Distributed Deep Learning on HPC for Infilling Holes in Spatial Precipitation Data

From: Johannes Meuer[1], Étienne Plésiat[1], Hannes Thiemann[1], Thomas Ludwig[1] and Christopher Kadow[1]
[1]German Climate Computing Center (DKRZ), Hamburg, Germany

DKRZ
www.dkrz.de

Introduction & Motivation

Spatial precipitation fields are recorded via weather **radars** and are used for predicting future climate states. In Germany, we have the **RADOLAN** data set [1], which combines spatial radar recordings and weather station measurements:

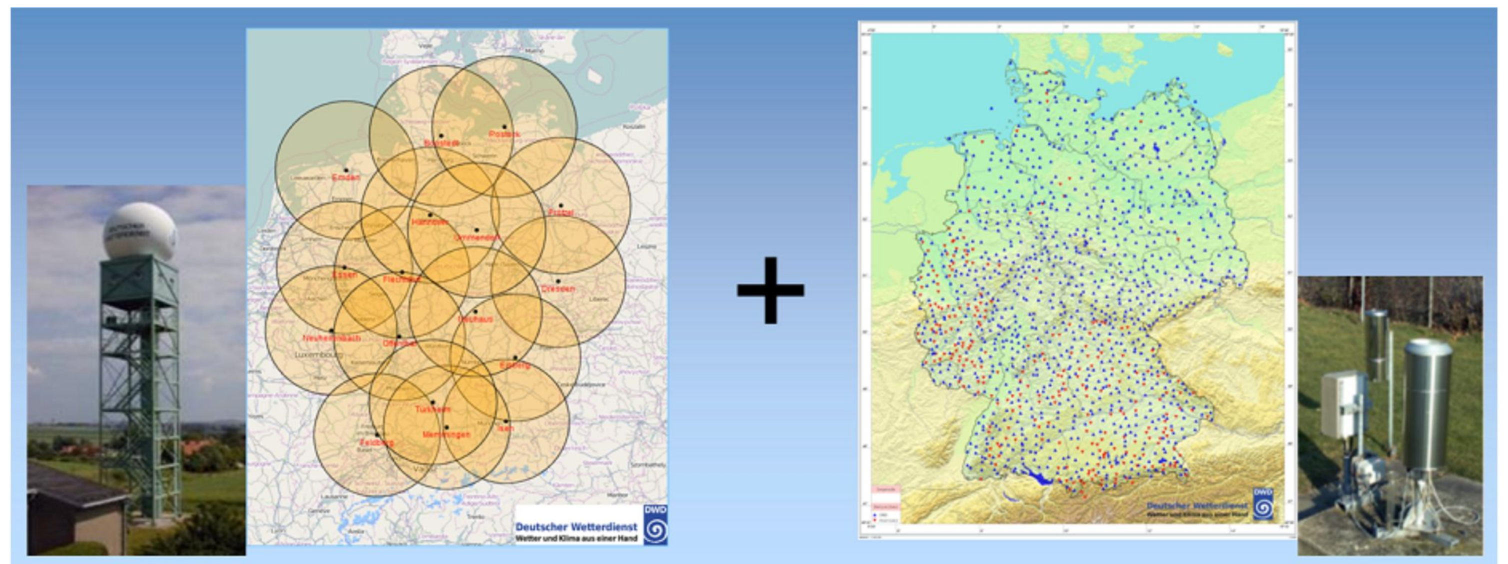

This highly resolved spatial precipitation data is used by other prediction models to create **nowcasts** or **severe weather warnings**. Since the predictions of such can be a **key factor** for policies, services, plans in the agricultural sector and many more, it is important to trust these predictions in order to take the right actions. However, **radar failures** occur due to blockage of radar beams, near-ground blind zones and other issues, which can lead to holes in the data:

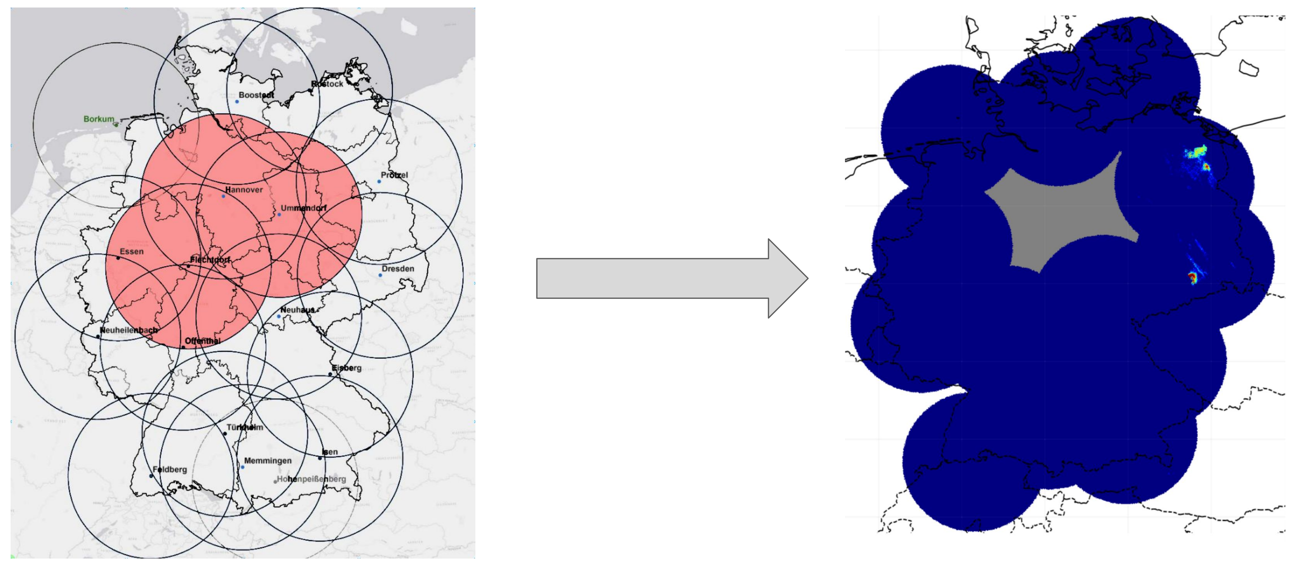

Radar coverage over Germany highlighting three radar failures that occurred simultaneously

Precipitation Field with a large missing value region resulting from the radar failures

These holes occur **very frequently** in the data set until today and affect the **accuracy** of prediction models that rely on this data. The following images show further examples of such that occurred in 2002:

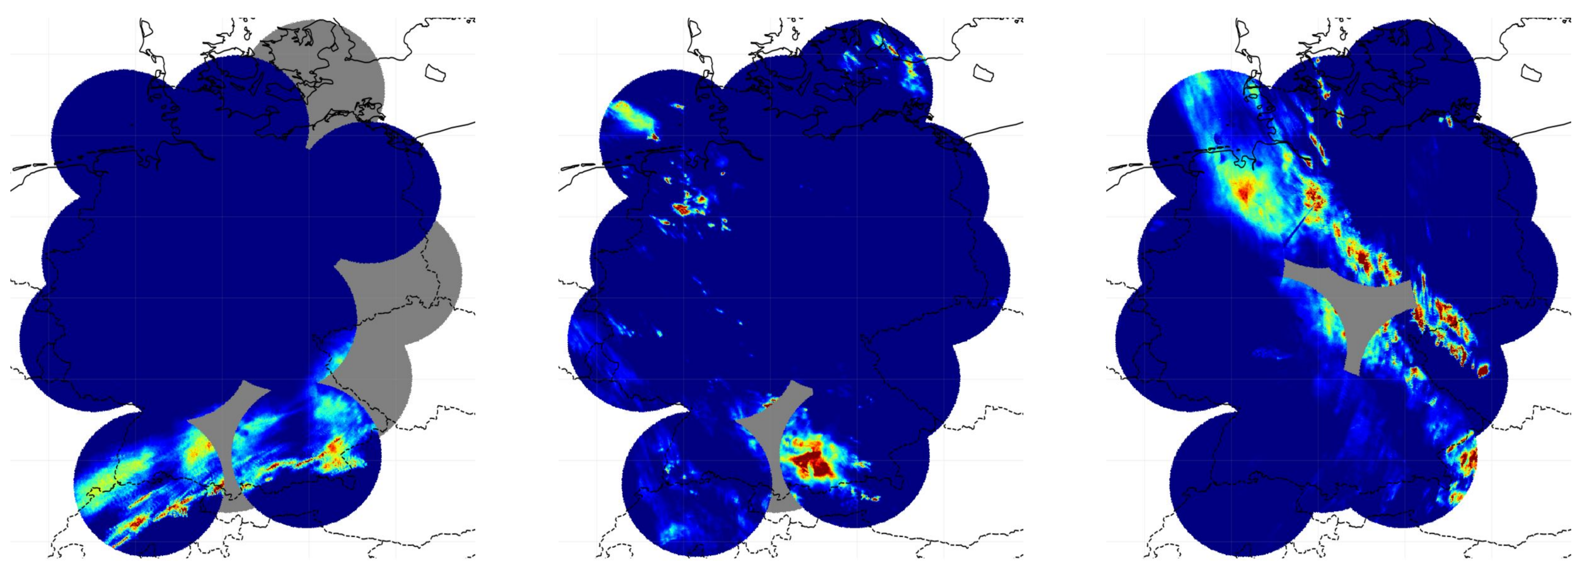

Our goal was to design a technique that is able to efficiently infill these holes in the **RADOLAN** data set in order to deal with radar failures and optimize the performance of the prediction models that use this data.

DKRZ DEUTSCHES KLIMARECHENZENTRUM

SC22

Johannes Meuer, meuer@dkrz.de

Methods

Image Inpainting is a method that was originally developed to **repair images** that were damaged by raindrops, to **optimize the quality** of old pictures or even **increase the resolution** of low quality images. Here, **Convolutional Neural Networks** have proven to produce astonishing results. Liu et al. [2] introduce **partial convolutions** that further enable infilling of irregular shaped holes:

Masked Image — Reconstructed Image — Original Image

For training, the method **masks** a complete image to simulate a corruption. The **U-shaped CNN** then performs **partial convolutions**, only considering the existing values in the image, and reconstructs the missing values. The reconstructed image is then compared to the original one.

This **CNN**, which has also proven to be sufficient for infilling missing climate data [3], is the baseline model of our approach. Since precipitation is **highly non-linear in time and space**, we further investigated two potential improvements and a final third investigation, that combines the previous two:

1. Considering additional **temporal** information in a **Convolutional Long Short-Term Memory** module [4]
2. Considering additional **atmospheric** information in an **Attention** module [5]

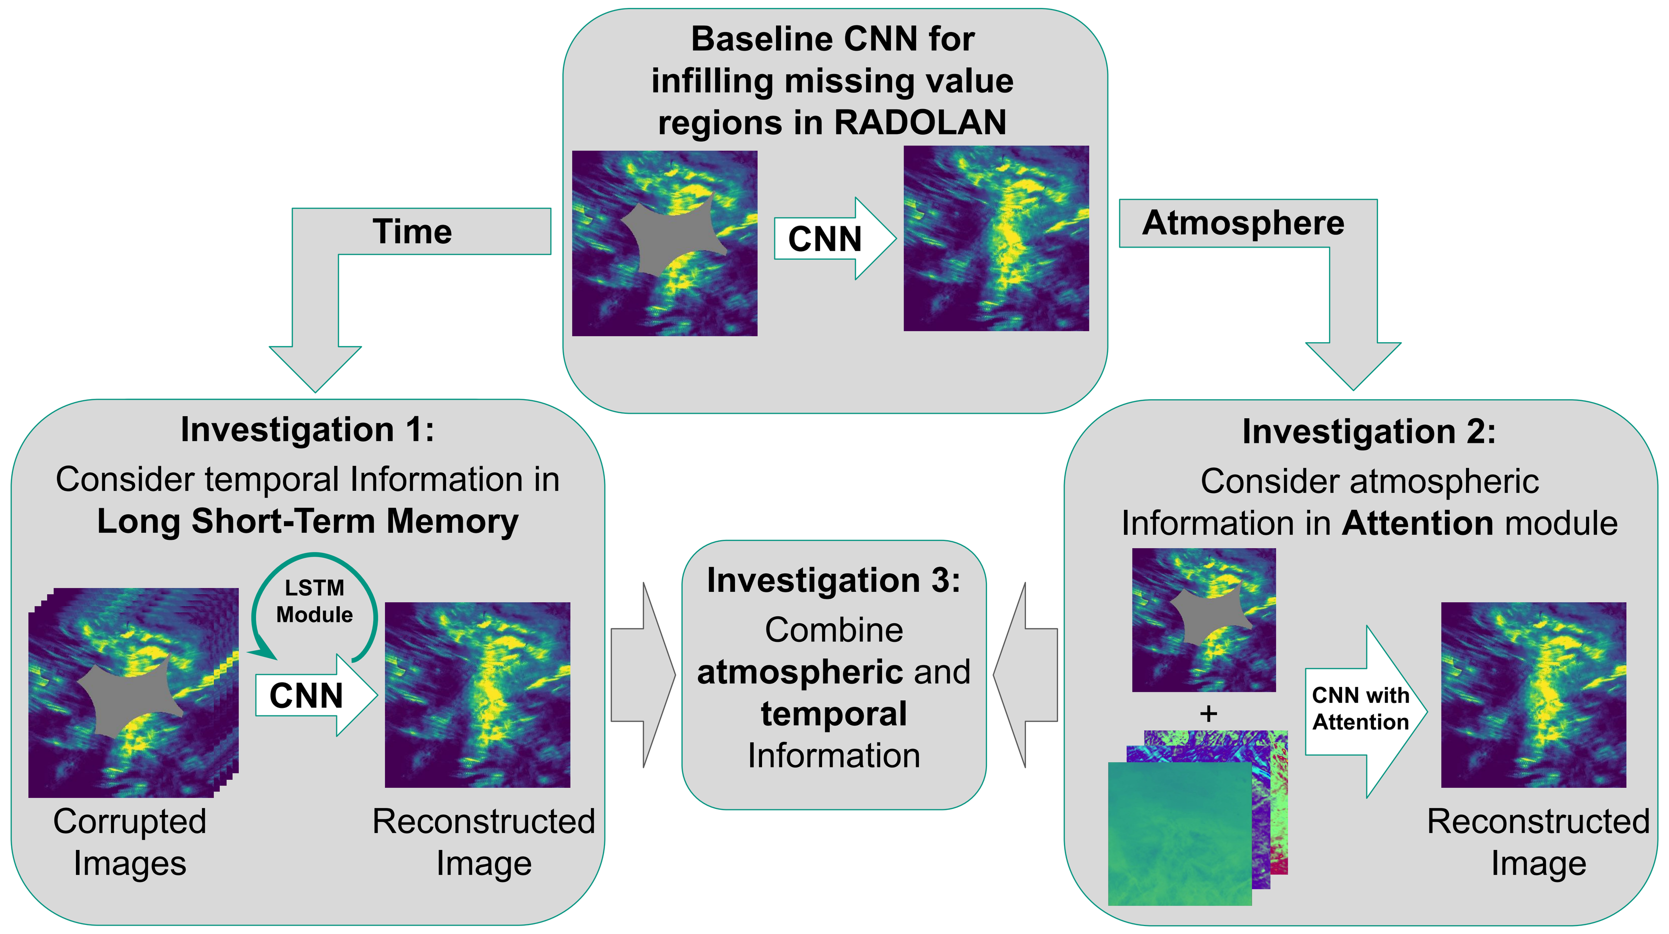

A general problem that we encountered was the amount of computational resources that the model required. Especially when considering the 1024x1024 sized images, containing the complete **RADOLAN** grid, we faced hardware limitations or extremely long training durations, even though the DKRZ provided high-end HPC hardware. Hence, we considered **distributing** the training batches on **multiple GPUs.** The **PyTorch** deep learning framework provides functionalities to efficiently parallelize the processing by splitting the input batch across multiple devices:

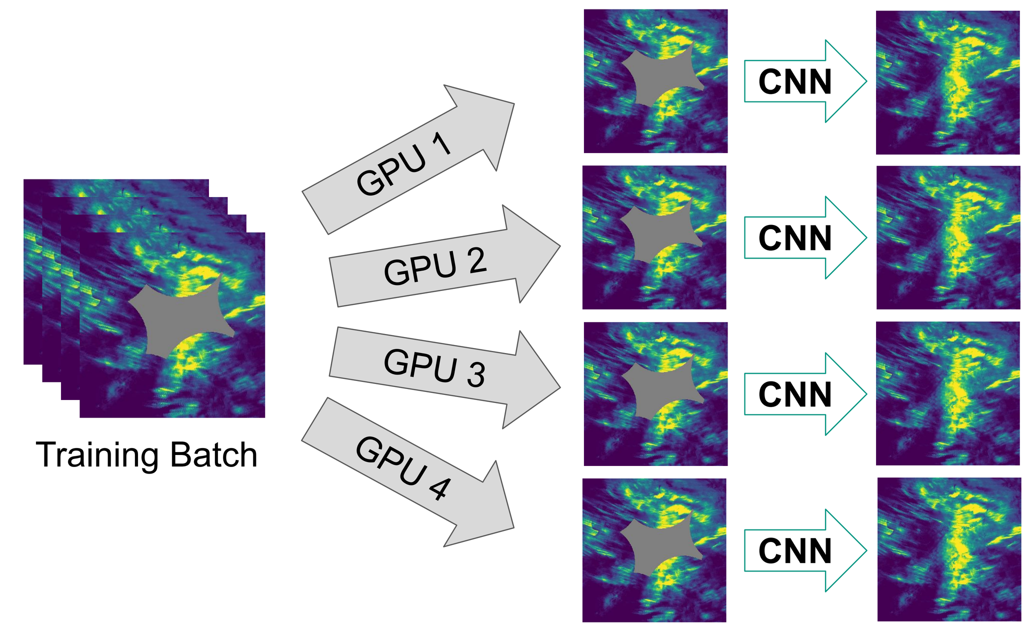

Results & Benchmarks

The results of the different investigations were evaluated on a large set of verification metrics. The **Baseline** model already produced quite good results. However, when considering additional **temporal or atmospheric** information, the accuracy of the results even increased. Finally, when **combining** the both improvements, we were able to achieve the overall best results:

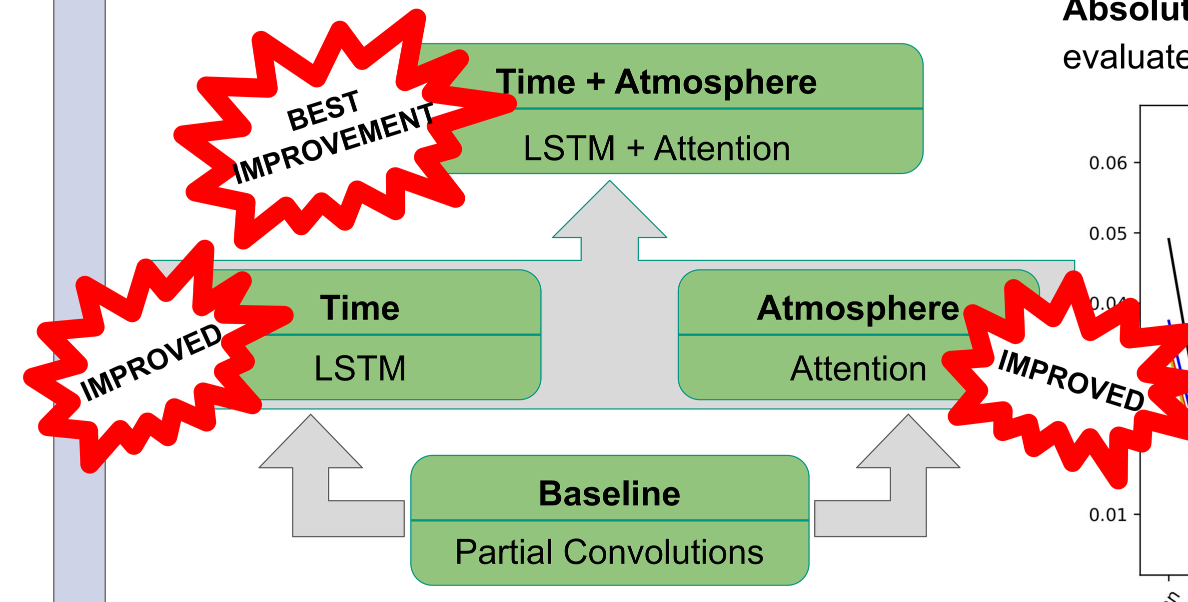

Absolute mean error in mm/h for the evaluated year 2012 as a monthly time series:

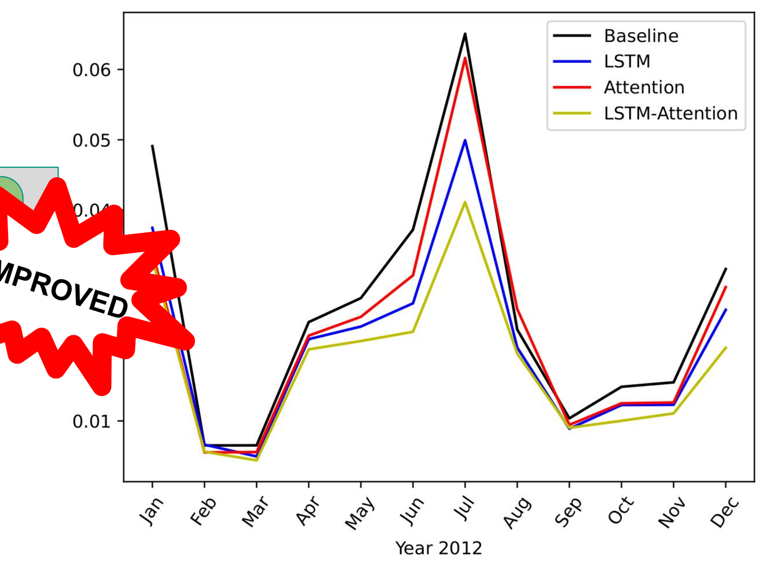

Model	RMSE in mm/h ↓	AME ↓ in mm/h	Temporal Corr. ↑	Spatial Corr. ↑
Baseline	0.3113	0.0827	0.9167	0.3753
LSTM	0.3024	0.0627	0.9532	0.3876
Attention	0.3107	0.0764	0.9222	0.3875
LSTM + Attention	**0.3004**	**0.0566**	**0.9575**	**0.4049**

The trained model was then applied to infill all incomplete grids in the **RADOLAN** data set, providing data with no missing values. Here are three representational samples:

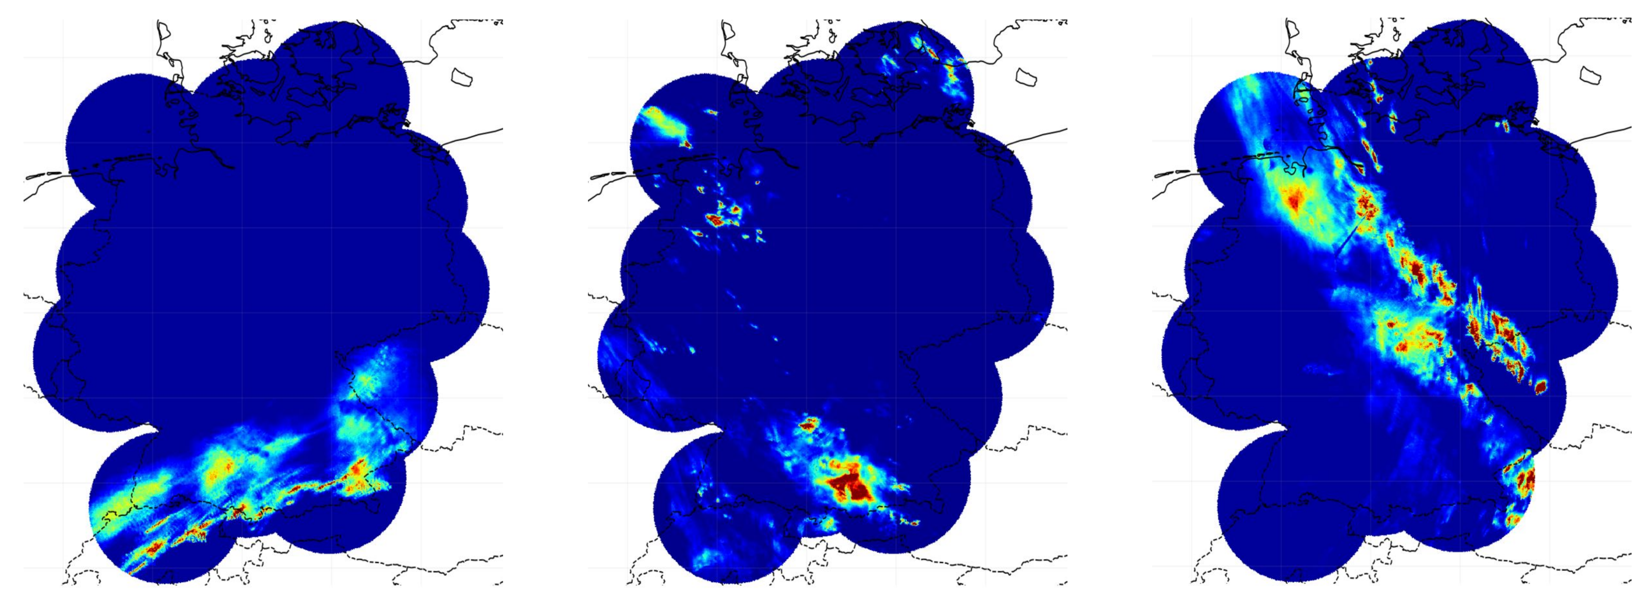

We further evaluated the training performance on different hardware components where we considered a scenario without **distributed deep learning** and one with. For the hardware components, the **Mistral** HPC system by the DKRZ provided a **NVIDIA Tesla M40** and a **NVIDIA Tesla V100** node. Furthermore, the new **Levante** HPC system by the DKRZ provided a **NVIDIA A100** node. The performance benchmarks are listed in the following table:

GPU Node	Graphical Memory	Speedup
NVIDIA Tesla M40	12 GB	1x
NVIDIA Tesla V100	32 GB	3x
NVIDIA A100 (non-distributed)	40 GB	6x
NVIDIA A100 (distributed)	40 GB	18x

References

[1] DWD: Radar-Online-Aneichung (RADOLAN). https://www.dwd.de/DE/leistungen/radolan/radolan.html
[2] LIU, Guilin ; REDA, Fitsum A. ; SHIH, Kevin J. ; WANG, Ting-Chun ; TAO, Andrew ; CATANZARO, Bryan: Image inpainting for irregular holes using partial convolutions. In: Proceedings of the European Conference on Computer Vision (ECCV), 2018, S. 85–100
[3] KADOW, Christopher ; HALL, David M. ; ULBRICH, Uwe: Artificial intelligence reconstructs missing climate information. In: Nature Geoscience 13 (2020), Nr. 6, S. 408–413
[4] SHI, Xingjian ; CHEN, Zhourong ; WANG, Hao ; YEUNG, Dit-Yan ; WONG, WaiKin ; WOO, Wang-chun: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In: Advances in neural information processing systems, 2015, S. 802–810
[5] HUANG, Junhao ; NIU, Dan ; ZANG, Zengliang ; CHEN, Xisong ; PAN, Xiaobin: RainfallNet: A Dual-Source of Spatial-Channel Attention Fusion Network for Precipitation Nowcasting. In: Journal of Physics: Conference Series Bd. 2050 IOP Publishing, 2021, S. 012008

Code available at: https://github.com/FREVA-CLINT/climatereconstructionAI