THE MAELSTROM PROJECT

Project overview and experience from the 1st MAELSTROM Boot Camp

2023-02-09 I **MICHAEL LANGGUTH**, BING GONG, MARTIN SCHULTZ

EuroHPC Joint Undertaking

Mitglied der Helmholtz-Gemeinschaft

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

JÜLICH Forschungszentrum

THE MAELSTROM PROJECT

Some key information

- Acronym for MAchinE Learning for Scalable meTeoROlogy and cliMate
- Interdisciplinary project of different universities, institutions and companies funded by EuroHPC Joint Undertaking
  → coordination by ECMWF
- Project period: April 2021 – March 2024

Overarching target

- Establish scalable machine learning (ML) applications in the weather and climate (W&C) domain

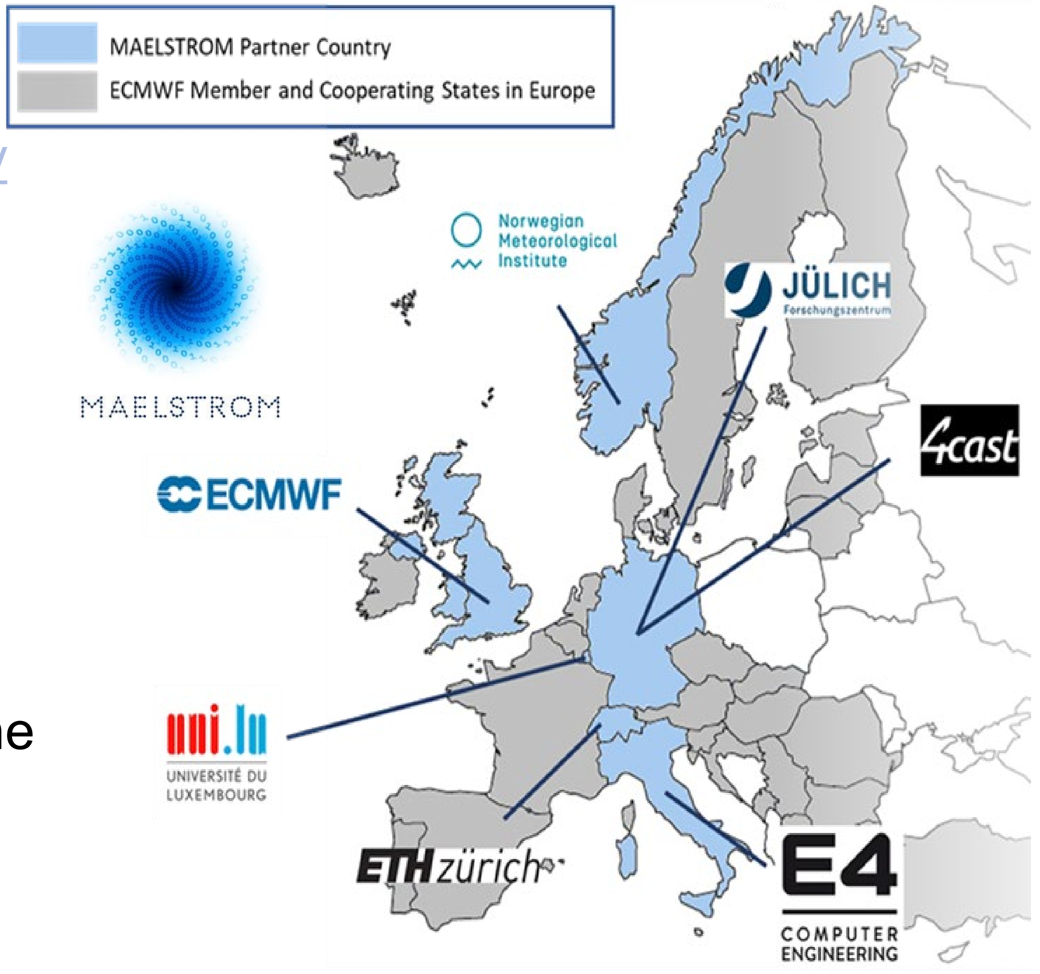

THE MAELSTROM PROJECT

Motivation

- Simulation of the atmosphere including its interactions with other Earth subsystems is inherently challenging:
    - High-dimensional Earth system
    - Highly non-linear (chaotic) atmospheric processes
    - Limited resolution realizable in numerical models
    - Errors und uncertainty in simulation data, but also in observations
- Successful application of Deep Learning (DL) in various domains: capable to learn highly non-linear mappings in a data-driven way

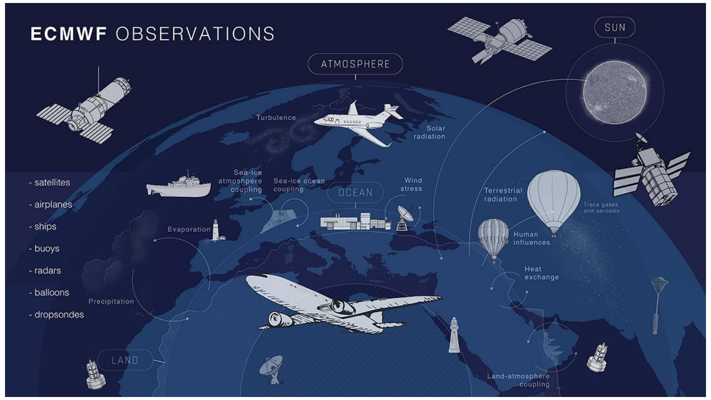

How can we tailor DL to meteorological problems?

THE MAELSTROM PROJECT

Motivation

- Recently substantial progress in DL applications for W&C
- Example: Deep neural networks for weather forecasting
    - FourCastNet by [Patha et al.] on 8th August 2022
    - PanguWeather by [Bi et al.] on 3th November 2022
    - GraphCast by [Lam et al.] on 24th December 2022
- However, myriad of options:

DL techniques (Conv Nets, GANs, Transformers etc.)
\+ **DL frameworks** (Tensorflow, Keras, PyTorch etc.)
\+ **bespoken hardware** (CPUs, GPUs, TPUs etc.)
= Obstacle for (early) scientists

Score card GraphCast vs. IFS

Model level and variable

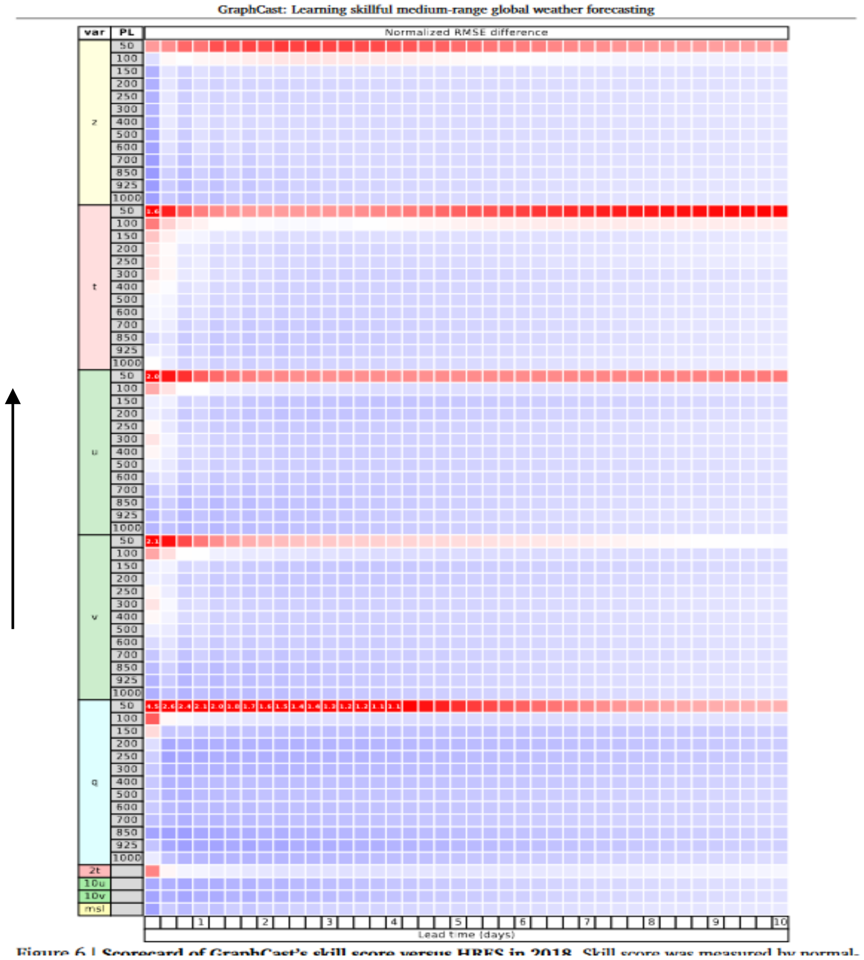

Forecast lead time

JÜLICH Forschungszentrum

THE MAELSTROM PROJECT

From the challenge to the approach

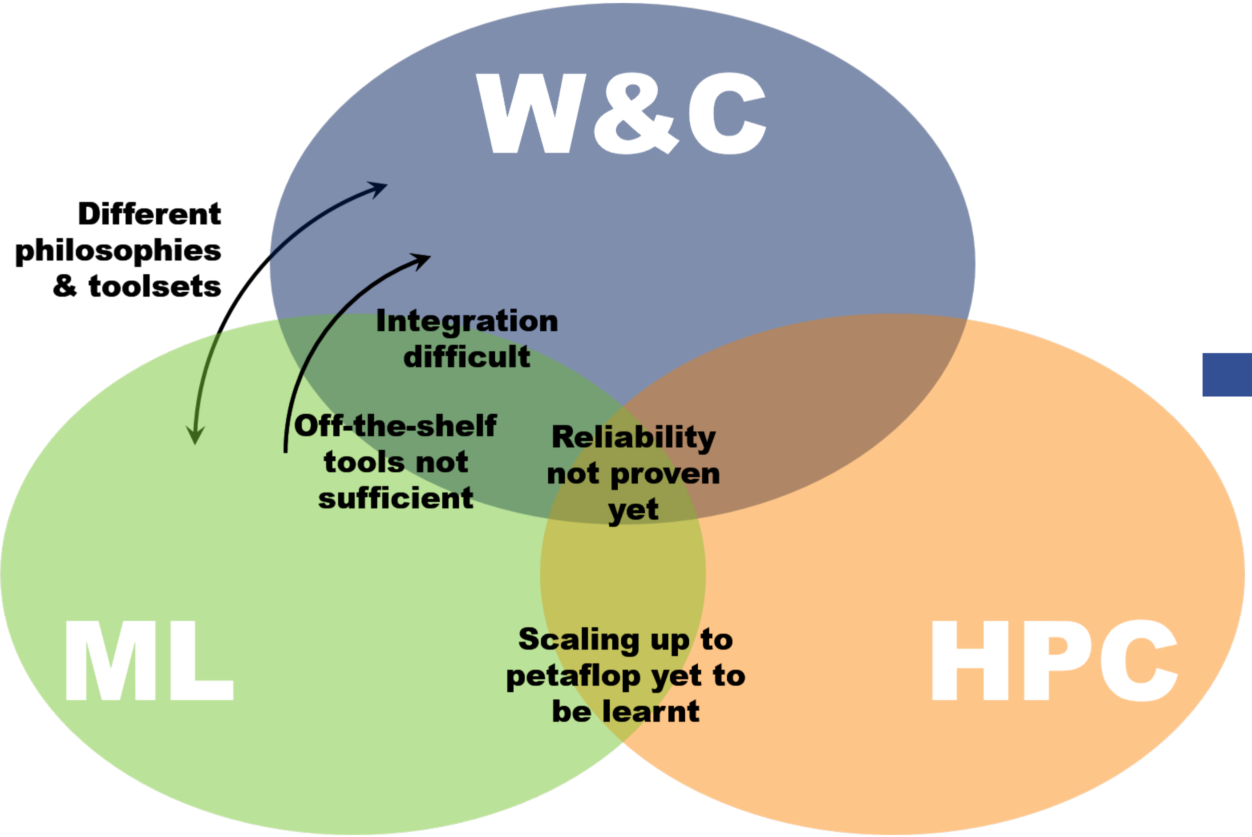

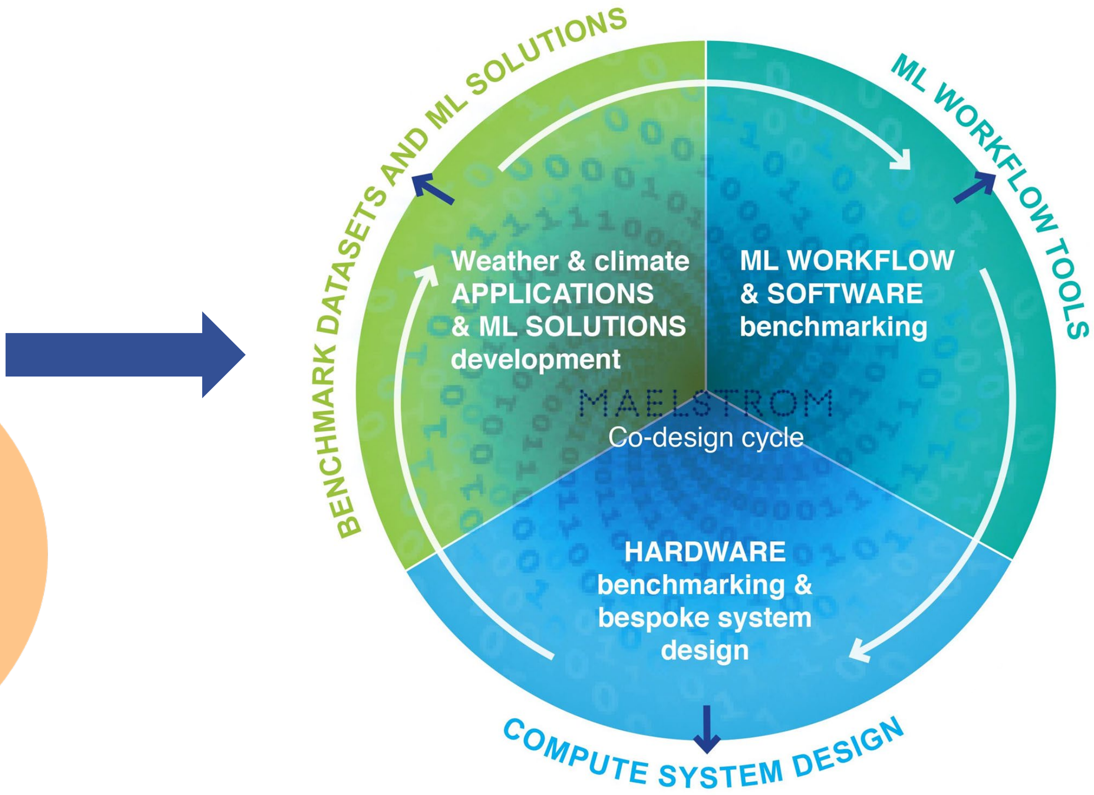

THE MAELSTROM PROJECT

The objectives in MAELSTROM

MAELSTROM

O1: To open W&C predictions as a new usage domain for machine learning applications that can exploit exaflop performance. **WP 1**

O2: To develop the optimal software environment to develop exascale-ready machine learning tools that can be used across the workflow of W&C predictions. **WP 2**

O3: To optimise compute system designs for machine learning applications for W&C predictions at the node and system level and to transfer this knowledge to other machine learning applications that will use future EuroHPC systems. **WP 3**

THE MAELSTROM PROJECT

MAELSTROM highlights – The current status

- First successful closing of co-design cycle:
  Dataset development (D1.1) → software development (D2.2) → hardware benchmark (D1.4) → Application development (D1.4)
- Achievements
  - All deliverables and milestones reached so far
  - Documentation and publication of MAELSTROM datasets
  - Usage of MAELSTROM applications as benchmark by industry (e.g. GRAPHCORE, Microsoft)
  - First hardware benchmarks
  - Successful dissemination
    1. MAELSTROM Dissemination Workshop with >200 participants (March ‘22)
    2. MAELSTROM Boot Camp with >30 participants on FZJ campus (September ‘22)

Check our deliverables here

THE MAELSTROM PROJECT

The MAELSTROM applications – an overview on WP1

A1: Blend citizen observations and numerical weather forecasts

A2: Incorporate social media data into prediction framework

A3: Neural network emulators for faster weather forecast models & data assimilation

A4: Improved ensemble predictions in forecast post-processing

A5: Improved local weather predictions in forecast post-processing

A6: Bespoke weather forecasts to support energy production in Europe

STATISTICAL DOWNSCALING WITH DL

The work in A5 of the MAELSTROM project - Motivation

MAELSTROM

- Improve local weather data = statistical downscaling
- High spatial variability in Earth Science data,
  but limited spatial resolution in numerical models
- Various adverse local effects
  - Local frost → loss in agriculture
  - Local precipitation events → flooding
- Increase in spatial resolution…
  - … is computationally costly
  - … imposes challenges on parameterization schemes

Protection measurement against nighty freeze in Neuwied. Photo: Rhein-Zeitung.

Objective:

- Develop generalizable and accurate neural networks for statistical downscaling
- Two-fold task for Earth Science data: Super-resolution + bias correction

JÜLICH Forschungszentrum

STATISTICAL DOWNSCALING WITH DL

The work in A5 of the MAELSTROM project - Approach

MAELSTROM

- Objective: Emulate a high-resolution reanalysis product
- Input: ERA5-reanalysis ($\Delta x_{ERA5} \cong 0.3°$),
  target: COSMO REA6-reanalysis ($\Delta x_{CREA6} \cong 0.055°$)
- The ERA5-dataset:
  - Global reanalysis dataset based on IFS model
  - Data coverage: 1979 – near real-time
- The COSMO REA6-dataset:
  - Regional reanalysis dataset based on *dynamical* downscaling with the COSMO model (ERA Interim)
  - Data coverage: Jan 1995 – Aug 2019
  - Added value against global ERA5-reanalysis dataset

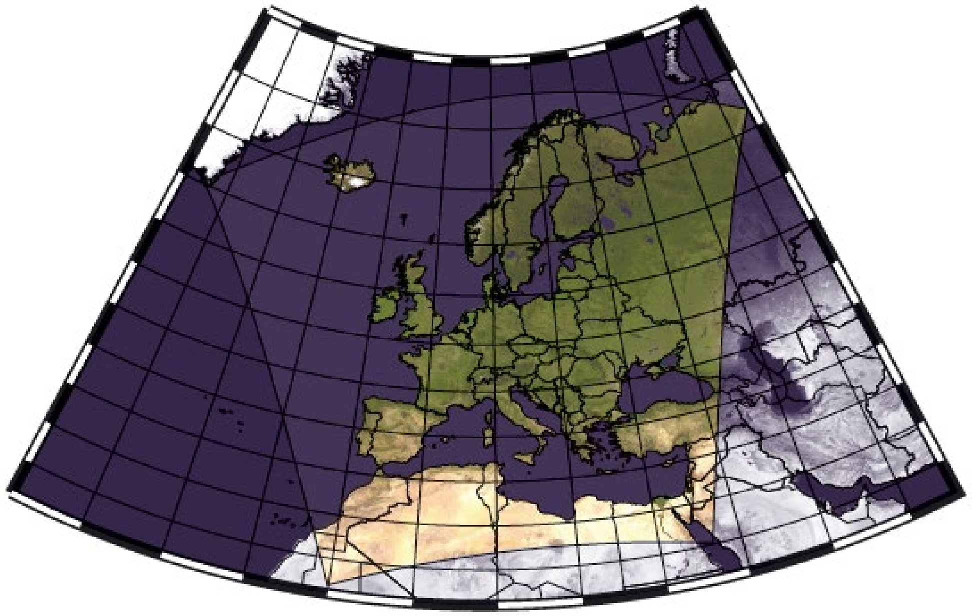

Model domain of the COSMO REA6 reanalysis dataset (from Bollmeyer, 2015a).

STATISTICAL DOWNSCALING WITH DL

The work in A5 of the MAELSTROM project - Approach

MAELSTROM

- Task: downscale 2m temperature (T2m)
- Predictors to encode planetary boundary layer state:
  T2m, 850 hPa- and 925 hPa temperature, 10m wind, boundary layer height, surface heat fluxes, surface topography
- Target region: 120x96 grid points over Central Europe
- Data published as [Tier 2-dataset][1]
- Test two neural network architectures:
  1) U-Net
  2) Wasserstein GAN (WGAN) with U-Net as generator

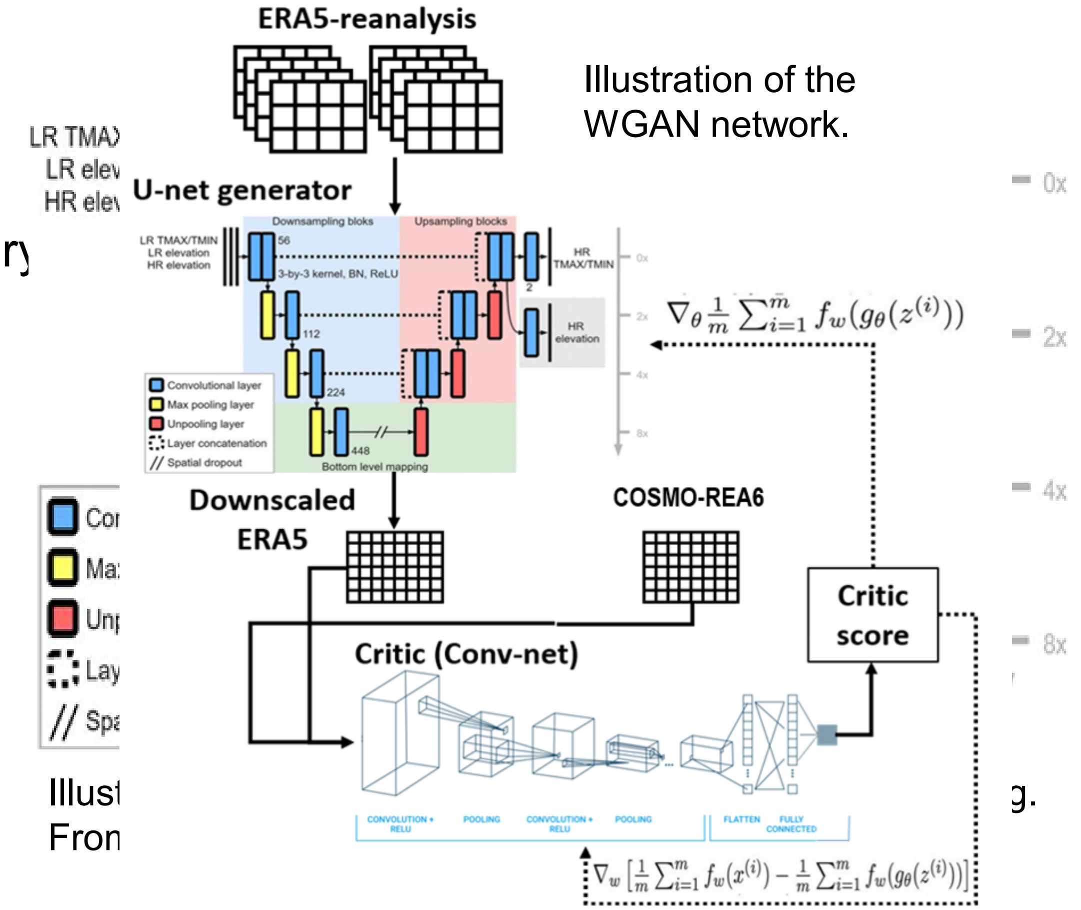

Illustration of the WGAN network.

[1] 25 years data (1/2 currently), hourly → ~100 GB (more predictors planned)

JÜLICH Forschungszentrum

STATISTICAL DOWNSCALING WITH DL

Results for Tier2-dataset

U-Net | WGAN

RMSE

Spatial variability

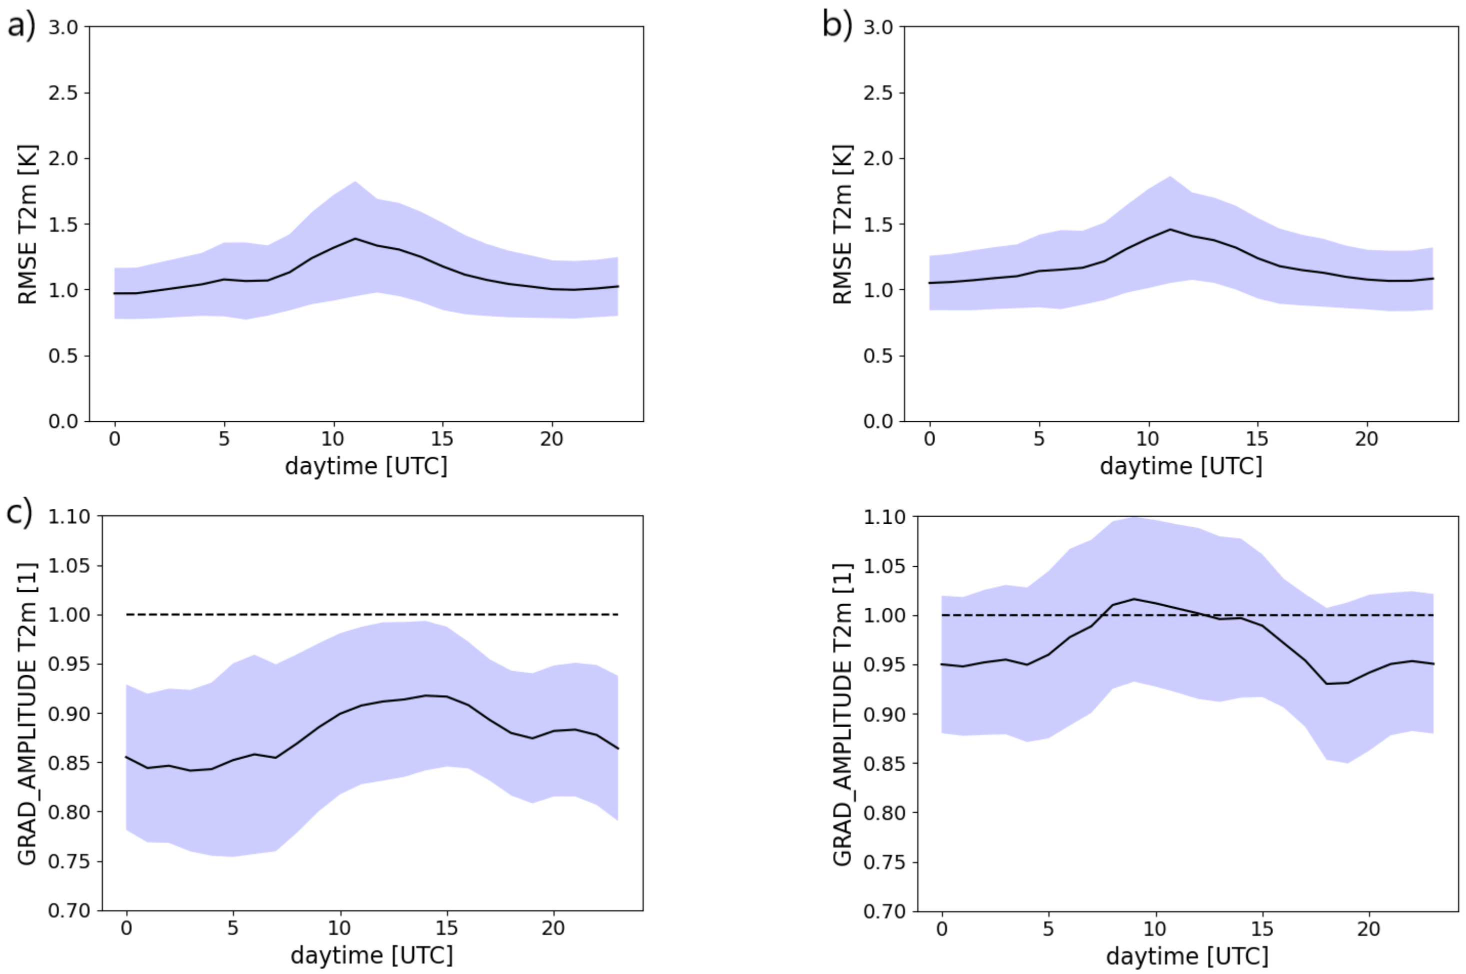

Better

Optimal value

STATISTICAL DOWNSCALING WITH DL

Outlook - T2m downscaling task

MAELSTROM

Open issues

- Increased errors around noon during summer
- Spatial variability underestimated at night

Next steps

- Add more informative predictors
- Test (deeper) state-of-the-art architectures adapted from computer vision (cf. next slides)
- Publish benchmark dataset
  - Based on ERA5- and COSMO REA6-reanalysis datasets
  - Various downscaling tasks
  - Baseline scores for architectures from literature

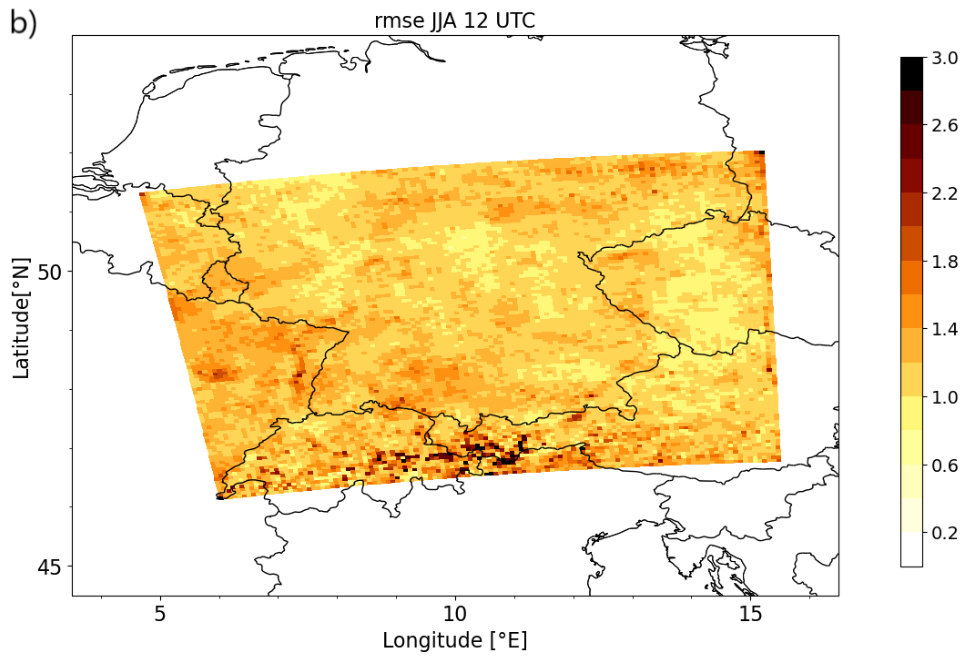

Spatial distribution of RMSE for the summer months JJA at 12 UTC on the test dataset (2017).

JÜLICH Forschungszentrum

STATISTICAL DOWNSCALING WITH DL

Outlook – precipitation downscaling task

MAELSTROM

- Recent progress in precipitation downscaling
- Reference studies, e.g. *Leinonen et al, 2021*, *Rasp and Price, 2022*, *Harris et al, 2022*:
  - Use GAN architectures
  - Map from NWP output to radar observations
  - Obtain promising results in downscaling + bias correction with sophisticated data handling
- Can we further improve with state-of-the-art model architectures adapted from computer vision?

Flooded Ahrtal
2021-07-15
Source:
info:bild.de

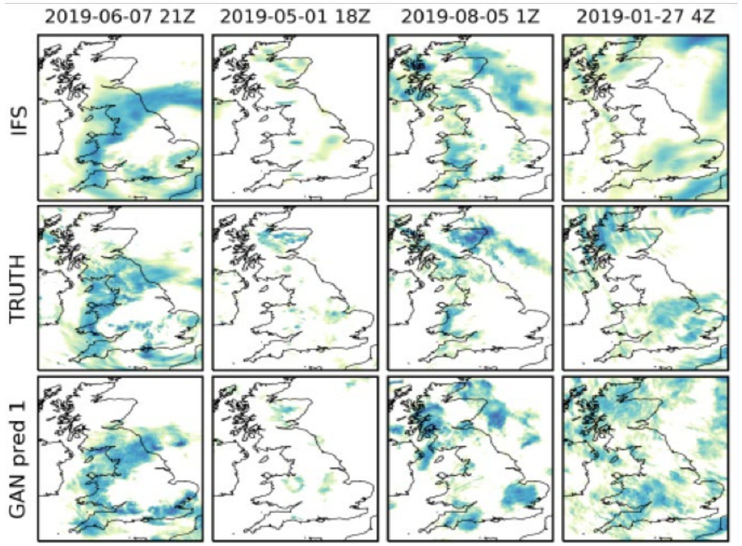

Continuation of work performed in DeepRain

SPONSORED BY THE
Federal Ministry of Education and Research

STATISTICAL DOWNSCALING WITH DL

Outlook – precipitation downscaling task

MAELSTROM

- Downscaling task: IFS short-range forecasts ($\Delta x_{IFS} = 0.1°$) to (remapped) RADKLIM observations ($\Delta x_{RAD} = 0.01°$)
- Test Swin-Transformer architecture and diffusion networks in addition to WGAN/U-Net baselines
- (very) preliminary results with diffusion network (with U-Net backbone):
    - Realistic spatial variability in data, but …
    - … strong underestimation of intensity

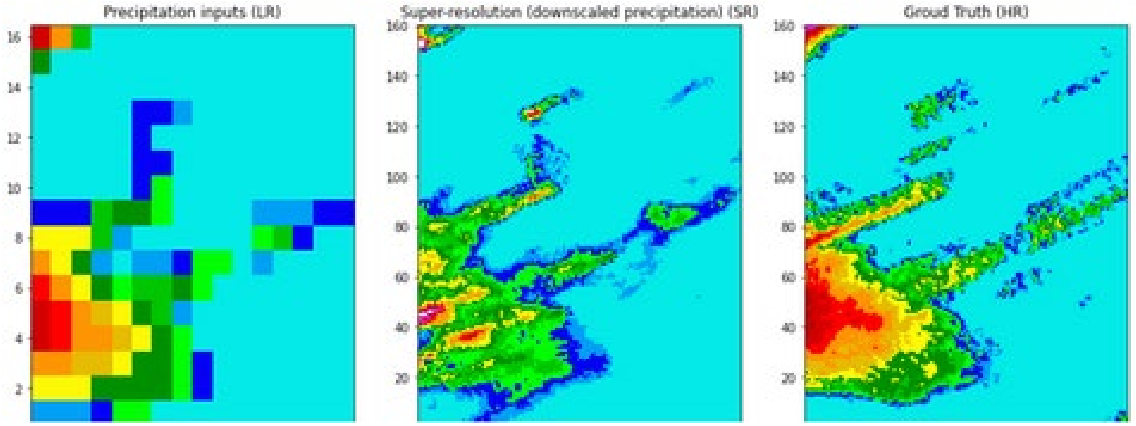

Example precipitation downscaling result: The coarse-grained input (left) was downscaled with a diffusion network (center). The ground truth RADKLIM data is displayed on the right.

STATISTICAL DOWNSCALING WITH DL

Outlook – and beyond?

- Yet, MAELSTROM applications do not require excessive scalability → ‘the MAELSTROM elephant’
- Inspiration from NLP: use transformer networks pre-trained on large corpus of unlabelled data (e.g. ERA5)
  - ✓ Learn abstraction of atmospheric dynamics with BERT-approach
  - ✓ General purpose for great variety of downstream applications

  → **AtmoRep** (in collaboration with Christian Lessig, University Magdeburg)

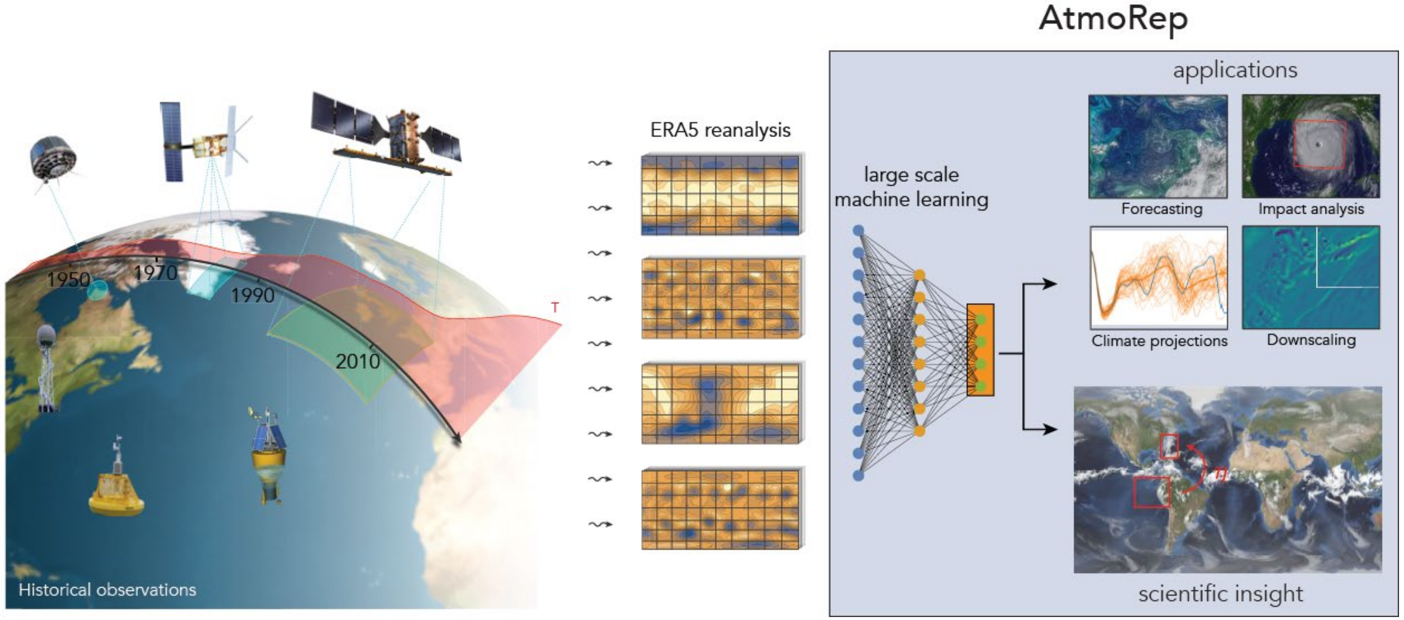

also part of *WestAI* service center

SPONSORED BY THE
Federal Ministry of Education and Research

MAELSTROM

JÜLICH Forschungszentrum

1ST MAELSTROM BOOTCAMP

Experience with a workshop on-site

MAELSTROM

- D3.4: Workshop (Boot Camp) for early dissemination of MAELSTROM outcomes
- Scope:
  - Target audience: Young researchers from Machine Learning and/or Meteorology
  - Train participants on developed DL applications in WP1
  - Application on HPC systems
- Event took place from 27th to 30th September at JSC
  - 32 participants (43 registers), 16 tutors
  - Two funded participants

Group photo of all participants and supervisors taken in front of the Rotunde on the FZJ campus.

JÜLICH Forschungszentrum

1ST MAELSTROM BOOTCAMP

Experience with a workshop on-site

MAELSTROM

- Boot Camp schedule:
  - Introductory lectures (1-1.5h) on MAELSTROM, DL, meteorology, HPC-systems and scalable DL
  - Hands-on training in applications (participants split up in teams for their favorite application)
  - Daily 5-min pitches from all teams → increase communication & networking (in addition to joint coffee breaks and dinners)
- Hands-on-training with Jupyter Notebooks on Jupyter-JSC (access to JSC‘s HPC-system)
  - ✓ Very interactive presentations
  - ✓ Convenient tool for interactive coding

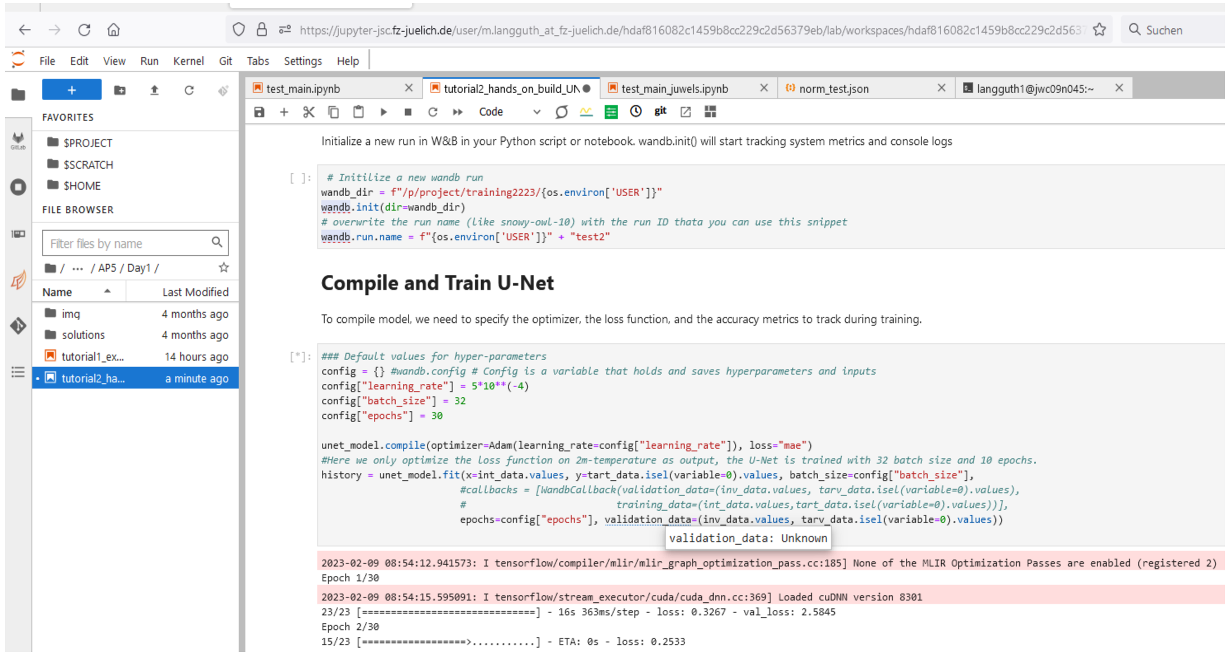

Example Jupyter Notebook from Application 5 used in the hands-on tutorials during the MAELSTROM Boot Camp.

1ST MAELSTROM BOOTCAMP

Experience with a workshop on-site

- Mostly positive feedback
  - Informative worshop material
  - Good balance between presentations and hands-on sessions
- Some ‚issues‘
  - Stability of JSC-Jupyter (fixed since October 22)
  - 4-day event rather short for the scope of the Boot Camp
- Boot Camp material will be published soon (~next week)

Join the 2nd Hackathon at ECMWF (Nov'22)
(check the MAELSTROM website)

Check deliverable 4.5

Questions?

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".