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

1st Johannes Meuer
Data Management
German Climate Computing Center
Hamburg, Germany
meuer@dkrz.de

2nd Étienne Pésiat
Data Management
German Climate Computing Center
Hamburg, Germany
plesiat@dkrz.de

3rd Hannes Thiemann
Data Management
German Climate Computing Center
Hamburg, Germany
thiemann@dkrz.de

4th Thomas Ludwig
Management
German Climate Computing Center
Hamburg, Germany
ludwig@dkrz.de

5th Christopher Kadow
Data Management
German Climate Computing Center
Hamburg, Germany
kadow@dkrz.de

***Abstract*—Missing climatological data is a general problem in climate research that leads to uncertainty of prediction models that rely on these data resources. So far, existing approaches for infilling missing precipitation data are mostly numerical or statistical techniques that require time consuming computations and are not suitable for large regions with missing data. Most recently, there have been several approaches to infill missing climate data with machine learning methods such as convolutional neural networks or generative adversarial networks. They have proven to perform well on infilling missing temperature or satellite data. However, these techniques consider only spatial variability in the data whereas precipitation data is much more variable in both space and time. Rainfall extremes with high amplitudes play an important role. We propose a convolutional inpainting network that additionally considers a memory module. One approach investigates the temporal variability in the missing data regions using a long-short term memory. An attention-based module has also been added to the technology to consider further atmospheric variables provided by reanalysis data. The model was trained and evaluated on the RADOLAN data set which is based on radar precipitation recordings and weather station measurements. Since the training of this high-resolved data set requires a large amount of computational resources, we apply distributed training on an HPC system to maximize the performance.**

***Index Terms*—Distributed-Deep-Learning, HPC, Machine-Learning, Image-Inpainting, Climate-Science, Convolutional-Neural-Network**

I. Introduction

The process of creating predictions can be a key factor for policies, services, plans in the agricultural sector and many more. Precipitation predictions, for example, can be used to assess water resources, floods or hazards, but also urban runoff volumes and duration or studying climate variability for long term trends [1]. In order to take the right actions based on predictions, one must rely on their precision which in return, whether it be a numerical weather prediction model or a machine learning approach, relies on the quality of provided data. Moreover, temporally and spatially continuous climate data often lack a required spatial or temporal resolution or values are even missing due to different circumstances [1]–[4]. Spatial precipitation data, for example, are collected via radars that face technical failures quite often due to blockage of radar beams, near-ground blind zones and other issues [2], [5]. The result is an incomplete, low-resolved or corrupted recording. Examples of this are illustrated in figure 1.

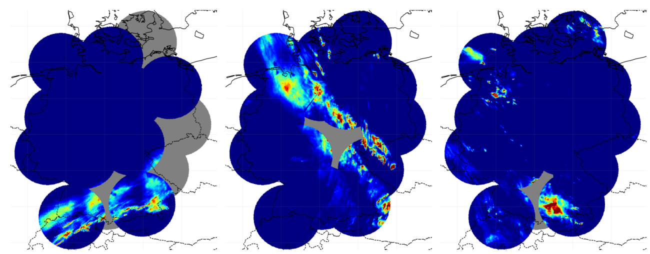

Fig. 1. Examples of corrupted radar recordings after radar failures.

These issues can lead to misinterpretations or uncertainty of predictions. However, collecting missing data can be a costly [6] or even impossible task when facing missing values from the past. Hence, methods to fill gaps in climate data have been elaborated, ranging from statistical approaches like spatial interpolation [7], [8] or linear regression [9], to machine learning approaches like support vector machines [10]. However, even when applying these techniques on HPC systems, they take up a lot of computation time or are not able to reliably infill larger regions. More advanced machine learning approaches, such as image inpainting, have been used to repair damages from images caused by raindrops, optimize quality of old pictures or even increase resolution of low quality images [11]–[13]. It has also been used to fill gaps in climate data: Shibata et al. [14] use image inpainting to reconstruct incomplete satellite images of sea surface temperatures and, most recently, [2] proposed an inpainting CNN to fill in gaps of corrupted radar images. Kadow et al. [3] apply an advanced image inpainting technique by Liu et al. [12] to reconstruct missing values in spatial surface temperature grids. The technique is able to reconstruct arbitrarily shaped missing value regions, which is especially useful when dealing with missing climate data. With this, Kadow et al. [3] are able to infill the temperature grids with astonishing results such as the reconstruction of an El-Niño from 1877 with only few grid points available for that year. Despite achieving good results, climate inpainting techniques do not consider temporal relationships in the data or are restricted to fix-shaped missing value regions.

Therefore, we propose an image inpainting CNN that is able to reconstruct arbitrarily shaped missing value regions with two improvements, that consider temporal relationships in the data and further atmospheric parameters.

II. Methods

The baseline precipitation inpainting network was derived from the implementation by Kadow et al. [3]. This network applies partial convolutions [12] in a U-shaped CNN [15]. For considering additional temporal information, we process sequences of precipitation input data with a convolutional LSTM approach [16] and apply a convolutional block attention module [17] for further atmospheric input parameters. The processing speed is optimized by applying functionalities of distributed deep learning provided by the PyTorch framework [18]. Here, the input training batch is split across multiple GPU devices. For the forward pass during training, the model is replicated on each device. Then, each input split is separately forwarded on the devices. For the backward pass, the gradients from the replicated model of each device are summed.

III. Results

The training of the network was performed via distributed deep learning on an A100 GPU, which was provided by the Levante HPC system at DKRZ. The baseline model already provided good results. However, the consideration of temporal and atmospheric information further improved the results and considering a combination of both, temporal and atmospheric information, gave the overall best performance. This observation is shown in table I that compares a set of verification metrics of the different models. The proposed models were trained on 28026 complete radar recordings ranging from 2007 to 2011 that were masked with a large region to simulate missing values. The illustrated metrics were calculated on predictions from 5885 radar fields from the year 2012. When applying distributed deep learning on the A100, we were able to speed up the training by a factor of 3 while maintaining the same prediction accuracy.

With the proposed method, we are able to efficiently infill real missing value regions from the RADOLAN precipitation data set. Figure 2 illustrates some predictions for the previously shown examples.

TABLE I
COMPARING THE ROOT MEAN SQUARE ERROR (RMSE IN mm/h), ABSOLUTE MEAN ERROR (AME IN mm/h), TEMPORAL CORRELATION ($r_{xy,Temp}$) AND SPATIAL CORRELATION $r_{xy,Spatial}$ OF THE PROPOSED MODELS.

Model	RMSE $\downarrow$	AME $\downarrow$	$r_{xy,Temp}$ $\uparrow$	$r_{xy,Spatial}$ $\uparrow$
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**

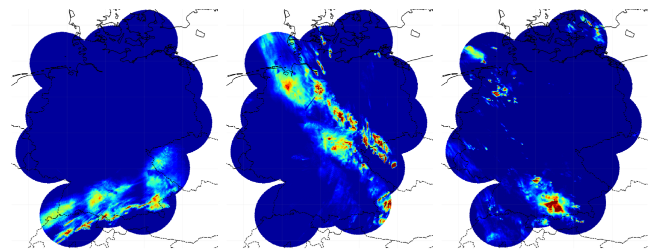

Fig. 2. Examples of infilled broken radar recordings.

IV. Future Work

With the provision of the open source software, there is the possibility for the community to put the results in context with other methods, apply them to weather services to improve predictions and potentially develop them further. The built-in components can now be adopted for other applications and other scientific predictions in space and time can be investigated. Super-resolution, for example, was investigated on the original partial convolutional network by Liu et al. [12] and could be applied to atmospheric images with the proposed modules. For example, we are currently investigating the application of super-resolution on ensembles of volcanic eruptions. Since the calculations to provide such high resolved data are very costly regarding time and hardware resources, this technique could enable faster predictions with lower computational resources. Another application scenario could be to generate radar images at regions where no radars exist, learning only from reanalysis data or weather station measurements. This could also enable to generate radar images from the past when no radars even exist yet.

References

[1] R. S. Teegavarapu, A. Aly, C. S. Pathak, J. Ahlquist, H. Fuelberg, and J. Hood, "Infilling missing precipitation records using variants of spatial interpolation and data-driven methods: use of optimal weighting parameters and nearest neighbour-based corrections," *International Journal of Climatology*, vol. 38, no. 2, pp. 776–793, 2018.

[2] A. Geiss and J. C. Hardin, "Inpainting radar missing data regions with deep learning," *Atmospheric Measurement Techniques Discussions*, pp. 1–28, 2021.

[3] C. Kadow, D. M. Hall, and U. Ulbrich, "Artificial intelligence reconstructs missing climate information," *Nature Geoscience*, vol. 13, no. 6, pp. 408–413, 2020.

[4] A. Barrios, G. Trincado, and R. Garreaud, "Alternative approaches for estimating missing climate data: application to monthly precipitation records in south-central chile," *Forest Ecosystems*, vol. 5, no. 1, pp. 1–10, 2018.

[5] DWD. Radar-online-aneichung (radolan). [Online]. Available: https://www.dwd.de/DE/leistungen/radolan/radolan.html

[6] R. J. Little, “Missing-data adjustments in large surveys,” *Journal of Business & Economic Statistics*, vol. 6, no. 3, pp. 287–296, 1988.

[7] T. M. Smith, R. W. Reynolds, R. E. Livezey, and D. C. Stokes, “Reconstruction of historical sea surface temperatures using empirical orthogonal functions,” *Journal of Climate*, vol. 9, no. 6, pp. 1403–1420, 1996.

[8] M. A. Oliver and R. Webster, “Kriging: a method of interpolation for geographical information systems,” *International Journal of Geographical Information System*, vol. 4, no. 3, pp. 313–332, 1990.

[9] R. L. Vislocky and J. M. Fritsch, “Generalized additive models versus linear regression in generating probabilistic mos forecasts of aviation weather parameters,” *Weather and Forecasting*, vol. 10, no. 4, pp. 669–680, 1995.

[10] T. Landot, S. Sgellari, C. Lima, and U. Lall, “In-filling missing historical daily rainfall data study,” *Final Report, South Florida Water Management District, Columbia University, New York, NY*, 2008.

[11] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, “Generative image inpainting with contextual attention,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5505–5514.

[12] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, “Image inpainting for irregular holes using partial convolutions,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 85–100.

[13] O. Elharrouss, N. Almaadeed, S. Al-Maadeed, and Y. Akbari, “Image inpainting: A review,” *Neural Processing Letters*, vol. 51, no. 2, pp. 2007–2028, 2020.

[14] S. Shibata, M. Iiyama, A. Hashimoto, and M. Minoh, “Restoration of sea surface temperature satellite images using a partially occluded training set,” in *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2018, pp. 2771–2776.

[15] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[16] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, “Convolutional lstm network: A machine learning approach for precipitation nowcasting,” in *Advances in neural information processing systems*, 2015, pp. 802–810.

[17] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, “Cbam: Convolutional block attention module,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 3–19.

[18] Facebook. Pytorch. [Online]. Available: https://pytorch.org/