TOWARDS DOWNSCALING GLOBAL AOD WITH MACHINE LEARNING

Josh Millar[1,2], **Paula Harder**[3], **Lilli Freischem**[1], **Philipp Weiss**[1], **Philip Stier**[1]
[1] University of Oxford, Oxford, United Kingdom
[2] Imperial College, London, United Kingdom
[3] Mila Quebec AI Institute, Quebec, Canada

ABSTRACT

Poor air quality represents a significant threat to human health, especially in urban areas. To improve forecasts of air pollutant mass concentrations, there is a need for high-resolution Aerosol Optical Depth (AOD) forecasts as proxy. However, current General Circulation Model (GCM) forecasts of AOD suffer from limited spatial resolution, making it difficult to accurately represent the substantial variability exhibited by AOD at the local scale. To address this, a deep residual convolutional neural network (ResNet) is evaluated for the GCM to local scale downscaling of low-resolution global AOD retrievals, outperforming a non-trainable interpolation baseline. We explore the bias correction potential of our ResNet using global reanalysis data, evaluating it against in-situ AOD observations. The improved resolution from our ResNet can assist in the study of local AOD variations.

1 INTRODUCTION

Aerosols, small atmospheric particles, can influence the global radiation balance as strongly as greenhouse gases and therefore play an important role in both local and global climate. $PM_{2.5}$ aerosol particles (i.e., those with diameter less than 2.5μm) are also the primary source of air pollutant responsible for causing significant health conditions, including respiratory and cardiovascular illnesses as well as lung cancer Beelen et al. (2013); Khomenko et al. (2021).

Aerosol Optical Depth (AOD) is the column-integrated light extinction over an atmospheric column i.e., a measure of the degree to which transmission of light through an atmospheric column is reduced due to its scattering or absorption by aerosols. This can be expressed as:

$$AOD_\lambda = \int_{z_1}^{z_2} e_\lambda(x, y, z, t)dz$$

where λ is a particular wavelength. Surface $PM_{2.5}$ mass concentrations have been found to be well-correlated with AOD Wang & Christopher (2003); Yang et al. (2019), leading to significant research directed towards estimating $PM_{2.5}$ mass concentrations from AOD retrievals. High-resolution (HR) GCM forecasts of AOD are therefore essential for understanding the effects of local-scale air pollution on the climate, as well as for other downstream health-related applications. However, GCM forecasts are limited in their spatial resolution by computational costs. Typically, these forecasts operate at a global spatial resolution of 80-250km Iles et al. (2020), which is much higher than the spatial scales over which AOD can vary, particularly in urban areas.

Image super-resolution (SR) is the problem of estimating a HR image from a lower-resolution counterpart image and has a wide range of applications from video surveillance Aakerberg et al. (2021) to medical diagnosis Chen et al. (2021). The problem is inherently ill-posed, and therefore challenging, since an infinite number of high-resolution images can be obtained from a single coarse-grained image. SR is a highly researched field, with diverse approaches ranging from early interpolation Zhang & Wu (2006); Keys (1981); Duchon (1979) and reconstruction Zhang et al. (2012); Dai et al. (2009); Liu et al. (2016) approaches to ML-based approaches.

The SR of climate variables is referred to as *downscaling*. Vandal et al. (2017) pioneered the application of ML for the task, utilizing a CNN to downscale satellite-observed precipitation data,

outperforming existing statistical methods. Consequently, ML has been established as the standard approach for the task. The SR-ResNet Ledig et al. (2017); Lim et al. (2017), built on residual blocks, has achieved state-of-the-art performance, in terms of restoration quality, for CNN approaches.

A limited number of studies have explored downscaling AOD . Wang et al. (2022) explored the downscaling of monthly-averaged 50km-resolution AOD, coupled with elevation, to finer 10km-resolution over a relatively large study area encompassing Southwest Asia. They utilise fully-connected neural networks and ensure the preservation of inherent within-scale temporal correlations. Li et al. (2020) utilise an SR-ResNet autoencoder to perform both gap-filling and 10x downscaling of weekly-averaged AOD in California, achieving impressive results. A recent study by Zhang et al. (2022a) utilizes a Generative Adversarial Network to gap-fill coverage over China.

Contributions
This study evaluates a CNN approach for the *spatial* downscaling of global AOD satellite observations. We also evaluate the application of the developed approach to global AOD reanalysis data, using comparison with in-situ AOD ground measurements to assess bias correction performance. To the best of our knowledge, we are the first to evaluate AOD downscaling performance on a global basis (i.e., its evaluation not restricted to a localised study area), as well as the first to apply machine learning (ML) towards AOD bias correction.

This study focuses on AOD at the reference wavelength 550nm, hereinafter denoted AOD_{550}.

2 DATA

MODIS Level 2 (MOD04_L2) Satellite Data
MODIS Level 2 "MOD04_L2" AOD_{550} satellite data MODIS Atmosphere Science Team (2017) is used for training the CNN framework. There are limited studies evaluating the global-scale performance of MODIS AOD; these studies generally report a low positive bias at global scale, but with large variability at regional scale Garrigues et al. (2022). A number of localised studies have shown MODIS AOD products to outperform other AOD satellite data-products when validated against ground-based measurements Payra et al. (2023); Handschuh et al. (2022); Garrigues et al. (2022).

The MOD04_L2 AOD data is produced at 10km (10x10 1-km at nadir) spatial resolution over both land and ocean. Different AOD retrieval algorithms are employed for land retrievals (Dark Target) and ocean retrievals (Deep Blue) in order to handle variations in surface reflectance Rubin & Collins (2014). Further details on each algorithm can be found in Levy et al. (2007).

We obtained MOD04_L2 AOD_{550} swaths for the period 24/02/2000 to 19/07/2023 from the Centre for Environmental Data Analysis (CEDA) archive. MOD04_L2 data is produced in sinusoidal projection. To facilitate later collocation of the data with validation datasets, the data is reprojected onto a fixed regular latitude-longitude grid. The HR MODIS data is coarsened using average pooling to generate the low-resolution (LR) input - HR target pairs required. The pooling results in LR inputs at approx. 0.9° resolution and HR targets at 0.09° resolution.

Various methods have been explored in the literature for gap-filling MODIS AOD data Zhang et al. (2022b); Kianian et al. (2021); Yang & Hu (2018). Here, we employed simple post-coarsening mean imputation, with crops consisting of over 25% missing values discarded pre-imputation.

CAMS Reanalysis
To evaluate bias correction efficacy, Copernicus Atmosphere Monitoring Service (CAMS) reanalysis data is used. CAMS AOD data is reported to have a slightly higher positive bias at global scale than MODIS MOD04_L2 Gueymard & Yang (2020). We obtained globaly daily CAMS AOD_{550} data for a shorter period 01/01/2009 to 31/12/2016 using the ECMWF web API. The data is produced on a fixed regular 0.75° grid.

AERONET Station Measurements
The standard methodology for the validation of AOD estimates is to compare with in-situ AERONET "ground-truth" measurement data. AERONET is a global network of multiwavelength sun-photometry ground instruments providing measurements of AOD at point location. The AOD retrieval algorithm is explained in Smirnov et al. (2000).

We obtained AERONET measurements for the shorter period 01/01/2009 to 31/12/2016 using the AERONET web API. The measurements are obtained for selected stations (that have been identified as well-maintained) to ensure a balanced global representation; a full list of the stations used, and their details, can be found in Table 2 in the Supplementary Material.

AERONET retrieves AOD measurements at eight wavelengths between 340nm and 1020nm; AOD_{550}, however, is not one of these and must be interpolated from neighbouring wavelengths Wang et al. (2017); Gueymard & Yang (2020) using the Ångström equation (and 440–870nm Ångström exponent). This assumes aerosol properties remain consistent within the measurement range.

The vertical profile of AOD is non-linear Stevens et al. (2017). Aerosols tend to be surface-concentrated i.e., AOD at elevated sites is typically lower than at sites nearer sea level within the same region. The pixel-mean surface elevation of each satellite/reanalysis grid cell and the elevation of the AERONET station within that cell may differ significantly, particularly in regions of complex terrain. Therefore, point-based AERONET measurements may not fully represent the aerosol distribution within a model/satellite grid cell Su et al. (2017); Song et al. (2018); Shi et al. (2019); Mukkavilli et al. (2019); He et al. (2019). This issue is exacerbated by lower grid cell resolution. To mitigate this, we employed a fixed 2.1km (i.e., the approx. top altitude of the boundary layer) scale-height correction when collocating AERONET measurements with gridded CAMS data.

3 METHODOLOGY

3.1 ARCHITECTURE

A modified SR-ResNet, based on the improved architecture presented by Lim et al. (2017) (with further modifications by Harder et al. (2023)), is developed. This architecture employs a "progressive upsampling" framework in which resolution is increased in stages within the network; this approach is more beneficial for large-factor (i.e., 10x) SR than pre- or post-upsampling approaches.

The network consists of two upsampling stages; the first stage performs an initial 5x upsampling with the second performing the further 2x upsampling required. The upsampling is performed by sub-pixel transpose convolutional layers. Each upsampling layer is followed by a set of residual blocks with skip connections. As per Lim et al. (2017), batch normalisation layers within the residual blocks are omitted. Unlike Lim et al., we maintained ReLU Agarap (2018) layers outside the residual blocks, including in the final network layer, to constrain network outputs to be non-negative.

3.2 EXPERIMENTAL SETUP

The ResNet is trained with a learning rate of $1e^{-3}$ and a batch size of 128. The Adam optimiser (Kingma & Ba (2017)) is used, with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. To deal with the inherent sparsity of the HR target images, a modified Mean Squared Error (MSE) loss function is used; this function calculates the MSE between each downscaled output and HR target image exclusively for the non-missing values in the target. We implemented early stopping to prevent overfitting using a monitor on validation loss; the number of training epochs varies but typically converges around 70-90 (approx. 5 hours on a single Nvidia A100 GPU).

Owing to its established usage as a baseline in SR Shi et al. (2016); Liu et al. (2021); Shocher et al. (2017) and downscaling studies Harder et al. (2023); Mu et al. (2020), Lanczos interpolation is used as a non-trainable baseline for evaluation. This being the first study to evaluate global AOD downscaling, there are unfortunately no existing ML baselines for further comparison.

4 RESULTS

The metric scores for each downscaling approach (i.e., ResNet and baseline) applied to LR MODIS AOD_{550} test set data are summarised in Table 1 below. As shown, our ResNet achieves significantly improved performance over the baseline in terms of KGE and MSE. The correlation scores on MODIS also show the ResNet-downscaled values to exhibit a stronger linear relationship with the HR target values than those downscaled using the baseline. This suggests the ResNet is more effective in capturing the inherent spatial patterns in the HR target data. However, the NMB scores also reveal a larger positive bias resulting from the ResNet's application, which is concerning for

bias correction performance as it suggests that the framework may have introduced/learned a biased representation. The visual improvement from our ResNet is evident in Figure 1.

Table 1: Metric scores for baseline and ResNet on coarsened LR MODIS test set data (evaluated against HR MODIS) and LR CAMS data (evaluated against AERONET). We report Kling-Gupta Efficiency (KGE) Gupta et al. (2009), MSE, Normalised Mean Bias (NMB), and (Pearson's) R. The best scores for each objective are highlighted in bold.

	MODIS (MOD04_L2)		**CAMS**		
	HR DS (baseline)	**HR DS (ResNet)**	**LR**	**HR DS (baseline)**	**HR DS (ResNet)**
KGE	0.571	**0.939**	0.646	**0.758**	0.738
MSE	0.008	**0.005**	0.075	**0.066**	0.069
NMB	**0.017**	0.039	0.203	**0.092**	0.098
R	0.920	**0.936**	**0.762**	0.758	0.752

The metric scores for the evaluation of reanalysis-bias correction performance, with each downscaling approach applied to CAMS data, are also summarized in Table 1. The improved performance of both approaches over the non-downscaled LR CAMS data is evident across various metrics, with the baseline showing marginally greater improvement than our CNN framework. However, both approaches also exhibit weaker correlations with the HR target values than the LR data. These bias correction results suggest that the ResNet is inadequately preserving the fidelity of the CAMS data to AERONET, despite being trained on observation data with a purportedly lower bias.

This could be attributed to the gap-filling technique employed, which may have introduced a biased representation of the MODIS training data, or inherent biases in the MODIS data itself. Fig 3 in the Supplementary Material shows the output of each approach applied to CAMS; despite its inferior numerical performance, the ResNet produces a slight visual improvement over the baseline.

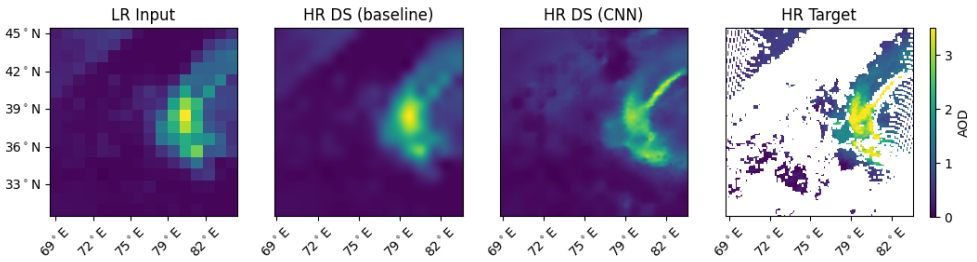

Figure 1: Example outputs, from baseline and ResNet (CNN), on a MOD04_L2 test image. The sparsity of the HR MODIS data is visible in the target image, while the LR image shows the results of the coarsening and gap-filling procedures.

5 Conclusion

In this study, we evaluated a ResNet-based approach for downscaling global AOD. The developed approach exhibited superior performance to an interpolation baseline when applied to MODIS satellite-observed AOD data, achieving KGE of 0.939 compared to 0.571 for the baseline. However, the evaluation also showed a larger positive bias resulting from the approach's application, which consequently did not outperform the baseline when applied to reanalysis data (KGE of 0.738 compared to 0.758 for the baseline). This study can be considered an exploratory work, highlighting the need for improved AOD satellite data as well as robust downscaling methods capable of better handling the sparsity typical of HR satellite imagery. Future work could explore incorporating physical constraining as part of the downscaling network Geiss et al. (2022); Harder (2022), regularization through the temporal recurrence mechanisms, and the unsupervised fine-tuning of ML-based downscaling approaches on reanalysis/GCM data. The latter could help mitigate distribution shift issues arising from using satellite data as a proxy when HR GCM data is unavailable.

REFERENCES

Andreas Aakerberg, Kamal Nasrollahi, and Thomas B. Moeslund. Real-World Super-Resolution of Face-Images from Surveillance Cameras. *CoRR*, abs/2102.03113, 2021. URL https://arxiv.org/abs/2102.03113.

Abien Fred Agarap. Deep Learning using Rectified Linear Units (ReLU). *CoRR*, abs/1803.08375, 2018. URL http://arxiv.org/abs/1803.08375.

Rob Beelen, Ole Raaschou-Nielsen, Zorana Andersen, Gudrun Weinmayr, Barbara Hoffmann, Kathrin Wolf, Evi Samoli, Paul Fischer, Mark Nieuwenhuijsen, Paolo Vineis, Wei Xun, Klea Katsouyanni, Konstantina Dimakopoulou, Anna Oudin, Bertil Forsberg, Lars Modig, Aki Havulinna, Timo Lanki, and Gerard Hoek. Effects of long-term exposure to air pollution on natural-cause mortality: An analysis of 22 European cohorts within the multicentre ESCAPE project. *Lancet*, 383, 12 2013. doi: 10.1016/S0140-6736(13)62158-3.

Zhen Chen, Xiaoqing Guo, Peter Y. M. Woo, and Yixuan Yuan. Super-Resolution Enhanced Medical Image Diagnosis With Sample Affinity Interaction. *IEEE Transactions on Medical Imaging*, 40 (5):1377–1389, 2021. doi: 10.1109/TMI.2021.3055290.

Shengyang Dai, Mei Han, Wei Xu, Ying Wu, Yihong Gong, and Aggelos K. Katsaggelos. SoftCuts: A Soft Edge Smoothness Prior for Color Image Super-Resolution. *IEEE Transactions on Image Processing*, 18(5):969–981, 2009. doi: 10.1109/TIP.2009.2012908.

Claude E. Duchon. Lanczos Filtering in One and Two Dimensions. *Journal of Applied Meteorology and Climatology*, 18(8):1016 – 1022, 1979. doi: https://doi.org/10.1175/1520-0450(1979)018⟨1016:LFIOAT⟩2.0.CO;2. URL https://journals.ametsoc.org/view/journals/apme/18/8/1520-0450_1979_018_1016_lfioat_2_0_co_2.xml.

S. Garrigues, S. Remy, J. Chimot, M. Ades, A. Inness, J. Flemming, Z. Kipling, I. Laszlo, A. Benedetti, R. Ribas, S. Jafariserajehlou, B. Fougnie, S. Kondragunta, R. Engelen, V.-H. Peuch, M. Parrington, N. Bousserez, M. Vazquez Navarro, and A. Agusti-Panareda. Monitoring multiple satellite aerosol optical depth (AOD) products within the Copernicus Atmosphere Monitoring Service (CAMS) data assimilation system. *Atmospheric Chemistry and Physics*, 22(22):14657–14692, 2022. doi: 10.5194/acp-22-14657-2022. URL https://acp.copernicus.org/articles/22/14657/2022/.

A. Geiss, S. J. Silva, and J. C. Hardin. Downscaling atmospheric chemistry simulations with physically consistent deep learning. *Geoscientific Model Development*, 15(17):6677–6694, 2022. doi: 10.5194/gmd-15-6677-2022. URL https://gmd.copernicus.org/articles/15/6677/2022/.

Christian A. Gueymard and Dazhi Yang. Worldwide validation of CAMS and MERRA-2 reanalysis aerosol optical depth products using 15 years of AERONET observations. *Atmospheric Environment*, 225:117216, 2020. ISSN 1352-2310. doi: https://doi.org/10.1016/j.atmosenv.2019.117216. URL https://www.sciencedirect.com/science/article/pii/S1352231019308556.

Hoshin V. Gupta, Harald Kling, Koray K. Yilmaz, and Guillermo F. Martinez. Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling. *Journal of Hydrology*, 377(1):80–91, 2009. ISSN 0022-1694. doi: https://doi.org/10.1016/j.jhydrol.2009.08.003. URL https://www.sciencedirect.com/science/article/pii/S0022169409004843.

Jana Handschuh, Thilo Erbertseder, Martijn Schaap, and Frank Baier. Estimating PM2.5 surface concentrations from AOD: A combination of SLSTR and MODIS. *Remote Sensing Applications: Society and Environment*, 26:100716, 2022. ISSN 2352-9385. doi: https://doi.org/10.1016/j.rsase.2022.100716. URL https://www.sciencedirect.com/science/article/pii/S2352938522000246.

Paula Harder. Generating physically-consistent high-resolution climate data with hard-constrained neural networks. In *AAAI 2022 Fall Symposium: The Role of AI in Responding to Climate Challenges*, 2022. URL https://www.climatechange.ai/papers/aaaifss2022/5.

Paula Harder, Venkatesh Ramesh, Alex Hernandez-Garcia, Qidong Yang, Prasanna Sattigeri, Daniela Szwarcman, Campbell Watson, and David Rolnick. Physics-Constrained Deep Learning for Climate Downscaling, 2023.

Lijie He, Aiwen Lin, Xinxin Chen, Hao Zhou, Zhigao Zhou, and Peipei He. Assessment of MERRA-2 Surface PM2.5 over the Yangtze River Basin: Ground-based Verification, Spatiotemporal Distribution and Meteorological Dependence. *Remote Sensing*, 11(4), 2019. ISSN 2072-4292. doi: 10.3390/rs11040460. URL https://www.mdpi.com/2072-4292/11/4/460.

C. E. Iles, R. Vautard, J. Strachan, S. Joussaume, B. R. Eggen, and C. D. Hewitt. The benefits of increasing resolution in global and regional climate simulations for European climate extremes. *Geoscientific Model Development*, 13(11):5583–5607, 2020. doi: 10.5194/gmd-13-5583-2020. URL https://gmd.copernicus.org/articles/13/5583/2020/.

R. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6):1153–1160, 1981. doi: 10.1109/TASSP.1981.1163711.

Sasha Khomenko, Marta Cirach, Evelise Pereira Barboza, Natalie Mueller, Jose Barrera-Gómez, David Rojas-Rueda, Kees de Hoogh, Gerard Hoek, and Mark Nieuwenhuijsen. Premature mortality due to air pollution in European cities: a health impact assessment. *The Lancet Planetary Health*, 5, 01 2021. doi: 10.1016/S2542-5196(20)30272-2.

Behzad Kianian, Yang Liu, and Howard H. Chang. Imputing Satellite-Derived Aerosol Optical Depth Using a Multi-Resolution Spatial Model and Random Forest for PM2.5 Prediction. *Remote Sensing*, 13(1), 2021. ISSN 2072-4292. doi: 10.3390/rs13010126. URL https://www.mdpi.com/2072-4292/13/1/126.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, 2017.

Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, 2017.

Robert C. Levy, Lorraine A. Remer, Shana Mattoo, Eric F. Vermote, and Yoram J. Kaufman. Second-generation operational algorithm: Retrieval of aerosol properties over land from inversion of Moderate Resolution Imaging Spectroradiometer spectral reflectance. *Journal of Geophysical Research: Atmospheres*, 112(D13), 2007. doi: https://doi.org/10.1029/2006JD007811. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2006JD007811.

Lianfa Li, Meredith Franklin, Mariam Girguis, Frederick Lurmann, Jun Wu, Nathan Pavlovic, Carrie Breton, Frank Gilliland, and Rima Habre. Spatiotemporal imputation of MAIAC AOD using deep learning with downscaling. *Remote Sensing of Environment*, 237:111584, 2020. ISSN 0034-4257. doi: https://doi.org/10.1016/j.rse.2019.111584. URL https://www.sciencedirect.com/science/article/pii/S0034425719306042.

Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced Deep Residual Networks for Single Image Super-Resolution, 2017.

Lu Liu, Wei Huang, Cheng Wang, Xuepan Zhang, and Bo Liu. SAR image super-resolution based on TV-regularization using gradient profile prior. In *2016 CIE International Conference on Radar (RADAR)*, pp. 1–4, 2016. doi: 10.1109/RADAR.2016.8059210.

Shuang Liu, Chengyi Xiong, Xiaodi Shi, and Zhirong Gao. Progressive face super-resolution with cascaded recurrent convolutional network. *Neurocomputing*, 449:357–367, 2021. ISSN 0925-2312. doi: https://doi.org/10.1016/j.neucom.2021.03.124. URL https://www.sciencedirect.com/science/article/pii/S0925231221005166.

MODIS Atmosphere Science Team. MODIS/Terra Aerosol 5-Min L2 Swath 10km, 2017. URL https://ladsweb.modaps.eosdis.nasa.gov/missions-and-measurements/products/MOD04_L2.

Bin Mu, Bo Qin, Shijin Yuan, and Xiaoyun Qin. A climate downscaling deep learning model considering the multiscale spatial correlations and chaos of meteorological events, Nov 2020. URL https://www.hindawi.com/journals/mpe/2020/7897824/.

S.K. Mukkavilli, A.A. Prasad, R.A. Taylor, J. Huang, R.M. Mitchell, A. Troccoli, and M.J. Kay. Assessment of atmospheric aerosols from two reanalysis products over Australia. *Atmospheric Research*, 215:149–164, 2019. ISSN 0169-8095. doi: https://doi.org/10.1016/j.atmosres.2018.08.026. URL https://www.sciencedirect.com/science/article/pii/S0169809518301236.

Swagata Payra, Ajay Sharma, Manoj Mishra, and Sunita Verma. Performance evaluation of MODIS and VIIRS satellite AOD products over the Indian subcontinent. *Frontiers in Environmental Science*, 11, 06 2023. doi: 10.3389/fenvs.2023.1158641.

J. I. Rubin and W. D. Collins. Global simulations of aerosol amount and size using MODIS observations assimilated with an Ensemble Kalman Filter. *Journal of Geophysical Research: Atmospheres*, 119(22):12,780–12,806, 2014. doi: https://doi.org/10.1002/2014JD021627. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JD021627.

Hanyu Shi, Zhiqiang Xiao, Xuchen Zhan, Han Ma, and Xiaodan Tian. Evaluation of MODIS and two reanalysis aerosol optical depth products over AERONET sites. *Atmospheric Research*, 220:75–80, 2019. ISSN 0169-8095. doi: https://doi.org/10.1016/j.atmosres.2019.01.009. URL https://www.sciencedirect.com/science/article/pii/S0169809518314674.

Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P. Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network, 2016.

Assaf Shocher, Nadav Cohen, and Michal Irani. "Zero-Shot" Super-Resolution using Deep Internal Learning, 2017.

A. Smirnov, B.N. Holben, T.F. Eck, O. Dubovik, and I. Slutsker. Cloud-Screening and Quality Control Algorithms for the AERONET Database. *Remote Sensing of Environment*, 73(3):337–349, 2000. ISSN 0034-4257. doi: https://doi.org/10.1016/S0034-4257(00)00109-7. URL https://www.sciencedirect.com/science/article/pii/S0034425700001097.

Zijue Song, Disong Fu, Xiaoling Zhang, Yunfei Wu, Xiangao Xia, Jianxin He, Xinlei Han, Renjian Zhang, and Huizheng Che. Diurnal and seasonal variability of PM2.5 and AOD in North China plain: Comparison of MERRA-2 products and ground measurements. *Atmospheric Environment*, 191:70–78, 2018. ISSN 1352-2310. doi: https://doi.org/10.1016/j.atmosenv.2018.08.012. URL https://www.sciencedirect.com/science/article/pii/S1352231018305296.

B. Stevens, S. Fiedler, S. Kinne, K. Peters, S. Rast, J. Müsse, S. J. Smith, and T. Mauritsen. MACv2-SP: a parameterization of anthropogenic aerosol optical properties and an associated Twomey effect for use in CMIP6. *Geoscientific Model Development*, 10(1):433–452, 2017. doi: 10.5194/gmd-10-433-2017. URL https://gmd.copernicus.org/articles/10/433/2017/.

Tianning Su, Jing Li, Chengcai Li, Alexis Kai-Hon Lau, Dongwei Yang, and Chuanyang Shen. An intercomparison of AOD-converted PM2.5 concentrations using different approaches for estimating aerosol vertical distribution. *Atmospheric Environment*, 166:531–542, 2017. ISSN 1352-2310. doi: https://doi.org/10.1016/j.atmosenv.2017.07.054. URL https://www.sciencedirect.com/science/article/pii/S1352231017305034.

Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew R. Michaelis, Ramakrishna R. Nemani, and Auroop R. Ganguly. DeepSD: Generating High Resolution Climate Change Projections through Single Image Super-Resolution. *CoRR*, abs/1703.03126, 2017. URL http://arxiv.org/abs/1703.03126.

Jun Wang and Sundar Christopher. Intercomparison between satellite-derived aerosol optical thickness and PM 2.5 mass: Implications for air quality studies. *Geophysical Research Letters*, 30, 01 2003.

Menglin Wang, Meredith Franklin, and Lianfa Li. Generating Fine-Scale Aerosol Data through Downscaling with an Artificial Neural Network Enhanced with Transfer Learning. *Atmosphere*, 13(2), 2022. ISSN 2073-4433. doi: 10.3390/atmos13020255. URL https://www.mdpi.com/2073-4433/13/2/255.

Wei Wang, Zengxin Pan, Feiyue Mao, Wei Gong, and Longjiao Shen. Evaluation of VIIRS Land Aerosol Model Selection with AERONET Measurements. *International Journal of Environmental Research and Public Health*, 14(9), 2017. ISSN 1660-4601. doi: 10.3390/ijerph14091016. URL https://www.mdpi.com/1660-4601/14/9/1016.

Jing Yang and Maogui Hu. Filling the missing data gaps of daily MODIS AOD using spatiotemporal interpolation. *Science of The Total Environment*, 633:677–683, 2018. ISSN 0048-9697. doi: https://doi.org/10.1016/j.scitotenv.2018.03.202. URL https://www.sciencedirect.com/science/article/pii/S0048969718309616.

Qianqian Yang, Qiangqiang Yuan, Linwei Yue, Tongwen Li, Huanfeng Shen, and Liangpei Zhang. The relationships between PM2.5 and aerosol optical depth (AOD) in mainland China: About and behind the spatio-temporal variations. *Environmental Pollution*, 248:526–535, 2019. ISSN 0269-7491. doi: https://doi.org/10.1016/j.envpol.2019.02.071. URL https://www.sciencedirect.com/science/article/pii/S0269749118345238.

Kaibing Zhang, Xinbo Gao, Dacheng Tao, and Xuelong Li. Single Image Super-Resolution With Non-Local Means and Steering Kernel Regression. *IEEE Transactions on Image Processing*, 21 (11):4544–4556, 2012. doi: 10.1109/TIP.2012.2208977.

Lei Zhang and Xiaolin Wu. An Edge-Guided Image Interpolation Algorithm via Directional Filtering and Data Fusion. *IEEE Transactions on Image Processing*, 15(8):2226–2238, 2006. doi: 10.1109/TIP.2006.877407.

Tao Zhang, Yuyu Zhou, Kaiguang Zhao, Zhengyuan Zhu, G. Asrar, and Zhao Xia. Gap-filling MODIS daily aerosol optical depth products by developing a spatiotemporal fitting algorithm. *GIScience Remote Sensing*, 59:762–781, 04 2022a. doi: 10.1080/15481603.2022.2060596.

Tao Zhang, Yuyu Zhou, Kaiguang Zhao, Zhengyuan Zhu, Ghassem R. Asrar, and Xia Zhao. Gap-filling MODIS daily aerosol optical depth products by developing a spatiotemporal fitting algorithm. *GIScience & Remote Sensing*, 59(1):762–781, 2022b. doi: 10.1080/15481603.2022.2060596. URL https://doi.org/10.1080/15481603.2022.2060596.

SUPPLEMENTARY MATERIAL

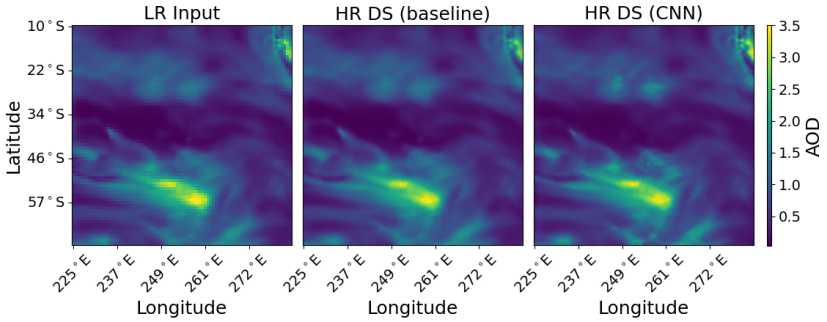

Figure 2: Example outputs, from baseline and ResNet (CNN), on a CAMS image.

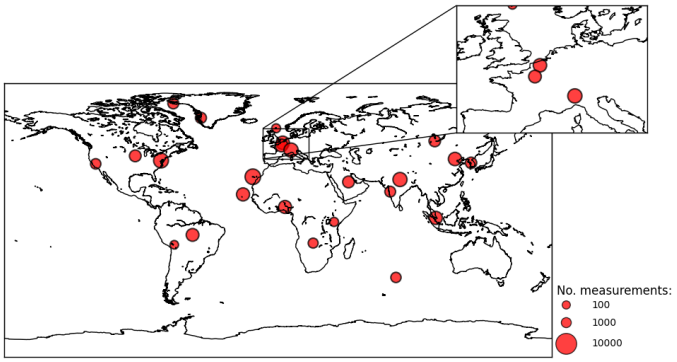

Figure 3: The AERONET stations selected for evaluation; the number of measurements obtained from each is proportional to its area. A list of stations and their details can be found below.

Table 2: List of AERONET stations used in this study.

Station	Coordinates		Elevation (m)
	Latitude	Longitude	
Alta_Floresta	-56.1	-9.9	277.0
Ames	-93.8	42.0	338.0
Amsterdam_Island	77.6	-37.8	30.0
Beijing	116.4	39.9	92.0
Capo_Verde	-22.9	16.7	60.0
Fresno	-119.8	36.8	0.0
GSFC	-76.8	38.9	87.0
Ilorin	4.3	8.3	350.0
Irkutsk	103.1	51.8	670.0
Ispra	8.6	45.8	235.0
Izana	-16.5	28.3	2391.0
Kangerlussuaq	-50.6	66.9	320.0
Kanpur	80.2	26.5	123.0
Lerwick_MO	-1.2	60.1	82.0
Lille	3.1	50.6	60.0
Mongu	23.2	-15.3	1107.0
Mount_Chacaltaya	-68.1	-16.4	5233.0
Nairobi	36.9	-1.3	1650.0
Paris	2.3	48.9	50.0
Pune	73.8	18.5	559.0
Seoul_SNU	126.9	37.5	116.0
Singapore	103.8	1.3	30.0
Solar_Village	46.4	24.9	764.0
Thule	-68.8	76.5	225.0