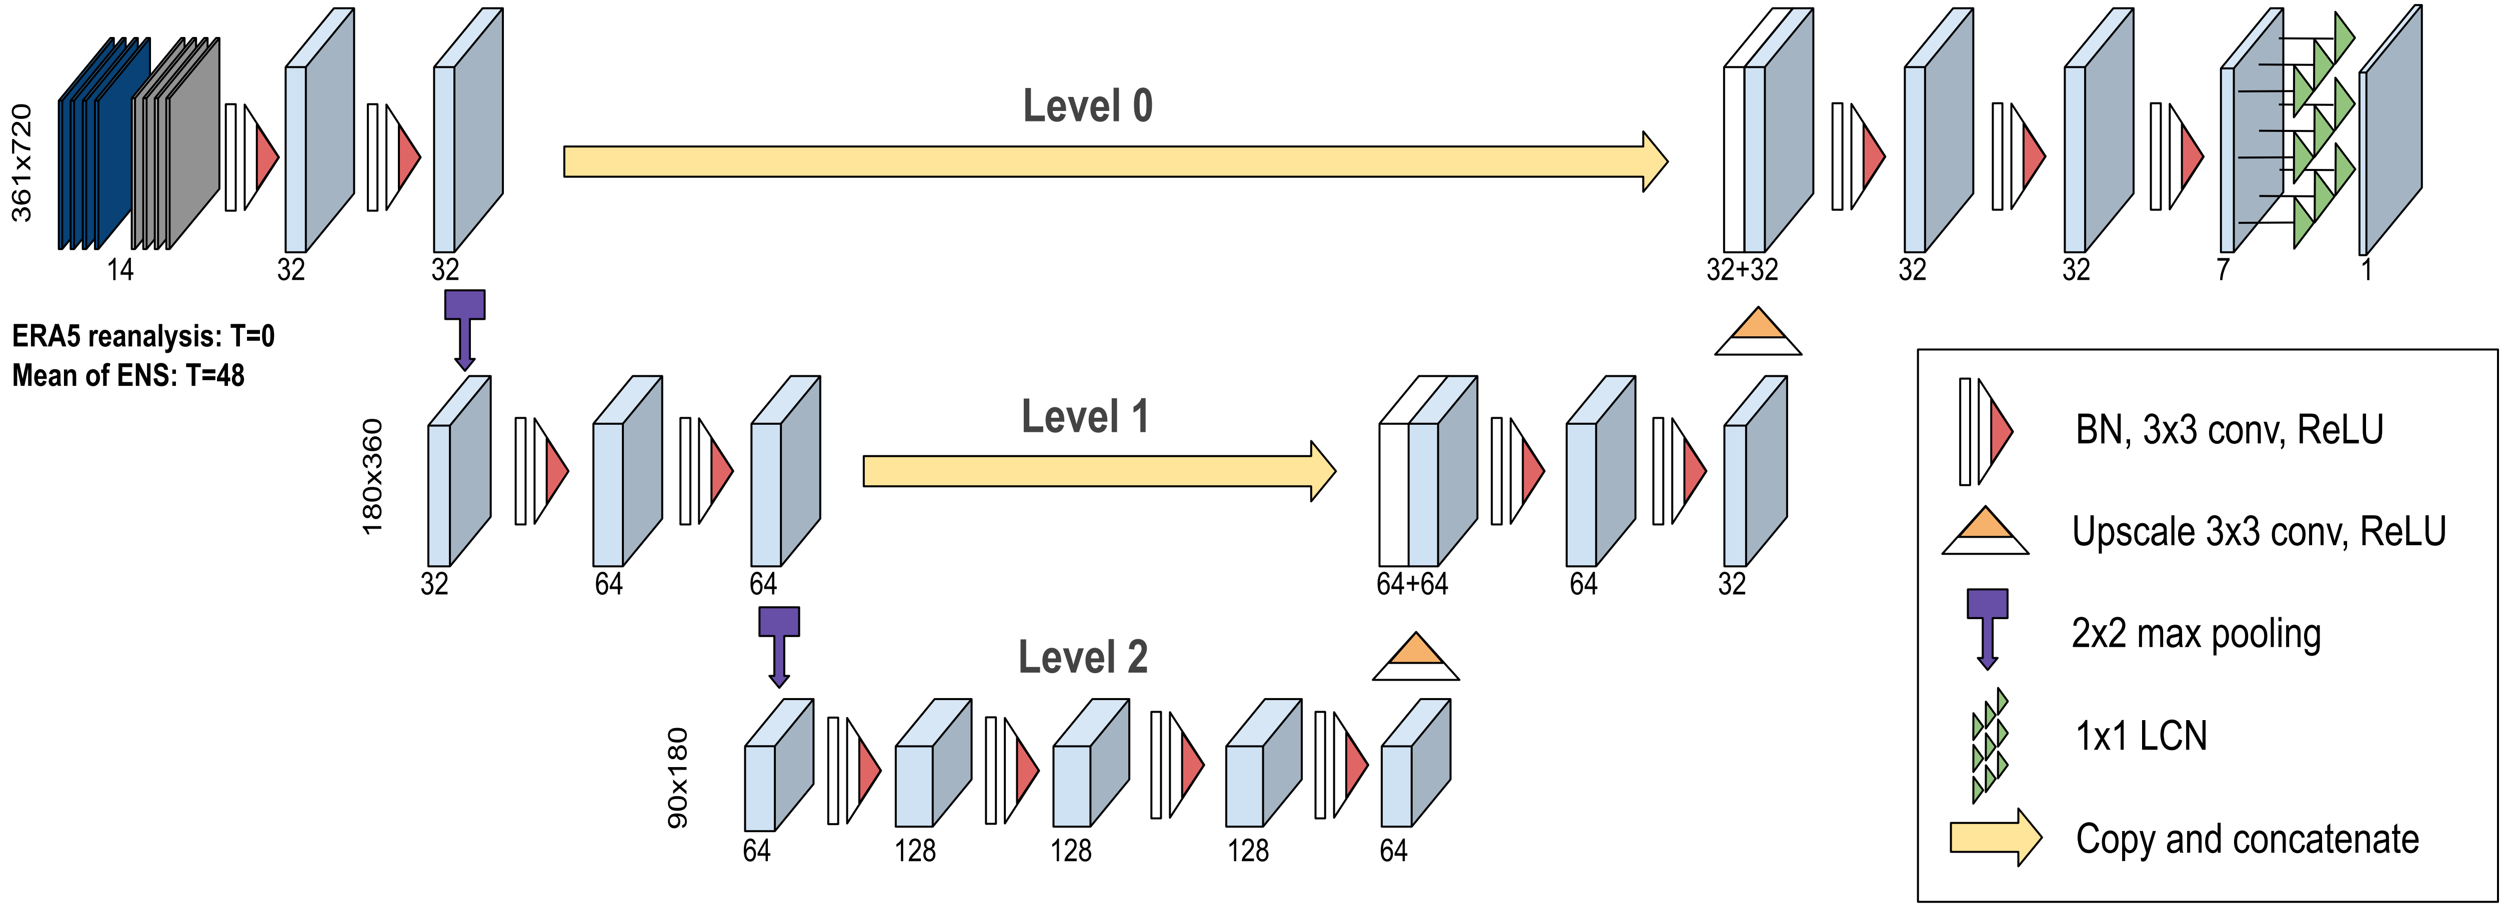

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Global RMSE results

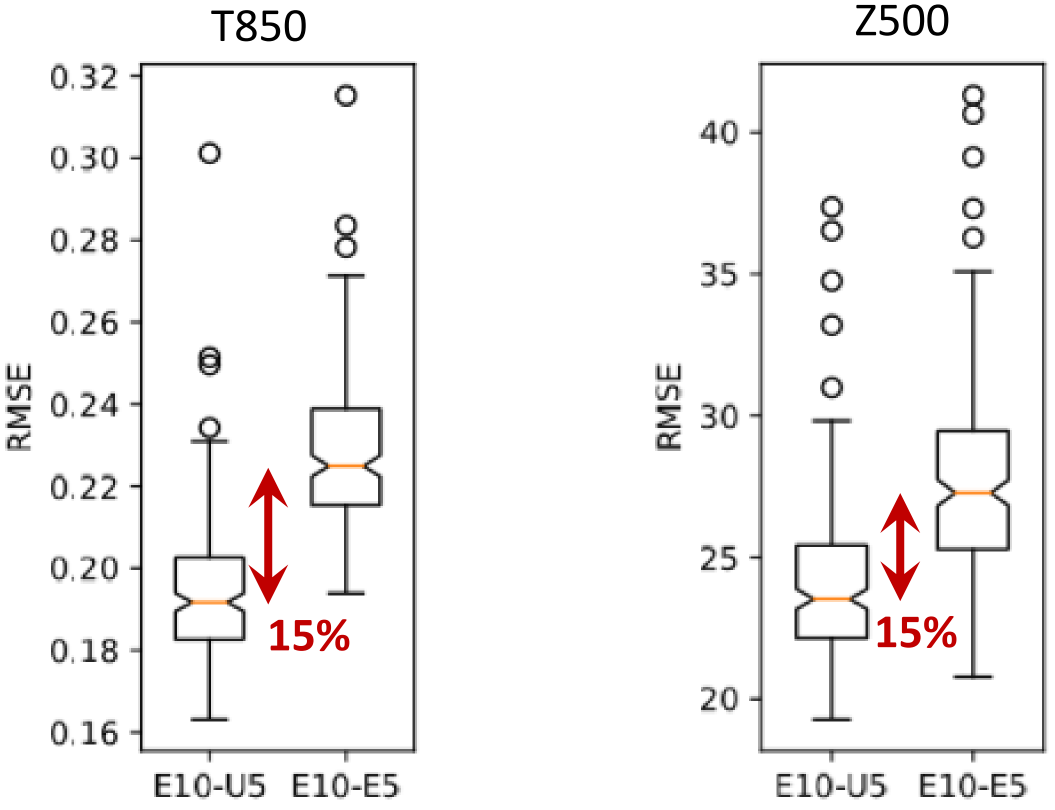

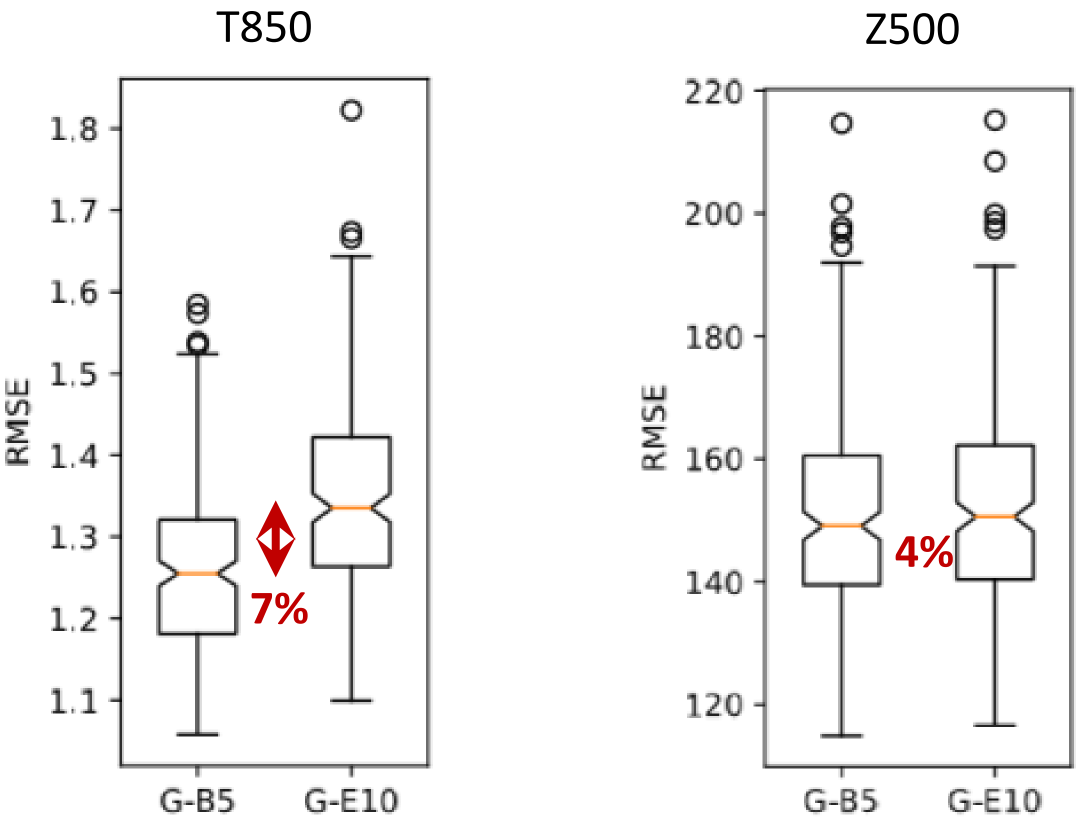

10 ensembles vs. UQ with 5 ensembles

10 ensembles vs. 5 ensembles

ERA5 (ground truth) vs. BC with 5 trajectories

ERA5 (ground truth) vs. 10 trajectories

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Extreme event: Tropical Cyclone Winston & Hurricane Matthews

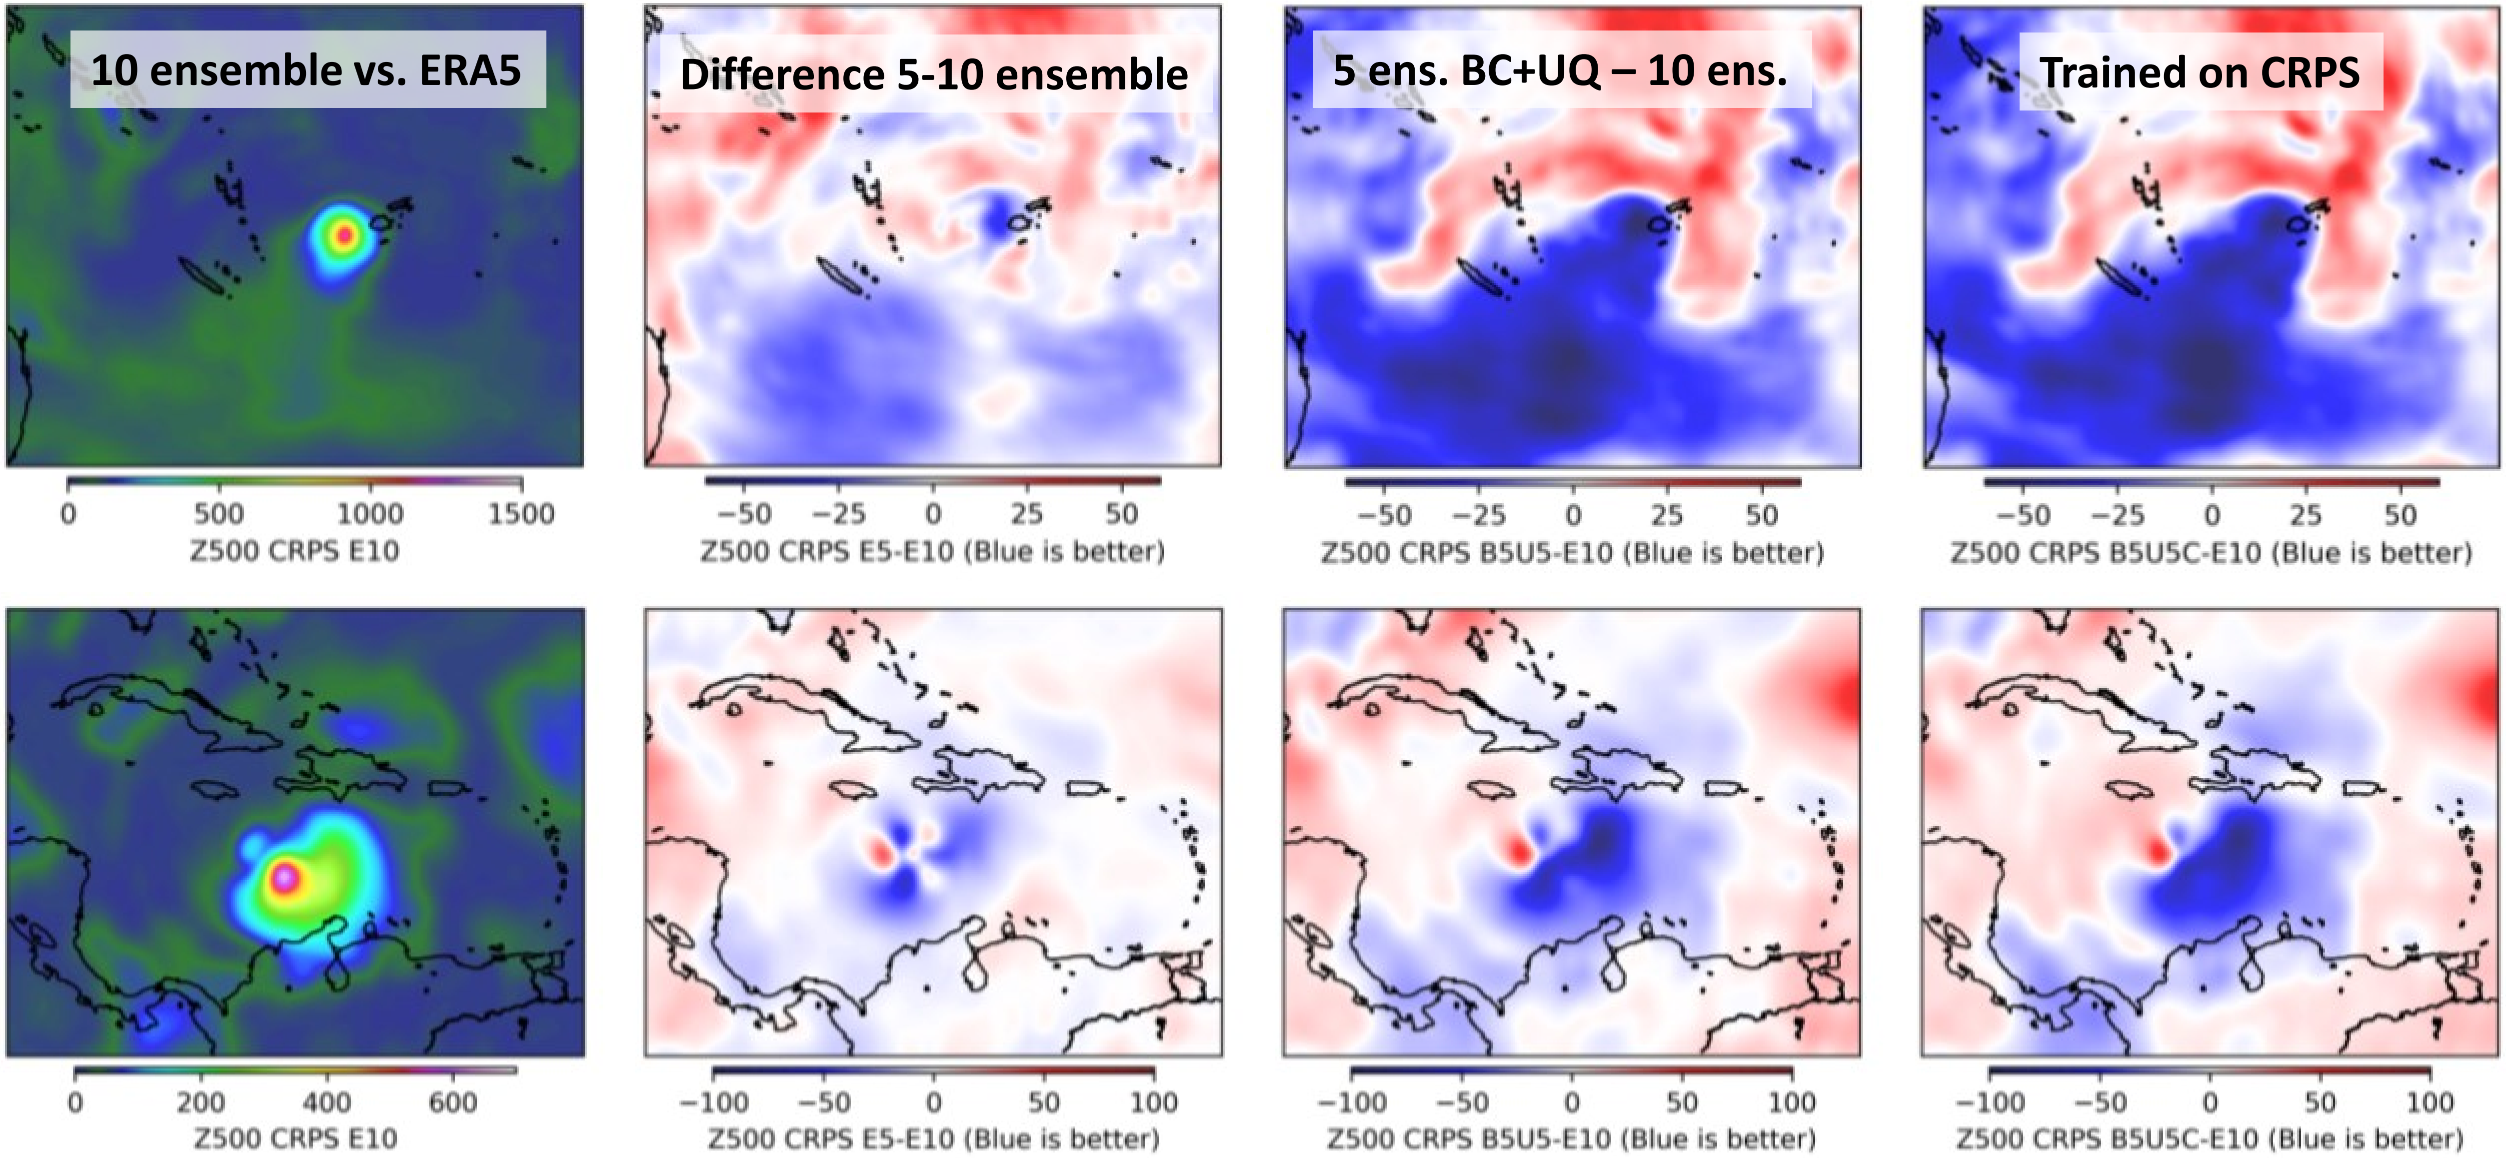

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Cold wave over Asia

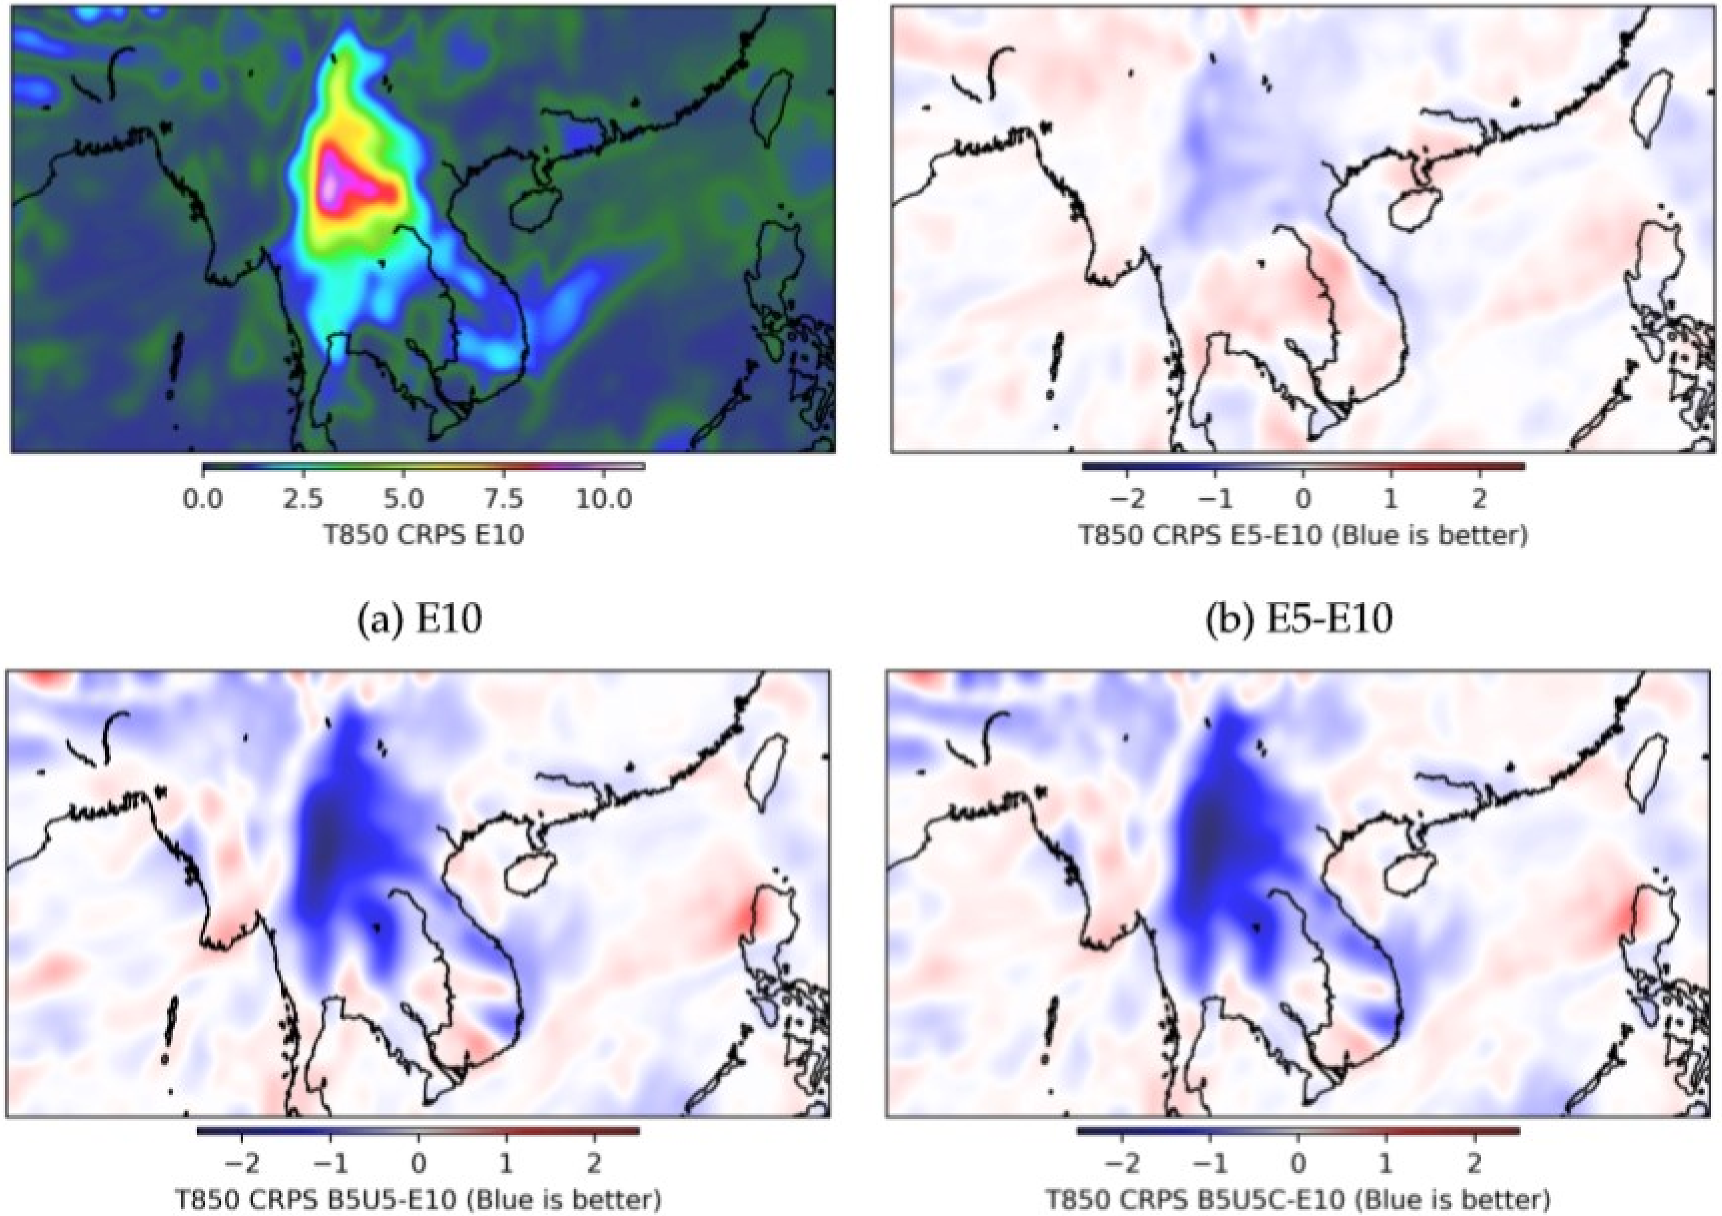

(a) E10

(b) E5-E10

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Intermediate Conclusions (Preliminary Study)

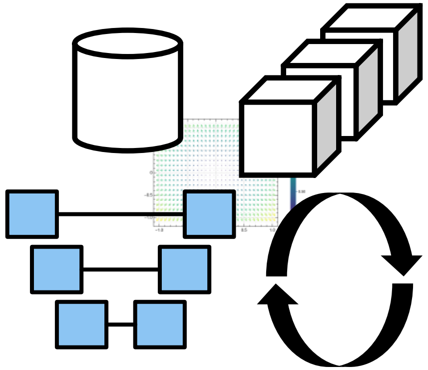

- **Simple Deep Learning can be used to accelerate forecast pipelines**
  - Take advantage of industry efforts to tune hardware and tool-chains
  - An informed approach is **necessary** to ensure improved results
- **Using Encoder-Decoder networks for predicting mean and StdDev in ensemble systems yields higher accuracy than using small ensemble statistics**
  - Fewer than half of the ensemble members are necessary
  - Accuracy improved with custom operators
- **Promising for increasing performance in large-scale settings**
  - Needs further investigation!
  - Join us/try yourself: https://github.com/spcl/deep-weather
- **Future directions:**
  - Larger datasets
  - Custom neural architectures for unstructured grids
  - Integrate into dace tool-chain for further optimization

P. Gronquist et al.: Deep Learning for Post-Processing Ensemble Weather Forecasts, RSTA'19, arXiv

Massive opportunity: Addressing the Climate Data Deluge (arXiv:2210.12538)

- **ECMWF soon produces 1PB/day of simulation data – we need to compress it!**

- **Key idea [1]: Overfit a Fourier Network to a block of (climate) data!**

[1]: Huang, Hoefler: "Compressing multidimensional weather and climate data into neural networks", ICLR'23 **top-5% paper**

Massive opportunity: Addressing the Climate Data Deluge (arXiv:2210.12538)

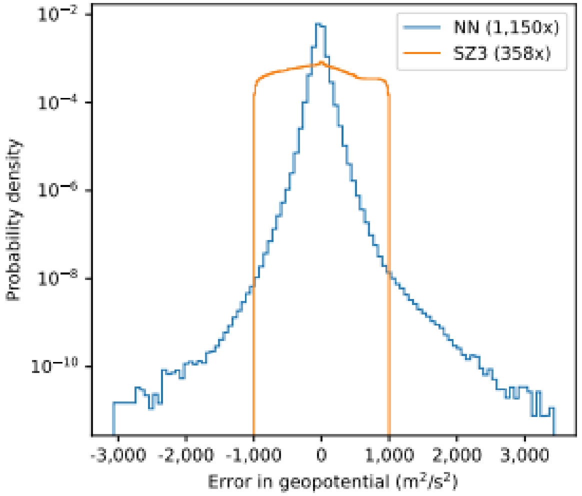

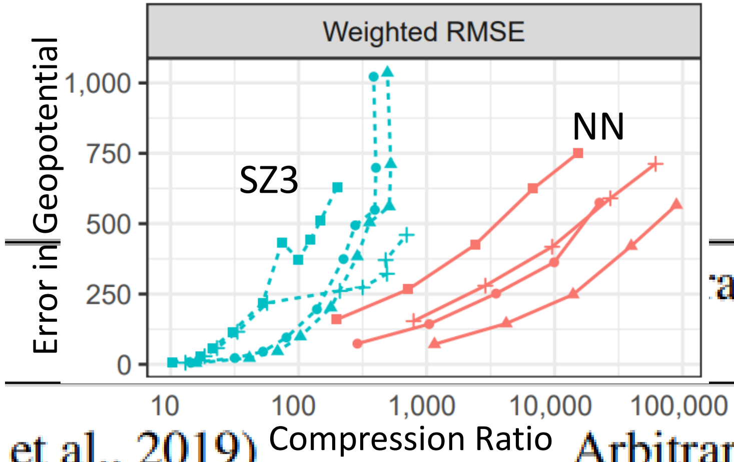

...ratio	Comp. speed	Decomp. speed Cont.	Rand.
et al., 2019) Arbitrary			
ZFP (Lindstrom, 2014) ≤ 10x			

SZ3 enforces strict error bounds… … which limits compression… … and introduces artifacts:

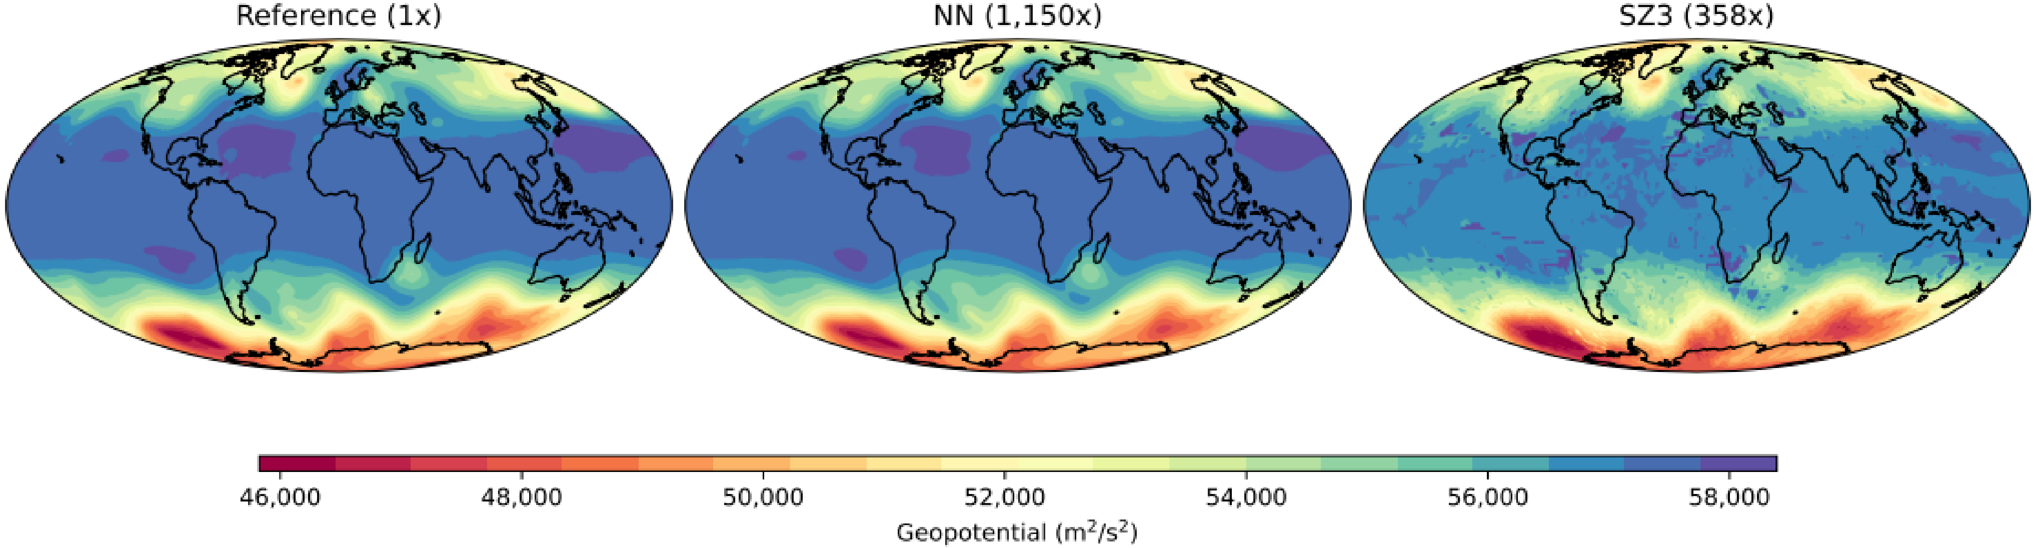

[1]: Huang, Hoefler: “Compressing multidimensional weather and climate data into neural networks”, ICLR’23 **top-5% paper**

Massive opportunity: Addressing the Climate Data Deluge (arXiv:2210.12538)

… but extreme events are dampened.

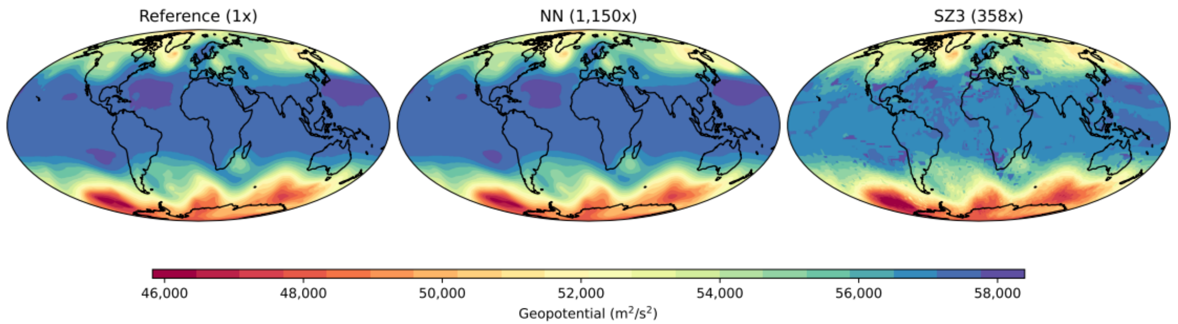

Hurricane Matthew (Oct. 5, 2016)

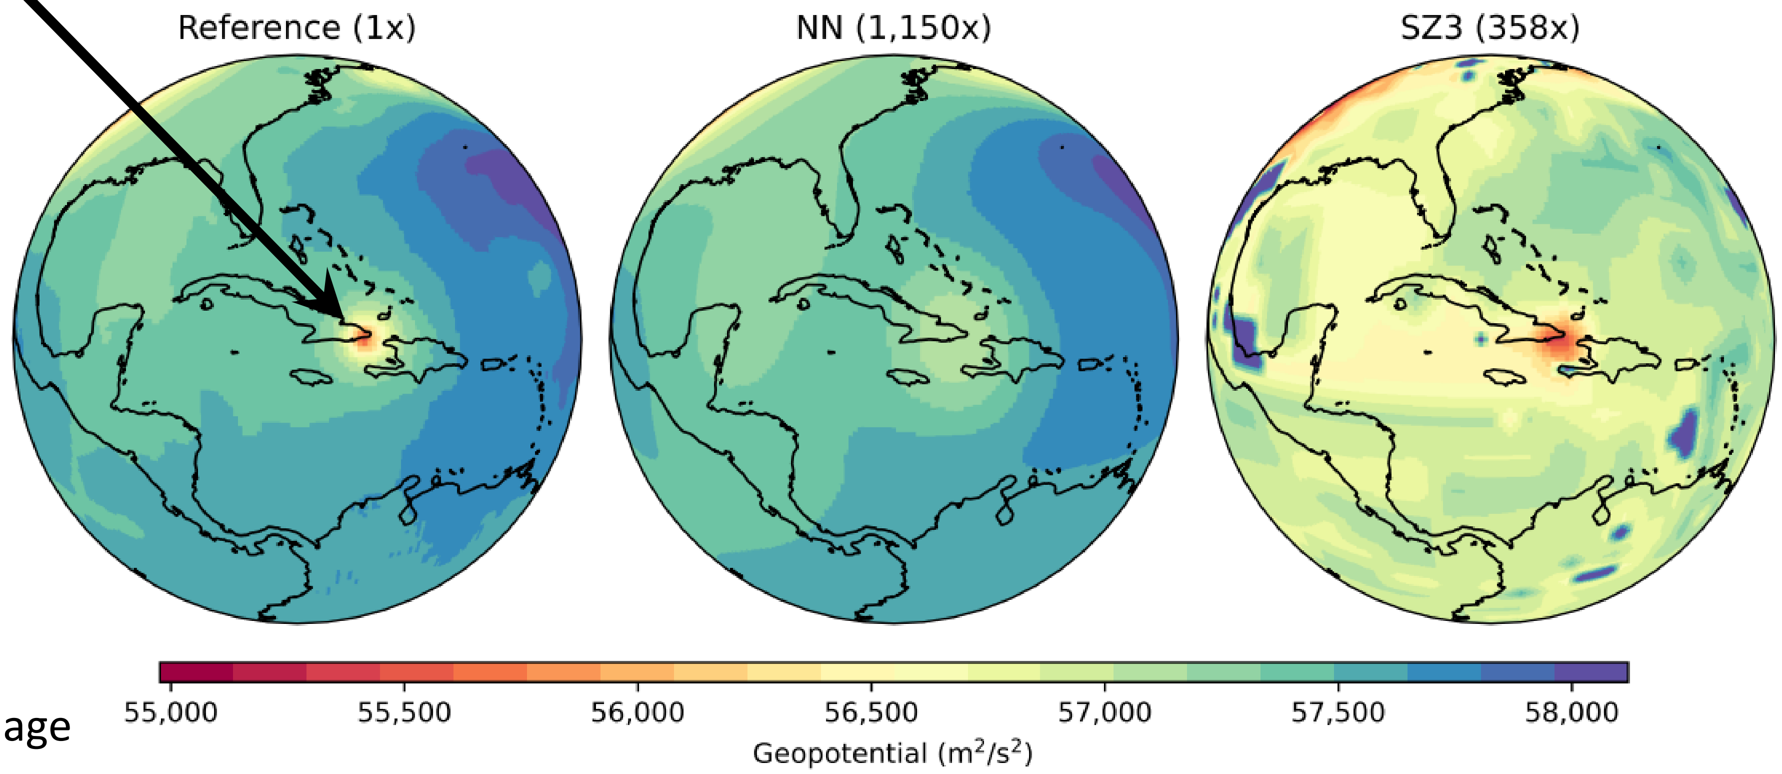

hit Cuba, Haiti, USA,
603 fatalities, $16.5bn damage

[1]: Huang, Hoefler: "Compressing multidimensional weather and climate data into neural networks", ICLR'23 **top-5% paper**

Meet the new wavelet basis – latest results from yesterday ☺

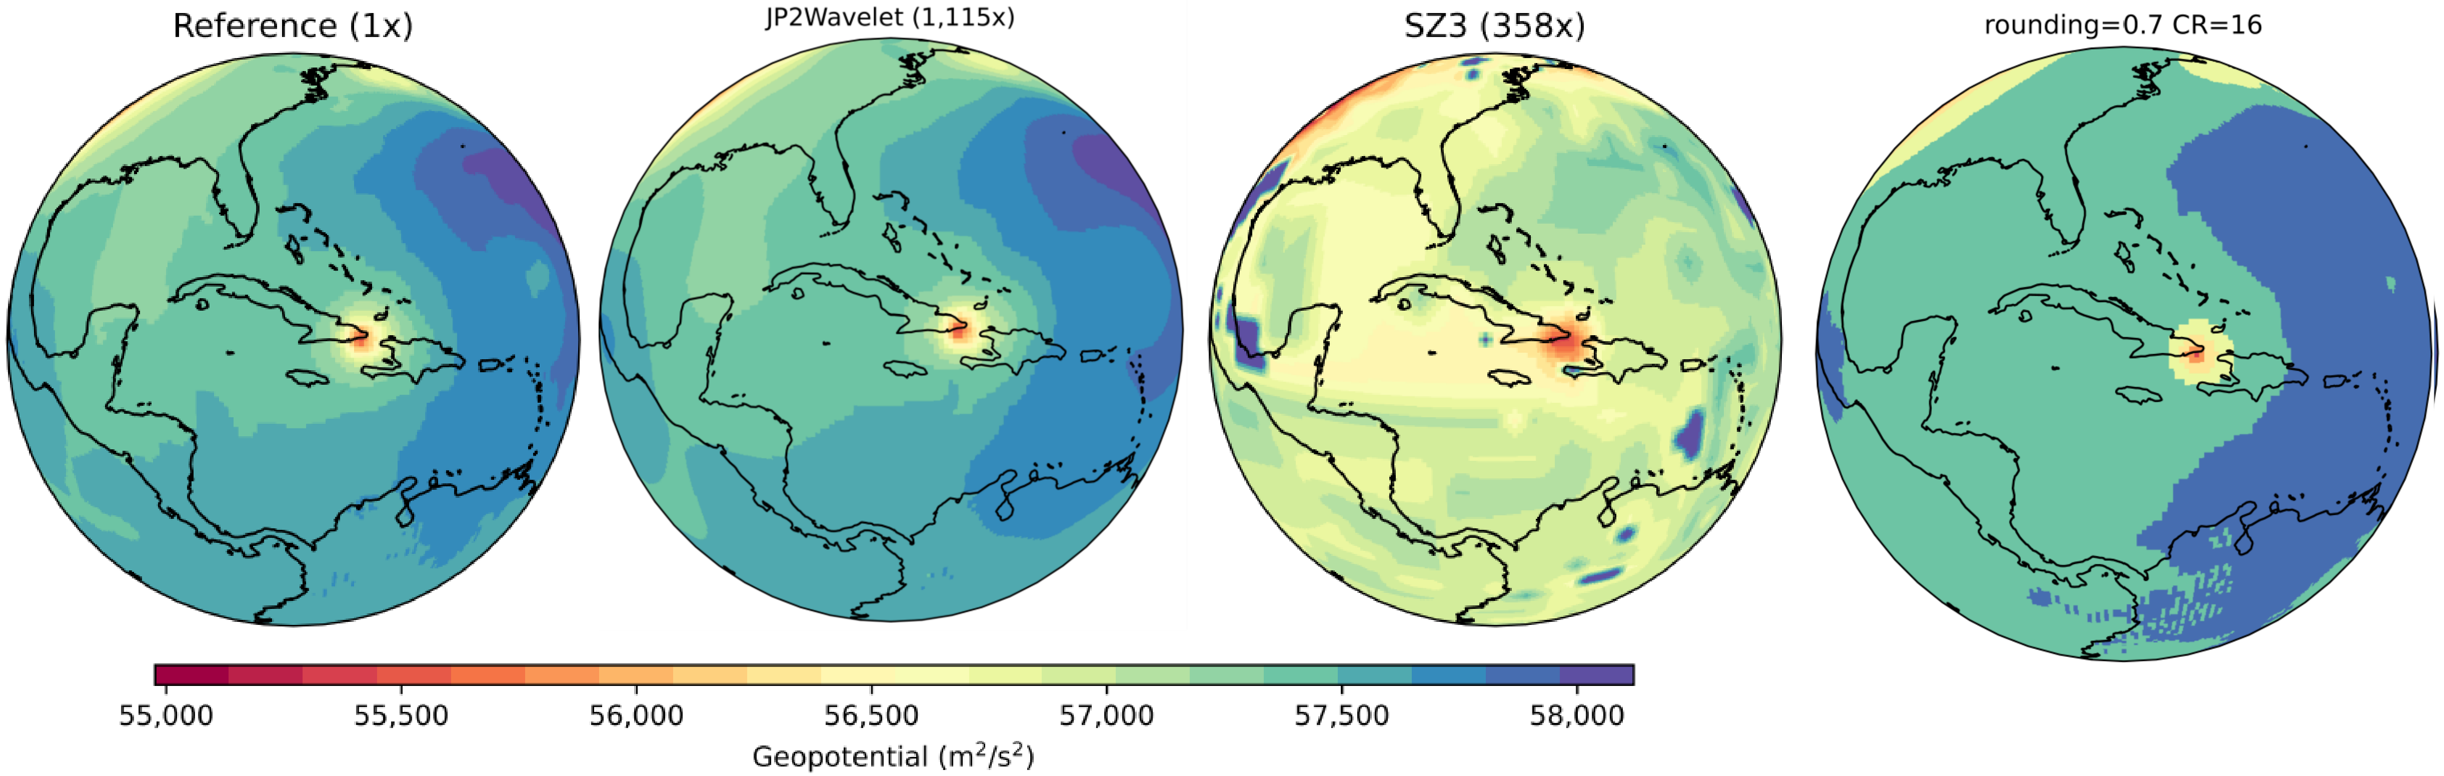

Compressing atmospheric data into its real information content

Milan Klöwer, Miha Razinger, Juan J. Dominguez, Peter D. Düben & Tim N. Palmer

Nature Computational Science **1**, 713–724 (2021) | Cite this article

Summary

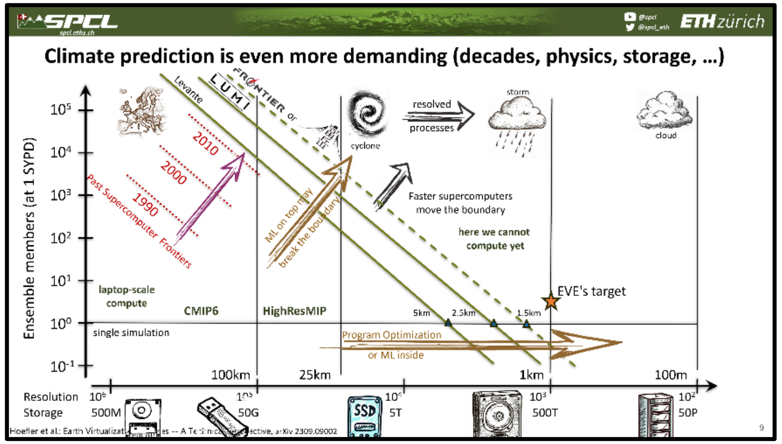

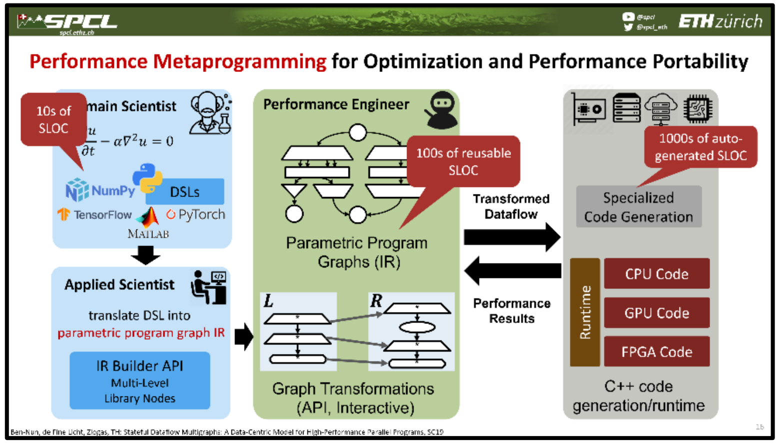

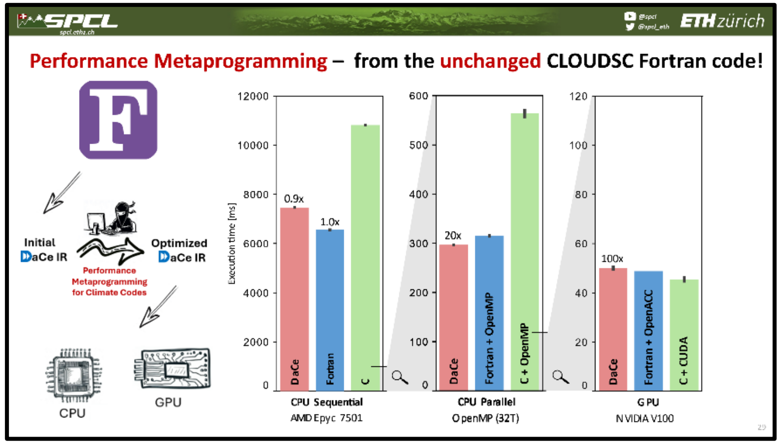

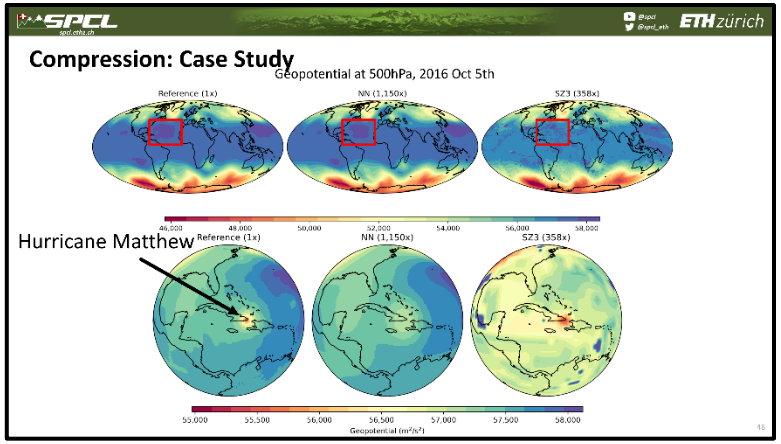

More of SPCL's research:

youtube.com/@spcl — 150+ Talks

twitter.com/spcl_eth — 1.2K+ Followers

github.com/spcl — 2K+ Stars

… or spcl.ethz.ch

Join us! We're looking for PhD students, postdocs, and academic visitors in Zurich!

http://spcl.inf.ethz.ch/Jobs/

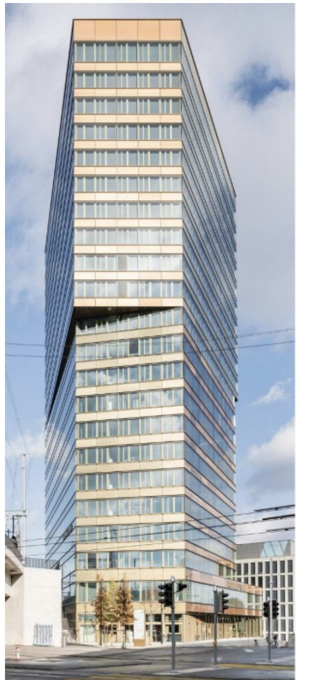

Bonus: Data Has Spatial Structure – Spatial Mixture of Experts

Weather and Climate grids have spatial structure!

Locality matters!

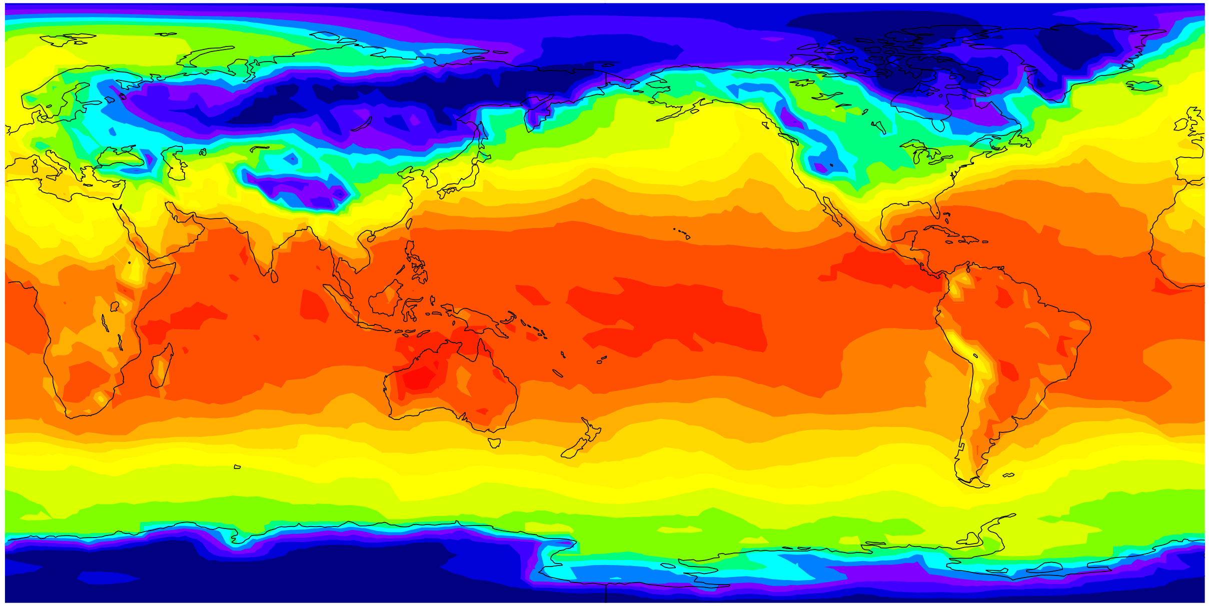

Convolution

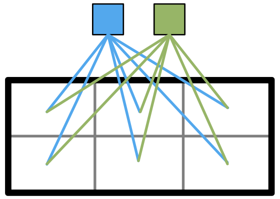

Locally-connected

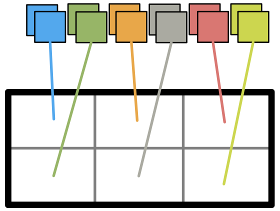

Low-rank locally-connected

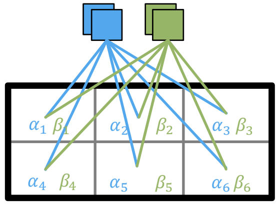

Mixture-of-Experts

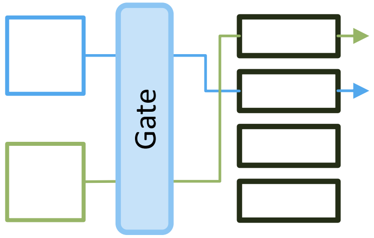

Spatial Mixture-of-Experts

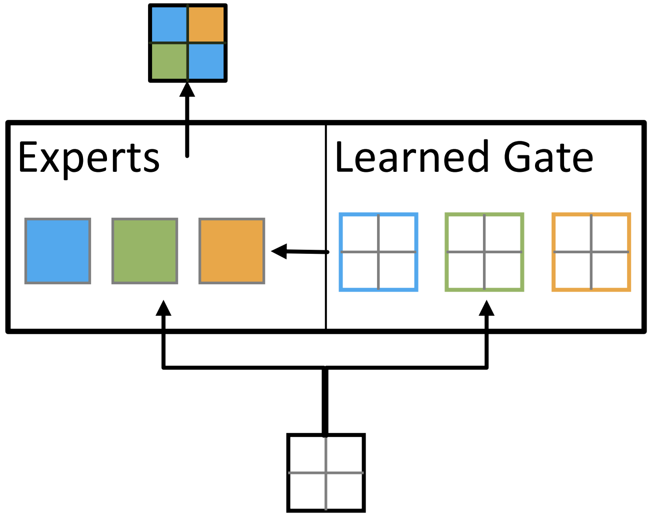

Dryden, Hoefler: Spatial Mixture-of-Experts, NeurIPS'22, arXiv

Spatial Mixture of Experts for Weather Prediction

Medium-range weather prediction WeatherBench

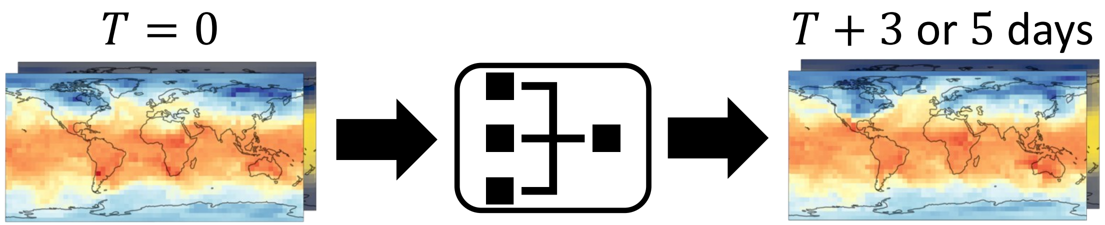

Model	Z500 [m²s⁻²] 3 days	Z500 [m²s⁻²] 5 days	T850 [K] 3 days	T850 [K] 5 days
Rasp & Thuerey	$316^{\pm 2.4}$	$563^{\pm 3.1}$	$1.80^{\pm 0.02}$	$2.84^{\pm 0.03}$
➥2× wide	$310^{\pm 2.0}$	$555^{\pm 2.8}$	$1.76^{\pm 0.03}$	$2.78^{\pm 0.01}$
LRLCN	$290^{\pm 1.4}$	$549^{\pm 1.9}$	$1.73^{\pm 0.03}$	$2.79^{\pm 0.01}$
ViT	$438^{\pm 2.8}$	$638^{\pm 3.1}$	$2.24^{\pm 0.04}$	$2.88^{\pm 0.03}$
SMoE	**$270^{\pm 2.0}$**	**$525^{\pm 2.0}$**	**$1.66^{\pm 0.02}$**	**$2.60^{\pm 0.01}$**

Ensemble post-processing ENS-10

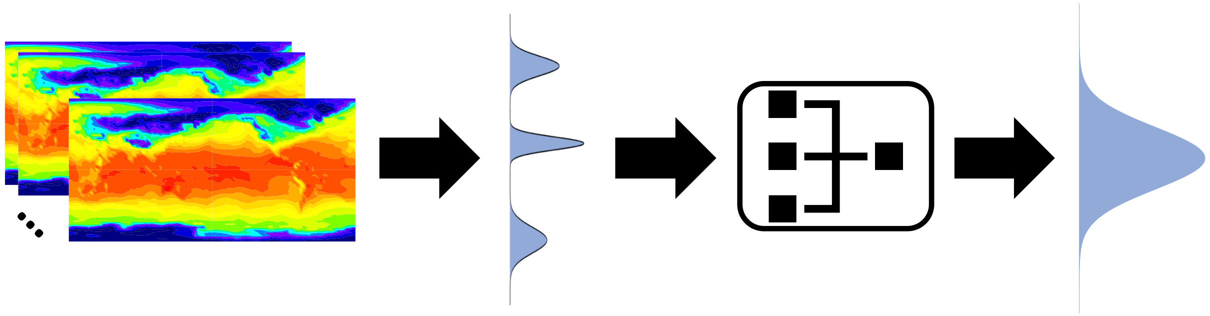

	Model	Z500 [m²s⁻²] 5 ens	Z500 [m²s⁻²] 10 ens	T850 [K] 5 ens	T850 [K] 10 ens	T2M [K] 5 ens	T2M [K] 10 ens
CRPS	EMOS	$79.12^{\pm 0.12}$	$78.80^{\pm 0.21}$	$0.721^{\pm 0.01}$	$0.706^{\pm 0.04}$	$0.720^{\pm 0.00}$	$0.711^{\pm 0.03}$
CRPS	U-Net	$76.54^{\pm 0.20}$	$76.18^{\pm 0.12}$	$0.685^{\pm 0.00}$	$0.670^{\pm 0.01}$	$0.657^{\pm 0.01}$	$0.644^{\pm 0.01}$
CRPS	SMoE	**$68.94^{\pm 0.14}$**	**$67.43^{\pm 0.12}$**	**$0.612^{\pm 0.01}$**	**$0.590^{\pm 0.02}$**	**$0.601^{\pm 0.02}$**	**$0.594^{\pm 0.02}$**
ECRPS	EMOS	$29.21^{\pm 0.18}$	$29.02^{\pm 0.13}$	$0.247^{\pm 0.00}$	$0.245^{\pm 0.02}$	$0.244^{\pm 0.00}$	$0.241^{\pm 0.02}$
ECRPS	U-Net	$27.78^{\pm 0.11}$	$27.55^{\pm 0.19}$	$0.230^{\pm 0.01}$	$0.229^{\pm 0.01}$	$0.225^{\pm 0.00}$	$0.220^{\pm 0.01}$
ECRPS	SMoE	**$23.79^{\pm 0.20}$**	**$23.10^{\pm 0.16}$**	**$0.207^{\pm 0.03}$**	**$0.197^{\pm 0.03}$**	**$0.199^{\pm 0.01}$**	**$0.190^{\pm 0.02}$**

Dryden, Hoefler: Spatial Mixture-of-Experts, NeurIPS'22, arXiv

Spatial Mixture-of-Experts

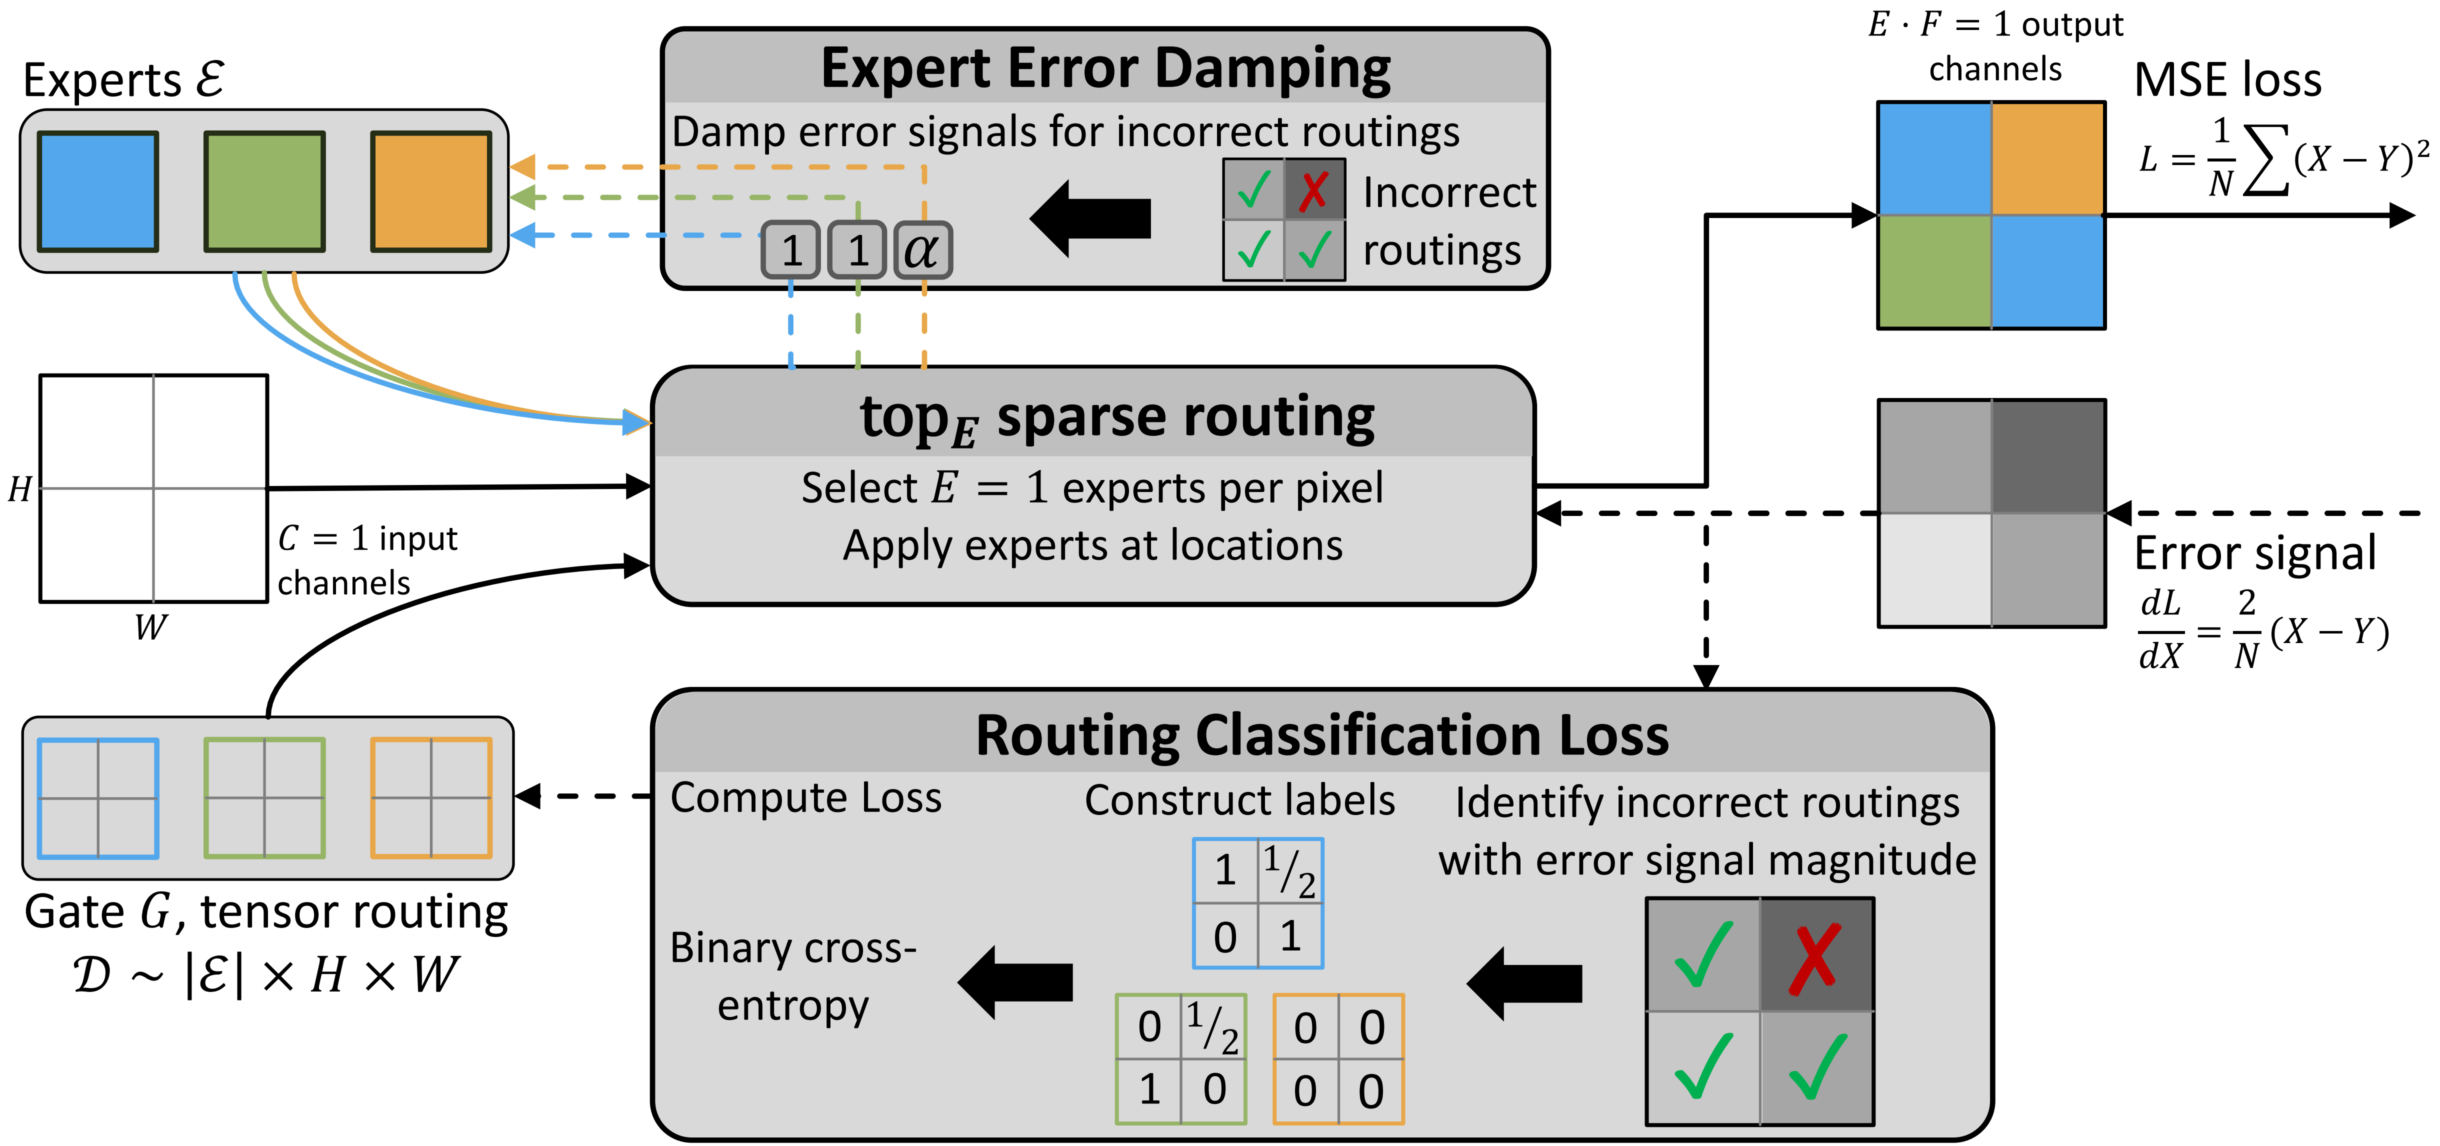

A first simple loop from CLOUDSC*

Data Parallelism

```
✓ ........ do  JK=1,KLEV
✓ ............. do  JL=1,KFDIA
                      ZQSM(JL,JK)=ZQSM(JL,JK)/(1.0-RE*ZQSM(JL,JK))
                    enddo
               enddo
```

Fully data parallel

Work	KLEV * KFDIA
Depth	1
Average Parallelism	KLEV * KFDIA

* examples are simplified for presentation purposes

A second more complex loop from CLOUDSC

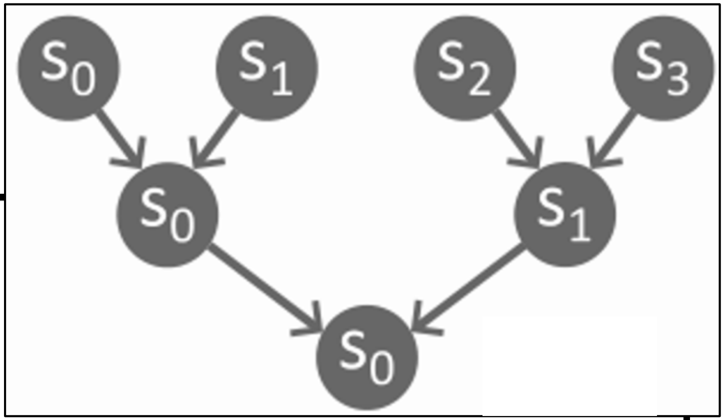

Data Parallelism

(array) accumulation prevents parallelization ☹

X ………. do JN=1, NSTEP-1
✓ ……………… do JL=1, KFDIA

```
do JN=1, NSTEP-1
   do JL=1, KFDIA
      ZQXN(JL, NSTEP) = ZQXN(JL, NSTEP)+ZQXN(JL, JN)
   enddo
enddo
```

Work	(NSTEP-1) * KFDIA	(NSTEP-1) * KFDIA
Depth	(NSTEP-1) * KFDIA	$\log_2$(NSTEP-1)
Average Parallelism	1	(NSTEP-1) * KFDIA / $\log_2$(NSTEP-1)

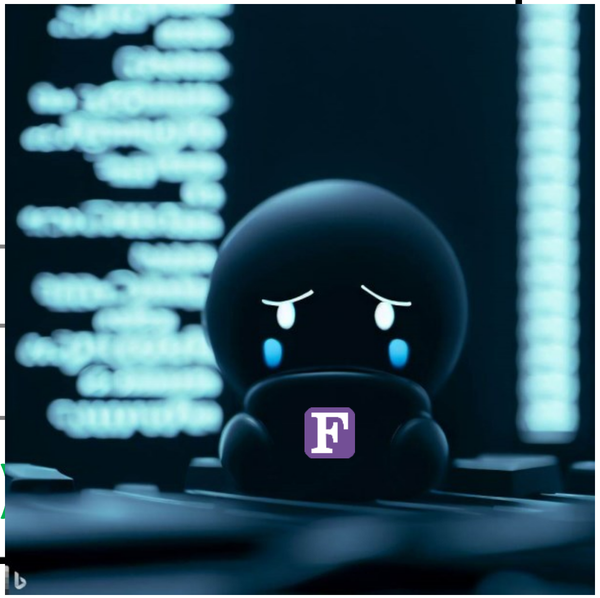

Now multiple realistic CLOUDSC loops

Data Parallelism

reuse of temporary variable prevents parallelization

L1

```
X ..... do JM=1,4
X ..... do JK=1,KLEV
X ..... do JL=1,KFDIA
  (1) if ZQX(JL,JK,JM)<RLMIN) then
  (2) ZQADJ=ZQX(JL,JK,JM)*ZQTMST
  (3) tend_q(JL,JK)=tend_q(JL,JK)+ZQADJ
  (4) tend_T(JL,JK)=tend_T(JL,JK)-RAL*Z
  (5) ZQX(JL,JK,JM)=0.0
```

L2

```
✓ ..... do JK=1,KLEV
✓ ..... do JL=1,KFDIA
  (6) ZQSM(JL,JK)=ZQSM(JL,JK)/(1.0-RE*ZQSM(JL
```

L3

```
✓ ..... do JK=1,KLEV
✓ ..... do JL=1,KFDIA
  (7) ZA(JL,JK)=MAX(0.0,MIN(1.0,ZA(JL,JK)))
  (8) ZLI(JL,JK)=ZQX(JL,JK,1)+ZQX(JL,JK,2)
  (9) if (ZLI(JL,JK)>RLMIN) then
  (10) ZLFRAC(JL,JK)=ZQX(JL,JK,1)/ZLI(JL,JK)
  else
  (11) ZLFRAC(JL,JK)=0.0
```

Order Constraints

→ Happens-before

1 → 2	Control	
2 → 3	RAW	
2 → 4	RAW	
2 → 5	WAR	

No order constraint: L2 || L1

5 → 8	RAW
8 → 9	RAW
9 → 10	Control
9 → 11	Control

L3 → L1

	L1		L2		L3		
Work	4 * KLEV * KFDIA * (1+4)		KLEV * KFDIA		KLEV * KFDIA * 4	L1	KLEV * KFDIA * 25
Depth	log2(4) * 1 *1 * (1+2)		1		1 * 1 * (2+1)	L2	8
Average Parallelism	KLEV * KFDIA * 10/3		KLEV * KFDIA		KLEV * KFDIA * 4/3	L3	KLEV * KFDIA * 25/8

Weather

Medium-range weather prediction WeatherBench

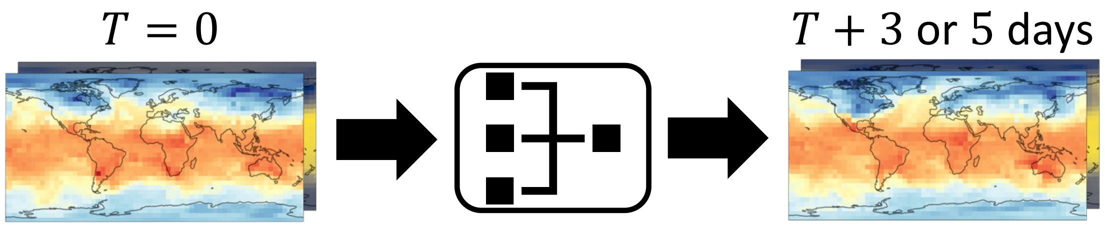

Model	Z500 [m^2s^{-2}]		T850 [K]	
	3 days	5 days	3 days	5 days
Rasp & Thuerey	$316^{\pm2.4}$	$563^{\pm3.1}$	$1.80^{\pm0.02}$	$2.84^{\pm0.03}$
➥2× wide	$310^{\pm2.0}$	$555^{\pm2.8}$	$1.76^{\pm0.03}$	$2.78^{\pm0.01}$
LRLCN	$290^{\pm1.4}$	$549^{\pm1.9}$	$1.73^{\pm0.03}$	$2.79^{\pm0.01}$
ViT	$438^{\pm2.8}$	$638^{\pm3.1}$	$2.24^{\pm0.04}$	$2.88^{\pm0.03}$
SMoE	**$270^{\pm2.0}$**	**$525^{\pm2.0}$**	**$1.66^{\pm0.02}$**	**$2.60^{\pm0.01}$**

[Rasp et al., 2020; Rasp & Thuerey, 2021; Ashkboos et al., 2022; Elsayed et al., 2020; Dosovitskiy et al., 2021]

Ensemble post-processing ENS-10

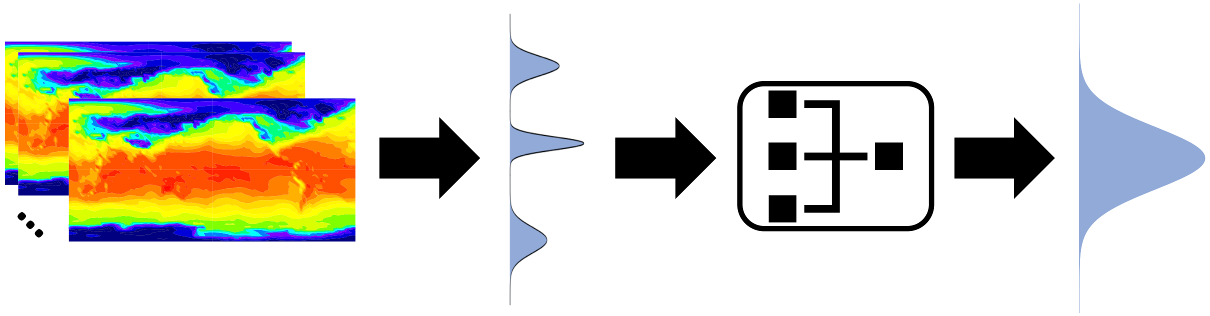

	Model	Z500 [m^2s^{-2}]		T850 [K]		T2M [K]	
		5 ens	10 ens	5 ens	10 ens	5 ens	10 ens
CRPS	EMOS	$79.12^{\pm0.12}$	$78.80^{\pm0.21}$	$0.721^{\pm0.01}$	$0.706^{\pm0.04}$	$0.720^{\pm0.00}$	$0.711^{\pm0.03}$
	U-Net	$76.54^{\pm0.20}$	$76.18^{\pm0.12}$	$0.685^{\pm0.00}$	$0.670^{\pm0.01}$	$0.657^{\pm0.01}$	$0.644^{\pm0.01}$
	SMoE	**$68.94^{\pm0.14}$**	**$67.43^{\pm0.12}$**	**$0.612^{\pm0.01}$**	**$0.590^{\pm0.02}$**	**$0.601^{\pm0.02}$**	**$0.594^{\pm0.02}$**
EECRPS	EMOS	$29.21^{\pm0.18}$	$29.02^{\pm0.13}$	$0.247^{\pm0.00}$	$0.245^{\pm0.02}$	$0.244^{\pm0.00}$	$0.241^{\pm0.02}$
	U-Net	$27.78^{\pm0.11}$	$27.55^{\pm0.19}$	$0.230^{\pm0.01}$	$0.229^{\pm0.01}$	$0.225^{\pm0.00}$	$0.220^{\pm0.01}$
	SMoE	**$23.79^{\pm0.20}$**	**$23.10^{\pm0.16}$**	**$0.207^{\pm0.03}$**	**$0.197^{\pm0.03}$**	**$0.199^{\pm0.01}$**	**$0.190^{\pm0.02}$**