

1 **DustNet (v1): Skillful neural network predictions**
2 **of dust aerosols over the Saharan Desert-Saharan**
3 **dust**

4 Trish E. Nowak^{1,3*}, Andy T. Augousti², Benno I. Simmons^{3†},
5 Stefan Siegert^{1†}

6 ^{1*}Mathematics and Statistics, University of Exeter, Exeter, EX4 QH, UK.

7 ²Department of Mechanical Engineering, Kingston University, London,
8 SW15 3DW, UK.

9 ³Centre for Ecology and Conservation, University of Exeter, Penryn,
10 TR10 9FE, UK.

11 *Corresponding author(s). E-mail(s): pn284@exeter.ac.uk;

12 Contributing authors: augousti@kingston.ac.uk;

13 bsimmons.research@gmail.com; s.siegert@exeter.ac.uk;

14 [†]These authors should be considered joint last author.

15 **Abstract**

16 Suspended in the atmosphere are millions of tonnes of mineral dust which
17 interacts with weather and climate. Accurate representation of mineral dust in
18 weather models is vital, yet remains challenging. Large scale weather models
19 use high power supercomputers and take hours to complete the forecast. Such
20 computational burden allows them to only include monthly climatological means
21 of mineral dust as input states inhibiting their forecasting accuracy. Here, we
22 introduce DustNet a simple, accurate and super fast forecasting model for 24-
23 hours (1-step) ahead predictions of aerosol optical depth (AOD). DustNet is a
24 custom-built 2-D Convolutional Neural Network (CNN) equipped with trans-
25 posed convolution layers. The model is trained on selected ERA5 meteorology
26 and past MODIS-AOD observational data as inputs. Our design of DustNet
27 ensures that the model trains in less than 8 minutes and creates predictions in
28 2.1 seconds on a desktop computer, without the need to utilize any Graphics
29 Processing Units (GPUs). Created by DustNet predictions outperform the
30 state-of-the-art physics-based model on coarse 1° x 1° resolution at 95% of
31 grid locations when compared to ground truth satellite data. The test results
32 show that the daily mean AOD over the entire study area highly correlates with

MODIS observational data, with Pearson's $r^2 = 0.91$. Our results demonstrate ~~show~~ DustNet's potential for fast and accurate AOD forecasting, which can easily be utilized by researchers without access to supercomputers or GPUs. ~~which could transform our understanding of dust's impacts on weather patterns.~~

Copyrights statement: For the purpose of open access, the author has applied a 'Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Keywords: 2D convolutional neural network, AOD, Saharan dust, forecast, spatiotemporal

1 Introduction

The Earth's atmosphere is loaded with approximately ~~26~~ ≈ 26 million tonnes of mineral dust - an atmospheric aerosol that represents the vast majority of mass burden in the atmosphere (Gliß et al., 2021; Kok et al., 2023). Each year, major sources emit over ~~5,000~~ $\approx 5,000$ million tonnes of dust globally (Kok et al., 2021b) and, although the majority of this material sinks at source, a substantial portion is transported over vast distances (Van Der Does et al., 2018). Once in the atmosphere, mineral dust interacts with the Earth systems and impacts weather, climate, human health and infrastructure, from fisheries to aviation (Shao et al., 2011; Knippertz and Stuut, 2014; Highwood and Ryder, 2014; Nenes et al., 2014; Miller et al., 2014; Jickells et al., 2014; Morman and Plumlee, 2014; Kok et al., 2023).

Despite its importance, representing atmospheric dust **aerosols** in weather and climate models is challenging (Parajuli et al., 2022; Kok et al., 2023). For example, physics-based Numerical Weather Prediction (NWP) and climate models struggle to fully represent the dust cycle with adequate emission, transport and generation (Evan et al., 2014; Kok et al., 2021a; Gliß et al., 2021; Zhao et al., 2022). ~~Instead, the Integrated Forecasting System (IFS) of the European Center for Medium-Range Weather Forecasting (ECMWF) creates predictions that use aerosol optical depth (AOD)~~ ~~Instead, the Integrated Forecasting System (IFS) of the European Center for Medium-Range Weather Forecasting (ECMWF) creates predictions that use aerosol optical depth (AOD)~~ based on monthly-mean climatological fields only (Bozzo et al., 2017). A limitation in computational resources is highlighted as one of the reasons for the lack of a dedicated aerosol scheme, since such a development would significantly increase the computational burden of the system (Mulcahy et al., 2014). The monthly mean AOD, developed by the Copernicus Atmosphere Monitoring Service (CAMS), provides a reasonable trade-off in global weather forecasting. **However, a more accurate representation of the AOD** would have significant benefits, such as large improvements in the representation of the summer monsoon circulation or precipitation patterns in the Sahel region (Bozzo et al., 2020; Balkanski et al., 2021).

Recent developments in the field of AI present a significant opportunity to overcome the computational burden of a dedicated physics-based aerosol scheme. Models

such as GraphCast, Pangu-Weather, and FourCastNet can now skillfully predict the main ERA5 variables and in many cases outperform the state-of-the-art NWP models (Lam et al., 2023; Bi et al., 2023; Pathak et al., 2022). To date, attempts to forecast atmospheric aerosols with neural network architectures have shown varying levels of success. “Satisfying” results were reported (Kang et al., 2019; Daoud et al., 2021) when applying a long-short-term memory (LSTM) architecture to local AOD forecasts. The application of a U-NET architecture revealed a skillful detection of classified ‘dust events’ at 67% precision rate (Sarafian et al., 2023). A lack of comparisons to the current physics-based forecasts, or inclusion of standardised skill metrics, makes direct comparison between AOD forecasting models nearly impossible.

Here, we present a unique application of 2D convolutional neural networks (CNN) to forecast atmospheric aerosol levels. We use our model (hereafter ‘DustNet’) to produce 24-hour spatial forecasts of AOD over North Africa. Computationally cheap and extremely fast, DustNet runs on a modestly configured laptop, rather than a high-power computer (HPC) - a fraction of the computational power required by traditional NWP models. The model trains in less than 8 minutes and predicts in 2.1 seconds. We compare the predictions of DustNet, and the corresponding daily CAMS forecasts, against the satellite-derived data using standard evaluation metrics, such as the root mean squared error (RMSE) and an accuracy correlation coefficient, to facilitate easy comparison with future AI models. The advantage of a smaller processing power requirement and rapid speed of prediction, combined with the accuracy of the forecast, makes our model a valuable complement to traditional AOD forecasting systems.

2 Methods

2.1 Study area

To effectively forecast dust aerosols, our study area encompasses the global principal dust generation source - the Sahara Desert - which is responsible for ~~~55%~~^{over 55%} of the 1,536 million tonnes of total global dust emitted annually (Ginoux et al., 2012). The region (~~Fig. 6 upper map~~) covers an area from 0° - 31°N and 20°W - 31°E (31 × 51 grid cells), with a longitudinal centre around the Bodélé Depression (16.5°N, 16.5°E). Located in northern Chad, this single location generates an estimated 6–18% of global dust emissions, which total to approximately 182 ± 65 million tonnes per year, the region is of major importance in models that seek to capture dust generation (Todd et al., 2007). To capture the seasonal south-westward dust transport across the Sahara and towards the Atlantic Ocean, our region includes additional grid cells to the south and west of the Bodélé Depression.

This choice allowed us to gain a sufficient amount of training data, with 51×31 grid cells providing 1,581 pixels for each training day, thereby ensuring robust model performance. By selecting this region, we were able to strike a balance between training efficiency, training speed, and prediction accuracy, making it possible to achieve effective dust aerosol forecasting. Furthermore, this approach enabled us to train the model on a traditional desktop computer without relying on cloud resources for data storage, making our approach more accessible and cost-effective. Additionally, the study region effectively captures dust aerosol generation and transport on selected

117 features, which is essential for accurate forecasting. Finally, by minimizing the area
118 to the Saharan Desert and consequently reducing the amount of chosen training fea-
119 tures, we were able to avoid adding different ocean and terrain processes, leading to
120 reduced model complexity without compromising performance.

121 2.2 Datasets

122 2.2.1 AOD data

123 We retrieved the AOD data from the Moderate Resolution Imaging Spectroradiometer
124 (MODIS) instrument located on board both Aqua and Terra spacecraft. With daily
125 temporal resolution over a period of 20 years starting from 1st January 2003 to 31st
126 December 2022, the AOD data yields 2×7305 files. We used the quality-controlled
127 level-3 data for AOD at 550nm. Choosing the combined mean of Dark Target and
128 Deep Blue algorithms provided a full coverage above bright and dark surfaces at a
129 horizontal resolution of $1^\circ \times 1^\circ$ (Hubanks et al., 2015). This choice provided a good
130 spatiotemporal coverage of AOD data above both land and ocean surfaces.

131 2.2.2 ERA5 data

132 Meteorological data comes from the fifth generation of European Centre for Medium-
133 Range Weather Forecast (ECMWF) atmospheric reanalysis project (ERA5) and
134 consists of 5 parameters: wind u component, wind v component, vertical velocity,
135 temperature and relative humidity. Each parameter was retrieved at 5 pressure levels
136 550hPa, 750hPa, 850hPa, 950hPa and 1000hPa. This choice provided us with 35 dis-
137 tinctive features representing atmospheric conditions from ground level to $\approx 5\text{km}$ in
138 vertical height. The ERA5 data is available on an hourly basis, but here we only chose
139 the data representing conditions for midday (12:00 UTC). This allows us to represent
140 the mid-point in atmospheric conditions between the Terra and Aqua satellite over-
141 passes above the equator (10:30am and 1:30pm respectively). To further match the
142 meteorological data with AOD, we chose a daily temporal resolution between 2003-
143 2022. The horizontal resolution of ERA5 data is $0.25^\circ \times 0.25^\circ$. To match this with the
144 AOD resolution of $1^\circ \times 1^\circ$, the data was regridded (see section Data pre-processing for
145 details). (details in Section 2.3.3).

146 2.2.3 Timestamps

147 We created timestamps using the NumPy package (version 1.23.0) in Python with
148 a daily temporal resolution over 20 years from 2003 to 2022 (7,305 days). We then
149 expanded the array dimensions through replication to match the exact spatial reso-
150 lution of atmospheric variables, resulting in a coverage of 31×51 grid cells for each
151 day. Then we multiplied the file to match the exact spatial resolution of atmospheric
152 variables and a coverage of 31×51 grid cells for each day.

2.2.4 CAMS forecast

We obtained daily ‘Total aerosol optical depth at 550nm’ forecast data from ‘CAMS global atmospheric composition forecasts’. CAMS forms a part of the ECMWF Integrated Forecasting System (IFS), and is a sophisticated numerical weather forecasting model (NWP) (Bozzo et al., 2017). During the AOD data assimilation process, CAMS utilises data from MODIS, among other satellites, together with data from ground-based observation stations. The model then uses physics and chemistry principles to forecast hourly AOD values on a single level for up to 5 days (120hr) ahead (Morcrette et al., 2009; Benedetti et al., 2009). For consistency, we only chose forecasts representing 12:00 UTC to capture the midpoint conditions between Aqua and Terra overpasses above the equator. The temporal extent choice was also matched to our predictions. Therefore, we initiated forecasts on midday 1st January 2020 until 30th December 2022 for 1095 days forecast between 2nd January 2020 and 31st December 2022. CAMS data is provided at a $0.4^\circ \times 0.4^\circ$ spatial resolution. To match with our data, we therefore used an identical approach as for the ERA5 datasets to regrid to a $1^\circ \times 1^\circ$ resolution (details in [Data pre-processing](#) section).

2.3 Data pre-processing

2.3.1 Data imputation

We combined data from the MODIS Aqua and Terra data sources at each individual location and time by labelling AOD data as missing whenever both sources were missing, using available data from one source if the other is missing, and averaging both sources whenever both are available. This data combination step reduces the total fraction of missing AOD values from 32.81% in Aqua and 30.89% in Terra to 19.89% in the combined data set. The remaining missing AOD values are imputed by spatial interpolation (individually for each time step) using Lattice Kriging (Hartman and Hössjer, 2008; Rue and Held, 2005) on four nearest neighbours with uniform weights. To validate the imputation method, we randomly held out 10% of the AOD data and compared them to their imputed values. The mean squared error of the imputed values is 0.005 which is less than 5.30% of the total variance of the AOD data. The MSE was found to be insensitive to the choice of the Kriging hyperparameter, with relative differences of less than 0.0003% over a wide range of values (see supplementary Fig.7). See Code and data availability section for links containing the full Python code for imputation.

2.3.2 AOD-lag

~~We use 5 preceding days of imputed AOD data as features to predict AOD on a given day. Hence, we had to remove the first 5 timestamps from the database as these did not have complete features available, resulting in a new total of 7,300.~~

2.3.3 ERA5-regridding

The ERA5 data (Hersbach et al., 2018) is supplied with a ~~horizontal~~[vertical](#) resolution of $0.25^\circ \times 0.25^\circ$ and thus needed regridding to match the AOD resolution. We processed

all meteorological data using Python version 3.8.13 and the Iris v 3.2.1 package. We used nearest-neighbour interpolation from the Iris package to convert each feature to a common $1^\circ \times 1^\circ$ resolution.

2.4 Feature engineering

To enhance the model’s predictive skill we incorporated two aspects of feature engineering: AOD lag and seasonal features. To account for temporal dependences, we use 5 preceding days of AOD data as features to predict AOD on a given day. Hence, we had to remove the first 5 timestamps from the database as these did not have complete features available, consequently reducing the total number of timestamps to 7,300. Additionally, we included trigonometric transformations of timestamps as seasonal features using sine:

$$x_{ijt}^{(42)} = \sin \left(2\pi \frac{t}{365.2425} \right), \quad (1)$$

and similarly using the cosine:

$$x_{ijt}^{(43)} = \cos \left(2\pi \frac{t}{365.2425} \right), \quad (2)$$

where t represents the day of the year. Timestamps are constant across space and allow the model to represent periodic variations on seasonal timescales. Thus, together with timestamps, our final total input consisted of 43 features

2.4.1 Combining and normalising

We combined the meteorological data with AOD data and engineered features into a single 4D NumPy array of shape 7300, 51, 31, 43, where the first dimension represents time, the second and third are longitude and latitude respectively, and features are stored along the last dimension. Let the x_{ijt} be the value of feature x at grid point i, j and time t . We normalised all features using min-max normalisation:

$$x_{ijt,norm} = \frac{(x_{ijt} - x_{min})}{(x_{max} - x_{min})} \quad (3)$$

where x_{min} and x_{max} are the overall minimum and maximum of a feature x over all grid points and timestamps in the training data.

2.4.2 Seasonal features

~~Our first 41 features contain atmospheric variables as described above. Additionally we included the sine and cosine of timestamps as seasonal features using:~~

$$x_{ijt}^{(42)} = \sin \left(2\pi \frac{t}{365.2425} \right) \quad (4)$$

219 ~~and similarly using the cosine:~~

$$x_{ijt}^{(43)} = \cos\left(2\pi \frac{t}{365.2425}\right) \quad (5)$$

220 ~~where t represents the day of the year. Timestamps are constant across space and~~
221 ~~allow the model to represent periodic variations on seasonal timescales. Thus, together~~
222 ~~with timestamps, our final total input consisted of 43 features.~~

223 2.4.3 Training, validation, test split

224 We split the data along the time dimension into 70%, 15% and 15% for training,
225 validation and test sets respectively. Splitting data with consecutive time steps yielded
226 better results than a random split. ~~The use of consecutive time steps ensures that each~~
227 ~~subset is composed of data points that are temporally distinct. This reduces the risk~~
228 ~~of autocorrelation and improves the model's ability to generalize to new, unseen data~~
229 ~~(Rasp et al., 2020).~~ Therefore, the training set covered 5,110 consecutive days from
230 6th January 2003 until 1st January 2017 (inclusive of both days). The validation set
231 took 1,095 consecutive days from 2nd January 2017 to 1st January 2020. Finally, we
232 set aside a test set, with 1,095 days of data from 2nd January 2020 to 31st December
233 2022. We made sure that the model never had access to the test set during the training
234 and validation processes and only after these were complete did we introduce the test
235 data and run our model to obtain predictions.

236 2.5 Designing CNN models

237 To find the best forecast of the daily AOD, we designed three CNN models based on
238 Hinton et al. (1995); LeCun et al. (2015); Goroshin et al. (2015). We used the end-
239 to-end open source machine learning platform TensorFlow 2, together with the Keras
240 high-level API (Chollet et al., 2015). Each model uses a different architecture based on
241 two-dimensional (2D) convolutions (hereafter Conv2D). In general, the Conv2D neural
242 network architecture enables regression problems in image analysis to be addressed
243 and is particularly effective at capturing spatial patterns in two-dimensional images.
244 The efficiency of Tensorflow allows for training and inference to be run on traditional
245 desktops or laptops rather than requiring HPC's. All models described hereafter were
246 run using Python version 3.10.10 on a MacBook Pro with an Apple M1 Pro and 32GB
247 RAM. ~~The models did not utilize any GPUs and can thus be replicated by users~~
248 ~~without access to a supercomputer.~~

249 We have chosen 'Adam' optimizer and the mean-squared-error (MSE) as a loss
250 function. ~~These options~~~~We also performed a series of diagnostic tests in order to choose~~
251 ~~the best optimizer and loss function. Tested loss functions included CoSine, Huber,~~
252 ~~LogCosh, Mean Absolute Percentage Error, Mean Absolute Error, Mean Squared~~
253 ~~Logarithmic Error and Mean Squared Error. We assessed the performance based on~~
254 ~~the lowest mean squared error, the speed of the overall training time and the time~~
255 ~~taken per step. The mean-squared-error (MSE) loss function together with the Adam~~
256 ~~optimizer offered optimal results in terms of training times and wereand was used~~
257 ~~for further analysis. For the Adam optimizer we used a learning rate of 0.001 and an~~

exponential decay rate of 0.9, which are default settings following (Kingma and Ba, 2014).

We determined the optimal size of the convolving window (kernel size) and the number of strides with a series of diagnostic tests. The results of these tests are presented in Table 1 with the optimal choice in bold based on minimising the mean squared error and the speed of the training time. The final design included a kernel size of (2,2) with a stride equal to 2, which produced the optimal MSE to training time ratio. We recognise that we have not tested every possible combination, thus it may be possible to achieve a better performing design.

Table 1 ~~Effects~~**Test results** of choosing different kernel sizes **on training time and MSE** for 2 models: Conv2D and U-NET. For simplicity, this test was run on a subset of data. The optimal choice is presented in bold font. Note that a small improvement in the MSE for a kernel size (3,3) was disregarded in favour of a much faster training time and time per step for kernel size (2,2).

	Conv2D			U-NET		
Kernel size	(5,5)	(3,3)	(2,2)	(5,5)	(3,3)	(2,2)
Training time	42min	23min	3min	1h41min	1h7min	23min
Time per step	28s	28s	12s	144s	140s	88s
MSE	0.00174	0.00133	0.00134	0.00175	0.00148	0.00151

We initially assigned 50 epochs to each training regime and monitored the performance using the mean squared error of training to validation loss. We also configured each model with Early Stopping and a patience of 4 epochs. This set up halts the training time when there is no improvement in validation loss after 4 consecutive iterations and prevents the model from over-fitting to training data (see [supplementary Fig. 8](#)). Our set-up saved the optimal ratio of training time versus validation loss and used the best performance to run predictions. Below, each model’s architecture is described in detail.

2.5.1 Conv2D model

For the first AOD prediction model we adapted a classical design of CNN. The Conv2D architecture, inspired by the visual system, applies filters (or convolutions) to capture spatial patterns in two-dimensional images (LeCun et al., 2015). The network performs feature extraction and learns representations at different scales. Such representations allow the network to identify relevant information and thus make predictions. Each of the hidden layers in our model was designed with a maximum of 264 and a minimum of 16 filters, as well as a 2×2 kernel size, which specifies the height and width of the 2D convolution window ([see model schematic in supplementary materials Fig. 9](#)). Learning of the complex representation is made possible by the non-linearity provided to the model by a correctly chosen activation function. Ramachandran et al. (2017) suggested an improvement to the popular ReLU activation function by proposing the Swish function. This method gained in popularity as it is capable of smoother output representation as well as more consistent performance (Rasamoelina et al., 2020). The Swish activation function proved to yield the best performance and thus we used

290 it throughout the model layers. An architecture constructed in this way provided
291 1,291,009,218,673 trainable parameters.

292 2.5.2 U-NET model

293 The architecture of our second model employed a U-NET like design, first proposed
294 by Ronneberger et al. (2015) for the purpose of biomedical image classification. The
295 model is characterised by its “U” shape design which employs both contracting and
296 expanding pathways to identify specific features within images. Here, we follow the
297 approach of Ayzel et al. (2020) who, inspired by U-NET, designed their RainNet model
298 for precipitation nowcasting. Thus, we also divided our model into two parts, encoder
299 and decoder, and utilised skip connections between both paths via concatenation layers
300 - unique features of the U-NET model. The encoder (or contracting) pathway of the
301 model included six Conv2D layers with Swish activation and a 2×2 kernel size, as well
302 as two MaxPooling2D layers with pool size $2 \times 21 \times 1$ (see supplementary Fig 10 for
303 model’s schematic drawing). The decoder (or expanding) pathway had five Conv2D
304 layers with two UpSampling2D and two Concatenate layers. The input layers were
305 bordered with a ZeroPadding2D layer which was cropped to the original size of 31×51
306 with Cropping2D in the output layer. Unlike the original U-NET network, our design
307 received 4-dimensional arrays of shape $7, 300 \times 31 \times 51 \times 42$ and generated an output
308 image of a shape of 31×51 for each prediction time step. Thus, the prediction generated
309 1905 images corresponding to dates from 2nd January 2020 to 31st December 2022.
310 The final U-NET model architecture provided 847,937 trainable parameters.

311 2.5.3 DustNet model

312 The last model design built upon the architecture of Conv2D and U-NET. This unique
313 design replaces the Concatenate layers with Transpose convolution layers, also
314 known as Deconvolutional Networks (Zeiler et al., 2010). Schematically represented in
315 Fig. 1 the input layer was first padded with a border of zeros (ZeroPadding2D) which
316 increased the input shape from $31 \times 51 \times 43$ to $40 \times 64 \times 43$. Zero padding enabled
317 the convolution to produce the same output size for multiple input sizes (Dumoulin
318 and Visin, 2016). We then applied the 2D convolving windows (Fig. 1 - pink arrows),
319 which moved over each padded input with a 2×2 kernel size and 2×2 strides which
320 allow upsampling. The first six layers of convolving, or contracting pathway, consisted
321 of double 64, 128 and 256 filters, where every second layer included strides. This
322 allowed the model to decrease the input size while increasing the number of channels
323 ($5 \times 8 \times 256$). The ‘deconvolutions’, or expanding pathways, were then applied by
324 adding six Conv2D Transpose layers, with reversed order of filters to the contracting
325 pathway. This allowed the model to decrease the input size (down to 5×8) while
326 increasing the amount of channels ($5 \times 8 \times 256$) trainable parameters. A ‘deconvolution’
327 was then applied by adding Conv2D Transpose layers. An advantage of transposed
328 convolution is its ability to efficiently upscale input data by applying inverse convolu-
329 tions. This enables the network to increase the size compared to the input and thus
330 generates high-resolution images at finer spatial scales (Zeiler et al., 2010). A 2D crop-
331 ping layer was then added to bring the width and height shape back to its initial input
332 size of 31×51 , while the final convolution matched the output with desired target size

333 of $31 \times 51 \times 1$. The final architecture allowed the model to create a total of 1,291,009
 334 trainable parameters. Since this design yielded the optimal results of predicting dust
 335 aerosols in comparison to baseline models, we called it DustNet.

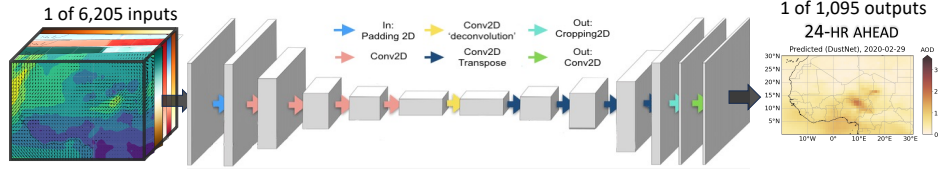


Fig. 1 Schematic representation of the DustNet model. Each of 6,205 inputs is first padded with a border of zeros using ZeroPadding2D (light blue arrow) to increase input sizedimensionality and allow the convolution windows to detect the borders. The features are then extracted by 2D convolution window (pink arrows) which decreases input sizedimensionality while increasing the number of trainable parameters. Then deconvolution is applied (yellow arrow) by including a 2D transpose network, which increases the size of the input (dark blue arrows) while maintaining connectivity between the layers. The output is then cropped back to match the initial input size (green arrow) and represents a 24-hr (1-step) ahead prediction.

336 2.5.4 Baseline models

337 We set the baselines as AOD climatological mean and persistence. The climatological
 338 means were calculated separately at each spatial location as the mean AOD over the
 339 training period. The climatological benchmark is constant in time. A time-varying
 340 baseline model is the persistence forecast, which uses the most recent observation of
 341 AOD as the 24-hour ahead prediction. Here, we used the values from the 1st day
 342 of calculated AOD lag from the reserved test set (values unseen by the model) to
 343 represent persistence. Both climatology and persistence act as null models, and a more
 344 sophisticated forecasting scheme should be able to outperform both in order to be
 345 considered useful.

346 2.6 Training CNN models

347 To train the models we used 17 years of daily data (2003-2019). We initiated the
 348 training on the first 15 years (70%) of data after which the models went into a self
 349 validation mode, for which we used the consecutive 2 years (15%) of data (see Section
 350 2.4.3 for full details on data split regime). The inputs included the value of the AOD
 351 over the previous 5 days and previous 1 day for each of 35 meteorological features
 352 (7 atmospheric variables at 5 pressure levels, see ERA5 data section in Methods).
 353 Regrided to a $1^\circ \times 1^\circ$ resolution over 31° of latitude by 51° of longitude, together
 354 with orography and the sine and cosine values of timestamps, the data produced
 355 a representative state consisting of 43 input features. Hence, for each of the 6,205
 356 training and validation days the models had access to 67,983 values.

2.7 Statistical analysis

2.7.1 CNN models evaluation

We evaluated each CNN model’s performance by assessing the training time, inference time taken per ‘time-step’, the MSE of predicted values in the test set, and the percentage improvement in the MSE above the climatology and persistence baseline models. ~~All three models were capable of producing an improved MSE above climatology and persistence baselines (see Table 2, with the best results indicated in bold font).~~ We then used the best performing model (DustNet) to visually evaluate its output against (unimputed) MODIS values. We inspected DustNet’s daily predictions for its ability to represent AOD spatially by mapping 28 consecutive days of predictions next to the corresponding data from MODIS (see ~~supplementary~~ Fig. 11). We looked for the model’s ability to capture the main dust generation sources, consistent AOD transport with prevailing winds, and correct distinctions of AOD accumulation between the ocean and land border.

To analyse the errors of the best performing model, we rearranged Equation 3 reverses normalisation of AOD predictions from each model:

$$y_{ijt,denorm} = y_{ijt,pred} (y_{max} - y_{min}) + y_{min} \quad (6)$$

where y_{pred} are the values predicted by the model, y_{max} is the maximum and y_{min} is the minimum AOD value from the training set. In this same manner, we used Equation 6 to reverse normalisation of the climatology and persistence predictions. We then assessed each CNN model by calculating the MSE between values predicted by the model using the de-normalised AOD denoted as \hat{A} , and the corresponding values from the test set (“true”) AOD value denoted as A . Here, we calculated a mean value along an axis of latitude N_{lat} and longitude N_{lon} , of our spatial coordinates at each prediction time step t , where $N_{lat}=31$, $N_{lon}=51$ and $N_t=1095$, using Equation 7:

$$MSE = \frac{1}{N_{lat}N_{lon}N_t} \sum_{i=1}^{N_{lat}} \sum_{j=1}^{N_{lon}} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt})^2 \quad (7)$$

We used this same process as described above to obtain the MSE for the climatology and persistence models. To ensure that model evaluation is only based on actually observed AOD values, all imputed AOD values were excluded from calculation of the MSE.

2.7.2 Validation of ~~results~~predictions

~~To validate our results, we compared our predictions with the ground-truth (not imputed) data from MODIS and the physics-based model (CAM5) fairly, we calculated the following metrics: mean bias error (MBE), RMSE, difference between RMSE’s ($\Delta RMSE$) and ACC. The metrics, defined below, follow a combination of notations from Bi et al. (2023) and Lam et al. (2023) adapted to spatial representation of temporally averaged values for each prediction day t ($N_t=1095$). All prediction values were~~

392 first de-normalised using Equation 6. Subsequently, we compared the model predic-
 393 tions (\hat{A}) with raw (unimputed) MODIS data (mean of Aqua and Terra) denoted as A .
 394 The climatological mean, denoted as A' , corresponds to the long-term average of AOD
 395 values from MODIS (2003-2022). To allow for comparison with the physics-based fore-
 396 cast, we tested the 24-hour (1-step ahead) predictions from CAMS using these same
 397 skill metrics, and compared them with the daily and seasonal results produced by the
 398 best performing model.

399 2.7.3 DustNet evaluation metrics

400 To compare predictions between the DustNet model, ground truth data from MODIS
 401 observations and the physics-based model (CAMS) fairly, we calculated the following
 402 metrics: mean bias error (MBE), RMSE, difference between RMSE's ($\Delta RMSE$)
 403 and ACC. The metrics, defined below, follow a combination of notations from Bi
 404 et al. (2023) and Lam et al. (2023) adapted to spatial representation of temporally
 405 averaged values for each prediction day t ($N_t=1095$). All prediction values were first
 406 de-normalised using Equation 6. Subsequently, we compared the model predictions
 407 (\hat{A}) with raw (unimputed) MODIS data (mean of Aqua and Terra) denoted as A . The
 408 climatological mean, denoted as A' , corresponds to the long-term average of AOD
 409 values from MODIS (2003-2022).

410 2.7.4 Spatial analysis

411 To analyse spatial characteristics of model performance, we calculated the temporal
 412 mean of model predictions ($N_t = 1095$) at each location (lat,lon). This allowed us to
 413 calculate mean bias error (MBE) between the predicted AOD (\hat{A}) and MODIS ground
 414 truth (A) for both DustNet and CAMS using Equation 8.

$$MBE_{spatial,ij} = \frac{1}{N_t} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt}) \quad (8)$$

415 We also calculated the spatial root mean square error ($RMSE_{spatial}$) for each model
 416 using Equation 9.

$$RMSE_{spatial,ij} = \sqrt{\frac{1}{N_t} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt})^2} \quad (9)$$

417 Calculating differences between RMSEs ($\Delta RMSE$) using Equation 10 allowed us to
 418 reveal specific locations at which predictions from one model outperformed the other.

$$\Delta RMSE_{spatial,ij} = RMSE_{spatial,ij}^{(CAMS)} - RMSE_{spatial,ij}^{(DustNet)} \quad (10)$$

419 Additionally, we calculated the spatial distribution of Anomaly Correlation Coefficient
 420 (ACC, Equation 11). Let \hat{A}' be the anomaly of predicted AOD values (\hat{A}), and A'

the anomaly of observed (ground truth A) AOD values, where the anomalies are the differences from MODIS climatology values, then:

$$ACC_{spatial,ij} = \frac{\sum_{t=1}^{N_t} [(\hat{A}'_{ijt} - \bar{A}'_{ijt}) \times (A'_{ijt} - \bar{A}'_{ijt})]}{\sqrt{\left[\sum_{t=1}^{N_t} (\hat{A}'_{ijt} - \bar{A}'_{ijt})^2\right] \times \left[\sum_{t=1}^{N_t} (A'_{ijt} - \bar{A}'_{ijt})^2\right]}} \quad (11)$$

The ACC is a common measure of skill which assesses the quality of prediction, and highlights anomalies between forecast and observed values. By subtracting the climatological mean from both, prediction and verification, the ACC measures the quality of prediction without giving misleadingly high results caused by seasonal variations. Refer to Fig. 2a—e and supplementary Fig. 12 for graphed representations of these calculations.

2.7.5 Temporal analysis

To analyse the model's predictions across different times, we calculated mean spatial AOD values for each prediction day. We also computed Pearson's correlation coefficients (r), associated p-values, and coefficient of determination (r^2) using the SciPy statistical package v.1.12 for each prediction day ($N=1095$) of spatially averaged data ($N_{lat}, N_{lon}=31, 51$). Corresponding results were calculated and plotted for both DustNet and CAMS forecasts with MODIS data. and plotted in supplementary Fig. ???. We have also adapted Equations 8 and 9 to temporal representation by using Equation 12 and 13.

$$MBE_{temporal,t} = \frac{1}{N_{lat}N_{lon}} \sum_{i=1}^{N_{lat}} \sum_{j=1}^{N_{lon}} (\hat{A}_{ijt} - A_{ijt}) \quad (12)$$

$$RMSE_{temporal,t} = \sqrt{\frac{1}{N_{lat}N_{lon}} \sum_{i=1}^{N_{lat}} \sum_{j=1}^{N_{lon}} (\hat{A}_{ijt} - A_{ijt})^2} \quad (13)$$

The graphed representations of temporal calculations can be seen in Fig. 6 and Fig. 13.

2.7.6 Justification of the selected points

In addition to spatial and temporal analyses, we focussed on four point locations to assess the model's performance at the local scale. The locations, shown in supplementary Fig. 14, were selected on the basis of a different aerosol type contributing to the total AOD, as well as prevailing meteorological conditions. We chose the region around the Bodélé Depression in Chad (16.5°N, 16.5°E) for its dust generation capability and consistency of high mineral dust loading (Washington et al., 2003). Nouadhibou in Mauritania (20.5°N, 17°W) is located at the edge of western Africa, where hot, dry Saharan air meets cool and moist Atlantic air (Carlson and Prospero, 1972). The temperature inversion creates a barrier for low horizontal flow of atmospheric dust, and instead forces an uplift of over 1.5km (Prospero and Carlson, 1972).

From this point atmospheric dust moves westward towards Central and South America at higher altitudes between 1.5km - 5km (Kaufman et al., 2005). To capture the transport of dust and fire smoke with southwestward winds towards South America (Kaufman et al., 2005) we chose a location over the Atlantic Ocean in the Gulf of Guinea (4°N, 4°W). For the fourth location, we chose the second largest city in Nigeria and the capital of Kano State (11.5°N, 8.5°E). Kano City is on a direct pathway of seasonal dust plumes known locally as the Harmattan season. During boreal winter the wind direction shifts to southwestward direction and transports the sand storms generated from the Bodélé Depression towards Kano, where they are associated with a large increase in air pollution (Anuforum, 2007; Schwanghart and Schütt, 2008; Sunnu et al., 2008).

2.7.7 Feature importance

We assessed feature importance using a perturbation-based method, where individual input channels were systematically altered to evaluate their contribution to model predictions. Specifically, each feature was zeroed out in turn, and the mean squared error (MSE) between the full prediction and the prediction with the altered input was calculated. This approach quantifies the sensitivity of the model's output to the absence of each feature, with higher MSE indicating greater importance. Results shown in supplementary Fig. 15 demonstrate that input channels corresponding to the features 'AOD 1 day lag' and 'vertical velocity at 850 hPa' exhibited the largest impact, indicating their relative importance for model predictions. Perturbation-based methods, such as this one, are widely used for assessing feature relevance in machine learning models due to their simplicity and interpretability (Covert et al., 2021; Molnar, 2022).

3 Results

3.1 DustNet model architecture and performance verification

To find the best deterministic AOD forecasting model, we compared three models. First, we adapted two leading CNN architectures, including 2-dimensional CNN and U-NET (Hinton et al., 1995; LeCun et al., 2015; Goroshin et al., 2015; Ayzel et al., 2020). We also custom-designed a 2D CNN with transposed layers (Zeiler et al., 2010). After comparing the performance of these three models (see Table 2 in the Methods section), we arrived at an optimal configuration for 24-hour dust aerosol forecasts and called our model DustNet. To train our DustNet model we used 17 years of daily AOD data (2003-2019) from the Moderate Resolution Imaging Spectroradiometer (MODIS) apparatus on board the Aqua and Terra satellites (see Methods section for full details). A schematic representation in Fig. 1 illustrates the inputs and output of the model. The inputs included the value of the AOD over the previous 5 days and previous 1 day for each of 35 meteorological features (7 atmospheric variables at 5 pressure levels, see ERA5 data section in Methods). Regrided to a $1^\circ \times 1^\circ$ resolution over 31° of latitude by 51° of longitude, together with orography and the sine and cosine values of timestamps, the data resulted in a representative state consisting of 67,983 values for each training day. The compiled model yielded nearly 1.3 million trainable

parameters and took 7min and 41s to complete the training process. Subsequently,
the the subsequent forecasts were produced in 2.1 seconds.

4 Performance verification

The comparative results of the three CNN models, shown in Table 2, demonstrate a clear advantage of the DustNet architecture in both computational efficiency and predictive accuracy. Developed in this study, DustNet achieves the shortest training time at 7 minutes and 41 seconds, nearly half that of the U-NET model. It also outperforms both U-NET and Conv2D in terms of mean squared error (MSE), achieving a value of 0.00153, which corresponds to a 53.68% improvement over the climatology baseline. Furthermore, DustNet generates forecasts in just 2.1 seconds, making it the fastest among the tested models. In contrast, Conv2D and U-NET require over 13 and 25 minutes, respectively, for training, while their resulting predictions show lower improvement over the climatological baseline. These findings indicate that DustNet is both more efficient and more accurate than the Conv2D and U-NET models, thereby demonstrating its skill in deterministic AOD forecasting.

To evaluate the resulting 24-hour (1-step ahead) predictions, we used 3 years of data (2020-2022), which were unseen by the model. Our initial baseline model included the climatological mean, which is often used in meteorological forecasts as a sensible default (Bozzo et al., 2020). The baseline tests revealed that DustNet improved (reduced) the mean squared error (MSE) by 53.68% in comparison to predictions based on the climatological mean (Table 2). The regimes used for training, validation and testing are included in section Training, validation, test split. To validate our results, we compared our predictions with the ground-truth (not imputed) data from MODIS, where the mean values between Aqua and Terra satellites, which record non-simultaneous measurements, provided the best representation of conditions around midday. To quantitatively assess the performance of the DustNet model against the ground truth we used two skill metrics: the root mean squared error (RMSE) and the anomaly correlation coefficient (ACC). To allow for comparison with the physics-based forecast, we tested the 24-hour predictions from CAMS using these same skill metrics, and compared them with the results produced by DustNet.

Added Table 2.

4.1 Performance of spatial forecast

We find that the DustNet model performs better in AOD forecasts than the physics-based CAMS model (Fig. 2). At nearly all spatial locations, DustNet predictions resulted in lower (better) RMSE values than CAMS during 2020-2022 (Fig. 2a and b). The greatest source of errors for both models was the most active dust source globally (Todd et al., 2007) — the Bodélé Depression (16.5°N, 16.5°E). Although this is the location of the highest error, here we show again that DustNet’s RMSE is nearly 50% lower than that produced by CAMS (0.62 *versus* 1.24 respectively). The Bodélé Depression is of global importance for two main reasons: (i) it is responsible for over 50% of the dust generated from the Sahara desert (Todd et al., 2007; Washington et al.,

Table 2 Normalised test results for three unique model architectures. Persistence and climatology baseline MSE’s of prediction to test data are presented below the table. The rows display results for total training time, time per iteration step and MSE for each kernel size of each model. The last column shows the percentage difference when compared to the climatological baseline.

CNN model	Training time	Time per step	MSE	Prediction time	Baseline ¹ improvement (%)
Conv2D	13min40s	34s	0.001895	4.1s	42.63%
U-NET	25min20s	53s	0.001691	4.9s	48.80%
DustNet	7min41s	17s	0.00153	2.1s	53.68%

Baseline MSE:

¹Climatology: 0.003303

²Persistence: 0.002992

2009; Jewell et al., 2021) and (ii) it was identified as the main source of minerals delivered seasonally to the Amazon basin (Koren et al., 2006; Jewell et al., 2021). A recent comparison of 14 physics-based models reveals their tendency to vastly underestimate the AOD forecast (ranging from -16% to -37%) in comparison to ground-based observations (Gliß et al., 2021). With nearly 40 million tonnes of dust emitted annually from the Bodélé Depression, lowering the forecasting error at this location, as achieved by DustNet, has the potential to vastly improve the forecasting of transported dust.

Overall, DustNet predictions outperformed CAMS forecasts on 95.26% of grid locations when comparing prediction errors [Fig. 2d](#)) and [e](#)).~~(Fig. 2e). In Fig. 2e~~), grid cells in the darkest brown colour indicate locations where the errors produced by CAMS were over 0.45 AOD higher than that of DustNet, with the maximum error difference reaching 1.24 AOD. These locations represent central Saharan desert and arid regions, indicating the AOD composed of mineral dust, and thereby the more skillful ability of DustNet to capture dust generation. Moreover, DustNet captures the high mean AOD over northern Nigeria (associated with the seasonal Harmattan haze (Anuforom, 2007; Sunnu et al., 2008; Schwanghart and Schütt, 2008)) more skillfully than CAMS (details in section Performance of seasonal-mean forecast below). However, there are two locations at which CAMS forecasts performed better than [Fig. 2d](#)) and [e](#)).~~DustNet (Fig. 2e)~~. Both of these locations are adjacent to the boundaries (SE and NW corners), beyond which DustNet was unable to obtain information on the processes during training, while the data used to generate the CAMS forecast was extracted from a larger region ([see section CAMS forecast for details](#)). Thus, the lack of information on processes at the boundaries may have affected the CAMS forecasts less than it affected DustNet. This, however, might be overcome by extending the study region for DustNet.

We also compare the ability of DustNet and CAMS to detect anomalies using the ACC, a quantitative metric used in previous similar studies (Lam et al., 2023; Bi et al., 2023). Here, DustNet also displays more skillful results than CAMS with a better (higher) ACC at [92.283%](#)~~92.28%~~ of grid cells shown in [Fig. 2d](#)) and [e](#)).~~Fig. 2d) and e~~). An ACC score above 60% is considered to be of value for forecasting purposes. The DustNet model surpasses this threshold at 79.89% of locations (white-yellow),

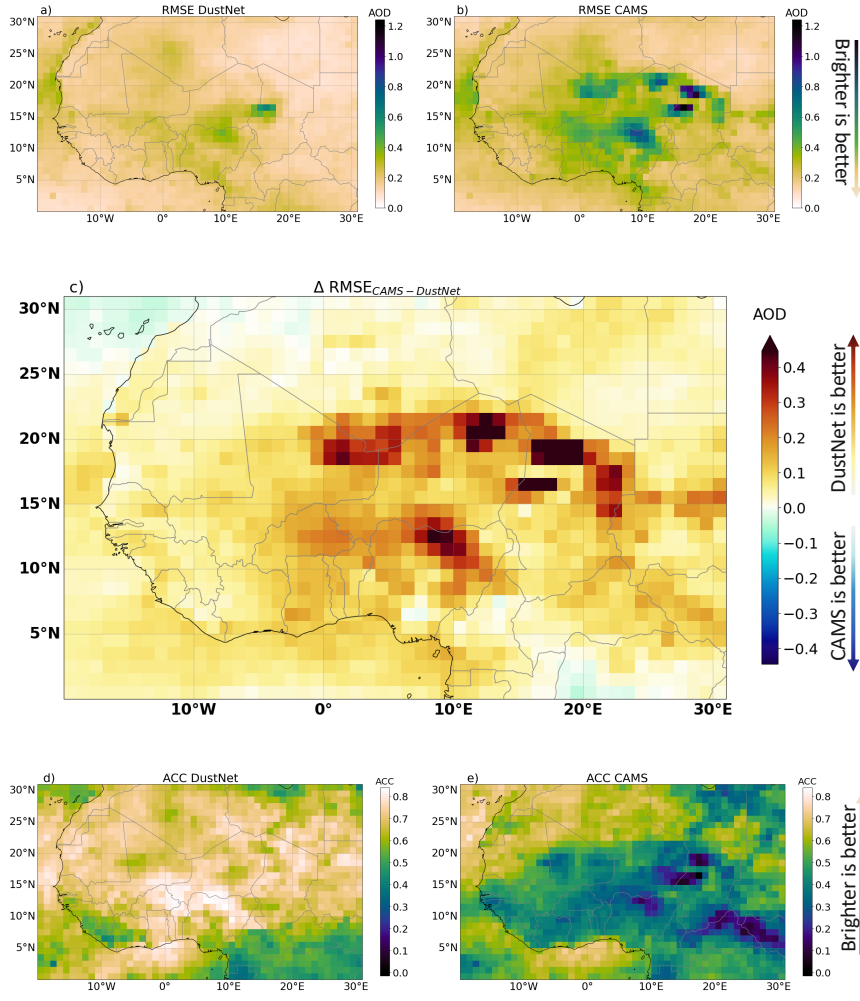


Fig. 2 Metrics indicating model performance. Results for 24-hour (1-step ahead) predictions of AOD values (mean across the daily prediction time 2020-2022, n=1095) compared with the ground truth data from MODIS. The RMSE for DustNet **a)** and CAMS **b)**, where the brighter the colour the smaller the error. Note, that the maximum error for DustNet is 0.62 AOD (medium green shades), while the maximum RMSE for CAMS reaches above 1.2 AOD (dark blue). In **c)** the difference in RMSE between CAMS and DustNet where all yellow to deep brown shades indicate the advantage of DustNet, while the blue shades indicate the advantage of CAMS. White grid cells indicate locations where both of the models performed equally when compared to the ground truth data. Note the lack of deeper blue shades and the dominance of yellow and brown grid cells where DustNet outperformed CAMS. **d)** and **e)** show the ACC for DustNet and CAMS respectively, where values above 0.6 (bright to white) indicate a valuable forecasting capability, while lower values (green to dark blue) indicate little to no predictive value. The ACC values in darkest blue indicate a misleading forecast.

564 indicating a better forecast value for a wider range of locations than CAMS (which
 565 had an ACC value above 60% at only 29.10% of the grid cells). Skillful detection of
 566 anomalies, combined with a high forecast value, indicates that the DustNet model
 567 could be a valuable addition to Earth System Models, where better representation
 568 of Saharan dust events leads to more realistic forecasts of precipitation and a better
 569 representation of the African monsoon (Anuforom, 2007; Düben, Peter and Modigliani,
 570 Umberto and Geer, Alan and Siemen, Stephan and Pappenberger, Florian and Bauer,
 571 Peter and Brown, Andy and Palkovic, Martin and Raoult, Baudouin and Wedi, Nils
 572 and Baousis, Vasileios, 2021; Balkanski et al., 2021).

573 Furthermore, we performed a comparative analysis of correlation coefficients
 574 between the forecasts and the ground-truth data. Figure 3 presents the daily correla-
 575 tion coefficients for two sets of comparisons: panel (a) displays the correlation between
 576 MODIS-derived AOD values and DustNet predictions, while panel (b) shows the correla-
 577 tion between MODIS and CAMS forecasts. Over the Saharan Desert, where mineral
 578 dust is the dominant contributor to AOD, DustNet exhibits a notably stronger correla-
 579 tion with MODIS, as indicated by the predominance of green shades. In contrast,
 580 CAMS demonstrates weaker correlations across the same region, evident in the pres-
 581 ence of white to brown shades, which aligns with previously identified dust generation
 582 zones highlighted in Figures 5 and 11. South of 8°N and east of 10°E, both models
 583 display weaker correlations. This outcome is expected, as DustNet’s training data did
 584 not incorporate information on black carbon or secondary organic aerosols, which can
 585 seasonally influence AOD in the equatorial regions of Central African forests (Jo et al.,
 586 2023).

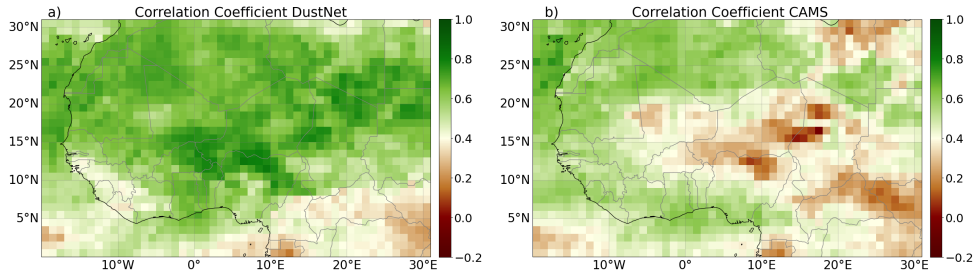


Fig. 3 Daily correlation coefficients between MODIS AOD observations and model predictions are shown for **a)** DustNet and **b)** CAMS. The maximum correlation for DustNet is 0.82, with a minimum of 0.16, while the maximum correlation for CAMS is 0.75, with a minimum of -0.04. Values with weaker correlations (≤ 0.4) are represented in white to brown shades, whereas stronger correlations (>0.4) are depicted in green. The predominance of green shades, particularly over the Sahara region, highlights the advantage of DustNet predictions over CAMS

587 When the data were spatially averaged over the study area for each test day,
 588 both DustNet and CAMS revealed high correlation with MODIS observations. How-
 589 ever, DustNet’s predictions exhibited stronger correlation with MODIS observations,
 590 achieving $r^2 = 0.91$ (see Fig. 4) in comparison to CAMS. This high correlation indi-
 591 cates that DustNet effectively captures the daily variability of AOD across the Sahara,

592 however with slight tendency to overestimate the high AOD values. In contrast, CAMS
 593 forecasts, while still highly correlated with MODIS ($r^2=0.71$), display a more frequent
 594 tendency to underestimate both low and high AOD values (Fig. 4b), and overestimate
 595 middle AOD values more frequently than DustNet. Both model results were highly
 596 significant, with $p - value \leq 0.0001$.

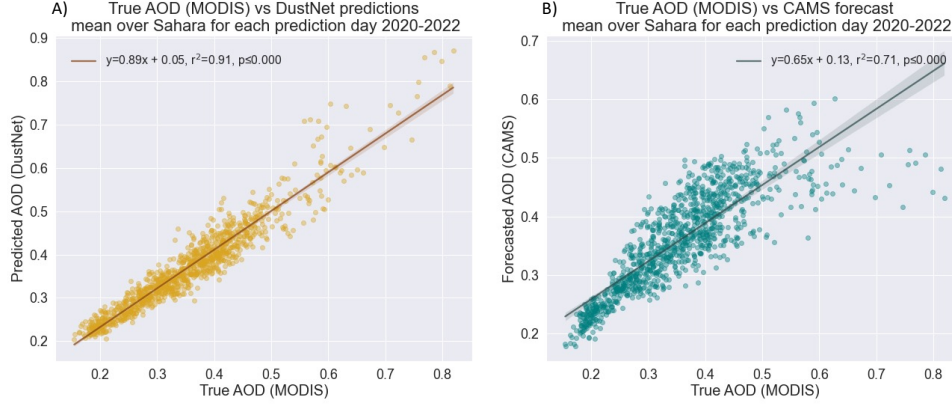


Fig. 4 Spatially averaged daily AOD (2020-2022, $n=1095$) regressed between model predictions and MODIS data. Linear regression with corresponding y equation, Pearson's r^2 and p values were calculated for daily spatial mean AOD over the Sahara for 2020 - 2022. Shown in **a**) AOD prediction results from DustNet correspond to MODIS data well with high $r^2 = 0.91$, and only a slight tendency to overestimate higher AOD. In **b**) the mean AOD forecasts from CAMS are shown to correspond with MODIS data well, $r^2 = 0.71$ though, with more frequent tendency to underestimate both low and high AOD values. Results from both predictions are highly significant with $p < 0.0001$.

597 4.2 Performance of seasonal-mean forecast

598 Saharan dust aerosols are highly seasonal in emission and transport direction (Anu-
 599 forum, 2007; Schwanghart and Schütt, 2008; Vandenbussche et al., 2020). Therefore,
 600 here we additionally compared the annual and seasonal means of DustNet predictions
 601 with MODIS and CAMS. Fig. 5a shows the annual mean AOD values of MODIS and
 602 the model predictions. DustNet is capable of producing more realistic predictions in
 603 comparison to MODIS than the mean annual forecasts from CAMS. This is confirmed
 604 by a highly significant correlation of the annual spatial mean AOD (DustNet: $r^2 =$
 605 0.91 ; CAMS: $r^2 = 0.71$, in [supplementary Fig. ??](#)). The DustNet model also captures
 606 the high AOD generated from the dustiest spot on Earth, the Bodélé Depression,
 607 more precisely than CAMS in both annual and all seasonal means (darkest colours
 608 in Fig. 5). Long-term comprehensive comparisons by Gliß et al. (2021) show that the
 609 forecasts produced by physics-based models tend to underestimate the AOD values
 610 compared to ground observations. While this underestimation of AOD is clear between
 611 5°N and 15°N , here we show that the CAMS forecast additionally tends to overesti-
 612 mate the AOD values around latitude 20°N over the Sahara during all the seasons of

the period 2020-2022 (Fig. 5, rightmost panel and [supplementary Fig. 16](#)). This could be attributed to the locations of most of the ground observation stations, concentrated along latitude 10°N (Gliß et al., 2021).

Moreover, we show that DustNet predictions capture the average seasonal displacement of AOD more skillfully than CAMS. The seasonal shift of Saharan dust by $\approx 10^\circ$ in latitude is consistent with past observations and studies (Prospero et al., 1981; Mbourou et al., 1997; Sunnu et al., 2008; Schepanski et al., 2017; Vandenbussche et al., 2020; Balkanski et al., 2021). Comparisons of AOD in Fig. 5b) and d) indicate that DustNet captures this shift more skillfully than CAMS. Associated with a seasonal change in wind direction and large plumes of transported dust, this phenomenon is locally well known as the Harmattan haze and is responsible for the high increase in air pollution (Anuforom, 2007; Schwanghart and Schütt, 2008; Sunnu et al., 2008). Previously noted mechanistic links between mineral dust and large-scale precipitation patterns, like the position of the Inter-tropical Convergence Zone (ITCZ) and the seasonal shift in the position of the West African monsoon, add to the importance of precise predictions of seasonal AOD displacement (Sunnu et al., 2008; Janicot et al., 2008; N'Datchoh et al., 2018; Balkanski et al., 2021). Additionally, seasonal means of the AOD, extracted from short forecast lead times of reanalysis models including CAMS, are used to validate other models including climate models (Zhao et al., 2022; O'Sullivan et al., 2020; Wu et al., 2020). Thus, achieving higher accuracy for the predictions of seasonal mean AOD forecasts with DustNet could improve the performance of current forecasting models.

The smoothness of predictions displayed by DustNet in comparison to CAMS is a characteristic of the regression algorithm used by deep learning models (explained in Bi et al., 2023).

4.3 Comparison of local predictions

We also test the ability of DustNet to provide accurate 24-hr predictions at four locations indicative of the main dust transport routes (see methods for details on locations). At all four locations, DustNet predictions align with satellite data (MODIS) better than forecasts produced by CAMS (Fig. 6, and Fig. 17 for correlations). This is especially evident at the Bodélé Depression, despite the site producing the highest prediction errors (see RMSE in Fig. 2a). The correlation between DustNet and MODIS is highly significant, with $r^2 = 0.62$, compared to CAMS which had $r^2 = 0.01$ (Fig. 6a) and Fig. 17a). DustNet also skillfully detects the daily and seasonal variability of the Bodélé Depression, demonstrating the ability of our model to skillfully capture dust generation at this location. Similarly, 24-hr DustNet predictions for Kano, the second most populous city in Nigeria, align better with MODIS ($r^2 = 0.74$) than forecasts from CAMS ($r^2 = 0.12$), whose predicted values stay close to the climatological mean (Fig. 6b and Fig. 17b).

During the first quarter ([Day of YearDOY](#) 0~ 90), the highest AOD values are present at the Bodélé Depression, Kano and the Gulf of Guinea (Fig. 6c). In Kano, the AOD values are just slightly lower than at the Bodélé and slightly lower in the Gulf of Guinea. Since both Kano and the Gulf of Guinea are positioned south-west from the Bodélé, their corresponding AOD values during quarter 1 indicate the Bodélé

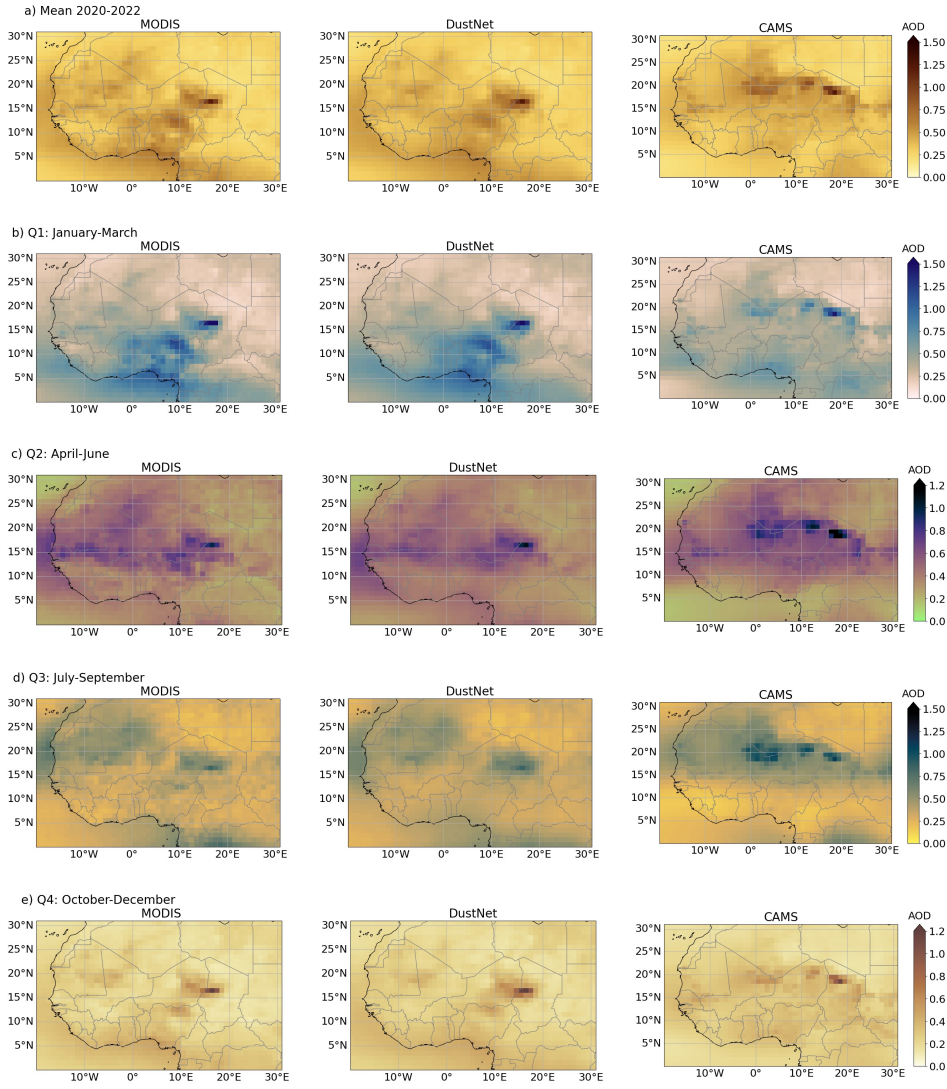


Fig. 5 Annual and quarterly means of daily AOD values for 2020-2022. All mean AOD values were calculated from daily 24-hour ahead predictions. The **left** column represents AOD values from MODIS observations, predictions from DustNet are in the **middle**, while forecasts from CAMS are in the **right** column. **Row a)** compares the 3-year annual mean AOD between the observations and models. In **row b)**, the 3-year mean of daily AOD for Q1: January - March is shown, noting the main generation site of the Bodélé Depression (dark blue) and the southwestward transport of mineral dust. In **row c)**, these same means are shown but for Q2: April - June. **Row d)** shows that both models, CAMS and DustNet, skillfully detected the northward shift of mean AOD transport during Q3: July - September. In **row e)**, the seasonal decrease in aerosol activity for Q4: October - December is skillfully captured by both models when compared to observations from MODIS. Note here the change in the colour bar range. Annual and quarterly mean AOD for 2020-2022. Mean AOD values calculated from 24-hr predictions. The left column represents AOD values from MODIS observations, while model predictions from DustNet are in the middle and from CAMS in the right column. **Row a)** compares the 3-year annual mean AOD between the observations and models, where DustNet skillfully captures the locations of the main dust events and the higher AOD around Nigeria and the Gulf of Guinea. In **row b)** the 3-year mean AOD for Q1: January - March, where the influence of the Harmattan wind has a visible effect on the mean AOD with a south-westward transport of mineral dust from the main generation site of the Bodélé Depression (dark blue). The effect of this transport is clearly picked up by our model. An increased AOD from biomass burning is also captured below 5°N. In **row c)** these same means are shown but for Q2: April - June where again the DustNet predictions skillfully capture the change in wind direction and westward aerosol transport in comparison to MODIS. **Row d)** shows that both models, CAMS and DustNet skillfully detected the northward shift of mean AOD transport during Q3: July-September. Here, CAMS forecasts tend to overestimate the AOD along the 20° latitude, but represent biomass burning-related AOD around equator more realistically than DustNet, whose smoother contours seem to overestimate the AOD below 10°N. In **row e)** the seasonal decrease in aerosol activity for Q4: October - December is skillfully captured by both models when compared to observations from MODIS. Here, DustNet captures the position of the Bodélé Depression more accurately than CAMS and shows the lack of aerosol generation from the eastern locations. Note here the change in the colour bar range.

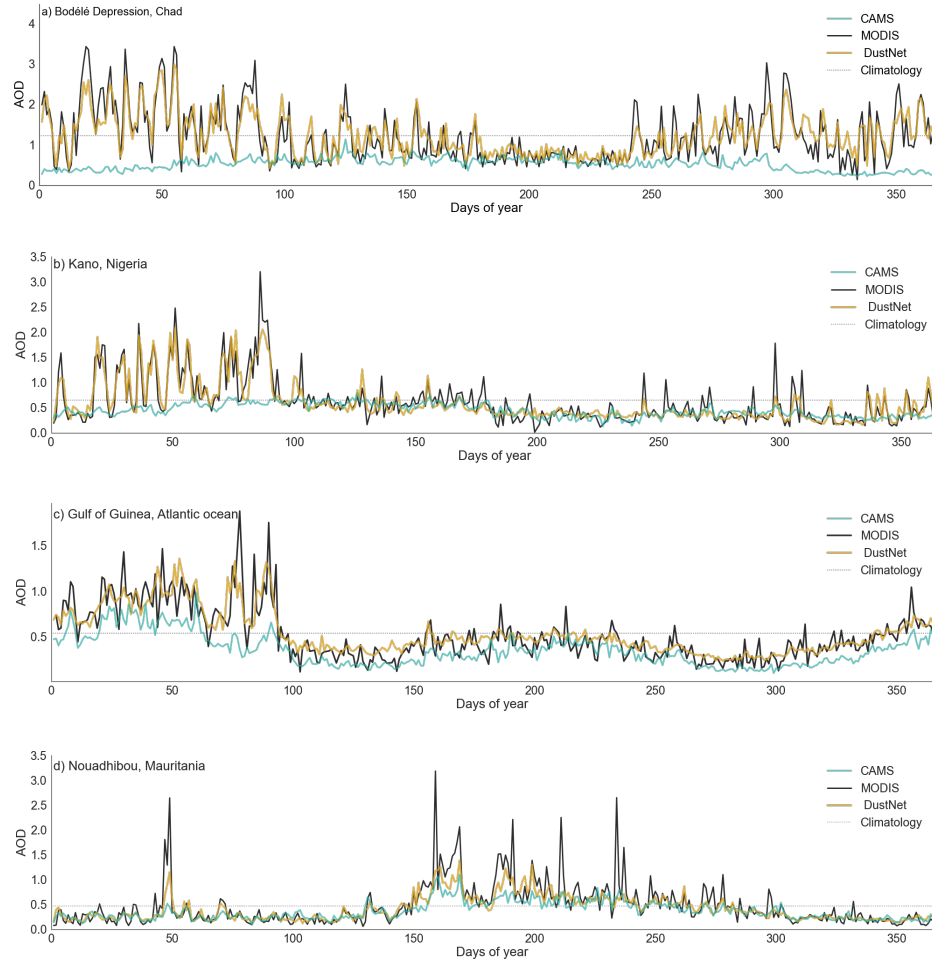


Fig. 6 Mean AOD predictions for each day of the year (2020-2022) at chosen locations. Shown are daily means (2020-2022) of AOD predictions from DustNet (golden line) and CAMS (light-sea-green line) as compared to MODIS (black line) and climatological mean (dotted line). At all four locations predictions from DustNet are closer to MODIS values than CAMS forecasts. An increase in AOD can be seen in the first 90 days of the year in **a)** the Bodélé Depression, with lower but still elevated values towards **b)** Kano and **c)** Gulf of Guinea. These elevated AOD values during quarter 1 are not observed in **d)** Nouadhibou, which is consistent with the south-western direction of the Harmattan wind. DustNet also predicts daily and seasonal AOD variability at each site more skillfully than CAMS, whose forecasts tend to stay closer to or below the climatological mean. Both models struggle to fully capture the highest AOD peaks recorded by MODIS at the westmost location - Nouadhibou, however the DustNet model replicates these peaks better than CAMS. ~~The background image, showing the position of the chosen locations (top), shows the December view of Blue Marble available from NASA <https://visibleearth.nasa.gov/collection/1484/blue-marble?page=4>.~~

Depression as a generation source (Schepanski et al., 2007; Jewell et al., 2021; Kok et al., 2021a). This also shows the ability of DustNet to capture generation and transport of AOD consistent with shifts in seasonal wind direction indicated in past studies (Schepanski et al., 2017; Schwanghart and Schütt, 2008; Anuforom, 2007; Sunnu et al., 2008).

During the third quarter (DOY 180~270), however, DustNet struggles to correctly capture the highest peaks in Kano and the Gulf of Guinea. The seasonal shift in meteorology and especially wind direction at these locations leads to an AOD composed of a mixture of aerosols, including sea-salt, and black carbon from biomass burning and industrial pollution (Anuforom, 2007; Mari et al., 2008; Knippertz et al., 2017). An area of future research could include information on vegetation and land cover during the training process, which would allow the model to distinguish between the ocean, Sahara Desert and central African forests. This would likely improve predictions for these regions and other aerosol species in general. The highest AOD values are also missed in Nouadhibou (Fig. 6d) during quarter 3 (DOY 180~260). However, here the seasonal increase in AOD points to a more localised origin, since dust generation at the Bodélé Depression is at its lowest with a daily AOD ≤ 1.0 . This finding is consistent with past analyses of boreal summertime dust generation, which point towards Western Sahara, Mauritania, Algeria and Mali as dust sources (Schepanski et al., 2007; Friese et al., 2017; Jewell et al., 2021; Kok et al., 2021a).

5 Discussion and future developments

The fast and skillful short-term predictions with DustNet present an opportunity for the forecasting community to incorporate a comprehensive aerosol scheme into future forecasts. The current coarse representation allows for quick testing and replication by professionals and enthusiasts alike. ~~DustNet also skillfully captures aspects of atmospheric processes such as dust generation, transport, and seasonal variations when compared to the satellite data. Despite DustNet not being explicitly trained to capture atmospheric processes such as dust generation, transport, or seasonal variations, these aspects are skillfully represented when compared to the satellite data.~~ Furthermore, skillful representation of atmospheric aerosols at specific locations opens a possibility for DustNet integration into more localised weather models.

The specific DustNet model architecture may be used for predicting other atmospheric particles, or indeed other environmental phenomena. However, this would require retraining the model using input features that represent the chosen particle or phenomenon. For example, to capture aerosols due to black carbon, features such as land cover types, vegetation, leaf area index, and forest fire locations should be considered. Similarly, when aiming to capture atmospheric aerosols due to sea-salt particles, features including wave height, energy flux into waves, peak wave period, and ocean surface stress should be taken into account. Moreover, DustNet model architecture may be used for predicting other spatio-temporal dynamics, such as phytoplankton concentrations from satellite-derived chlorophyll-a data, by substituting input variables with relevant meteorological and ocean state data.

While DustNet outperforms CAMS in short-term forecasts, it is not without limitations. Although the model is trained on 43 features, only one - terrain - represented the ground conditions. Thus, incorporating additional information could be beneficial in capturing more nuanced or indeed wider interactions. For example, the generation of dust depends not only on the atmospheric conditions, but also on the soil moisture, type and mineral composition from which atmospheric dust derives (Knippertz et al., 2017; Van Der Does et al., 2018). Soil type and mineralogy impact the dust interactions with other atmospheric particles, and wider Earth systems by delivering essential minerals to the oceans and rain-forests (Kok et al., 2023; Jickells et al., 2014; Koren et al., 2006). Information on ground vegetation and cover can also play a role in determining dust generation locations and transport, especially over forests and in urban areas.

Additionally, DustNet’s predictions at the northern and southeastern locations of the region boundaries are visibly weaker than at the centre (Fig. 2c and 2d). The predominant wind and transport directions of the atmospheric dust during this study are confirmed as west and southwest (Fig. 5, especially 5b and 5c), which indicates that the northern and southeastern areas may be governed by processes not included in the feature selection of this study. This is not surprising, since the Mediterranean Sea is directly to the north of our study region, while the Congolian rainforest covers grids directly to the south and southwest of the boundaries. These indicate the potential for more skillful forecasts with a broader study area, which, together with additional features, could capture more nuanced processes above the oceans and rainforests. ~~at the study region boundaries are visibly weaker than at the centre. This highlights the potential for a more skillful forecasts with a broader study area.~~

Likewise, the daily predictions of extreme AOD values at specific point locations (especially in Nouadhibu, Fig. 6d) can fall short of the values captured by the satellites. Together with the deterministic nature of the model, DustNet’s predictions lack the probability distribution with the length of the tail for the extreme values.

Addressing these limitations is crucial for future advancements. Rather than increasing the model’s training time or epochs, we propose expanding the training data with diverse geographical information. This approach would capture nuanced interactions of atmospheric dust with Earth’s systems. The inclusion of data from broader environmental disciplines, expanding study locations, and extending lead-time predictions are important next steps. Thus, a multidisciplinary approach can further enhance models capabilities and contribute to a range of specialised AI models with skillful predictions.

6 Conclusions

This study introduces a novel application of neural networks to improve the prediction of dust aerosols over the Saharan Desert, the world’s most significant source of atmospheric dust. Dust aerosols play a critical role in global climate systems, air quality, and ecosystems, yet traditional models often struggle with accuracy and speed due to the complex nature of dust dynamics and computational burden.

The research employs machine learning to bridge these gaps, offering a method that is both efficient and accurate. By training the DustNet model on satellite-based and reanalysis datasets, the research demonstrates significant improvements in capturing spatial and temporal variability of dust emissions. The results show that the neural network can produce skillful predictions while requiring fewer computational resources than conventional models.

Moreover, the framework is designed for accessibility and reproducibility, utilising open-source tools and emphasising transparency to facilitate broader adoption within the scientific community. This work not only advances the predictive capabilities for dust aerosols but also serves as a template for applying machine learning to other challenging atmospheric problems. Its potential implications span across atmospheric research, and practical applications like air quality management.

Acknowledgements We acknowledge NASA for producing, maintaining and releasing the MODIS AOD data which was used for training and comparison in this study. For these same reasons we acknowledge Copernicus Atmospheric Monitoring Service and ECMWF for their open release of CAMS AOD data. We would also like to acknowledge the reviewers of this paper: the Anonymous Referee #1 and Prof Narendra Ojha – Referee #2, whose comments contributed to better communication of our results and an overall improvement of this paper.

Funding This work was supported by the UKRI Centre for Doctoral Training in Environmental Intelligence, Engineering and Physical Sciences Research Council Grant Reference: EP/S022074/1.

Author contributions Conceptualization: T.E.N., S.S., B.I.S., A.T.A. Methodology: T.E.N., S.S. Investigation: T.E.N. Visualization: T.E.N. Supervision: S.S., B.I.S., A.T.A. Writing—original draft: T.E.N. Writing—review and editing: T.E.N., S.S., B.I.S., A.T.A.

Competing interests The authors declare no competing interests.

Code and data availability The full code for each model (DustNet, U-NET and Conv2D) with structured input data were deposited Nowak et al. (2024a) and are available from Zenodo at <https://zenodo.org/records/10722953>. The repository includes all results from the DustNet model (output data), and Python code to replicate all statistical analysis to reproduce each figure included in this article. Pre-processed ERA5 and AOD data are deposited as NumPy files in Zenodo together with Python imputation code at <https://zenodo.org/records/10593152> (Nowak et al., 2024b).

Reanalysis of atmospheric features were downloaded from the Copernicus Climate Data Store collection ‘ERA5 hourly data on pressure levels from 1940 to present’. Unprocessed datasets are available from Copernicus Climate Change Services (C3S) Climate Data Store (CDS) at <https://cds.climate.copernicus.eu/cdsapp/>. Pre-processed ERA5 data is also included in the aforementioned Zenodo repository.

The AOD at 550nm Level 3 daily data for combined Dark Target and Deep Blue algorithms were retrieved from Moderate Resolution Imaging Spectroradiometer (MODIS) on both Aqua and Terra spacecraft. Both datasets are available from NASA’s Atmosphere Archive & Distribution System (LAADS) Distributed Active

784 Archive Center (DAAC). Both MOD08.D3 and MYD08.D3 files can be retrieved
 785 from <https://ladsweb.modaps.eosdis.nasa.gov/search/>. Pre-processed AOD data is
 786 also included in the aforementioned Zenodo repository.

787 The forecast of AOD was downloaded from the Atmosphere Data Store of Coper-
 788 nicus Atmosphere Monitoring Service (CAMS). The total aerosol optical depth at
 789 550nm from the Global atmospheric composition forecast for midday run with a 24hr
 790 lead-time can be obtained from <https://ads.atmosphere.copernicus.eu/#!/home>.

791 7 References

792 Anuforom, A.C. 2007. Spatial distribution and temporal variability of Harmattan dust
 793 haze in sub-Sahel West Africa. *Atmospheric Environment* 41(39): 9079–9090 .

794 Ayzel, G., T. Scheffer, and M. Heistermann. 2020. RainNet v1. 0: a convolu-
 795 tional neural network for radar-based precipitation nowcasting. *Geoscientific Model*
 796 *Development* 13(6): 2631–2644 .

797 Balkanski, Y., R. Bonnet, O. Boucher, R. Checa-Garcia, and J. Servonnat. 2021, July.
 798 Better representation of dust can improve climate models with too weak an African
 799 monsoon. *Atmospheric Chemistry and Physics* 21(14): 11423–11435. <https://doi.org/10.5194/acp-21-11423-2021> .

801 Benedetti, A., J.J. Morcrette, O. Boucher, A. Dethof, R.J. Engelen, M. Fisher,
 802 H. Flentje, N. Huneeus, L. Jones, J.W. Kaiser, S. Kinne, A. Mangold, M. Razinger,
 803 A.J. Simmons, and M. Suttie. 2009, July. Aerosol analysis and forecast in the
 804 European Centre for Medium-Range Weather Forecasts Integrated Forecast Sys-
 805 tem: 2. Data assimilation. *Journal of Geophysical Research* 114(D13). <https://doi.org/10.1029/2008jd011115> .

807 Bi, K., L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian. 2023, July. Accurate medium-
 808 range global weather forecasting with 3D neural networks. *Nature* 619(7970):
 809 533–538. <https://doi.org/10.1038/s41586-023-06185-3> .

810 Bozzo, A., A. Benedetti, J. Flemming, Z. Kipling, and S. Rémy. 2020, March. An
 811 aerosol climatology for global models based on the tropospheric aerosol scheme in the
 812 Integrated Forecasting System of ECMWF. *Geoscientific Model Development* 13(3):
 813 1007–1034. <https://doi.org/10.5194/gmd-13-1007-2020> .

814 Bozzo, A., S. Remy, A. Benedetti, J. Flemming, P. Bechtold, M. Rodwell, and J.J.
 815 Morcrette 2017. Implementation of a CAMS-based aerosol climatology in the IFS.
 816 Technical report, European Centre for Medium-Range Weather Forecasts Reading,
 817 UK.

818 Carlson, T.N. and J.M. Prospero. 1972. The large-scale movement of Saharan air
 819 outbreaks over the northern equatorial Atlantic. *Journal of Applied Meteorology*
 820 *and Climatology* 11(2): 283–297 .

- 821 Chollet, F. et al. 2015. Keras.
- 822 Covert, I., S. Lundberg, and S.I. Lee. 2021. Explaining by removing: A unified frame-
823 work for model explanation. *Journal of Machine Learning Research* 22(209): 1–90
824 .
- 825 Daoud, N., M. Eltahan, and A. Elhennawi 2021. Aerosol optical depth forecast over
826 global dust belt based on LSTM, CNN-LSTM, CONV-LSTM and FFT algorithms.
827 In *IEEE EUROCON 2021-19th International Conference on Smart Technologies*,
828 pp. 186–191. IEEE.
- 829 Dumoulin, V. and F. Visin. 2016. A guide to convolution arithmetic for deep learning.
830 *arXiv preprint arXiv:1603.07285* .
- 831 Düben, Peter and Modigliani, Umberto and Geer, Alan and Siemen, Stephan and
832 Pappenberger, Florian and Bauer, Peter and Brown, Andy and Palkovic, Martin and
833 Raoult, Baudouin and Wedi, Nils and Baousis, Vasileios. 2021, 01/2021. Machine
834 learning at ECMWF: A roadmap for the next 10 years.
- 835 Evan, A.T., C. Flamant, S. Fiedler, and O. Doherty. 2014. An analysis of aeolian dust
836 in climate models. *Geophysical Research Letters* 41(16): 5996–6001 .
- 837 Friese, C.A., J.A. Van Hateren, C. Vogt, G. Fischer, and J.B.W. Stuut. 2017. Seasonal
838 provenance changes in present-day Saharan dust collected in and off Mauritania.
839 *Atmospheric Chemistry and Physics* 17(16): 10163–10193 .
- 840 Ginoux, P., J.M. Prospero, T.E. Gill, N.C. Hsu, and M. Zhao. 2012, August. Global-
841 scale attribution of anthropogenic and natural dust sources and their emission rates
842 based on MODIS Deep Blue aerosol products. *Reviews of Geophysics* 50(3). <https://doi.org/10.1029/2012rg000388> .
- 843
- 844 Gliß, J., A. Mortier, M. Schulz, E. Andrews, Y. Balkanski, S.E. Bauer, A.M.K. Bene-
845 dictow, H. Bian, R. Checa-Garcia, M. Chin, P. Ginoux, J.J. Griesfeller, A. Heckel,
846 Z. Kipling, A. Kirkevåg, H. Kokkola, P. Laj, P.L. Sager, M.T. Lund, C.L. Myhre,
847 H. Matsui, G. Myhre, D. Neubauer, T. van Noije, P. North, D.J.L. Olivié, S. Rémy,
848 L. Sogacheva, T. Takemura, K. Tsigaridis, and S.G. Tsyro. 2021, January. Aero-
849 Com phase III multi-model evaluation of the aerosol life cycle and optical properties
850 using ground- and space-based remote sensing as well as surface in situ observa-
851 tions. *Atmospheric Chemistry and Physics* 21(1): 87–128. [https://doi.org/10.5194/](https://doi.org/10.5194/acp-21-87-2021)
852 [acp-21-87-2021](https://doi.org/10.5194/acp-21-87-2021) .
- 853 Goroshin, R., J. Bruna, J. Tompson, D. Eigen, and Y. LeCun 2015. Unsuper-
854 vised learning of spatiotemporally coherent metrics. In *2015 IEEE International*
855 *Conference on Computer Vision (ICCV)*, pp. 4086–4093.
- 856 Hartman, L. and O. Hössjer. 2008. Fast kriging of large data sets with Gaussian
857 Markov random fields. *Computational Statistics & Data Analysis* 52(5): 2331–2349

- Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, et al. 2018. ERA5 hourly data on single levels from 1979 to present. *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*. <https://doi.org/10.24381/cds.bd0915c6> .
- Highwood, E.J. and C.L. Ryder 2014. *Radiative Effects of Dust*, pp. 267–286. Dordrecht: Springer Netherlands.
- Hinton, G.E., P. Dayan, B.J. Frey, and R.M. Neal. 1995. The "wake-sleep" algorithm for unsupervised neural networks. *Science* 268(5214): 1158–1161 .
- Hubanks, P., S. Platnick, M. King, and B. Ridgway. 2015. MODIS Atmosphere L3 gridded product algorithm theoretical basis document (atbd) & users guide. *ATBD reference number ATBD-MOD-30*, NASA 125: 585 .
- Janicot, S., C.D. Thorncroft, A. Ali, N. Asencio, G. Berry, O. Bock, B. Bourlès, G. Caniaux, F. Chauvin, A. Deme, et al. 2008. Large-scale overview of the summer monsoon over West Africa during the AMMA field experiment in 2006. In *Annales Geophysicae*, Volume 26, pp. 2569–2595. Copernicus Publications Göttingen, Germany.
- Jewell, A.M., N. Drake, A.J. Crocker, N.L. Bakker, T. Kunkelova, C.S. Bristow, M.J. Cooper, J.A. Milton, P.S. Breeze, and P.A. Wilson. 2021. Three North African dust source areas and their geochemical fingerprint. *Earth and Planetary Science Letters* 554: 116645 .
- Jickells, T., P. Boyd, and K.A. Hunter 2014. *Biogeochemical Impacts of Dust on the Global Carbon Cycle*, pp. 359–384. Dordrecht: Springer Netherlands.
- Jo, D.S., S. Tilmes, L.K. Emmons, S. Wang, and F. Vitt. 2023. A new simplified parameterization of secondary organic aerosol in the Community Earth System Model Version 2 (CESM2; CAM6. 3). *Geoscientific Model Development Discussions* 2023: 1–24 .
- Kang, S., N. Kim, and B.D. Lee 2019. Fine dust forecast based on recurrent neural networks. In *2019 21st International Conference on Advanced Communication Technology (ICACT)*, pp. 456–459. IEEE.
- Kaufman, Y., I. Koren, L. Remer, D. Tanré, P. Ginoux, and S. Fan. 2005. Dust transport and deposition observed from the Terra-Moderate Resolution Imaging Spectroradiometer (MODIS) spacecraft over the Atlantic Ocean. *Journal of Geophysical Research: Atmospheres* 110(D10) .
- Kingma, D.P. and J. Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .

- Knippertz, P., A.H. Fink, A. Deroubaix, E. Morris, F. Tocquer, M.J. Evans, C. Flamant, M. Gaetani, C. Lavaysse, C. Mari, et al. 2017. A meteorological and chemical overview of the DACCIWA field campaign in West Africa in June–July 2016. *Atmospheric Chemistry and Physics* 17(17): 10893–10918 .
- Knippertz, P. and J.B.W. Stuut. 2014. *A key player in the Earth system*. Dordrecht: Springer Netherlands.
- Kok, J.F., A.A. Adebisi, S. Albani, Y. Balkanski, R. Checa-Garcia, M. Chin, P.R. Colarco, D.S. Hamilton, Y. Huang, A. Ito, et al. 2021a. Contribution of the world’s main dust source regions to the global cycle of desert dust. *Atmospheric Chemistry and Physics* 21(10): 8169–8193 .
- Kok, J.F., A.A. Adebisi, S. Albani, Y. Balkanski, R. Checa-Garcia, M. Chin, P.R. Colarco, D.S. Hamilton, Y. Huang, A. Ito, et al. 2021b. Improved representation of the global dust cycle using observational constraints on dust properties and abundance. *Atmospheric Chemistry and Physics* 21(10): 8127–8167 .
- Kok, J.F., T. Storelvmo, V.A. Karydis, A.A. Adebisi, N.M. Mahowald, A.T. Evan, C. He, and D.M. Leung. 2023, January. Mineral dust aerosol impacts on global climate and climate change. *Nature Reviews Earth & Environment* 4(2): 71–86. <https://doi.org/10.1038/s43017-022-00379-5> .
- Koren, I., Y.J. Kaufman, R. Washington, M.C. Todd, Y. Rudich, J.V. Martins, and D. Rosenfeld. 2006. The Bodélé Depression: a single spot in the Sahara that provides most of the mineral dust to the Amazon forest. *Environmental Research Letters* 1(1): 014005 .
- Lam, R., A. Sanchez-Gonzalez, M. Willson, P. Wernsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merosé, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, and P. Battaglia. 2023, nov. Learning skillful medium-range global weather forecasting. *Science*: eadi2336 .
- LeCun, Y., Y. Bengio, and G. Hinton. 2015, May. Deep learning. *Nature* 521(7553): 436–444. <https://doi.org/10.1038/nature14539> .
- Mari, C., G. Cailley, L. Corre, M. Sauniois, J.L. Attié, V. Thouret, and A. Stohl. 2008. Tracing biomass burning plumes from the Southern Hemisphere during the AMMA 2006 wet season experiment. *Atmospheric Chemistry and Physics* 8(14): 3951–3961 .
- Mbouroou, G., J. Bertrand, and S. Nicholson. 1997. The diurnal and seasonal cycles of wind-borne dust over Africa north of the equator. *Journal of Applied Meteorology and Climatology* 36(7): 868–882 .
- Miller, R.L., P. Knippertz, C. Pérez García-Pando, J.P. Perlwitz, and I. Tegen 2014. *Impact of Dust Radiative Forcing upon Climate*, pp. 327–357. Dordrecht: Springer

931 Netherlands.

932 Molnar, C. 2022. *Interpretable Machine Learning* (2 ed.), Chapter Chapter 10: Neural
933 Network Interpretation. online: GitHub.

934 Morcrette, J.J., O. Boucher, L. Jones, D. Salmond, P. Bechtold, A. Beljaars,
935 A. Benedetti, A. Bonet, J. Kaiser, M. Razinger, et al. 2009. Aerosol analysis and fore-
936 cast in the European Centre for medium-range weather forecasts integrated forecast
937 system: Forward modeling. *Journal of Geophysical Research: Atmospheres* 114 (D6)
938 .

939 Morman, S.A. and G.S. Plumlee 2014. *Dust and Human Health*, pp. 385–409.
940 Dordrecht: Springer Netherlands.

941 Mulcahy, J.P., D.N. Walters, N. Bellouin, and S.F. Milton. 2014, May. Impacts
942 of increasing the aerosol complexity in the Met Office global numerical weather
943 prediction model. *Atmospheric Chemistry and Physics* 14 (9): 4749–4778. <https://doi.org/10.5194/acp-14-4749-2014> .

945 N’Datchoh, E., I. Diallo, A. Konaré, S. Silué, K. Ogunjobi, A. Diedhiou, and
946 M. Doumbia. 2018. Dust induced changes on the West African summer monsoon
947 features. *International Journal of Climatology* 38(1): 452–466 .

948 Nenes, A., B. Murray, and A. Bougiatioti 2014. *Mineral Dust and its Microphysical*
949 *Interactions with Clouds*, pp. 287–325. Dordrecht: Springer Netherlands.

950 Nowak, T.E., A.T. Augousti, B.I. Simmons, and S. Siebert. 2024a. Dustnet - structured
951 data and python code to reproduce the model, statistical analysis and figures.

952 Nowak, T.E., A.T. Augousti, B.I. Simmons, and S. Siebert. 2024b. Pre-processed daily
953 era5 and modis aod data (2003 - 2022) ready for use in ai/ml forecasting.

954 O’Sullivan, D., F. Marengo, C.L. Ryder, Y. Pradhan, Z. Kipling, B. Johnson,
955 A. Benedetti, M. Brooks, M. McGill, J. Yorks, et al. 2020. Models transport Saharan
956 dust too low in the atmosphere: a comparison of the MetUM and CAMS forecasts
957 with observations. *Atmospheric Chemistry and Physics* 20(21): 12955–12982 .

958 Parajuli, S.P., Q. Jin, and D. Francis. 2022, November. Editorial: Atmospheric dust:
959 How it affects climate, environment and life on Earth? *Frontiers in Environmental*
960 *Science* 10. <https://doi.org/10.3389/fenvs.2022.1058052> .

961 Pathak, J., S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mar-
962 dani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath,
963 and A. Anandkumar. 2022. FourCastNet: A Global Data-driven High-resolution
964 Weather Model using Adaptive Fourier Neural Operators.

- 965 Prospero, J., R. Glaccum, and R. Nees. 1981. Atmospheric transport of soil dust from
966 Africa to South America. *Nature* 289(5798): 570–572 .
- 967 Prospero, J.M. and T.N. Carlson. 1972. Vertical and areal distribution of Saharan
968 dust over the western equatorial North Atlantic Ocean. *Journal of Geophysical*
969 *Research* 77(27): 5255–5265 .
- 970 Ramachandran, P., B. Zoph, and Q.V. Le. 2017. Searching for activation functions.
971 *arXiv preprint arXiv:1710.05941* .
- 972 Rasamoelina, A.D., F. Adjailia, and P. Sinčák 2020. A review of activation function for
973 artificial neural network. In *2020 IEEE 18th World Symposium on Applied Machine*
974 *Intelligence and Informatics (SAMI)*, pp. 281–286. IEEE.
- 975 Rasp, S., P.D. Dueben, S. Scher, J.A. Weyn, S. Mouatadid, and N. Thuerey. 2020.
976 WeatherBench: a benchmark data set for data-driven weather forecasting. *Journal*
977 *of Advances in Modeling Earth Systems* 12(11): e2020MS002203 .
- 978 Ronneberger, O., P. Fischer, and T. Brox 2015. U-NET: Convolutional net-
979 works for biomedical image segmentation. In *Medical Image Computing and*
980 *Computer-Assisted Intervention–MICCAI 2015: 18th International Conference,*
981 *Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18, pp. 234–241.
982 Springer.
- 983 Rue, H. and L. Held. 2005. *Gaussian Markov random fields: theory and applications*.
984 New York: Chapman and Hall/CRC press.
- 985 Sarafian, R., D. Nissenbaum, S. Raveh-Rubin, V. Agrawal, and Y. Rudich. 2023,
986 March. Deep multi-task learning for early warnings of dust events implemented for
987 the Middle East. *npj Climate and Atmospheric Science* 6(1). [https://doi.org/10.](https://doi.org/10.1038/s41612-023-00348-9)
988 [1038/s41612-023-00348-9](https://doi.org/10.1038/s41612-023-00348-9) .
- 989 Schepanski, K., B. Heinold, and I. Tegen. 2017. Harmattan, Saharan heat low,
990 and West African monsoon circulation: modulations on the Saharan dust out-
991 flow towards the North Atlantic. *Atmospheric Chemistry and Physics* 17(17):
992 10223–10243 .
- 993 Schepanski, K., I. Tegen, B. Laurent, B. Heinold, and A. Macke. 2007. A new Saha-
994 ran dust source activation frequency map derived from MSG-SEVIRI IR-channels.
995 *Geophysical Research Letters* 34(18) .
- 996 Schwanghart, W. and B. Schütt. 2008. Meteorological causes of Harmattan dust in
997 West Africa. *Geomorphology* 95(3–4): 412–428 .
- 998 Shao, Y., K.H. Wyrwoll, A. Chappell, J. Huang, Z. Lin, G.H. McTainsh, M. Mikami,
999 T.Y. Tanaka, X. Wang, and S. Yoon. 2011. Dust cycle: An emerging core theme in
1000 Earth system science. *Aeolian Research* 2(4): 181–204 .

- 1001 Sunnu, A., G. Afeti, and F. Resch. 2008. A long-term experimental study of the
1002 Saharan dust presence in West Africa. *Atmospheric Research* 87(1): 13–26 .
- 1003 Todd, M.C., R. Washington, J.V. Martins, O. Dubovik, G. Lizcano, S. M’bainayel, and
1004 S. Engelstaedter. 2007. Mineral dust emission from the Bodélé Depression, northern
1005 Chad, during BoDEx 2005. *Journal of Geophysical Research: Atmospheres* 112(D6)
1006 .
- 1007 Van Der Does, M., P. Knippertz, P. Zschenderlein, R. Giles Harrison, and J.B.W.
1008 Stuut. 2018. The mysterious long-range transport of giant mineral dust particles.
1009 *Science advances* 4(12): eaau2768 .
- 1010 Vandenbussche, S., S. Callewaert, K. Schepanski, and M. De Mazière. 2020. North
1011 African mineral dust sources: new insights from a combined analysis based on 3D
1012 dust aerosol distributions, surface winds and ancillary soil parameters. *Atmospheric
1013 Chemistry and Physics* 20(23): 15127–15146 .
- 1014 Washington, R., C. Bouet, G. Cautenet, E. Mackenzie, I. Ashpole, S. Engelstaedter,
1015 G. Lizcano, G.M. Henderson, K. Schepanski, and I. Tegen. 2009. Dust as a tipping
1016 element: the Bodélé Depression, Chad. *Proceedings of the National Academy of
1017 Sciences* 106(49): 20564–20571 .
- 1018 Washington, R., M. Todd, N.J. Middleton, and A.S. Goudie. 2003. Dust-storm
1019 source areas determined by the total ozone monitoring spectrometer and surface
1020 observations. *Annals of the Association of American Geographers* 93(2): 297–313 .
- 1021 Wu, C., Z. Lin, and X. Liu. 2020. The global dust cycle and uncertainty in CMIP5
1022 (Coupled Model Intercomparison Project phase 5) models. *Atmospheric Chemistry
1023 and Physics* 20(17): 10401–10425 .
- 1024 Zeiler, M.D., D. Krishnan, G.W. Taylor, and R. Fergus 2010. Deconvolutional net-
1025 works. In *2010 IEEE Computer Society Conference on computer vision and pattern
1026 recognition*, pp. 2528–2535. IEEE.
- 1027 Zhao, A., C.L. Ryder, and L.J. Wilcox. 2022, February. How well do the CMIP6 models
1028 simulate dust aerosols? *Atmospheric Chemistry and Physics* 22(3): 2095–2119.
1029 <https://doi.org/10.5194/acp-22-2095-2022> .

1030 Appendix

1031 Changed the occurrence of Appendix Figures and numbers

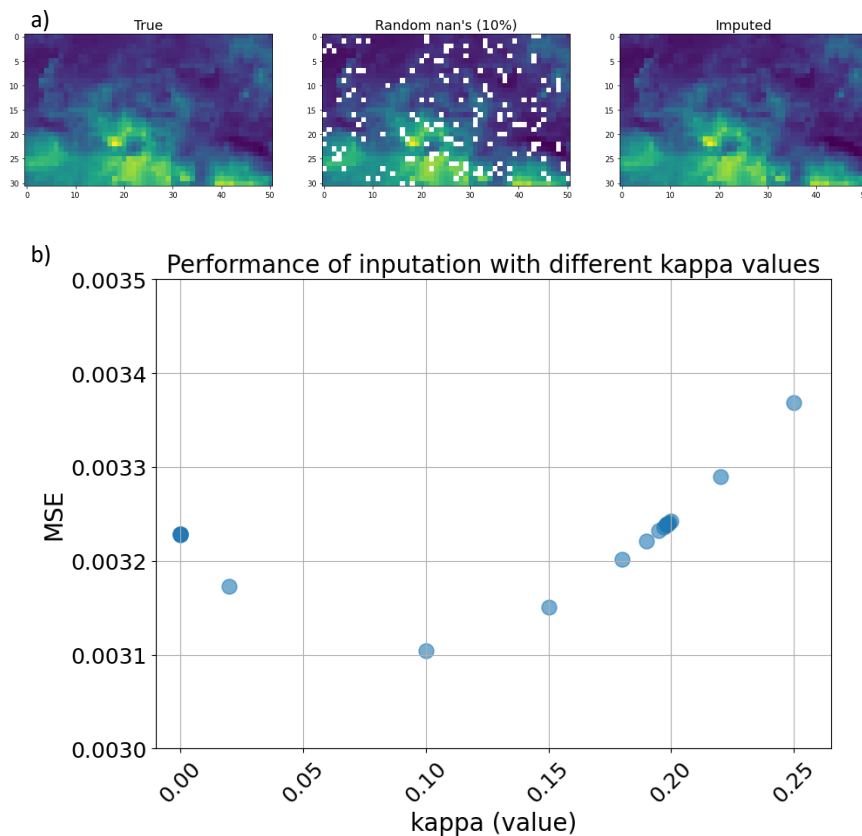


Fig. 7 Performance validation of imputation code. Three visual inspections of imputation in **a)** show the "true" values on the left, randomly assigned 10% of missing values in the middle, and imputed with Lattice Kriging method values in the right panel. In **b)** the results of MSEs between "true" and imputed values are displayed for different kappa values (Kriging hyperparameter) used during imputation. We tested 24 different kappa values ranging between 0.000001 and 0.1998) The MSE appears to be insensitive to changes in kappa values producing only marginal improvements in the overall MSE.

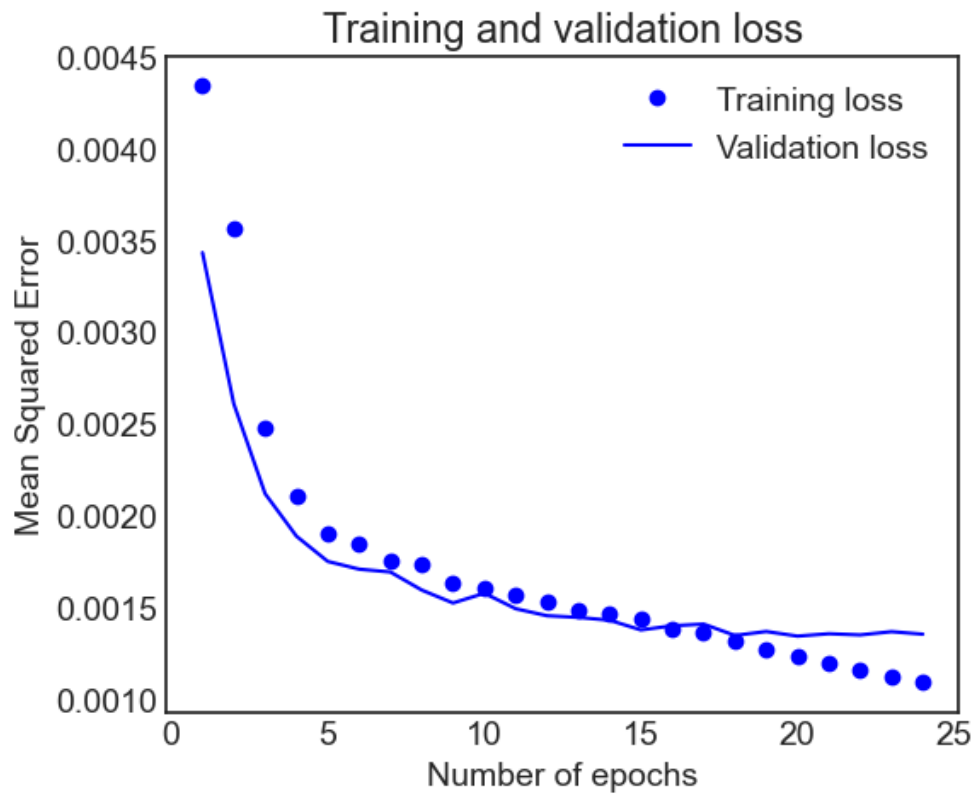


Fig. 8 Training and validation loss for the optimal model - DustNet. The model's architecture ensures Early Stopping is performed following the 4th iteration without any improvement in validation loss. Here, stopping occurred after 24 epochs and the model with the lowest ratio of training to validation loss was saved and used for predictions.

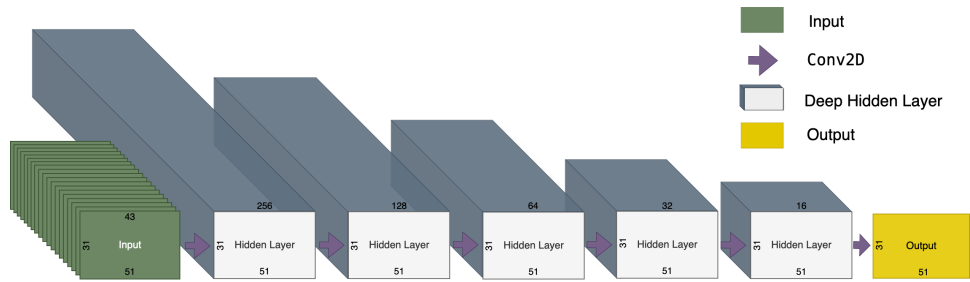


Fig. 9 Schematic representation of simple Conv2D model. **From left:** the Input Layer of shape (31,51,43) is represented in green. Following, are the 5 Hidden Layers of this same width and height as Inputs, but with different depths. The depths (number of hidden connections) were set in decreasing order to 256, 128, 64, 32 and 16. The last 2D Convolution with depth 1 created the output which shape was matching our target AOD.

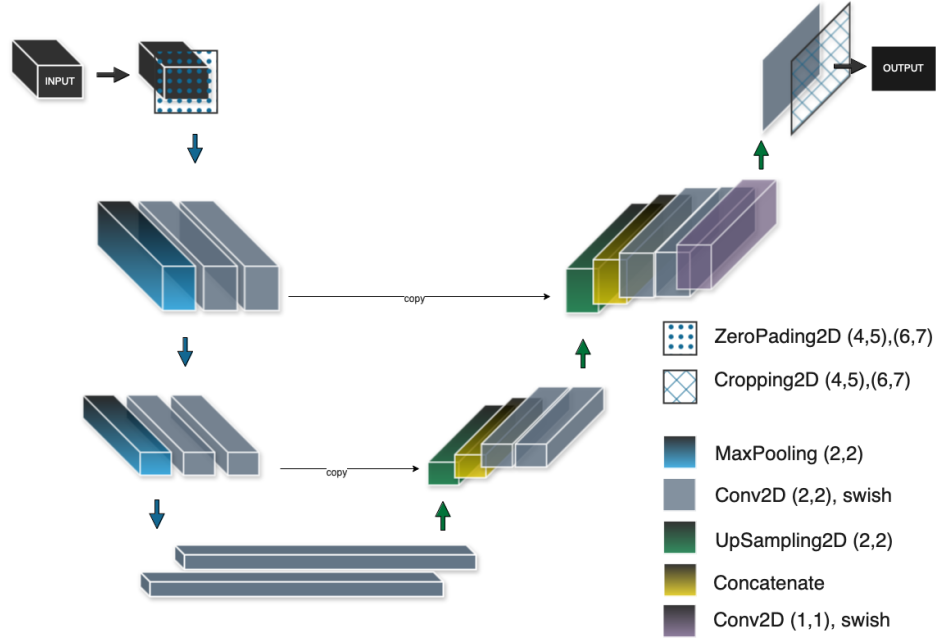


Fig. 10 Illustrative sketch of U-NET model architecture with individual blocks representing model layers. The input layer (31,51,43) is first padded with 2D zeros layer which increases the height and width of input shape (40,64,43). The encoding pathway (blue arrows down) includes 2 successive layers of Conv2D which increase the depth of the input size ($40 \times 64 \times 64$). Following MaxPooling layer decreases the first two dimensions, while Conv2D increases the third ($20 \times 32 \times 128$). After the second MaxPooling and double Conv2D the input is reshaped to (10×128). The decoding pathway (green arrows up) includes 2D Upsampling and Concatenation which now increases the width, height and depth to ($20 \times 32 \times 384$). Following 2 layers of Conv2D decrease the depth while UpSampling and Concatenation increases the shape to ($40 \times 64 \times 192$). The last two layers of Conv2D decreased the depth to ($40, 64, 64$) while its final layer brought the depth down to ($40 \times 64 \times 1$). Last layer, Cropping2D ensured the output matched the target size of ($31 \times 51 \times 1$).

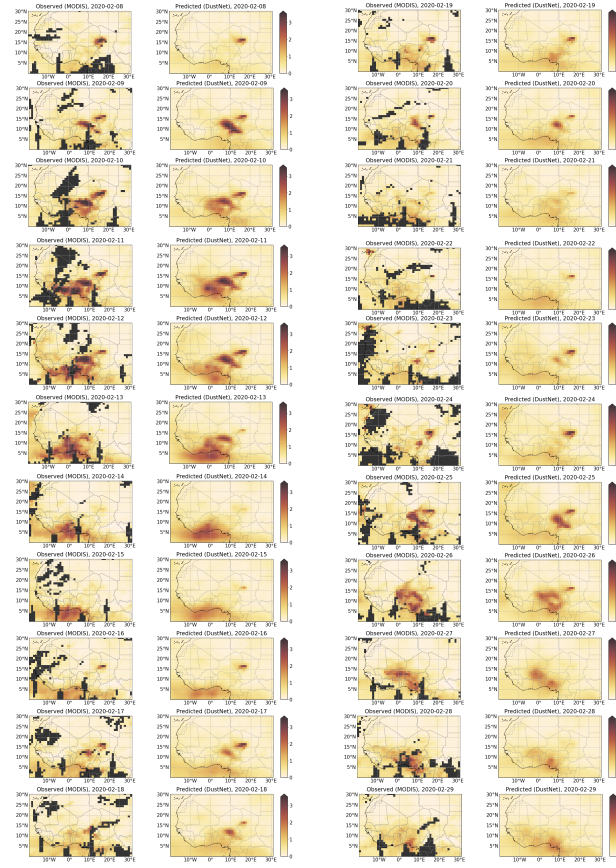


Fig. 11 Comparison of daily AOD values as observed by MODIS (mean of Aqua and Terra) (**left panel** in both columns) and corresponding DustNet predictions (**right panel**) for selected continuous 3 weeks (22 days), from 8th - 29th February 2020. The dark grey colour in the MODIS maps represents missing values. Despite an initial assumption of heavy reliance on the past 5 days of AOD during training, DustNet presents a skillful ability to predict the next time-step (24-hr) which visibly differs from the last 5 days. This is evident on 13th -14th Feb and 21st Feb, where the AOD values start to decrease despite an increasing past trend. Similarly, prediction of an increasing AOD from 22nd - 26th Feb was captured, despite the previous 5 days of decreasing AOD. The south-western direction of aerosol transport during boreal winter is also skillfully captured (10th - 14th Feb), as is the position of the Bodélé Depression during dust generation episodes (22nd - 26th Feb), but without overly relying on this location as a constant dust source (27th - 29th Feb).

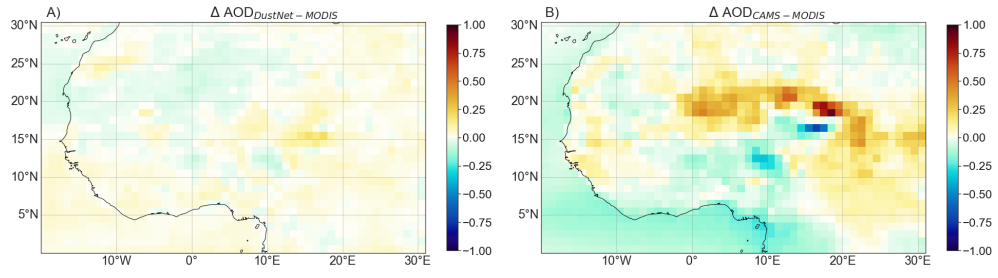


Fig. 12 Bias of daily predictions for **a)** DustNet and **b)** for CAMS with respect to MODIS data (n=1095). The lighter the shade, the lower the bias. Note that the maximum bias produced by DustNet is 0.21 while the maximum bias for CAMS is 0.93. The areas of AOD over-prediction in comparison to MODIS are shaded in yellow-brown, while under-predicted AOD is shaded in blue.

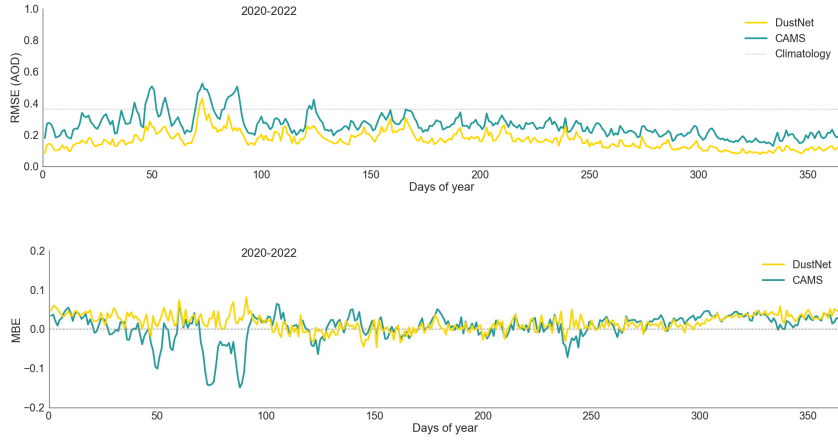


Fig. 13 The **top panel** presents the study area mean RMSE calculated from daily AOD values predicted by DustNet (yellow) and CAMS (cyan) are shown. At all time-steps the DustNet model predictions show smaller (better) errors than those produced by CAMS. The **lower panel** shows the temporal mean bias errors (MBE) from the DustNet predictions (yellow) and CAMS (cyan). Here, the DustNet bias fluctuates close to zero more often than the bias produced by CAMS.

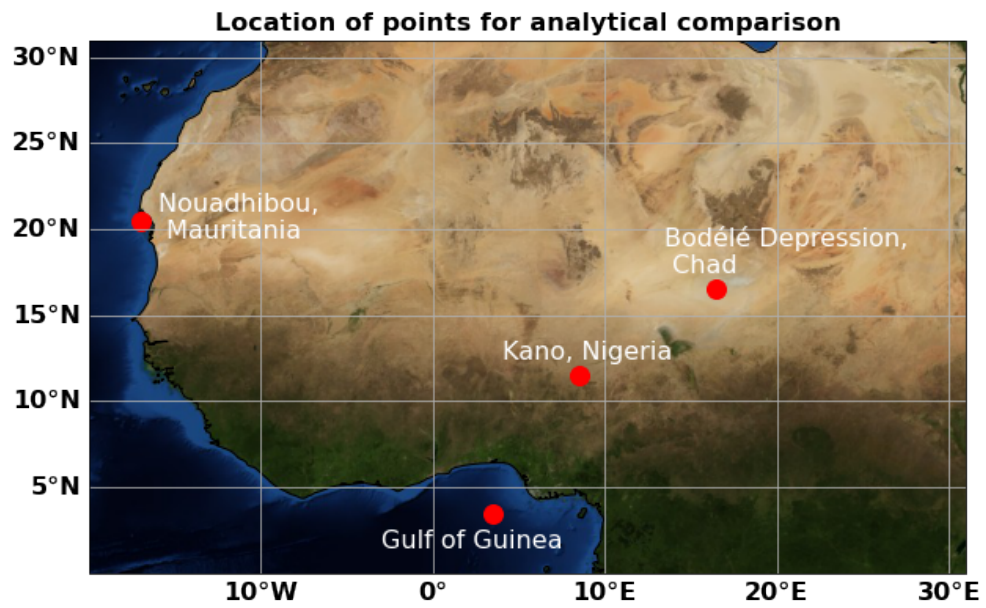


Fig. 14 Locations of selected grid points used to assess the model's predictive accuracy on a local scale ($1^\circ \times 1^\circ$ resolution grid size). The background image for the December view of Blue Marble is available from NASA <https://visibleearth.nasa.gov/collection/1484/blue-marble?page=4>

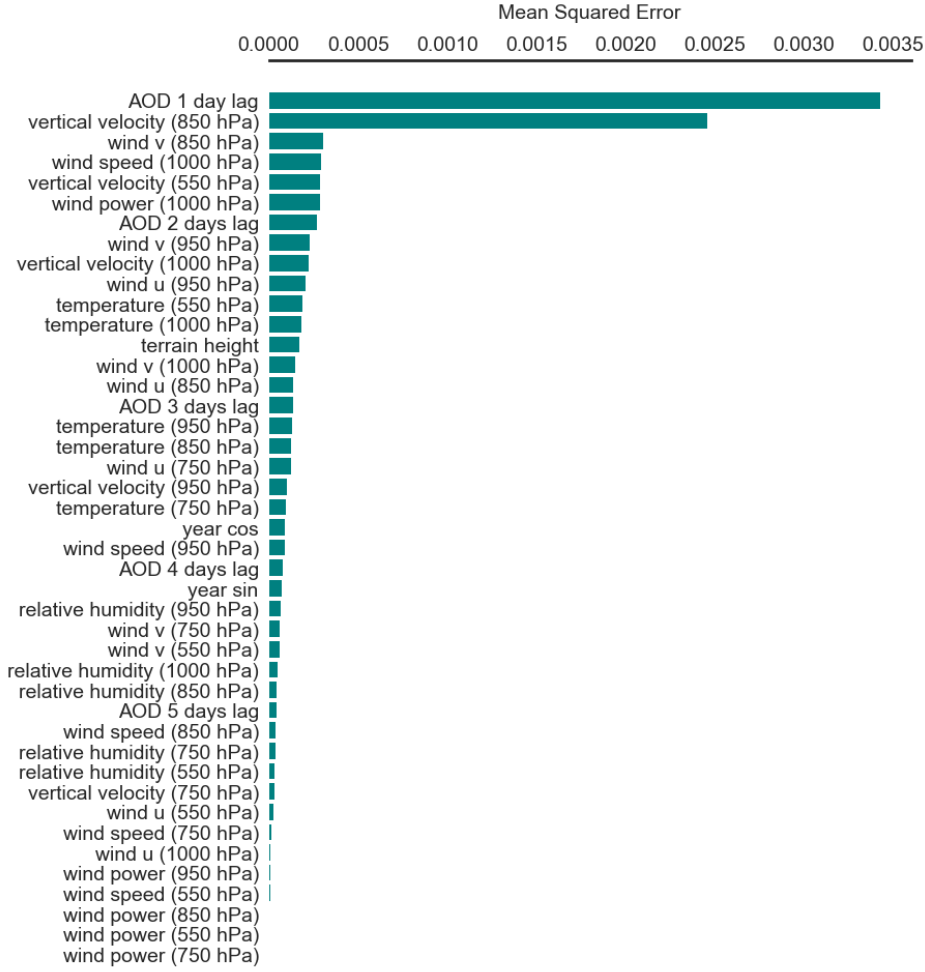


Fig. 15 Feature importance analysis for the DustNet model based on Mean Squared Error (MSE). Features with the highest MSE are deemed most critical for the pre-trained DustNet model's 24-hour (1-step ahead) predictions. The bar chart highlights that model performance exhibited the highest MSE when the 'AOD 1-day lag' feature values were removed (zeroed-out). Similarly, the removal of 'vertical velocity at 850 hPa' values during the prediction run also resulted in increased MSE. These findings indicate that the input channels associated with 'AOD 1-day lag' and 'vertical velocity at 850 hPa' are particularly influential, underscoring their importance for accurate model predictions.

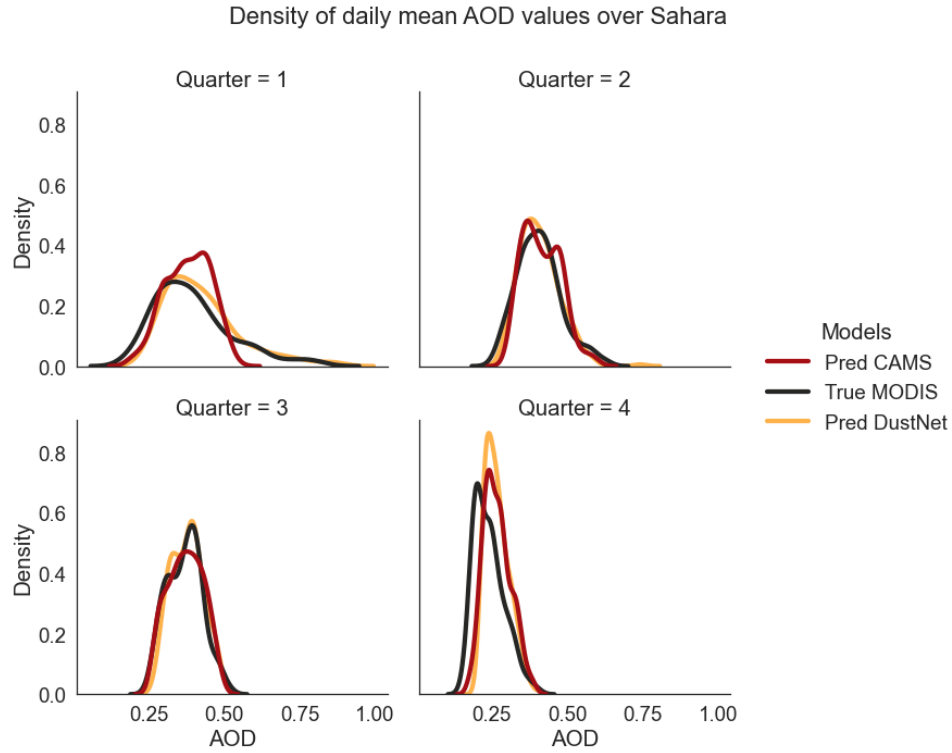


Fig. 16 Seasonal mean distribution of daily **predicted and true** AOD values. The data was averaged over the study region for the testing period of 2020-2022, and shows CAMS forecasts (red) DustNet predictions (yellow) and ground-true MODIS (black). The long tail, indicative of higher AOD values, is clearly missing in the CAMS distribution (red) for Quarter 1: January - March, while the lower AOD values are overestimated. The opposite is true for Quarter 4: October - December, where lower AOD values tend to be underestimated by both CAMS and DustNet in comparison to MODIS. Both models forecast fairly well during Quarter 2 and 3, although DustNet captures the bimodal distribution of AOD in Quarter 3 more skillfully than CAMS.

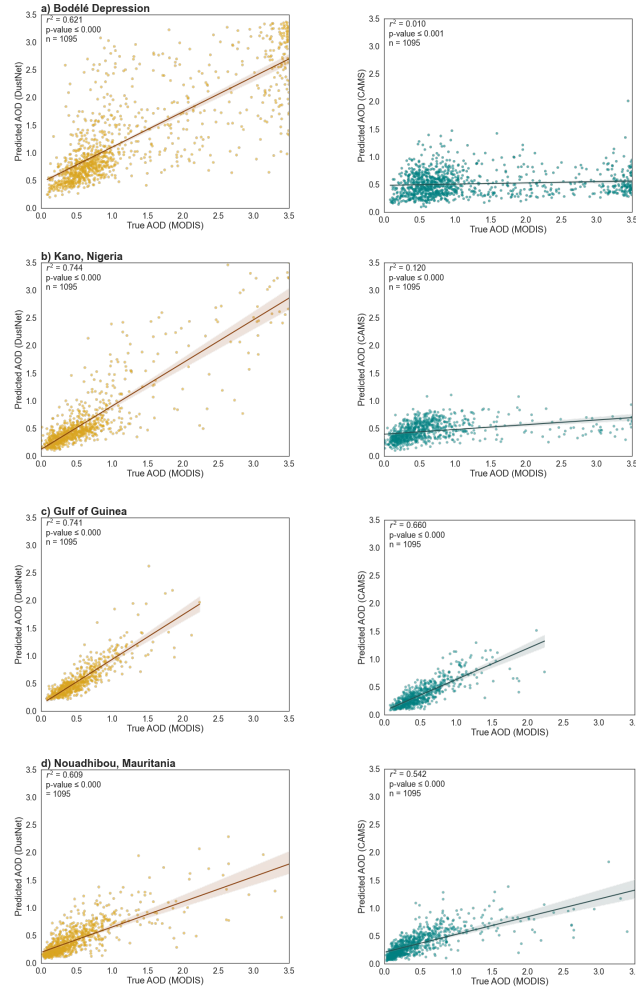


Fig. 17 Scatter plot relationship between predicted mean AOD values (2020-2022) and MODIS data. Results for DustNet (**left panel**) and forecasts from CAMS (**right panel**) at four chosen locations show better agreement of DustNet predictions with MODIS data at each location. In **a)** the Bodélé Depression, Chad - highest source of dust in the Sahara - DustNet is significantly better than CAMS; **b)** Kano, Nigeria - the second most populous province **c)** Gulf of Guinea - over the ocean; and **d)** Nouadhibou, Mauritania - coastal location.

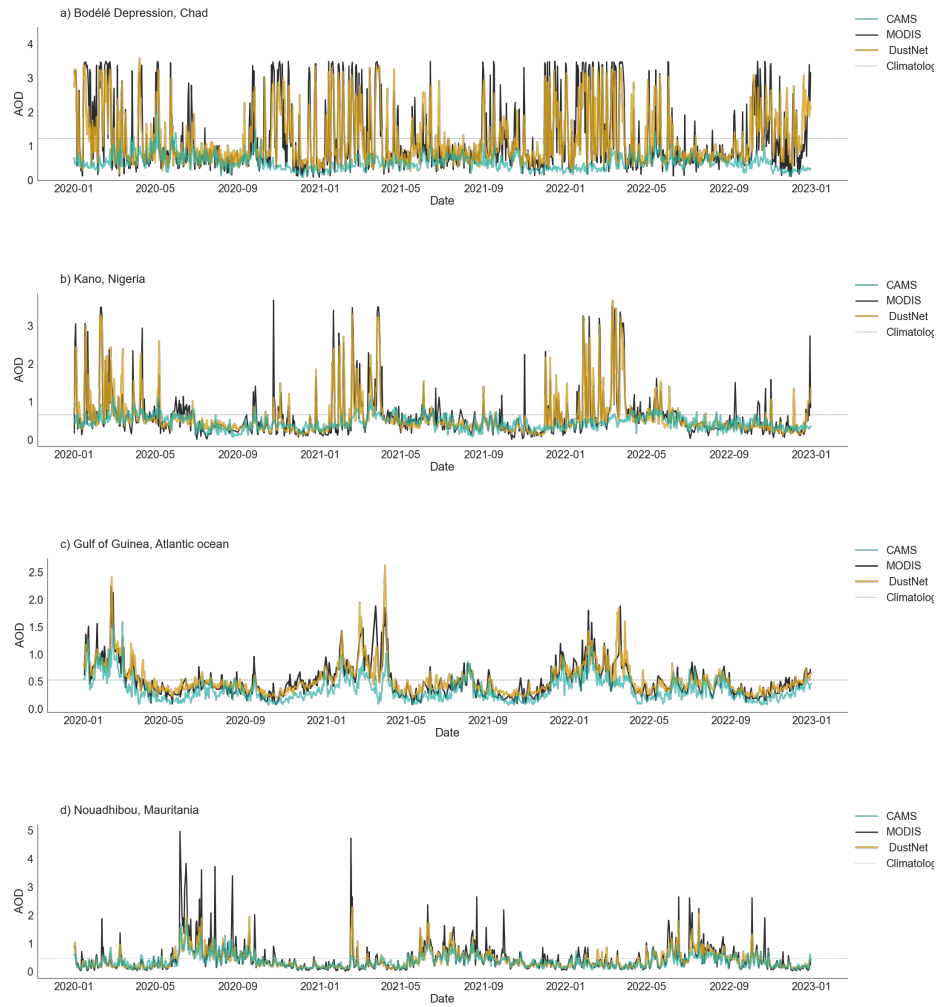


Fig. 18 Same as Fig. 4, but for daily data.