

Skillful neural network predictions of dust aerosols over the Saharan Desert–Saharan dust

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

¹*Mathematics and Statistics, University of Exeter, Exeter, EX4 QH, UK.

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

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

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

Contributing authors: augousti@kingston.ac.uk;

bsimmons.research@gmail.com; s.siebert@exeter.ac.uk;

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

Abstract

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

34 easily be utilized by researchers without access to supercomputers or GPUs.
35 ~~which could transform our understanding of dust's impacts on weather patterns.~~

36
37 **Copyrights statement:** For the purpose of open access, the author has applied
38 a ‘Creative Commons Attribution (CC BY) licence to any Author Accepted
39 Manuscript version arising from this submission.

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

42 1 Introduction

43 The Earth’s atmosphere is loaded with ≈ 26 million tonnes of mineral dust - an atmo-
44 spheric aerosol that represents the vast majority of mass burden in the atmosphere
45 [1, 2]. Each year, major sources emit $\approx 5,000$ million tonnes of dust globally [3] and,
46 although the majority of this material sinks at source, a substantial portion is trans-
47 ported over vast distances [4]. Once in the atmosphere, mineral dust interacts with the
48 Earth systems and impacts weather, climate, human health and infrastructure, from
49 fisheries to aviation [2, 5–11].

50 Despite its importance, representing atmospheric dust **aerosols** in weather and cli-
51 mate models is challenging [2, 12]. For example, physics-based Numerical Weather
52 Prediction (NWP) and climate models struggle to fully represent the dust cycle
53 with adequate emission, transport and generation [1, 13–15]. Instead, the Integrated
54 Forecasting System (IFS) of the European Center for Medium-Range Weather Fore-
55 casting (ECMWF) creates predictions that use aerosol optical depth (AOD) based on
56 monthly-mean climatological fields only [16]. A limitation in computational resources
57 is highlighted as one of the reasons for the lack of a dedicated aerosol scheme, since such
58 a development would significantly increase the computational burden of the system
59 [17]. The monthly mean AOD, developed by the Copernicus Atmosphere Monitoring
60 Service (CAMS), provides a reasonable trade-off in global weather forecasting. How-
61 ever, a more accurate representation of the AOD would have significant benefits, such
62 as large improvements in the representation of the summer monsoon circulation or
63 precipitation patterns in the Sahel region [18, 19].

64 Recent developments in the field of AI present a significant opportunity to over-
65 come the computational burden of a dedicated physics-based aerosol scheme. Models
66 such as GraphCast, Pangu-Weather, and FourCastNet can now skillfully predict the
67 main ERA5 variables and in many cases outperform the state-of-the-art NWP mod-
68 els [20–22]. To date, attempts to forecast atmospheric aerosols with neural network
69 architectures have shown varying levels of success. “Satisfying” results were reported
70 [23, 24] when applying a long-short-term memory (LSTM) architecture to local AOD
71 forecasts. The application of a U-NET architecture revealed a skillful detection of
72 classified ‘dust events’ at 67% precision rate [25]. A lack of comparisons to the cur-
73 rent physics-based forecasts, or inclusion of standardised skill metrics, makes direct
74 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 Results

2.1 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 [26–29]. We also custom-designed a 2D CNN with transposed layers [30]. After comparing the performance of these three models (see 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 forecasts were produced in 2.1 seconds.

To evaluate the resulting 24-hour 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 [18]. 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. 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




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.

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.

2.2 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 [31] — 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 [31–33] and (ii) it was identified as the main source of minerals delivered seasonally to the Amazon basin [33, 34]. 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 [1]. 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. 2c). In Fig. 2c), 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 [35–37]) 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 DustNet (Fig. 2c). Both of these locations are adjacent to the boundaries (SE and NW

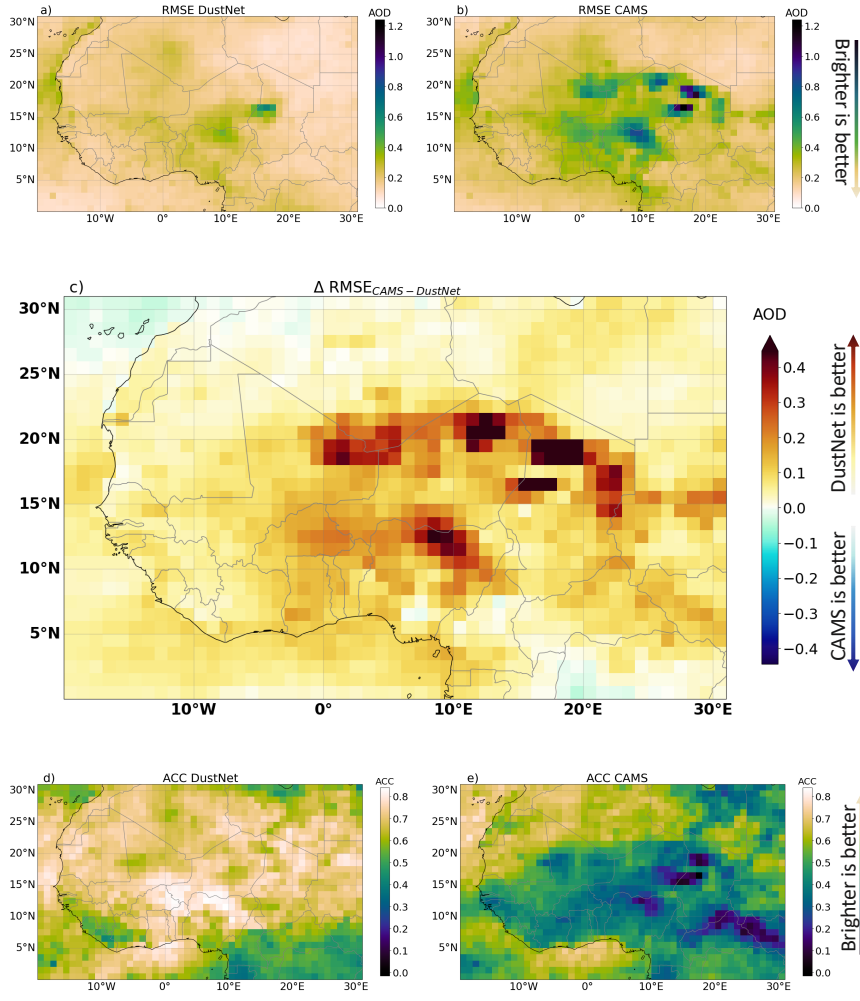


Fig. 2 Metrics indicating model performance. Results for 24-hour (1-step ahead) predictions of AOD values (mean across the daily prediction time 2020-2022, $n=1095$) compared with the ground truth data from MODIS. The RMSE for DustNet (a) and CAMS (b), where the brighter the colour the smaller the error. 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.

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 [20, 21] (see section Methods for details). Here, DustNet also displays more skillful results than CAMS with a better (higher) ACC at 92.283% of grid cells shown in 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), indicating a better forecast value for a wider range of locations than CAMS (which had an ACC value above 60% at only 29.10% of the grid cells). Skillful detection of anomalies, combined with a high forecast value, indicates that the DustNet model could be a valuable addition to Earth System Models, where better representation of Saharan dust events leads to more realistic forecasts of precipitation and a better representation of the African monsoon [19, 35, 38].

2.3 Performance of seasonal-mean forecast

Saharan dust aerosols are highly seasonal in emission and transport direction [35, 37, 39]. Therefore, here we additionally compared the annual and seasonal means of DustNet predictions with MODIS and CAMS. Fig. 3a shows the annual mean AOD values of MODIS and the model predictions. DustNet is capable of producing more realistic predictions in comparison to MODIS than the mean annual forecasts from CAMS. This is confirmed by a highly significant correlation of the annual spatial mean AOD (DustNet: $r^2 = 0.91$; CAMS: $r^2 = 0.71$, in Fig. 5). The DustNet model also captures the high AOD generated from the dustiest spot on Earth, the Bodélé Depression, more precisely than CAMS in both annual and all seasonal means (darkest colours in Fig. 3). Long-term comprehensive comparisons [1] show that the forecasts produced by physics-based models tend to underestimate the AOD values compared to ground observations. While this underestimation of AOD is clear between 5°N and 15°N, here we show that the CAMS forecast additionally tends to overestimate the AOD values around latitude 20°N over the Sahara during all the seasons of the period 2020-2022 (Fig. 3, rightmost panel and Fig. 6). This could be attributed to the locations of most of the ground observation stations, concentrated along latitude 10°N [1].

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 [19, 36, 39–42]. Comparisons of AOD in Fig. 3b) 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 [35–37]. Previously noted mechanistic links between mineral dust and large-scale precipitation patterns, like the

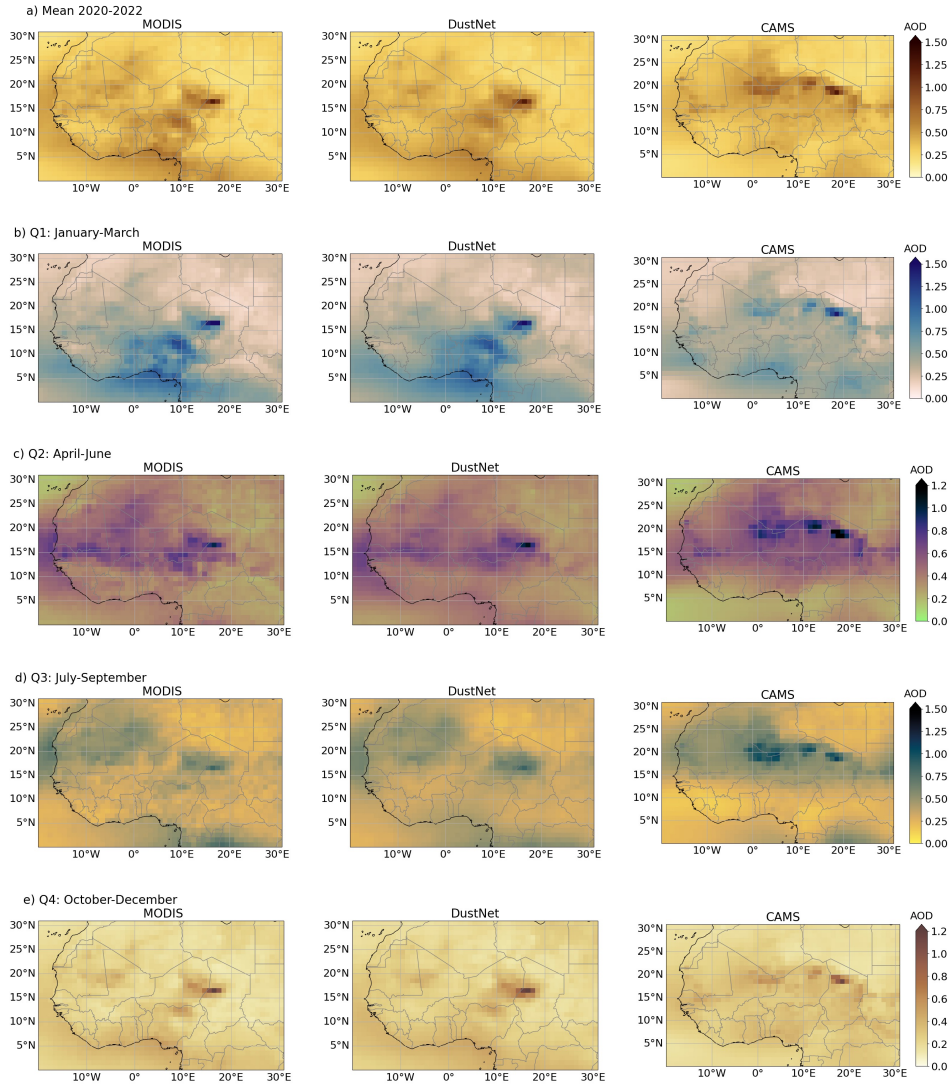


Fig. 3 Annual and quarterly means of daily AOD values for 2020-2022. All mean AOD values were calculated from daily 24-hour ahead predictions. The **left** column represents AOD values from MODIS observations, predictions from DustNet are in the **middle**, while forecasts from CAMS are in the **right** 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.

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 [19, 36, 43, 44]. 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 [15, 45, 46]. 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 [21]).

2.4 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. 4, and Fig. 8 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. 4a) and Fig. 8a). 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. 4b and Fig. 8b).

During the first quarter (Day of Year DOY 0~90), the highest AOD values are present at the Bodélé Depression, Kano and the Gulf of Guinea (Fig. 4c). 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 Depression as a generation source [14, 33, 47]. 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 [35–37, 42].

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 [35, 48, 49]. 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. 4d) 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

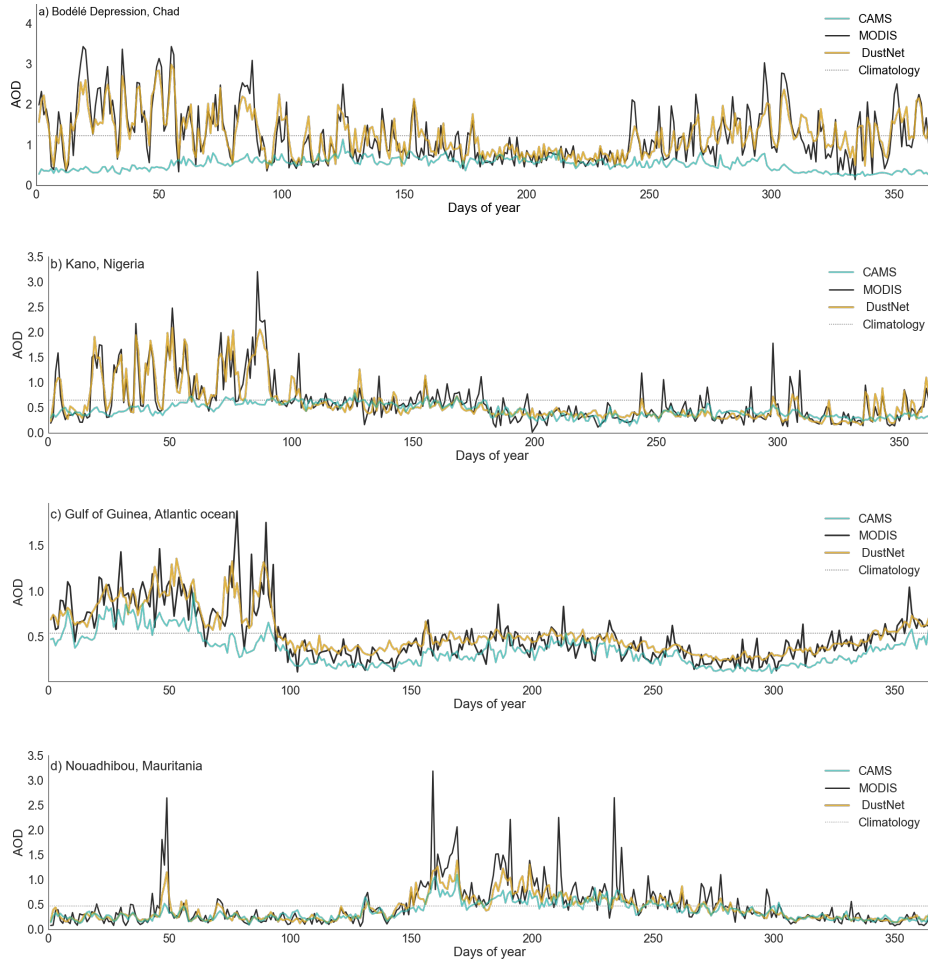


Fig. 4 Mean AOD predictions for each day of the year (2020-2022) at chosen locations. Shown are daily means (2020-2022) of AOD predictions from DustNet (golden line) and CAMS (light-sea-green line) as compared to MODIS (black line) and climatological mean (dotted line). At all four locations predictions 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.](https://visibleearth.nasa.gov/collection/1484/blue-marble?page=4)

dust generation, which point towards Western Sahara, Mauritania, Algeria and Mali as dust sources [14, 33, 47, 50].

3 Discussion

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 adapted for predicting other atmospheric particles. However, this would require retraining the model using input features that represent the chosen particle. 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.

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, soils mineral composition from which atmospheric dust derives [4, 49]. 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 [2, 10, 34]. 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. 3, especially 3b and 3c), 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. 4d) 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.

4 Methods

4.1 Study area

To effectively forecast dust aerosols, our study area encompasses the global principal dust generation source - the Sahara Desert - which is responsible for $\sim 55\%$ of the 1,536 million tonnes of total global dust emitted annually [51]. The region (Fig. 4 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 [31]. 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 features, which is essential for accurate forecasting. Finally, by minimizing the area to the Saharan Desert and consequently reducing the amount of chosen training features, we were able to avoid adding different ocean and terrain processes, leading to reduced model complexity without compromising performance.

316 4.2 Datasets

317 4.2.1 AOD data

318 We retrieved the AOD data from the Moderate Resolution Imaging Spectroradiometer
319 (MODIS) instrument located on board both Aqua and Terra spacecraft. With daily
320 temporal resolution over a period of 20 years starting from 1st January 2003 to 31st
321 December 2022, the AOD data yields 2×7305 files. We used the quality-controlled
322 level-3 data for AOD at 550nm. Choosing the combined mean of Dark Target and Deep
323 Blue algorithms provided a full coverage above bright and dark surfaces at a horizontal
324 resolution of $1^\circ \times 1^\circ$ [52]. This choice provided a good spatiotemporal coverage of AOD
325 data above both land and ocean surfaces.

326 4.2.2 ERA5 data

327 Meteorological data comes from the fifth generation of European Centre for Medium-
328 Range Weather Forecast (ECMWF) atmospheric reanalysis project (ERA5) and
329 consists of 5 parameters: wind u component, wind v component, vertical velocity,
330 temperature and relative humidity. Each parameter was retrieved at 5 pressure levels
331 550hPa, 750hPa, 850hPa, 950hPa and 1000hPa. This choice provided us with 35 dis-
332 tinctive features representing atmospheric conditions from ground level to $\approx 5\text{km}$ in
333 vertical height. The ERA5 data is available on an hourly basis, but here we only chose
334 the data representing conditions for midday (12:00 UTC). This allows us to represent
335 the mid-point in atmospheric conditions between the Terra and Aqua satellite over-
336 passes above the equator (10:30am and 1:30pm respectively). To further match the
337 meteorological data with AOD, we chose a daily temporal resolution between 2003-
338 2022. The horizontal resolution of ERA5 data is $0.25^\circ \times 0.25^\circ$. To match this with the
339 AOD resolution of $1^\circ \times 1^\circ$, the data was regridded (see section [Data pre-processing](#) for
340 details).

341 4.2.3 Timestamps

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

348 4.2.4 CAMS forecast

349 We obtained daily ‘Total aerosol optical depth at 550nm’ forecast data from ‘CAMS
350 global atmospheric composition forecasts’. CAMS forms a part of the ECMWF Inte-
351 grated Forecasting System (IFS), and is a sophisticated numerical weather forecasting
352 model (NWP) [16]. During the AOD data assimilation process, CAMS utilises data
353 from MODIS, among other satellites, together with data from ground-based observa-
354 tion stations. The model then uses physics and chemistry principles to forecast hourly
355 AOD values on a single level for up to 5 days (120hr) ahead [53, 54]. 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).

4.3 Data pre-processing

4.3.1 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 [55, 56] 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.9). See [Code and data availability](#) section for links containing the full Python code for imputation.

4.3.2 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.

4.3.3 ERA5 regriding

The ERA5 data [57] is supplied with a ~~horizontal~~vertical resolution of $0.25^\circ \times 0.25^\circ$ and thus needed regriding 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.

4.3.4 Combining and normalising

We combined the meteorological data with AOD data into a single 4D NumPy array of shape 7300, 51, 31, 41, 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

394 normalised all features using min-max normalisation:

$$x_{ijt,norm} = \frac{(x_{ijt} - x_{min})}{(x_{max} - x_{min})} \quad (1)$$

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

397 4.3.5 Seasonal features

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

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

400 and similarly using the cosine:

$$x_{ijt}^{(43)} = \cos\left(2\pi \frac{t}{365.2425}\right), \quad (3)$$

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

404 4.3.6 Training, validation, test split

405 We split the data along the time dimension into 70%, 15% and 15% for training,
406 validation and test sets respectively. Splitting data with consecutive time steps yielded
407 better results than a random split. The use of consecutive time steps ensures that
408 each subset is composed of data points that are temporally distinct. This reduces the
409 risk of autocorrelation and improves the model's ability to generalize to new, unseen
410 data [58]. Therefore, the training set covered 5,110 consecutive days from 6th January
411 2003 until 1st January 2017 (inclusive of both days). The validation set took 1,095
412 consecutive days from 2nd January 2017 to 1st January 2020. Finally, we set aside a
413 test set, with 1,095 days of data from 2nd January 2020 to 31st December 2022. We
414 made sure that the model never had access to the test set during the training and
415 validation processes and only after these were complete did we introduce the test data
416 and run our model to obtain predictions.

417 4.4 Designing CNN models

418 To find the best forecast of the daily AOD, we designed three CNN models based on
419 [26–28]. We used the end-to-end open source machine learning platform TensorFlow
420 2, together with the Keras high-level API [59]. Each model uses a different architec-
421 ture based on two-dimensional (2D) convolutions (hereafter Conv2D). In general, the
422 Conv2D neural network architecture enables regression problems in image analysis

to be addressed and is particularly effective at capturing spatial patterns in two-dimensional images. The efficiency of Tensorflow allows for training and inference to be run on traditional desktops or laptops rather than requiring HPC's. All models described hereafter were run using Python version 3.10.10 on a MacBook Pro with an Apple M1 Pro and 32GB RAM. The models did not utilize any GPUs and can thus be replicated by users without access to a supercomputer.

We have chosen 'Adam' optimizer and the mean-squared-error (MSE) as a loss function. These options We also performed a series of diagnostic tests in order to choose the best optimizer and loss function. Tested loss functions included CoSine, Huber, LogCosh, Mean Absolute Percentage Error, Mean Absolute Error, Mean Squared Logarithmic Error and Mean Squared Error. We assessed the performance based on the lowest mean squared error, the speed of the overall training time and the time taken per step. The mean-squared-error (MSE) loss function together with the Adam optimizer offered optimal results in terms of training times and were and was used 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 [60].

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 a it may be possible to achieve a better performing design.

Table 1 Test results of choosing different kernel sizes for 2 models: Conv2D.T 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).

Kernel size	Conv2D			U-NET		
	(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. 12). 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.

4.4.1 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 [27]. 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. 15). 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., [61] 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 [62]. The Swish activation function proved to yield the best performance and thus we used it throughout the model layers. An architecture constructed in this way provided 1,291,009 trainable parameters.

4.4.2 U-NET inspired model

The architecture of our second model employed a U-NET like design, first proposed by [63] for the purpose of biomedical image classification. The model is characterised by its “U” shape design which employs both contracting and expanding pathways to identify specific features within images. Here, we follow the approach of [29] who, inspired by U-NET, designed their RainNet model for precipitation nowcasting. Thus, we also divided our model into two parts, encoder and decoder, and utilised skip connections between both paths via concatenation layers - unique features of the U-NET model. The encoder (or contracting) pathway of the model included six Conv2D layers with Swish activation and a 2×2 kernel size, as well as two MaxPooling2D layers with pool size $2 \times 2 \times 1$ (see supplementary Fig 16 for model’s schematic drawing). The decoder (or expanding) pathway had five Conv2D layers with two UpSampling2D and two Concatenate layers. The input layers were bordered with a ZeroPadding2D layer which was cropped to the original size of 31×51 with Cropping2D in the output layer. Unlike the original U-NET network, our design received 4-dimensional arrays of shape $7,300 \times 31 \times 51 \times 42$ and generated an output image of a shape of 31×51 for each prediction time step. Thus, the prediction generated 1905 images corresponding to dates from 2nd January 2020 to 31st December 2022.

4.4.3 DustNet model

The last model design built upon the architecture of Conv2D and U-NET. This unique design replaces the Concatenate layers with Transpose convolution layers, also known as Deconvolutional Networks [30]. Schematically represented in Fig. 1 the input layer was first padded with a border of zeros (ZeroPadding2D) which increased the input shape from $31 \times 51 \times 43$ to $40 \times 64 \times 43$. Zero padding enabled the convolution to produce the same output size for multiple input sizes [64]. We then applied the 2D convolving windows (Fig. 1 - pink arrows) which moved over each padded input

with a 2×2 kernel size and 2×2 strides which allow upsampling. This allowed the model to decrease the input size (~~down to 5×8~~) while increasing the amount of **channels ($5 \times 8 \times 256$)**~~trainable parameters~~. A ‘deconvolution’ was then applied by adding Conv2D Transpose layers. An advantage of transposed convolution is its ability to efficiently upscale input data by applying inverse convolutions. This enables the network to increase the size compared to the input and thus generates high-resolution images at finer spatial scales [30]. A 2D cropping layer was then added to bring the **width and height shape** back to its initial input size of 31×51 , **while the final convolution matched the output with desired target size of $31 \times 51 \times 1$** . The final architecture allowed the model to create a total of 1,286,913 trainable parameters. Since this design yielded the optimal results of predicting dust aerosols in comparison to baseline models, we called it DustNet.

4.4.4 Baseline models

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

4.5 Statistical analysis

4.5.1 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. 10). 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.

In Table 2 replaced: **U-NET 25min20s & 53s & 0.001691 & 4.9s & 48.80%**~~51min15s & 194s & 0.001904 & 18.6s & 42.36%~~

To analyse the errors of the best performing model, we rearranged Equation 1 reverses normalisation of AOD predictions from each model:

$$y_{ijt,denorm} = y_{ijt,pred} (y_{max} - y_{min}) + y_{min} \quad (4)$$

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

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 4 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 $N_{\text{lat}}=31$, $N_{\text{lon}}=51$ and $N_t=1095$, using Equation 5:

$$MSE = \frac{1}{N_{\text{lat}}N_{\text{lon}}N_t} \sum_{i=1}^{N_{\text{lat}}} \sum_{j=1}^{N_{\text{lon}}} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt})^2 \quad (5)$$

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.

4.5.2 DustNet evaluation metrics

To compare predictions between the DustNet model, ground truth data from MODIS observations and the physics-based model (CAMS) fairly, we calculated the following metrics: mean bias error (MBE), RMSE, difference between RMSE’s (ΔRMSE) and ACC. The metrics, defined below, follow a combination of notations from [21] and [20] adapted to spatial representation of temporally averaged values for each prediction day t ($N_t=1095$). All prediction values were first de-normalised using Equation 4. Subsequently, we compared the model predictions (\hat{A}) with raw (unimputed) MODIS data (mean of Aqua and Terra) denoted as A . The climatological mean, denoted as A' , corresponds to the long-term average of AOD values from MODIS (2003-2022).

4.5.3 Spatial analysis

To analyse spatial characteristics of model performance, we calculated the temporal mean of model predictions ($N_t = 1095$) at each location (lat,lon). This allowed us to calculate mean bias error (MBE) between the predicted AOD (\hat{A}) and MODIS ground truth (A) for both DustNet and CAMS using Equation 6.

$$MBE_{spatial,ij} = \frac{1}{N_t} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt}) \quad (6)$$

We also calculated the spatial root mean square error ($RMSE_{spatial}$) for each model using Equation 7.

$$RMSE_{spatial,ij} = \sqrt{\frac{1}{N_t} \sum_{t=1}^{N_t} (\hat{A}_{ijt} - A_{ijt})^2} \quad (7)$$

Calculating differences between RMSEs ($\Delta RMSE$) using Equation 8 allowed us to reveal specific locations at which predictions from one model outperformed the other.

$$\Delta RMSE_{spatial,ij} = RMSE_{spatial,ij}^{(CAMS)} - RMSE_{spatial,ij}^{(DustNet)} \quad (8)$$

Additionally, we calculated the spatial distribution of Anomaly Correlation Coefficient (ACC, Equation 9). 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:

$$ACC_{spatial,ij} = \frac{\sum_{t=1}^{N_t} [(\hat{A}'_{ijt} - \bar{A}'_{ijt}) \times (A'_{ijt} - \bar{A}'_{ijt})]}{\sqrt{\left[\sum_{t=1}^{N_t} (\hat{A}'_{ijt} - \bar{A}'_{ijt})^2 \right] \times \left[\sum_{t=1}^{N_t} (A'_{ijt} - \bar{A}'_{ijt})^2 \right]}} \quad (9)$$

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. 7 for graphed results of these calculations.

4.5.4 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 for both DustNet and CAMS

forecasts with MODIS data and are plotted in supplementary Fig. 5. We have also adapted Equations 6 and 7 to temporal representation by using Equation 10 and 11.

$$MBE_{temporal,t} = \frac{1}{N_{lat}N_{lon}} \sum_{i=1}^{N_{lat}} \sum_{j=1}^{N_{lon}} (\hat{A}_{ijt} - A_{ijt}) \quad (10)$$

$$RMSE_{temporal,t} = \sqrt{\frac{1}{N_{lat}N_{lon}} \sum_{i=1}^{N_{lat}} \sum_{j=1}^{N_{lon}} (\hat{A}_{ijt} - A_{ijt})^2} \quad (11)$$

The graphed results of temporal calculations can be seen in Fig. 4 and supplementary Fig. 13.

4.5.5 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 Figure 11, 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 [65]. 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 [66]. The temperature inversion creates a barrier for low horizontal flow of atmospheric dust, and instead forces an uplift of over 1.5km [67]. From this point atmospheric dust moves westward towards Central and South America at higher altitudes between 1.5km - 5km [68]. To capture the transport of dust and fire smoke with southwestward winds towards South America [68] 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 [35–37].

4.5.6 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. 17 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 ([69, 70]).

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 [71] 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> [72].

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 Archive Center (DAAC). Both MOD08_D3 and MYD08_D3 files can be retrieved from <https://ladsweb.modaps.eosdis.nasa.gov/search/>. Pre-processed AOD data is also included in the aforementioned Zenodo repository.

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

5 References

- [1] Gliß, J., Mortier, A., Schulz, M., Andrews, E., Balkanski, Y., Bauer, S.E., Benedictow, A.M.K., Bian, H., Checa-Garcia, R., Chin, M., Ginoux, P., Griesfeller, J.J., Heckel, A., Kipling, Z., Kirkevåg, A., Kokkola, H., Laj, P., Sager, P.L.,

- 660 Lund, M.T., Myhre, C.L., Matsui, H., Myhre, G., Neubauer, D., Noije, T., North,
661 P., Olivié, D.J.L., Rémy, S., Sogacheva, L., Takemura, T., Tsigaridis, K., Tsyro,
662 S.G.: AeroCom phase III multi-model evaluation of the aerosol life cycle and opti-
663 cal properties using ground- and space-based remote sensing as well as surface
664 in situ observations. *Atmospheric Chemistry and Physics* **21**(1), 87–128 (2021)
665 <https://doi.org/10.5194/acp-21-87-2021>
- 666 [2] Kok, J.F., Storelvmo, T., Karydis, V.A., Adebisi, A.A., Mahowald, N.M., Evan,
667 A.T., He, C., Leung, D.M.: Mineral dust aerosol impacts on global climate and
668 climate change. *Nature Reviews Earth & Environment* **4**(2), 71–86 (2023) [https:](https://doi.org/10.1038/s43017-022-00379-5)
669 [//doi.org/10.1038/s43017-022-00379-5](https://doi.org/10.1038/s43017-022-00379-5)
- 670 [3] Kok, J.F., Adebisi, A.A., Albani, S., Balkanski, Y., Checa-Garcia, R., Chin, M.,
671 Colarco, P.R., Hamilton, D.S., Huang, Y., Ito, A., *et al.*: Improved representation
672 of the global dust cycle using observational constraints on dust properties and
673 abundance. *Atmospheric Chemistry and Physics* **21**(10), 8127–8167 (2021)
- 674 [4] Van Der Does, M., Knippertz, P., Zschenderlein, P., Giles Harrison, R., Stuut,
675 J.-B.W.: The mysterious long-range transport of giant mineral dust particles.
676 *Science advances* **4**(12), 2768 (2018)
- 677 [5] Shao, Y., Wyrwoll, K.-H., Chappell, A., Huang, J., Lin, Z., McTainsh, G.H.,
678 Mikami, M., Tanaka, T.Y., Wang, X., Yoon, S.: Dust cycle: An emerging core
679 theme in Earth system science. *Aeolian Research* **2**(4), 181–204 (2011)
- 680 [6] Knippertz, P., Stuut, J.-B.W.: A Key Player in the Earth System, pp. 267–410.
681 Springer, Dordrecht (2014). https://doi.org/10.1007/978-94-017-8978-3_1 . [http:](http://dx.doi.org/10.1007/978-94-017-8978-3_1)
682 [//dx.doi.org/10.1007/978-94-017-8978-3_1](http://dx.doi.org/10.1007/978-94-017-8978-3_1)
- 683 [7] Highwood, E.J., Ryder, C.L.: Radiative Effects of Dust, pp. 267–286. Springer,
684 Dordrecht (2014). https://doi.org/10.1007/978-94-017-8978-3_11 . [http://dx.doi.](http://dx.doi.org/10.1007/978-94-017-8978-3_11)
685 [org/10.1007/978-94-017-8978-3_11](http://dx.doi.org/10.1007/978-94-017-8978-3_11)
- 686 [8] Nenes, A., Murray, B., Bougiatioti, A.: Mineral Dust and its Microphysical Inter-
687 actions with Clouds, pp. 287–325. Springer, Dordrecht (2014). [https://doi.org/10.](https://doi.org/10.1007/978-94-017-8978-3_12)
688 [1007/978-94-017-8978-3_12](https://doi.org/10.1007/978-94-017-8978-3_12) . http://dx.doi.org/10.1007/978-94-017-8978-3_12
- 689 [9] Miller, R.L., Knippertz, P., Pérez García-Pando, C., Perlwitz, J.P., Tegen,
690 I.: Impact of Dust Radiative Forcing upon Climate, pp. 327–357. Springer,
691 Dordrecht (2014). https://doi.org/10.1007/978-94-017-8978-3_13 . [http://dx.doi.](http://dx.doi.org/10.1007/978-94-017-8978-3_13)
692 [org/10.1007/978-94-017-8978-3_13](http://dx.doi.org/10.1007/978-94-017-8978-3_13)
- 693 [10] Jickells, T., Boyd, P., Hunter, K.A.: Biogeochemical Impacts of Dust on the
694 Global Carbon Cycle, pp. 359–384. Springer, Dordrecht (2014). [https://doi.org/](https://doi.org/10.1007/978-94-017-8978-3_14)
695 [10.1007/978-94-017-8978-3_14](https://doi.org/10.1007/978-94-017-8978-3_14) . http://dx.doi.org/10.1007/978-94-017-8978-3_14
696 [14](http://dx.doi.org/10.1007/978-94-017-8978-3_14)

- [11] Morman, S.A., Plumlee, G.S.: Dust and Human Health, pp. 385–409. Springer, Dordrecht (2014). https://doi.org/10.1007/978-94-017-8978-3_15 . http://dx.doi.org/10.1007/978-94-017-8978-3_15
- [12] Parajuli, S.P., Jin, Q., Francis, D.: Editorial: Atmospheric dust: How it affects climate, environment and life on Earth? *Frontiers in Environmental Science* **10** (2022) <https://doi.org/10.3389/fenvs.2022.1058052>
- [13] Evan, A.T., Flamant, C., Fiedler, S., Doherty, O.: An analysis of aeolian dust in climate models. *Geophysical Research Letters* **41**(16), 5996–6001 (2014)
- [14] Kok, J.F., Adebisi, A.A., Albani, S., Balkanski, Y., Checa-Garcia, R., Chin, M., Colarco, P.R., Hamilton, D.S., Huang, Y., Ito, A., *et al.*: Contribution of the world’s main dust source regions to the global cycle of desert dust. *Atmospheric Chemistry and Physics* **21**(10), 8169–8193 (2021)
- [15] Zhao, A., Ryder, C.L., Wilcox, L.J.: How well do the CMIP6 models simulate dust aerosols? *Atmospheric Chemistry and Physics* **22**(3), 2095–2119 (2022) <https://doi.org/10.5194/acp-22-2095-2022>
- [16] Bozzo, A., Remy, S., Benedetti, A., Flemming, J., Bechtold, P., Rodwell, M., Morcrette, J.-J.: Implementation of a CAMS-based aerosol climatology in the IFS. Technical report, European Centre for Medium-Range Weather Forecasts Reading, UK (2017)
- [17] Mulcahy, J.P., Walters, D.N., Bellouin, N., Milton, S.F.: Impacts of increasing the aerosol complexity in the Met Office global numerical weather prediction model. *Atmospheric Chemistry and Physics* **14**(9), 4749–4778 (2014) <https://doi.org/10.5194/acp-14-4749-2014>
- [18] Bozzo, A., Benedetti, A., Flemming, J., Kipling, Z., Rémy, S.: An aerosol climatology for global models based on the tropospheric aerosol scheme in the Integrated Forecasting System of ECMWF. *Geoscientific Model Development* **13**(3), 1007–1034 (2020) <https://doi.org/10.5194/gmd-13-1007-2020>
- [19] Balkanski, Y., Bonnet, R., Boucher, O., Checa-Garcia, R., Servonnat, J.: Better representation of dust can improve climate models with too weak an African monsoon. *Atmospheric Chemistry and Physics* **21**(14), 11423–11435 (2021) <https://doi.org/10.5194/acp-21-11423-2021>
- [20] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Meroze, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., Battaglia, P.: Learning skillful medium-range global weather forecasting. *Science*, 2336 (2023)
- [21] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., Tian, Q.: Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **619**(7970), 533–538

- (2023) <https://doi.org/10.1038/s41586-023-06185-3>
- [22] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., Anandkumar, A.: FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators (2022)
- [23] Kang, S., Kim, N., Lee, B.-D.: Fine dust forecast based on recurrent neural networks. In: 2019 21st International Conference on Advanced Communication Technology (ICACT), pp. 456–459 (2019). IEEE
- [24] Daoud, N., Eltahan, M., Elhennawi, A.: Aerosol optical depth forecast over global dust belt based on LSTM, CNN-LSTM, CONV-LSTM and FFT algorithms. In: IEEE EUROCON 2021-19th International Conference on Smart Technologies, pp. 186–191 (2021). IEEE
- [25] Sarafian, R., Nissenbaum, D., Raveh-Rubin, S., Agrawal, V., Rudich, Y.: Deep multi-task learning for early warnings of dust events implemented for the Middle East. *npj Climate and Atmospheric Science* **6**(1) (2023) <https://doi.org/10.1038/s41612-023-00348-9>
- [26] Hinton, G.E., Dayan, P., Frey, B.J., Neal, R.M.: The "wake-sleep" algorithm for unsupervised neural networks. *Science* **268**(5214), 1158–1161 (1995)
- [27] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015) <https://doi.org/10.1038/nature14539>
- [28] Goroshin, R., Bruna, J., Tompson, J., Eigen, D., LeCun, Y.: Unsupervised learning of spatiotemporally coherent metrics. In: 2015 IEEE International Conference on Computer Vision (ICCV), pp. 4086–4093 (2015). <https://doi.org/10.1109/ICCV.2015.465>
- [29] Ayzel, G., Scheffer, T., Heistermann, M.: RainNet v1. 0: a convolutional neural network for radar-based precipitation nowcasting. *Geoscientific Model Development* **13**(6), 2631–2644 (2020)
- [30] Zeiler, M.D., Krishnan, D., Taylor, G.W., Fergus, R.: Deconvolutional networks. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 2528–2535 (2010). IEEE
- [31] Todd, M.C., Washington, R., Martins, J.V., Dubovik, O., Lizcano, G., M’bainayel, S., Engelstaedter, S.: Mineral dust emission from the Bodélé Depression, northern Chad, during BoDEx 2005. *Journal of Geophysical Research: Atmospheres* **112**(D6) (2007)
- [32] Washington, R., Bouet, C., Cautenet, G., Mackenzie, E., Ashpole, I., Engelstaedter, S., Lizcano, G., Henderson, G.M., Schepanski, K., Tegen, I.: Dust as

- 770 a tipping element: the Bodélé Depression, Chad. Proceedings of the National
771 Academy of Sciences **106**(49), 20564–20571 (2009)
- 772 [33] Jewell, A.M., Drake, N., Crocker, A.J., Bakker, N.L., Kunkelova, T., Bristow,
773 C.S., Cooper, M.J., Milton, J.A., Breeze, P.S., Wilson, P.A.: Three North African
774 dust source areas and their geochemical fingerprint. Earth and Planetary Science
775 Letters **554**, 116645 (2021)
- 776 [34] Koren, I., Kaufman, Y.J., Washington, R., Todd, M.C., Rudich, Y., Martins, J.V.,
777 Rosenfeld, D.: The Bodélé Depression: a single spot in the Sahara that provides
778 most of the mineral dust to the Amazon forest. Environmental Research Letters
779 **1**(1), 014005 (2006)
- 780 [35] Anuforom, A.C.: Spatial distribution and temporal variability of Harmattan dust
781 haze in sub-Sahel West Africa. Atmospheric Environment **41**(39), 9079–9090
782 (2007)
- 783 [36] Sunnu, A., Afeti, G., Resch, F.: A long-term experimental study of the Saharan
784 dust presence in West Africa. Atmospheric Research **87**(1), 13–26 (2008)
- 785 [37] Schwanghart, W., Schütt, B.: Meteorological causes of Harmattan dust in West
786 Africa. Geomorphology **95**(3-4), 412–428 (2008)
- 787 [38] Düben, Peter and Modigliani, Umberto and Geer, Alan and Siemen, Stephan
788 and Pappenberger, Florian and Bauer, Peter and Brown, Andy and Palkovic,
789 Martin and Raoult, Baudouin and Wedi, Nils and Baousis, Vasileios: Machine
790 learning at ECMWF: A roadmap for the next 10 years. ECMWF (2021). <https://doi.org/10.21957/ge7ckgm> . <https://www.ecmwf.int/node/19877>
- 791
- 792 [39] Vandenbussche, S., Callewaert, S., Schepanski, K., De Mazière, M.: North African
793 mineral dust sources: new insights from a combined analysis based on 3D dust
794 aerosol distributions, surface winds and ancillary soil parameters. Atmospheric
795 Chemistry and Physics **20**(23), 15127–15146 (2020)
- 796 [40] Prospero, J., Glaccum, R., Nees, R.: Atmospheric transport of soil dust from
797 Africa to South America. Nature **289**(5798), 570–572 (1981)
- 798 [41] Mbourou, G., Bertrand, J., Nicholson, S.: The diurnal and seasonal cycles of wind-
799 borne dust over Africa north of the equator. Journal of Applied Meteorology and
800 Climatology **36**(7), 868–882 (1997)
- 801 [42] Schepanski, K., Heinold, B., Tegen, I.: Harmattan, Saharan heat low, and West
802 African monsoon circulation: modulations on the Saharan dust outflow towards
803 the North Atlantic. Atmospheric Chemistry and Physics **17**(17), 10223–10243
804 (2017)
- 805 [43] Janicot, S., Thorncroft, C.D., Ali, A., Asencio, N., Berry, G., Bock, O., Bourlès,

- 806 B., Caniaux, G., Chauvin, F., Deme, A., *et al.*: Large-scale overview of the sum-
 807 mer monsoon over West Africa during the AMMA field experiment in 2006. In:
 808 Annales Geophysicae, vol. 26, pp. 2569–2595 (2008). Copernicus Publications
 809 Göttingen, Germany
- 810 [44] N'Datchoh, E., Diallo, I., Konaré, A., Silué, S., Ogunjobi, K., Diedhiou, A.,
 811 Doumbia, M.: Dust induced changes on the West African summer monsoon
 812 features. *International Journal of Climatology* **38**(1), 452–466 (2018)
- 813 [45] O'Sullivan, D., Marengo, F., Ryder, C.L., Pradhan, Y., Kipling, Z., Johnson,
 814 B., Benedetti, A., Brooks, M., McGill, M., Yorks, J., *et al.*: Models transport
 815 Saharan dust too low in the atmosphere: a comparison of the MetUM and CAMS
 816 forecasts with observations. *Atmospheric Chemistry and Physics* **20**(21), 12955–
 817 12982 (2020)
- 818 [46] Wu, C., Lin, Z., Liu, X.: The global dust cycle and uncertainty in CMIP5 (Coupled
 819 Model Intercomparison Project phase 5) models. *Atmospheric Chemistry and*
 820 *Physics* **20**(17), 10401–10425 (2020)
- 821 [47] Schepanski, K., Tegen, I., Laurent, B., Heinold, B., Macke, A.: A new Saharan
 822 dust source activation frequency map derived from MSG-SEVIRI IR-channels.
 823 *Geophysical Research Letters* **34**(18) (2007)
- 824 [48] Mari, C., Cailley, G., Corre, L., Saunio, M., Attié, J.-L., Thouret, V., Stohl,
 825 A.: Tracing biomass burning plumes from the Southern Hemisphere during the
 826 AMMA 2006 wet season experiment. *Atmospheric Chemistry and Physics* **8**(14),
 827 3951–3961 (2008)
- 828 [49] Knippertz, P., Fink, A.H., Deroubaix, A., Morris, E., Tocquer, F., Evans, M.J.,
 829 Flamant, C., Gaetani, M., Lavaysse, C., Mari, C., *et al.*: A meteorological and
 830 chemical overview of the DACCIIWA field campaign in West Africa in June–July
 831 2016. *Atmospheric Chemistry and Physics* **17**(17), 10893–10918 (2017)
- 832 [50] Friese, C.A., Van Hateren, J.A., Vogt, C., Fischer, G., Stuut, J.-B.W.: Seasonal
 833 provenance changes in present-day Saharan dust collected in and off Mauritania.
 834 *Atmospheric Chemistry and Physics* **17**(16), 10163–10193 (2017)
- 835 [51] Ginoux, P., Prospero, J.M., Gill, T.E., Hsu, N.C., Zhao, M.: Global-scale attri-
 836 bution of anthropogenic and natural dust sources and their emission rates based
 837 on MODIS Deep Blue aerosol products. *Reviews of Geophysics* **50**(3) (2012)
 838 <https://doi.org/10.1029/2012rg000388>
- 839 [52] Hubanks, P., Platnick, S., King, M., Ridgway, B.: MODIS Atmosphere L3 grid-
 840 ded product algorithm theoretical basis document (atbd) & users guide. ATBD
 841 reference number ATBD-MOD-30, NASA **125**, 585 (2015)
- 842 [53] Morcrette, J.-J., Boucher, O., Jones, L., Salmond, D., Bechtold, P., Beljaars,

- 843 A., Benedetti, A., Bonet, A., Kaiser, J., Razinger, M., et al.: Aerosol analy-
844 sis and forecast in the European Centre for medium-range weather forecasts
845 integrated forecast system: Forward modeling. *Journal of Geophysical Research:*
846 *Atmospheres* **114**(D6) (2009)
- 847 [54] Benedetti, A., Morcrette, J.-J., Boucher, O., Dethof, A., Engelen, R.J., Fisher,
848 M., Flentje, H., Huneeus, N., Jones, L., Kaiser, J.W., Kinne, S., Mangold, A.,
849 Razinger, M., Simmons, A.J., Suttie, M.: Aerosol analysis and forecast in the
850 European Centre for Medium-Range Weather Forecasts Integrated Forecast Sys-
851 tem: 2. Data assimilation. *Journal of Geophysical Research* **114**(D13) (2009)
852 <https://doi.org/10.1029/2008jd011115>
- 853 [55] Hartman, L., Hössjer, O.: Fast kriging of large data sets with Gaussian Markov
854 random fields. *Computational Statistics & Data Analysis* **52**(5), 2331–2349
855 (2008)
- 856 [56] Rue, H., Held, L.: *Gaussian Markov Random Fields: Theory and Applications*,
857 pp. 40–48. Chapman and Hall/CRC press, New York (2005). Chap. 2.4 Numerical
858 methods for sparse matrices. <https://doi.org/10.1201/9780203492024>
- 859 [57] Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J.,
860 Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al.: ERA5 hourly data on single
861 levels from 1979 to present. Copernicus Climate Change Service (C3S) Climate
862 Data Store (CDS) (2018) <https://doi.org/10.24381/cds.bd0915c6>
- 863 [58] Rasp, S., Dueben, P.D., Scher, S., Weyn, J.A., Mouatadid, S., Thuerey, N.: Weath-
864 erBench: a benchmark data set for data-driven weather forecasting. *Journal of*
865 *Advances in Modeling Earth Systems* **12**(11), 2020–002203 (2020)
- 866 [59] Chollet, F., et al.: Keras. GitHub (2015). <https://github.com/fchollet/keras>
- 867 [60] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint
868 arXiv:1412.6980 (2014)
- 869 [61] Ramachandran, P., Zoph, B., Le, Q.V.: Searching for activation functions. arXiv
870 preprint arXiv:1710.05941 (2017)
- 871 [62] Rasamoelina, A.D., Adjailia, F., Sinčák, P.: A review of activation function for
872 artificial neural network. In: 2020 IEEE 18th World Symposium on Applied
873 Machine Intelligence and Informatics (SAMI), pp. 281–286 (2020). IEEE
- 874 [63] Ronneberger, O., Fischer, P., Brox, T.: U-NET: Convolutional networks for
875 biomedical image segmentation. In: Medical Image Computing and Computer-
876 Assisted Intervention–MICCAI 2015: 18th International Conference, Munich,
877 Germany, October 5–9, 2015, Proceedings, Part III 18, pp. 234–241 (2015).
878 Springer

- [64] Dumoulin, V., Visin, F.: A guide to convolution arithmetic for deep learning. arXiv preprint arXiv:1603.07285 (2016)
- [65] Washington, R., Todd, M., Middleton, N.J., Goudie, A.S.: Dust-storm source areas determined by the total ozone monitoring spectrometer and surface observations. *Annals of the Association of American Geographers* **93**(2), 297–313 (2003)
- [66] Carlson, T.N., Prospero, J.M.: The large-scale movement of Saharan air outbreaks over the northern equatorial Atlantic. *Journal of Applied Meteorology and Climatology* **11**(2), 283–297 (1972)
- [67] Prospero, J.M., Carlson, T.N.: Vertical and areal distribution of Saharan dust over the western equatorial North Atlantic Ocean. *Journal of Geophysical Research* **77**(27), 5255–5265 (1972)
- [68] Kaufman, Y., Koren, I., Remer, L., Tanré, D., Ginoux, P., Fan, S.: 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) (2005)
- [69] Covert, I., Lundberg, S., Lee, S.-I.: Explaining by removing: A unified framework for model explanation. *Journal of Machine Learning Research* **22**(209), 1–90 (2021)
- [70] Molnar, C.: Chapter 10: Neural Network Interpretation. *Interpretable Machine Learning*, 2nd edn. Github, online (2022). <https://christophm.github.io/interpretable-ml-book>
- [71] Nowak, T.E., Augousti, A.T., Simmons, B.I., Siegert, S.: DustNet - structured data and Python code to reproduce the model, statistical analysis and figures. zenodo (2024). <https://doi.org/10.5281/zenodo.10631953> . <https://zenodo.org/records/10722953>
- [72] Nowak, T.E., Augousti, A.T., Simmons, B.I., Siegert, S.: Pre-processed daily ERA5 and MODIS AOD data (2003 - 2022) ready for use in AI/ML forecasting. zenodo (2024). <https://doi.org/10.5281/zenodo.10593151> . <https://zenodo.org/records/10593152>
- [73] Jo, D.S., Tilmes, S., Emmons, L.K., Wang, S., Vitt, F.: 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 (2023)

Supplementary Figures

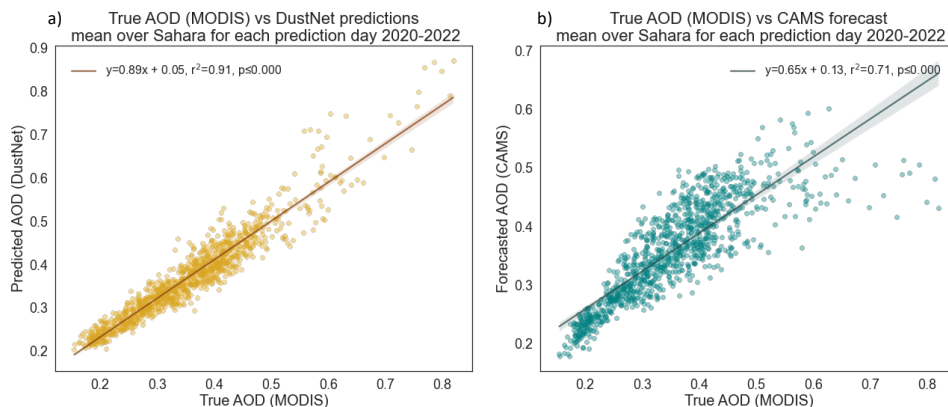


Fig. 5 Fig. A1: Spatially averaged daily AOD (2020-2022, n=1095) regressed between model predictions and MODIS data. ~~Daily spatial mean AOD (2020-2022) regressed between model predictions and MODIS data.~~ Linear regression with corresponding y equation, Pearson's r^2 and p values were calculated for 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 $p < 0.0001$.




Fig. 6 Fig. A2: Seasonal mean distribution of daily 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. 7 Fig. A3: Bias of predictions **a)** DustNet and **b)** CAMS with respect to MODIS data. Note that the maximum bias produced by DustNet is 0.21 while the maximum bias for CAMS is 0.93.

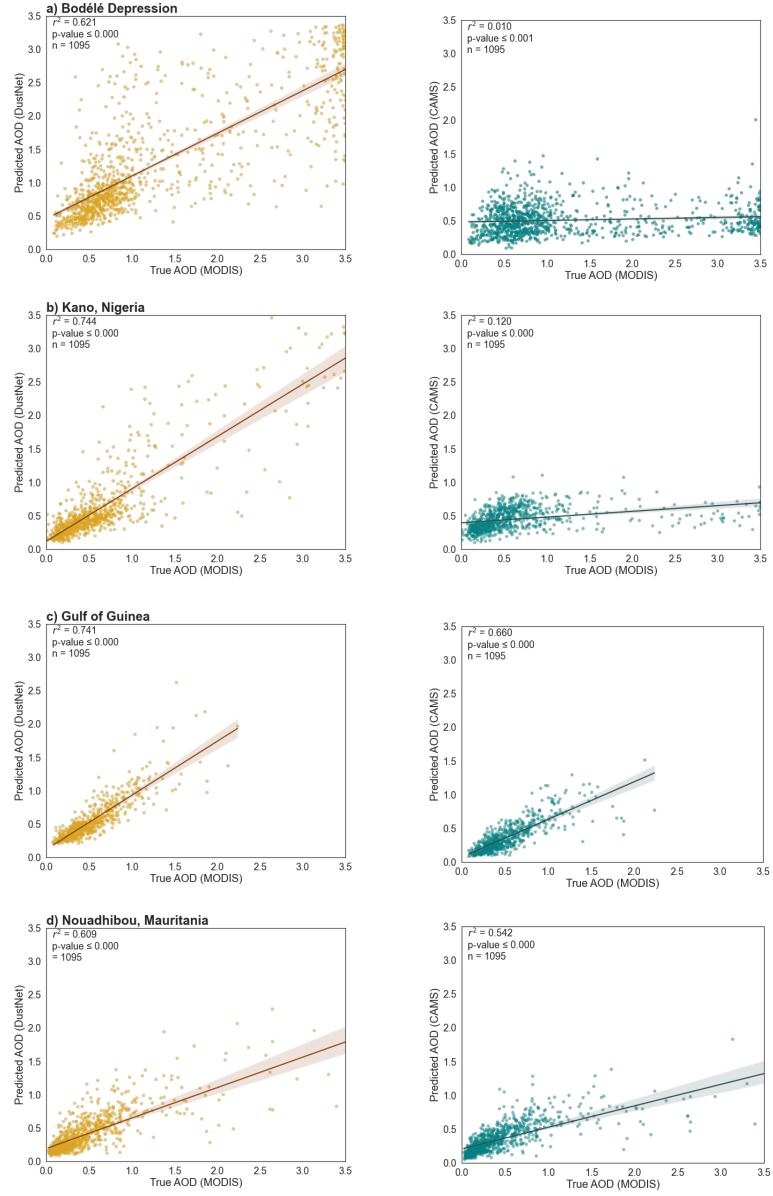


Fig. 8 Fig. A4: Scatter plot relationship between predicted mean AOD values (2020-2022) and MODIS data. Results for DustNet (**left panel**) and forecasts from CAM5 (**right panel**) at four chosen locations show better agreement of DustNet predictions with MODIS data at each location. In **a)** the Bodélé Depression, Chad - highest source of dust in the Sahara - DustNet is significantly better than CAM5; **b)** Kano, Nigeria - the second most populous province **c)** Gulf of Guinea - over the ocean; and **d)** Nouadhibou, Mauritania - coastal location.




Fig. 9 Fig. A5: 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.

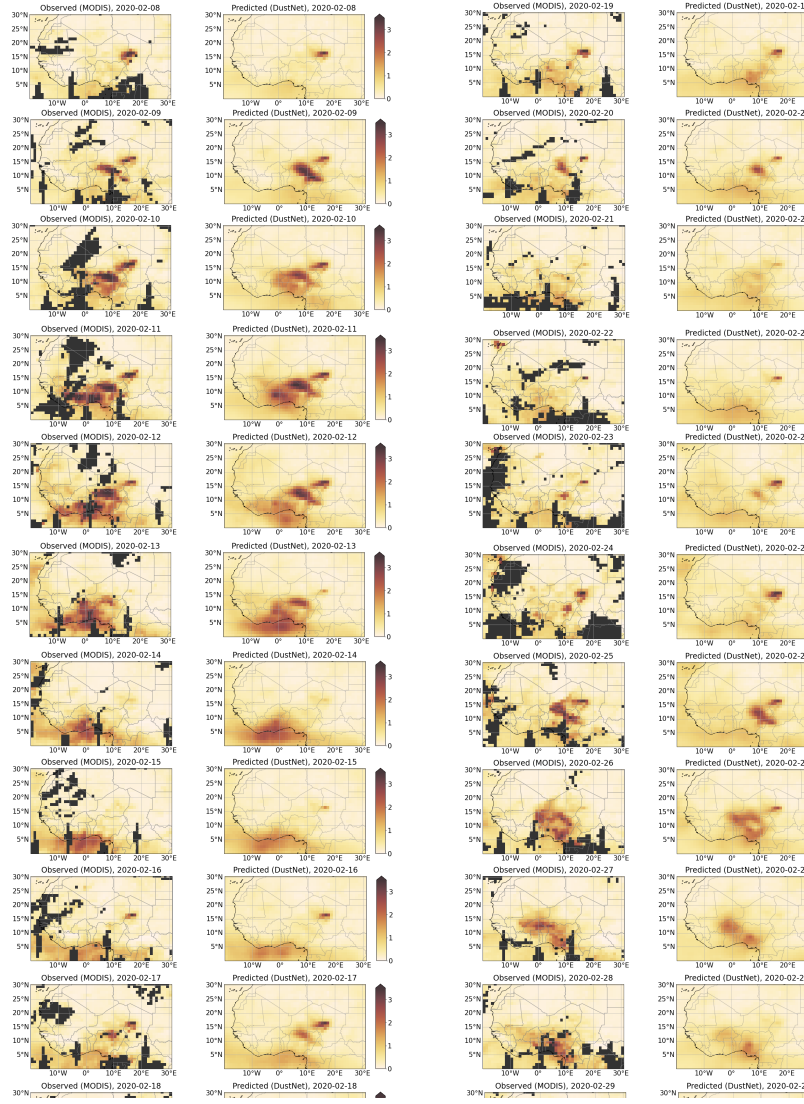


Fig. 10 Fig. A6: Comparison of daily AOD values as observed by MODIS (mean of Aqua and Terra) (**left panel** in both columns) and corresponding DustNet predictions (**right panel**) for selected continuous 3 weeks (22 days), from 8th - 29th February 2020. The dark grey colour in the MODIS maps represents 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).

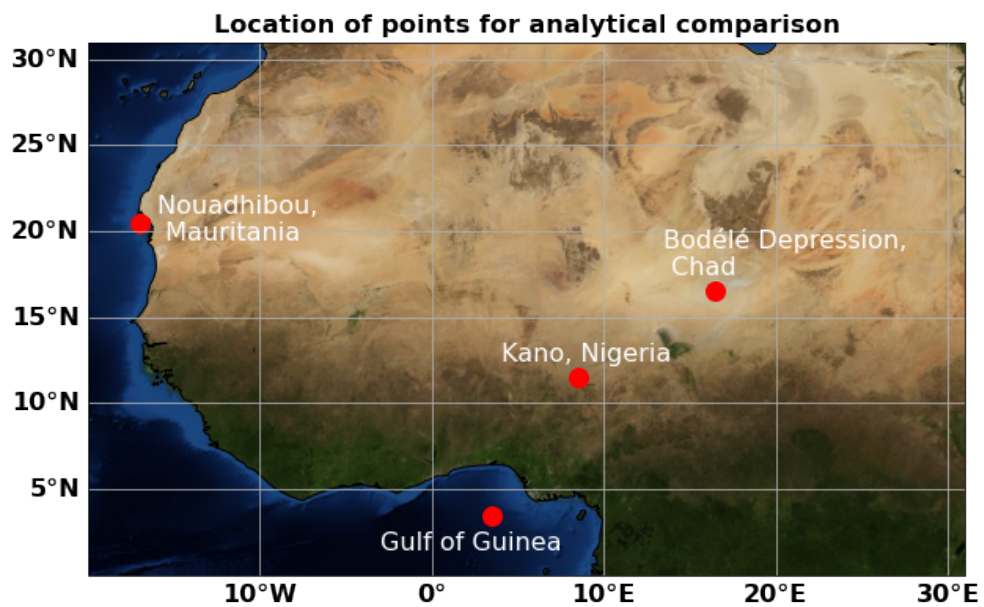


Fig. 11 Fig. A7: Locations of selected grid points used to asses the model's predictive accuracy on a local scale ($1^\circ \times 1^\circ$ grid size). The background image for the December view of Blue Marble is available from NASA <https://visibleearth.nasa.gov/collection/1484/blue-marble?page=4>




Fig. 12 Fig. A8: 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. 13 Fig. A9: In the **top panel** the temporal mean RMSE values for AOD predicted by DustNet (yellow) and CAMS (cyan) are shown. The dashed line represents the climatological mean from MODIS. 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.

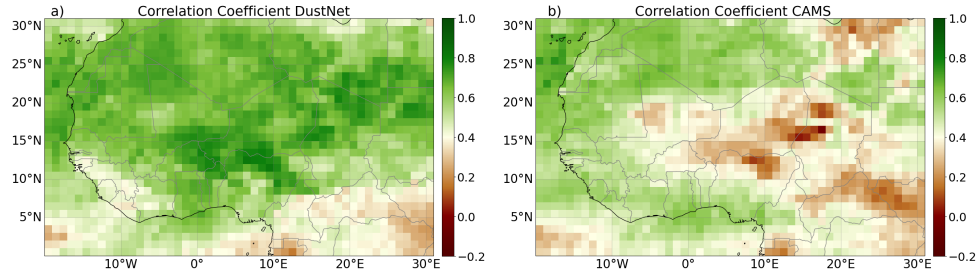


Fig. 14 Fig. A10: Daily MODIS correlation coefficient with predictions from **a)** DustNet and **b)** CAMS. The maximum correlation for DustNet is 0.82 and minimum is 0.16, while the maximum correlation for CAMS is 0.75 and minimum is -0.04. Values with weaker (≤ 0.4) correlation are represented in white to brown, while values with stronger >0.4 correlation are shown in green shades. With predominantly green shades, especially over the Sahara region, the advantage of DustNet predictions is clearly visible.

914 Represented in Figure 14 is the daily correlation coefficient between the ground
 915 true values from MODIS and DustNet predictions in panel a), and between MODIS
 916 and CAMS in panel b). The area of the Saharan Desert, where mineral dust is the main
 917 contributor to the AOD values, is the stronghold of DustNet-produced predictions. In
 918 comparison, CAMS displays a lower (weaker) correlation with MODIS over areas in
 919 the Saharan Desert, which were highlighted in previous figures (Fig. 3 and Fig. 10)
 920 as dust generation regions. The regions south of 8°N and east of 10°E display weaker
 921 correlation for both models. This is not surprising, since DustNet's training regime did
 922 not include any data indicative of black carbon or secondary organic aerosols, which
 923 seasonally dominate the AOD over the equatorial region of Central African forests
 924 ([73]).




Fig. 15 Fig. A11: 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. 16 Fig. A12: 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 Max-Pooling 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. 17 Fig. A13: DustNet model's feature importance based on MSE. The features with the highest MSE indicates the highest importance for the pre-trained DustNet model in creating 24-hours (1-step ahead) predictions. The bar chart indicates that the model predictions result in highest MSE when the information on 'AOD 1 day lag' feature values are removed (zeroed-out). Similarly, when feature values for 'vertical velocity (850hPa)' were removed from prediction run, the model resulted in higher MSE. This suggests that the pre-trained DustNet model relies on 'yesterdays' AOD values and vertical velocity at 850hPa the most during prediction run.

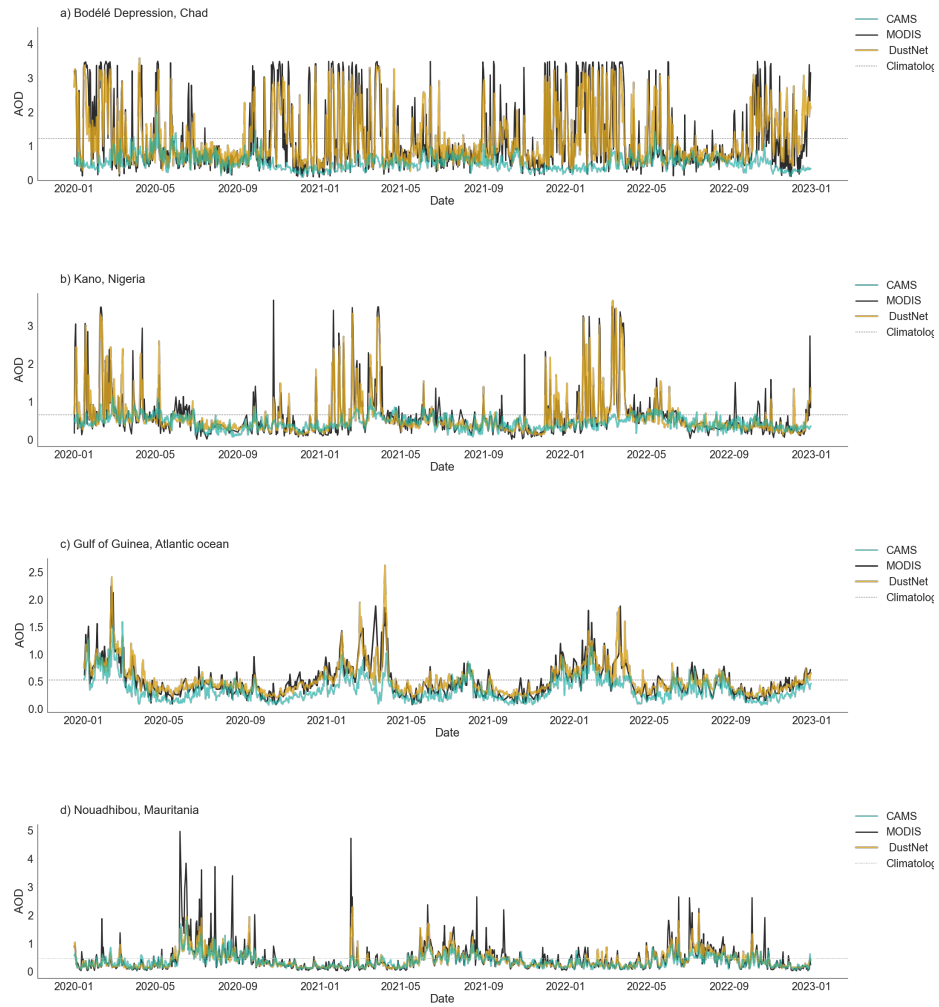


Fig. 18 Fig. A14: Same as Fig. 4, but for daily data.