

Climate model-informed deep learning of global soil moisture distribution

Klaus Klingmüller¹ and Jos Lelieveld^{1,2}

¹Max Planck Institute for Chemistry, Hahn-Meitner-Weg 1, 55128 Mainz, Germany

²The Cyprus Institute, P.O. Box 27456, 1645 Nicosia, Cyprus

Correspondence: Klaus Klingmüller (k.klingmueller@mpic.de)

Abstract. We present a deep neural network (DNN) that produces accurate predictions of observed surface soil moisture, ~~based on~~ applying meteorological data from a climate model. The network was trained on daily satellite retrievals of soil moisture from the European Space Agency (ESA) Climate Change Initiative (CCI). The predictors precipitation, temperature and humidity were simulated with the ECHAM/MESSy atmospheric chemistry-climate model (EMAC). Our evaluation shows that predictions of the trained DNN are highly correlated with the observations, both, spatially and temporally, and free of bias. This offers an alternative for parametrisation schemes in climate models, especially in simulations that use, but may not focus on soil moisture, which we illustrate with the threshold wind speed for mineral dust emissions. Moreover, the DNN can provide proxies for missing values in satellite observations to produce realistic, comprehensive, high resolution global datasets. As the approach presented here could be similarly used for other variables and observations, the study is a proof of concept for basic but expedient machine learning techniques in climate modelling, which may motivate additional applications.

1 Introduction

Since decades, global climate and atmospheric-chemistry models rely on supercomputers with a traditional cluster architecture, utilising many powerful compute nodes in parallel. The individual nodes are typically very potent by themselves, using general purpose central processing unit (CPU) cores with large memory. They are required to serve the needs of diverse algorithms representing many different physical and chemical processes. The implementations of these algorithms quite often have a source code legacy with a history of multiple decades. This creates a dependency on the system type which may limit progress in model studies once the CPU performance increase slows down, typical for the recent past. In fact, due to challenges in the continued miniaturisation of semiconductors, such a slowdown is ongoing, which motivates the search for other performance gains (Leiserson et al., 2020).

Progress has been made to accelerate weather, climate and atmospheric-chemistry models with general purpose graphics processing units (GPUs) (Yashiro et al., 2016; Alvanos and Christoudias, 2017; Sun et al., 2018; Fuhrer et al., 2018; Müller et al., 2019), but a wider utilisation is pending. Other high performance computing applications such as lattice quantum chromodynamics with typically smaller codebase and less legacy code swiftly exploited the computational resources of GPUs (Egri et al., 2007; Clark et al., 2010). A discipline which not only uses the full potential of GPUs but also builds its present success

25 on the advent of GPU computing, is machine learning, particularly deep learning (Krizhevsky et al., 2012). Strong commercial
interest even boosted the development of specialised machine learning hard- and software (e.g. Abadi et al., 2015; Jouppi
et al., 2017; Markidis et al., 2018). Atmospheric and climate models can benefit from this development by making use of the
computational capabilities of the new hardware to accelerate existing algorithms (Hatfield et al., 2019), or by applying machine
learning techniques to complement existing, physical process-based and in particular empirical parametrisations (Chevallier
30 et al., 2000).

Given the success of deep learning in many different disciplines (Schmidhuber, 2015; LeCun et al., 2015; Silver et al.,
2016) it seems likely that atmospheric modelling and climate science in general can benefit from this methodology, not only in
terms of performance gains, and promising results have been presented lately (e.g. Kadow et al., 2020). However, reservations
about machine learning exist. In contrast to physical models, where the simulation result is deduced from the laws of physics
35 and physical parameters, trained models often represent black boxes where the rules by which they compute the output tend
to be non-transparent and cannot easily be modified to represent varied physical conditions. On the other hand, complex
conventional, phenomenologically derived parametrisations at times also lack a clear physical meaning and their parameters are
often tuned to obtain realistic results, which essentially is a non-systematic form of training. Moreover, methods for interpreting
machine learning models are emerging (Montavon et al., 2018; Kohoutová et al., 2020). Aside from that, depending on the
40 application, it may be irrelevant whether a process is implemented as a black box or not if the scientific focus is on other
processes. For instance, the present study was motivated by the need for reliable soil moisture data to develop a mineral dust
emission scheme.

Soil moisture near the surface has a significant impact on the emissions of mineral dust (Klingmüller et al., 2016) and
it is generally of great importance for weather and climate. It has been identified by the Global Observing System for Cli-
45 mate (GCOS) as an essential climate variable (ECV, Bojinski et al., 2014) and, for example, soil moisture can greatly impact
mesoscale convective systems (Klein and Taylor, 2020). While detailed parametrisations of soil moisture exist (e.g., Ekici
et al., 2014), many models, such as the ECHAM/MESSy atmospheric chemistry-climate model (EMAC) (Jöckel et al., 2006)
still use relatively simple models.

On the other hand, because moisture at the surface affects the dielectric properties of the soil, it can be well retrieved
50 remotely using microwaves, and an extensive global daily dataset covering the past four decades is provided by the European
Space Agency (ESA) Climate Change Initiative (CCI) (Dorigo et al., 2017; Gruber et al., 2017, 2019).

To make use of the ESA CCI surface soil moisture data in climate models, it cannot be imported directly, because the
daily subsets have substantial gaps depending on the local retrieval conditions at overpass time. Moreover, merely importing
observations would limit model applications to hindcasting. Therefore we pursued an alternative approach and use the satellite
55 data for supervised training of a deep neural network (DNN) to predict soil moisture based on modelled meteorological data.
In doing so, we explored the potential of machine learning to complement physical models using an introductory example,
which demonstrates the useful results which can be achieved with limited technical effort, and which might instigate further
applications.

This article is structured as follows: the datasets used are presented in section 2, the DNN is introduced and evaluated in section 3, applications of the DNN are discussed in section 4 before we draw conclusions in section 5.

2 Data

We used results from a 10 year simulation with the ECHAM/MESSy atmospheric chemistry-climate model (EMAC) (Jöckel et al., 2006) version 2.52, covering the years 2006 to 2015. The exact configuration is described by Klingmüller et al. (2020). Horizontally, the setup employs a Gaussian T63 grid with approximately 1.9° spacing. EMAC assimilates observational data by nudging temperature, vorticity and divergence above the boundary layer to meteorological reanalysis data of the European Centre for Medium-Range Weather Forecasts (ECMWF) and by using the sea surface temperature from the same dataset.

The EMAC output variables we considered are precipitation, surface temperature and humidity, which are preprocessed to daily average values. In addition, we made use of the static ecosystem rooting depth map of the online dust emission scheme, originally from Schenk and Jackson (2009).

The EMAC data was complemented by daily volumetric soil moisture (i.e. the ratio of water relative to soil volume) observations from the ESA CCI Soil Moisture Product Release v04.5 (Dorigo et al., 2017; Gruber et al., 2017, 2019). This dataset is representative for the soil moisture in the topmost few centimetres of soil (down to about 5 cm). We used the dataset combining retrievals from active and passive spaceborne microwave instruments, which we aggregated from the original spatial grid with 0.25° spacing to the Gaussian T63 grid of the EMAC results.

We subdivided the 10 year period covered by the EMAC simulation into a training period of 8 years from 2006 to 2013 and a test period of 2 years from 2014 to 2015. This choice maximises the length of the training period while keeping more than one year for testing, allowing to identify interannual variations during the testing period. Moreover, using training and testing periods in chronological order, all predictions in the testing period represent forecasts. The test period was exclusively used to evaluate the DNN after training. Every third year of the training period (2008 and 2011) was used for validating and monitoring the training procedure.

3 DNN model

The basic concept of our approach is to relate the observed soil moisture to relevant quantities modelled by the global climate model. We applied a simple DNN architecture which operates on one grid cell at a time. As a consequence, the DNN can easily be integrated as a submodel in global climate models such as EMAC. The meteorological variables used as predictors are selected based on availability in the climate simulation output and physical relevance: the rain rate represents the source of soil moisture, whereas surface temperature and atmospheric humidity control soil drying. To account for cumulative effects, in addition to the actual daily mean rain rate, surface temperature and specific humidity, we provided the DNN with the corresponding values lagged by 1 day, the mean of values with 2 to 3 days lag, and the mean of 4 to 7 days lagged values. According to preliminary tests, considering longer lags did not yield a significant improvement, but it might become relevant if

90 the overall performance can be enhanced in future versions. To account for regional characteristics such as soil properties, the DNN uses longitude and latitude, encoded in the triple $\sin(\text{lon})$, $\cos(\text{lon})$, $\sin(\text{lat})$ as well as the local rooting depth. Likewise, to account for seasonality, for example by vegetation variations, we supplied the DNN with the time of year encoded in $\sin(2\pi t/a)$ and $\cos(2\pi t/a)$, where t/a is the time in years. In total this amounts to 18 input variables which had to be mapped to one output variable, the surface soil moisture.

95 For that purpose, we employed a generic DNN of linearly stacked densely connected layers as illustrated in Fig. 1. Four hidden layers of 512 units with rectified linear activation are followed by the output layer with a single unit and linear activation. There is no strict formula for the number of layers and units per layer, but too small DNNs are not capable of learning complex rules, whereas too large DNNs are more at risk of overfitting and require more computational resources. Our values are a compromise that is proving to work well, but systematic optimisation could possibly yield better results. To generalise the DNN and prevent overfitting during training, we regularised by applying a 10 % dropout rate to the hidden layers (Hinton et al., 2012) and stopped the training process as soon as the validation loss was no longer improving. Before training, all input and output variables were normalised independently to have a mean of 0 and a standard deviation of 1. Accordingly, before using the DNN for predictions, the same transformations had to be applied to the input data and an inverse transformation to the output.

105 We implemented the DNN and performed the training and inference using the TensorFlow library (Abadi et al., 2015) with the Keras interface (Chollet et al., 2015, 2017) for R (R Core Team, 2019). The training took about ~~1.5 hours~~ 40 minutes on an Nvidia Tesla V100 GPU accelerator. The computationally much less demanding predictions were evaluated on common desktop hardware.

To assess the overall predictive power of the DNN, we compared the volumetric soil moisture calculations with the corresponding observations for all grid cells and days during the test period (2014 to 2015) where retrievals are available. Figure 2 shows the scatter of the 1247041 observation-prediction pairs. It demonstrates a remarkably high quality of the predictions, which are strongly correlated with the observations, with a Pearson correlation coefficient of 0.92, and do not show any bias, resulting in a small root-mean-square error of 0.033 and a mean absolute error of 0.024, being an order of magnitude smaller than the average data values. Note that the Gaussian grid has more grid cells per area at higher latitudes, giving those regions more weight in this comparison, but the number of relevant grid cells in polar regions is small.

115 While Fig. 2 combines the effect of the spatial and temporal variability, Fig. 3 shows the spatial correlation for each day separately, using all grid cells with observations during that day. The correlation coefficient attains high values around 0.92 throughout the test period and rarely drops below 0.9. Considering that the training takes place before the test period, it is noteworthy that there is no significant decline of the correlation over the two years. Essentially, at any time the DNN yields a realistic representation over the globe, predicting low soil moisture in arid regions and high soil moistures in wet areas.

120 Equally important is a realistic representation of the temporal variation for each grid cell. Due to the strong temporal variability of weather and in particular precipitation, this is more challenging and we do not expect equally high correlation coefficients as for the spatial correlation. Nevertheless, Fig. 4 shows that the temporal correlation coefficients are above-globally high with a mean of almost 0.7, which appears to be a good performance when compared with state of the art model data




Figure 1. DNN architecture and hyperparameters. The input variables on the left are processed by four fully connected hidden layers of 512 units each to yield the soil moisture on the right. All four hidden layers use a rectified linear unit function activation and during the training a dropout rate of 10 % to avoid overfitting. The output layer with only one unit applies a linear activation function to obtain the final result. The DNN operates on normalised variables with ~~stadard~~ standard deviation 1 and mean 0.

125 (Dorigo et al., 2017). The correlation coefficient is larger than 0.5 ~~virtually globally~~. The only notable exceptions are the in almost all grid cells except those representing extremely dry soils of the Sahara, Rub' al Khali, Taklamakan and Gobi deserts. Related to this, a slight overestimation by the DNN towards the lower end of the soil moisture range can be identified in Fig. 2. Training and evaluation of the DNN are challenging in this range because satellite retrievals are both sparse and uncertain (see, e.g., Fig. S1 in the supplement), therefore, depending on the application, some additional efforts that focus on these

130 regions might be required. For dust emissions, the soil moisture in these regions is too low to significantly influence the results, while it is more relevant in semi-arid regions, where observations and predictions are more reliable (Fig. 4 and Fig. S2 in the supplement).

4 Applications

Volumetric soil moisture




Figure 2. Comparison of observed and predicted volumetric surface soil moisture considering all daily grid-cell values available globally during the years 2014 and 2015. The colours represent the density distribution of the scatter, outliers are represented by dots. Pearson correlation coefficient r , root-mean-square error (RMSE), mean absolute error (MAE) and the number of data points n are provided in the upper left corner.

Spatial correlation


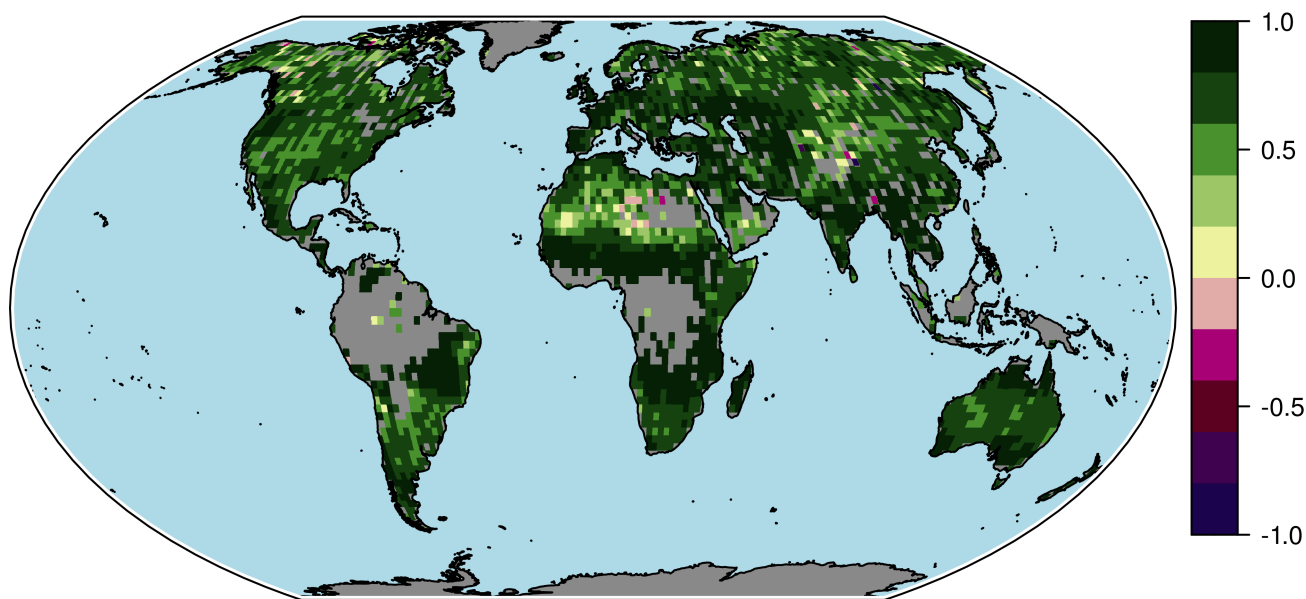


Figure 3. Spatial correlation of predicted and observed volumetric soil moisture throughout the test period.

Temporal correlation



Analysis of the temporal correlation of observed and predicted volumetric soil moisture in the individual grid cells during the years 2014 and 2015. In stippled regions the correlation coefficient r of the time series exceeds 0.5. Observations are available for regions shaded in grey.

Figure 4. Correlation coefficient of the observed and predicted volumetric soil moisture time series in the individual grid cells during the test years 2014 and 2015. In the grey regions no observations are available during this period.

The DNN operates on a single grid cell at a time and therefore can easily be incorporated in climate models, for example in the submodel core layer of EMAC. This provides realistic soil moisture values to other submodels, which can optionally replace the current parametrisation of soil water. For instance, this could be advantageous for mineral dust emission schemes which should account for reduced emissions from moist soils but so far have been limited by the inadequate representation of soil moisture in the topmost surface layer by the physically modelled soil water (Klingmüller et al., 2018).

Considering the grid cell centred at 49.4°N 7.5°E, the upper panel of Fig. 5 exemplifies how closely the predicted volumetric soil moisture time series resembles the observations. As noted above, there is no bias in the predictions and the seasonal cycle is well represented. Moreover, the short-term variations of the predictions show a clear similarity to those of the observations. Both have a comparable amplitude and frequency, and characteristic features in the observed time series are reproduced by the predictions, e.g., in July 2014, October 2014, December 2014/January 2015 or March 2015–2015 (Fig. S3 in the supplement). These features occur irregularly and are not repeated year after year, which demonstrates that the DNN did not simply learn one representative climatology but utilises information from the meteorological data provided by the climate model. We reiterate that the nudged EMAC meteorology is close but not identical to the real atmosphere, hence small deviations are expected.

For comparison, the realistic representation of the temporal soil moisture variation in each grid cell, the meteorological predictors are essential. The climatological model obtained by training the DNN omitting the meteorological predictors still reaches an overall correlation coefficient of close to 0.9, but the global mean of temporal correlation coefficient drops substantially from close to 0.7 to below 0.5 (Fig. S4 in the supplement). In addition to the interannual variations, in particular the variability on time scales shorter than 1 month is lost by ignoring the meteorology (Fig. S5 in the supplement). The relevance of the individual meteorological predictors for the temporal variations differs regionally. For example in Europe, the DNN predictions are sensitive to all meteorological predictors, but in arid and semi-arid regions of the Middle East they are not sensitive to reducing the rain (Fig. S6 to S8 in the supplement). Here, all variability is inferred from temperature and humidity. This does not necessarily imply rain to be the least important driver of soil moisture variability in the region, rather it is the least reliable predictor due to the uncertainty in modelled precipitation.

The meteorological predictors are also required to make the DNN sensitive to climatic changes. Increasing temperature or decreasing humidity promotes soil drying, whereas increasing rain enhances the soil moisture, which is reproduced by the DNN (Fig. S6 to S8 in the supplement). Therefore, if climatic changes are represented in the meteorological predictors, the DNN soil moisture will respond to these changes. The response might be limited because in the present configuration the DNN can substitute meteorological information with the knowledge of season and location. Consequently, for applications in long-term climate projections, training and evaluation on longer time scales and possibly refinements of the DNN architecture are advisable.

The spatial coordinate predictors are crucial for a realistic spatial variation of the soil moisture. They are the parameters encoding information on the wide range of surface properties including soil and vegetation types, and substantially improve the DNN. Less relevant are the seasonal predictors, but if provided, the DNN makes use of them so that the DNN predictions




Figure 5. Time series of the observed and predicted daily volumetric soil moisture values and the EMAC soil water in the grid cell centred at 49.4°N 7.5°E [in Europe](#).

[are sensitive to variations of the time of year \(Figs. S6 to S8 in the supplement\) and the temporal correlation is improved \(Fig. S4 in the supplement\).](#)

170 **4 Applications**

[The DNN operates on a single grid cell at a time and therefore can easily be incorporated in climate models, for example in the submodel core layer of EMAC. This provides realistic soil moisture values to other submodels, which can optionally replace the current parametrisation of soil hydrology. For instance, this could be advantageous for mineral dust emission schemes which should account for reduced emissions from wet soils but so far have been limited by the inadequate representation of soil moisture in the topmost surface layer by the physically modelled soil water \(Klingmüller et al., 2018\).](#)

175

[The lower panel of Fig. 5 shows the corresponding time series of the physically modelled EMAC soil water. Evidently, this time series is largely unrelated to the observed surface soil moisture in the upper panel: the short-term variability is smaller](#)

whereas the long-term variability is much larger, showing a strong decline during summer 2015 which is not present in the observations. Regardless of the question whether this decline reflects a true decline of the water content in deeper soil layers or whether it is only a model artefact, it is obviously impossible to map the EMAC soil water to a realistic representation of the observed surface soil moisture in the panel above. However, it is the latter which is required for parametrisations such as the mineral dust emission scheme. Because the DNN presented here fulfils this requirement, we propose to use it as improvement over the EMAC computed soil water and as viable alternative to more sophisticated physical soil moisture models. The algorithm cannot replace physically-based soil moisture representations for first-principle process studies, but is an accurate substitute for parametrisations that depend on limited empirical information.

Mineral dust emissions are predominantly caused by saltation bombardment where saltating particles on impact with the surface eject finer dust sediments or disintegrate to finer particles themselves. To activate and sustain a horizontal flux of saltating particles, the surface friction velocity of the air has to exceed a threshold which depends on the soil properties. Soil moisture increases this threshold, thereby reducing the dust emissions. We studied this effect using the parametrisation presented by Fécan et al. (1999),

$$\frac{u_{*t}}{u_{*td}} = \sqrt{1 + 1.21(w - (0.0014\phi_{\text{clay}}^2 + 0.17\phi_{\text{clay}}))^{0.68}}, \quad (1)$$

where u_{*t} is the threshold surface friction velocity, u_{*td} the corresponding threshold for dry soil, w the gravimetric soil moisture in percent and ϕ_{clay} the soil clay fraction in percent. The equation is applied if the soil moisture exceeds $0.0014\phi_{\text{clay}}^2 + 0.17\phi_{\text{clay}}$, representing the minimum soil moisture required to induce an increase in the threshold. Like Astitha et al. (2012), we combined this soil moisture dependency with the threshold surface friction velocity parametrisation of Marticorena and Bergametti (1995), applied to a saltation particle diameter of 60 μm . The full equation for the threshold surface friction velocity is reproduced in appendix A. We evaluated the equation for air density $\rho_{\text{air}} = 1.2 \text{ kg m}^{-3}$ using the clay fraction distribution from Shangguan et al. (2014). To convert the volumetric to gravimetric soil moisture we assumed a soil bulk density of 1600 kg m^{-3} . For comparison, we converted the EMAC soil water from water volume per area to gravimetric soil moisture using the same bulk density and assuming the water to be evenly distributed over the soil column defined by the rooting depth.

We focused our analysis on Mesopotamia and the Arabian Peninsula where a significant correlation of soil moisture and dust emissions was reported (Klingmüller et al., 2016), and considered the regional average of the threshold surface friction velocity over the territory of Iraq, Israel, Jordan, Kuwait, Lebanon, Oman, Qatar, Saudi Arabia, State of Palestine, Syria, United Arab Emirates and Yemen. The threshold surface friction velocity during the test period is shown in Fig. 6 calculated using the observed, predicted and EMAC-calculated soil moisture. The results based on the observed and predicted soil moisture show good agreement and a strong seasonal cycle, ~~whereas the result based-~~ During summer, the threshold calculated based on the DNN predictions tends to be slightly higher than that based on the observed values indicating an effect of the aforementioned challenges in hyper arid regions. Nevertheless, the difference between both results is much smaller than the interannual variations and the DNN result is in the range of the short term variations of the observation based result which result from the varying retrieval coverage. In contrast, the result based on the EMAC soil water has little variability, and is therefore inconsistent with the other two results, irrespective of the precise conversion factor used to obtain the gravimetric

surface soil moisture. We conclude that the value of the EMAC soil water, which represents the total water including moisture in deeper soil layers, is limited in this context, but so far it has been the only estimate available for EMAC simulations.

Fig. 6 additionally shows results from a recent dataset of Pu et al. (2020) who have retrieved a climatological monthly global distribution of the threshold in terms of the wind speed at 10 m altitude based on satellite and reanalysis data. Two versions of the data represent different assumptions used to identify dust events based on the dust optical depth (DOD). The seasonal cycle is similar to that of the predicted threshold surface friction velocity with a comparable relative amplitude and a similar but slightly shifted phase. Assuming a logarithmic wind profile, the predicted surface friction velocity threshold can be converted to the corresponding 10 m wind speed $u_{10m,t} = u_{*t} \ln(10m/z_0)/\kappa$, where $\kappa \approx 0.4$ is the Von Kármán constant and z_0 the surface roughness. ~~Consistently~~ Consistent with Astitha et al. (2012) we used the surface roughness $z_0 = 0.0001$ m. According to the logarithmic profile, the threshold that Eq. (A1) yields for dry soils, i.e., the minimal threshold $u_{*td} = 0.26 \text{ m s}^{-1}$, corresponds to the 10 m wind speed $u_{10m,td} = 7.5 \text{ m s}^{-1}$. This value is higher than most of the climatological values, however, the latter were derived based on 6-hourly wind speeds, whereas the parametrisation we used is meant for and applied to instantaneous surface friction velocities which vary on shorter time-scales (only limited by the model time step, e.g., 12 min in our T63 simulation and even shorter for higher spatial resolutions) and reach higher peak values. Therefore the threshold definitions differ, not allowing direct comparisons of the absolute values. Additionally, because the retrieval of the climatology does not account for dust transport, it may regionally underestimate the threshold.

The substantial variations of the threshold surface friction velocity obtained based on the observed soil moisture emphasize the relevance of soil moisture for dust emissions. The soil moisture predicted by the DNN is sufficiently realistic to reproduce these variations and to account for this important effect in global climate model simulations.

In addition to the incorporation in global climate models, another application of the DNN is the reprocessing of remote sensing data. Based on meteorological input data, the DNN predicts the global daily soil moisture distribution consistent with the observations. In contrast to the observational datasets that have substantial gaps in regions and time periods where conditions do not allow retrievals, the meteorological input data does not have any missing values, and consequently the same applies to the predicted soil moisture. The latter can therefore also be used to consistently fill the gaps in the observations to obtain a complete daily global soil moisture dataset. Figure 7 shows the global distribution of the observed and predicted soil moisture on two example days from the training period, one during winter on the northern hemisphere (15. January 2012), the other during summer (15. July 2012). Observations outside the training period can be processed as well (Fig. S9 in the supplement) but to obtain optimal results the training should include the period of interest. Regardless of the extensive regions without observations, the prediction yields global values within a reasonable range, closely resembling the observations where available. ~~Only in the~~ Note that in regions or seasons where no or only few observations are available throughout the training period, the predictions have to be interpreted with caution. This applies to the rain forests, central deserts and regions permanently or seasonally frozen or covered by snow. In the complete soil moisture distribution the seasonal variations become apparent, promoting the DNN predictions for use in further studies such as trend analysis. Moreover, in contrast to the incomplete observations, the optimized predictions can straightforwardly be assimilated into climate models.

Dust emission threshold wind speed over the Arabian Peninsula and Mesopotamia

Based on observed soil moisture — DOD threshold 0.2 —
 Based on predicted soil moisture — DOD threshold 0.5 —
 Based on EMAC soil moisture —

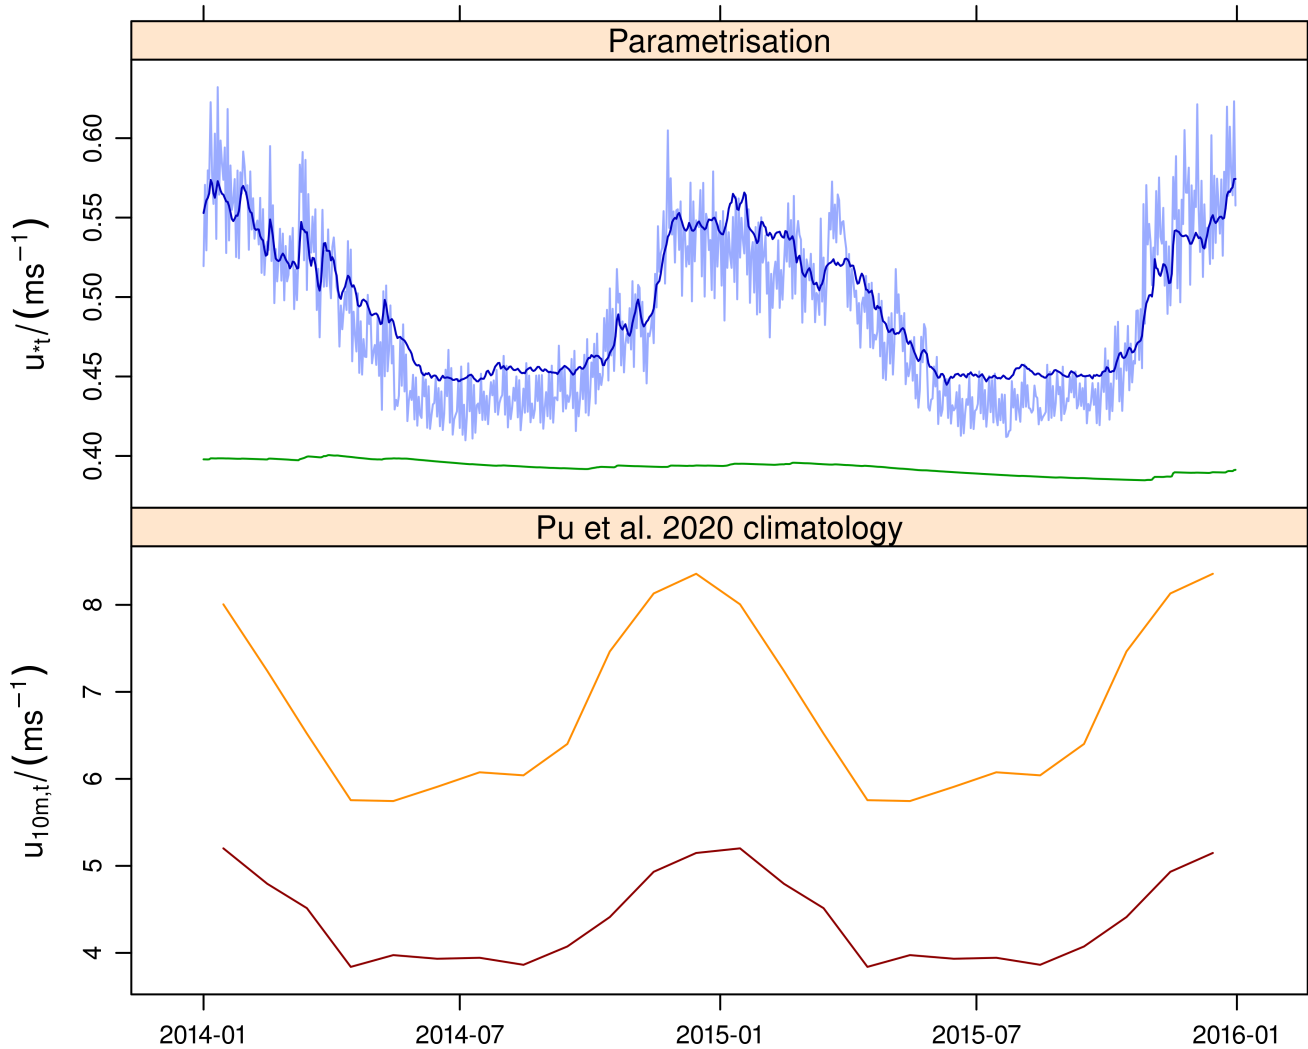


Figure 6. The surface friction velocity threshold above which dust is emitted, averaged over the Arabian Peninsula and Mesopotamia. The surface friction velocity threshold u_{*t} is computed using the parametrisation provided in the appendix A. The threshold in terms of the wind speed at 10 m altitude $u_{10m,t}$ represents the climatological data set retrieved by Pu et al. (2020).

Volumetric soil moisture

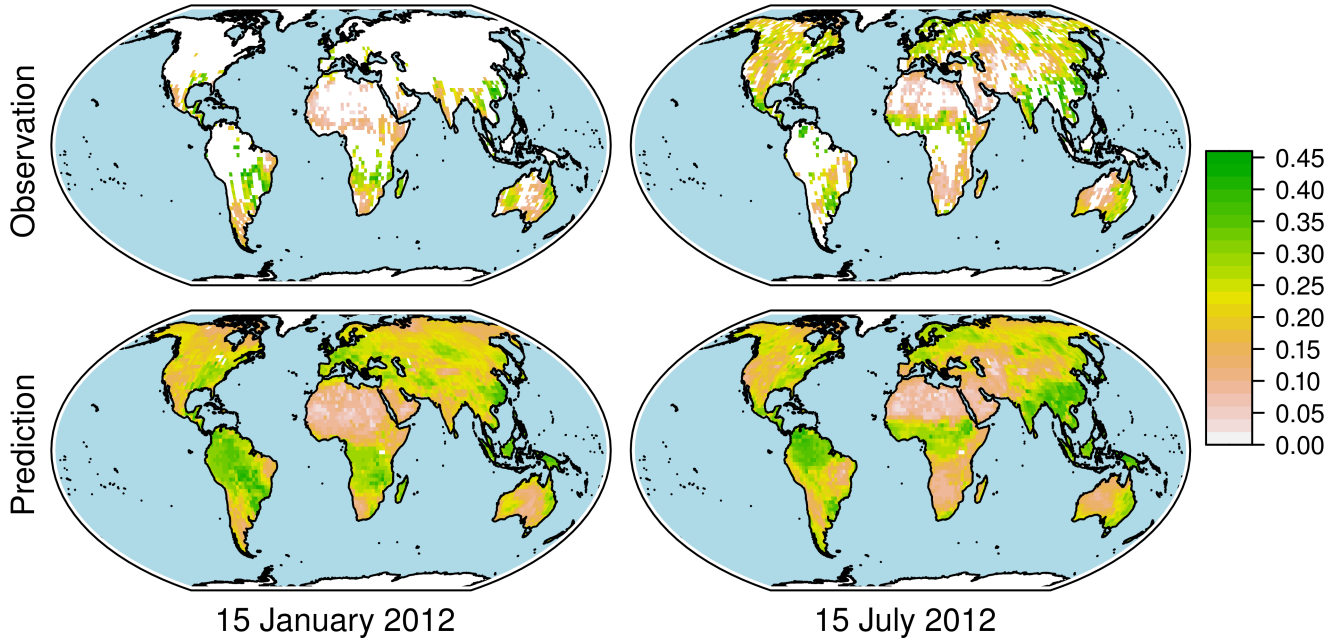


Figure 7. Global distribution of the observed and predicted volumetric soil moisture on a northern hemisphere winter (15. January 2012) and summer (15 July 2012) day in the training period.

5 Conclusions

We have presented a machine learning model which relates soil moisture to meteorological conditions. Informed by a climate model, this DNN is able to accurately predict satellite based surface soil moisture observations, as demonstrated by our evaluation. Using the example of the threshold wind speed for mineral dust emissions we showed that the DNN predictions can be used for improved representations of surface soil moisture-dependent processes within climate models.

The DNN in its present form should be regarded as a proof of concept, and there is room for improvement. The current DNN architecture, the simple stack of several densely connected layers, is very generic. While it is generally quite powerful, it is not tailored to our specific application and other concepts might be considered as well. Convolutional neural networks could exploit the spatial relationship of neighbouring grid cells and recurrent neural networks might more optimally account for the causal relationship of the soil moisture at successive days including long-term accumulative effects. The causal relation is partly addressed in our implementation by the consideration of lagged meteorological variables, representing a temporal convolution. However, the prediction is not informed about the conditions prior to one week in the past. This apparently works well for the surface soil moisture, but is probably not sufficient for additional applications, such as those that require information about the moisture in deeper soil layers. The hyperparameters of the DNN, including the number of layers, the number of units per layer and the selection of input variables, are chosen to be appropriate for the problem, but have not been systematically optimised.

We conclude that there are various pathways for future developments that may enhance the DNN performance. Nevertheless, the present DNN setup can already be beneficial for applications such as online mineral dust emission schemes in climate models. Therefore, the trained DNN will be implemented as an EMAC submodel using the MESSy interface.

265 The CPU time required for the inferencing using the trained DNN is negligible compared to the total computational demand
of a global climate model. Other applications of DNNs within climate models may be more demanding, in particular if they
process three-dimensional data instead of only two-dimensional surface data and if they have to be called more often than
once per day. In this case, GPUs or specialised inferencing hardware can be employed to evaluate the trained model, which
by design efficiently utilises such accelerators. In this way, climate models can benefit from GPUs that are available in many
supercomputers without the need for porting complex algorithms to GPUs, and from the rapid development of machine learning
270 hardware.

For hindcasting applications, an alternative to implementing the trained soil moisture DNN into climate models is to import the consistent and comprehensive global soil moisture prediction from the DNN at runtime. Because such a dataset is also of general use, it appears to be promising to repeat the procedure presented here with high resolution meteorological reanalysis data.

275 Overall, our example demonstrates that machine learning models informed by data from traditional, physical process-based climate models can perform well in learning and predicting observational data. In return, they can complement process parametrisations in climate models, especially when the parametrisations rely on limited empirical data. ~~On top of that~~Furthermore, they may help climate models to efficiently utilise recent hardware architectures.

Code and data availability. The data and DNN parameters used in this study are archived at the German Climate Computing Centre (DKRZ) and available from the corresponding author KK until they are deposited in the public Edmond Open Research Data Repository of the Max Planck Society. The ECHAM climate model is available to the scientific community under the MPI-M Software License Agreement (<https://mpimet.mpg.de/en/science/modeling-with-icon/code-availability>, last access: 11 November 2020, MPI-M, 2020). The Modular Earth Submodel System (MESSy) is continuously further developed and applied by a consortium of institutions. The usage of MESSy and access to the source code are licensed to all affiliates of institutions which are members of the MESSy Consortium. Institutions can become a
285 member of the MESSy Consortium by signing the MESSy Memorandum of Understanding. More information can be found on the MESSy Consortium Website (<https://www.messy-interface.org>, last access: 11 November 2020, MESSy, 2020).

Appendix A: Surface friction velocity threshold for dust emissions

The full equation for the threshold surface friction velocity u_{*t} used in section 4 is

$$\begin{aligned}
 u_{*t} = & 0.129 \sqrt{\frac{D_p}{\rho_{\text{air}}} \left(\rho_p g + \frac{0.006 g \sqrt{\text{cm}} / \text{s}^2}{D_p^{5/2}} \right)} \\
 & \times \begin{cases} \frac{1}{\sqrt{1.928 B^{0.092} - 1}} & B < 10 \\ (1 - 0.0858 e^{-0.0617(B-10)}) & B \geq 10 \end{cases} \\
 & \times \left(1 - \frac{\ln \frac{z_o}{z_{os}}}{\ln \left(0.35 \left(\frac{10 \text{cm}}{z_{os}} \right)^{0.8} \right)} \right)^{-1} \\
 & \times \sqrt{1 + 1.21 \max(0, (w - (0.0014 \phi_{\text{clay}}^2 + 0.17 \phi_{\text{clay}})))^{0.68}}, \tag{A1}
 \end{aligned}$$

where

$D_p = 60 \mu\text{m}$	saltation particle diameter
ρ_{air}	air density
$\rho_p = 2.65 \text{ g/cm}^3$	particle density
$g = 9.80665 \text{ m/s}^2$	gravitational acceleration
$B = \frac{u_{*t} D_p}{\nu}$	friction Reynolds number, initially $B = 1331(D_p/\text{cm})^{1.56} + 0.38$
$\nu = 0.157 \cdot 10^{-4} \text{ m}^2/\text{s}$	kinematic viscosity of air
$z_o = 0.01 \text{ cm}$	surface roughness length
$z_{os} = 0.00333 \text{ cm}$	local roughness length of the uncovered surface
w	gravimetric soil moisture in %
ϕ_{clay}	clay fraction in %

Author contributions. KK conceived the study, implemented the DNN and wrote the manuscript supported by JL. Both authors discussed the results and finalised the article.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. We acknowledge financial support from the MaxWater Initiative of the Max Planck Society.

300 References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous
305 Systems, <https://www.tensorflow.org/>, software available from tensorflow.org, 2015.
- Alvanos, M. and Christoudias, T.: GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESSy (EMAC) Earth system model (version 2.52), *Geoscientific Model Development*, 10, 3679–3693, <https://doi.org/10.5194/gmd-10-3679-2017>, <https://gmd.copernicus.org/articles/10/3679/2017/>, 2017.
- Astitha, M., Lelieveld, J., Abdel Kader, M., Pozzer, A., and de Meij, A.: Parameterization of dust emissions in the global atmospheric
310 chemistry-climate model EMAC: impact of nudging and soil properties, *Atmospheric Chemistry and Physics*, 12, 11 057–11 083, <https://doi.org/10.5194/acp-12-11057-2012>, <https://acp.copernicus.org/articles/12/11057/2012/>, 2012.
- Bojinski, S., Verstraete, M., Peterson, T. C., Richter, C., Simmons, A., and Zemp, M.: The Concept of Essential Climate Variables in Support of Climate Research, Applications, and Policy, *Bulletin of the American Meteorological Society*, 95, 1431–1443, <https://doi.org/10.1175/BAMS-D-13-00047.1>, <https://doi.org/10.1175/BAMS-D-13-00047.1>, 2014.
- 315 Chevallier, F., Morcrette, J.-J., Chéruy, F., and Scott, N. A.: Use of a neural-network-based long-wave radiative-transfer scheme in the ECMWF atmospheric model, *Quarterly Journal of the Royal Meteorological Society*, 126, 761–776, <https://doi.org/10.1002/qj.49712656318>, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49712656318>, 2000.
- Chollet, F., Allaire, J., et al.: R Interface to Keras, <https://github.com/rstudio/keras>, 2017.
- Chollet, F. et al.: Keras, <https://keras.io>, 2015.
- 320 Clark, M., Babich, R., Barros, K., Brower, R., and Rebbi, C.: Solving lattice QCD systems of equations using mixed precision solvers on GPUs, *Computer Physics Communications*, 181, 1517 – 1528, <https://doi.org/https://doi.org/10.1016/j.cpc.2010.05.002>, <http://www.sciencedirect.com/science/article/pii/S0010465510001426>, 2010.
- Dorigo, W., Wagner, W., Albergel, C., Albrecht, F., Balsamo, G., Brocca, L., Chung, D., Ertl, M., Forkel, M., Gruber, A., Haas, E., Hamer, P. D., Hirschi, M., Ikonen, J., de Jeu, R., Kidd, R., Lahoz, W., Liu, Y. Y., Miralles, D., Mistelbauer, T., Nicolai-Shaw, N.,
325 Parinussa, R., Pratola, C., Reimer, C., van der Schalie, R., Seneviratne, S. I., Smolander, T., and Lecomte, P.: ESA CCI Soil Moisture for improved Earth system understanding: State-of-the art and future directions, *Remote Sensing of Environment*, 203, 185 – 215, <https://doi.org/https://doi.org/10.1016/j.rse.2017.07.001>, <http://www.sciencedirect.com/science/article/pii/S0034425717303061>, *earth Observation of Essential Climate Variables*, 2017.
- Egri, G. I., Fodor, Z., Hoelbling, C., Katz, S. D., Nógrádi, D., and Szabó, K. K.: Lattice QCD as a video game, *Computer Physics Communications*, 177, 631 – 639, <https://doi.org/https://doi.org/10.1016/j.cpc.2007.06.005>, <http://www.sciencedirect.com/science/article/pii/S0010465507003025>, 2007.
- 330 Ekici, A., Beer, C., Hagemann, S., Boike, J., Langer, M., and Hauck, C.: Simulating high-latitude permafrost regions by the JSBACH terrestrial ecosystem model, *Geoscientific Model Development*, 7, 631–647, <https://doi.org/10.5194/gmd-7-631-2014>, <https://gmd.copernicus.org/articles/7/631/2014/>, 2014.
- 335 Fuhrer, O., Chadha, T., Hoefler, T., Kwasniewski, G., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, C., Schulthess, T. C., and Vogt, H.: Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0, Geo-

- scientific Model Development, 11, 1665–1681, <https://doi.org/10.5194/gmd-11-1665-2018>, <https://gmd.copernicus.org/articles/11/1665/2018/>, 2018.
- 340 Fécan, F., Marticorena, B., and Bergametti, G.: Parametrization of the increase of the aeolian erosion threshold wind friction velocity due to soil moisture for arid and semi-arid areas, 17, 149–157, <https://doi.org/10.1007/s00585-999-0149-7>, 1999.
- Gruber, A., Dorigo, W. A., Crow, W., and Wagner, W.: Triple Collocation-Based Merging of Satellite Soil Moisture Retrievals, *IEEE Transactions on Geoscience and Remote Sensing*, 55, 6780–6792, 2017.
- Gruber, A., Scanlon, T., van der Schalie, R., Wagner, W., and Dorigo, W.: Evolution of the ESA CCI Soil Moisture climate data records and their underlying merging methodology, *Earth System Science Data*, 11, 717–739, <https://doi.org/10.5194/essd-11-717-2019>, <https://essd.copernicus.org/articles/11/717/2019/>, 2019.
- 345 Hatfield, S., Chantry, M., Düben, P., and Palmer, T.: Accelerating High-Resolution Weather Models with Deep-Learning Hardware, in: *Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '19*, Association for Computing Machinery, New York, NY, USA, <https://doi.org/10.1145/3324989.3325711>, <https://doi.org/10.1145/3324989.3325711>, 2019.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R.: Improving neural networks by preventing co-adaptation of feature detectors, *CoRR*, abs/1207.0580, <http://arxiv.org/abs/1207.0580>, 2012.
- 350 Jöckel, P., Tost, H., Pozzer, A., Brühl, C., Buchholz, J., Ganzeveld, L., Hoor, P., Kerkweg, A., Lawrence, M. G., Sander, R., Steil, B., Stiller, G., Tanarhte, M., Taraborrelli, D., van Aardenne, J., and Lelieveld, J.: The atmospheric chemistry general circulation model ECHAM5/MESSy1: consistent simulation of ozone from the surface to the mesosphere, *Atmospheric Chemistry and Physics*, 6, 5067–5104, <https://doi.org/10.5194/acp-6-5067-2006>, 2006.
- 355 Joutsi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., Boyle, R., Cantin, P.-I., Chao, C., Clark, C., Coriell, J., Daley, M., Dau, M., Dean, J., Gelb, B., Ghaemmaghani, T. V., Gottipati, R., Gulland, W., Hagmann, R., Ho, C. R., Hogberg, D., Hu, J., Hundt, R., Hurt, D., Ibarz, J., Jaffey, A., Jaworski, A., Kaplan, A., Khaitan, H., Killebrew, D., Koch, A., Kumar, N., Lacy, S., Laudon, J., Law, J., Le, D., Leary, C., Liu, Z., Lucke, K., Lundin, A., MacKean, G., Maggiore, A., Mahony, M., Miller, K., Nagarajan, R., Narayanaswami, R., Ni, R., Nix, K., Norrie, T., Omernick, M., Penukonda, N., Phelps, A., Ross, J., Ross, M., Salek, A., Samadiani, E., Severn, C., Sizikov, G., Snelham, M., Souter, J., Steinberg, D., Swing, A., Tan, M., Thorson, G., Tian, B., Toma, H., Tuttle, E., Vasudevan, V., Walter, R., Wang, W., Wilcox, E., and Yoon, D. H.: In-Datacenter Performance Analysis of a Tensor Processing Unit, in: *ISCA '17: The 44th Annual International Symposium on Computer Architecture*, ACM, New York, NY, USA, <https://doi.org/10.1145/3079856.3080246>, <https://doi.org/10.1145/3079856.3080246>, 2017.
- 360 Kadow, C., Hall, D. M., and Ulbrich, U.: Artificial intelligence reconstructs missing climate information, *Nature Geoscience*, 13, 408–413, <https://doi.org/10.1038/s41561-020-0582-5>, <https://doi.org/10.1038/s41561-020-0582-5>, 2020.
- Klein, C. and Taylor, C. M.: Dry soils can intensify mesoscale convective systems, *Proceedings of the National Academy of Sciences*, 117, 21 132–21 137, <https://doi.org/10.1073/pnas.2007998117>, 2020.
- Klingmüller, K., Pozzer, A., Metzger, S., Stenchikov, G. L., and Lelieveld, J.: Aerosol optical depth trend over the Middle East, *Atmospheric Chemistry and Physics*, 16, 5063–5073, <https://doi.org/10.5194/acp-16-5063-2016>, <https://acp.copernicus.org/articles/16/5063/2016/>, 2016.
- 370 Klingmüller, K., Metzger, S., Abdelkader, M., Karydis, V. A., Stenchikov, G. L., Pozzer, A., and Lelieveld, J.: Revised mineral dust emissions in the atmospheric chemistry–climate model EMAC (MESSy 2.52 DU_Astitha1 KKDU2017 patch), *Geoscientific Model Development*, 11, 989–1008, <https://doi.org/10.5194/gmd-11-989-2018>, <https://gmd.copernicus.org/articles/11/989/2018/>, 2018.

- Klingmüller, K., Karydis, V. A., Bacer, S., Stenchikov, G. L., and Lelieveld, J.: Weaker cooling by aerosols due to dust–pollution interactions, *Atmospheric Chemistry and Physics*, 20, 15 285–15 295, <https://doi.org/10.5194/acp-20-15285-2020>, <https://acp.copernicus.org/articles/20/15285/2020/>, 2020.
- Kohoutová, L., Heo, J., Cha, S., Lee, S., Moon, T., Wager, T. D., and Woo, C.-W.: Toward a unified framework for interpreting machine-learning models in neuroimaging, *Nature Protocols*, 15, 1399–1435, <https://doi.org/10.1038/s41596-019-0289-5>, <https://doi.org/10.1038/2Fs41596-019-0289-5>, 2020.
- 380 Krizhevsky, A., Sutskever, I., and Hinton, G. E.: ImageNet Classification with Deep Convolutional Neural Networks, in: *Advances in Neural Information Processing Systems*, edited by Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q., vol. 25, pp. 1097–1105, Curran Associates, Inc., <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>, 2012.
- LeCun, Y., Bengio, Y., and Hinton, G.: Deep learning, 521, 436–444, <https://doi.org/10.1038/nature14539>, <http://www.pubmed.org/26017442>, 2015.
- 385 Leiserson, C. E., Thompson, N. C., Emer, J. S., Kuszmaul, B. C., Lamson, B. W., Sanchez, D., and Schardl, T. B.: There’s plenty of room at the Top: What will drive computer performance after Moore’s law?, *Science*, 368, <https://doi.org/10.1126/science.aam9744>, <https://science.sciencemag.org/content/368/6495/eaam9744>, 2020.
- Markidis, S., Chien, S. W. D., Laure, E., Peng, I. B., and Vetter, J. S.: NVIDIA Tensor Core Programmability, Performance Precision, in: *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pp. 522–531, <https://doi.org/10.1109/IPDPSW.2018.00091>, 2018.
- 390 Marticorena, B. and Bergametti, G.: Modeling the atmospheric dust cycle: 1. Design of a soil-derived dust emission scheme, 100, 16 415, <https://doi.org/10.1029/95JD00690>, 1995.
- MESSy: The Modular Earth Submodel System, <https://www.messy-interface.org>, last access: 11 November 2020, 2020.
- Montavon, G., Samek, W., and Müller, K.-R.: Methods for interpreting and understanding deep neural networks, *Digital Signal Processing*, 73, 1 – 15, <https://doi.org/https://doi.org/10.1016/j.dsp.2017.10.011>, <http://www.sciencedirect.com/science/article/pii/S1051200417302385>, 2018.
- MPI-M: Code availability, <https://mpimet.mpg.de/en/science/modeling-with-icon/code-availability>, last access: 11 November 2020, 2020.
- Müller, A., Deconinck, W., Kühnlein, C., Mengaldo, G., Lange, M., Wedi, N., Bauer, P., Smolarkiewicz, P. K., Diamantakis, M., Lock, S.-J., Hamrud, M., Saarinen, S., Mozdzyński, G., Thiemert, D., Gllinton, M., Bénard, P., Voitus, F., Colavolpe, C., Marguinaud, P., Zheng, Y., Van Bever, J., Degrauwe, D., Smet, G., Termonia, P., Nielsen, K. P., Sass, B. H., Poulsen, J. W., Berg, P., Osuna, C., Fuhrer, O., Clement, V., Baldauf, M., Gillard, M., Szmelter, J., O’Brien, E., McKinstry, A., Robinson, O., Shukla, P., Lysaght, M., Kulczewski, M., Ciznicki, M., Piątek, W., Ciesielski, S., Błażewicz, M., Kurowski, K., Procyk, M., Sychala, P., Bosak, B., Piotrowski, Z. P., Wyszogrodzki, A., Raffin, E., Mazauric, C., Guibert, D., Douriez, L., Vigouroux, X., Gray, A., Messmer, P., Macfaden, A. J., and New, N.: The ESCAPE project: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale, *Geoscientific Model Development*, 12, 4425–4441, <https://doi.org/10.5194/gmd-12-4425-2019>, <https://gmd.copernicus.org/articles/12/4425/2019/>, 2019.
- 400 Pu, B., Ginoux, P., Guo, H., Hsu, N. C., Kimball, J., Marticorena, B., Malyshev, S., Naik, V., O’Neill, N. T., Pérez García-Pando, C., Paireau, J., Prospero, J. M., Shevliakova, E., and Zhao, M.: Retrieving the global distribution of the threshold of wind erosion from satellite data and implementing it into the Geophysical Fluid Dynamics Laboratory land–atmosphere model (GFDL AM4.0/LM4.0), *Atmospheric Chemistry and Physics*, 20, 55–81, <https://doi.org/10.5194/acp-20-55-2020>, <https://doi.org/10.5194/2FACP-20-55-2020>, 2020.
- 410 R Core Team: R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, <https://www.R-project.org/>, 2019.

- Schenk, H. J. and Jackson, R. B.: ISLSCP II Ecosystem Rooting Depths, <https://doi.org/10.3334/ORNLDAAAC/929>, http://daac.ornl.gov/cgi-bin/dsviewer.pl?ds_id=929, 2009.
- Schmidhuber, J.: Deep learning in neural networks: An overview, *Neural Networks*, 61, 85 – 117, <https://doi.org/https://doi.org/10.1016/j.neunet.2014.09.003>, <http://www.sciencedirect.com/science/article/pii/S0893608014002135>, 2015.
- Shangguan, W., Dai, Y., Duan, Q., Liu, B., and Yuan, H.: A global soil data set for earth system modeling, *Journal of Advances in Modeling Earth Systems*, 6, 249–263, <https://doi.org/https://doi.org/10.1002/2013MS000293>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2013MS000293>, 2014.
- 420 Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D.: Mastering the game of Go with deep neural networks and tree search, *Nature*, 529, 484–489, <https://doi.org/10.1038/nature16961>, <https://doi.org/10.1038%2Fnature16961>, 2016.
- Sun, J., Fu, J. S., Drake, J. B., Zhu, Q., Haidar, A., Gates, M., Tomov, S., and Dongarra, J.: Computational Benefit of
425 GPU Optimization for the Atmospheric Chemistry Modeling, *Journal of Advances in Modeling Earth Systems*, 10, 1952–1969, <https://doi.org/10.1029/2018MS001276>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001276>, 2018.
- Yashiro, H., Terai, M., Yoshida, R., Iga, S.-i., Minami, K., and Tomita, H.: Performance Analysis and Optimization of Non-hydrostatic ICosahedral Atmospheric Model (NICAM) on the K Computer and TSUBAME2.5, in: Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '16, Association for Computing Machinery, New York, NY, USA,
430 <https://doi.org/10.1145/2929908.2929911>, <https://doi.org/10.1145/2929908.2929911>, 2016.