

APPLICATION PAPER  

# Pollution tracker: Finding industrial sources of aerosol emission in satellite imagery

Peter Manshausen<sup>1</sup> , Duncan Watson-Parris<sup>1,2</sup> , Lena Wagner<sup>3</sup>, Pirmin Maier<sup>3</sup>, Sybrand J. Muller<sup>3</sup>, Gernot Ramming<sup>3</sup> and Philip Stier<sup>1</sup>

<sup>1</sup>Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, Oxford, United Kingdom

<sup>2</sup>Scripps Institution of Oceanography and Halicioğlu Data Science Institute, University of California San Diego, La Jolla, CA, USA

<sup>3</sup>GAF AG, Munich, Germany

**Corresponding author:** Peter Manshausen; Email: [peter.manshausen@physics.ox.ac.uk](mailto:peter.manshausen@physics.ox.ac.uk)

**Received:** 09 May 2023; **Accepted:** 30 May 2023

**Keywords:** aerosols; air pollution; classification; clouds; computer vision; remote sensing

## Abstract

The effects of anthropogenic aerosol, solid or liquid particles suspended in the air, are the biggest contributor to uncertainty in current climate perturbations. Heavy industry sites, such as coal power plants and steel manufacturers, large sources of greenhouse gases, also emit large amounts of aerosol in a small area. This makes them ideal places to study aerosol interactions with radiation and clouds. However, existing data sets of heavy industry locations are either not public, or suffer from reporting gaps. Here, we develop a supervised deep learning algorithm to detect unreported industry sites in high-resolution satellite data, using the existing data sets for training. For the pipeline to be viable at global scale, we employ a two-step approach. The first step uses 10 m resolution data, which is scanned for potential industry sites, before using 1.2 m resolution images to confirm or reject detections. On held-out test data, the models perform well, with the lower resolution one reaching up to 94% accuracy. Deployed to a large test region, the first stage model yields many false positive detections. The second stage, higher resolution model shows promising results at filtering these out, while keeping the true positives, improving the precision to 42% overall, so that human review becomes feasible. In the deployment area, we find five new heavy industry sites which were not in the training data. This demonstrates that the approach can be used to complement existing data sets of heavy industry sites.

## Impact Statement

Aerosols from industrial air pollution play an important role in both public health and climate. Complete and freely accessible data on large aerosol sources are of great interest, but not currently available. Here, we present a scoping study to fill the gaps in existing data sets by automatically detecting heavy industry in satellite data. This approach can be rolled out globally to complete the available data.

## 1. Climate Change and the Need for Pollution Source Databases

Anthropogenic aerosol, such as black carbon or sulfate, plays an important role for the earth's energy budget. Aerosol can interact directly with solar radiation by absorbing or scattering light. More subtly, it

  This research article was awarded Open Data and Open Materials badges for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

can change the properties of clouds to make them more reflective by providing additional condensation nuclei and increasing the number of cloud droplets, which in turn sets off a series of further adjustments. Together, aerosol-induced changes in cloud properties are the largest source of uncertainty in current estimates of man-made climate forcing (Masson-Delmotte *et al.*, 2021). This is mainly because the different effects are difficult to untangle and experiments are infeasible in nature on the large scales required.

Research has therefore relied heavily on so-called opportunistic experiments (Christensen *et al.*, 2021), where an aerosol source is known and localized, such that polluted and unpolluted clouds can be directly compared. Examples of such sources include volcanoes (Malavelle *et al.*, 2017), ships (Coakley and Walsh, 2002), and heavy industry (Toll *et al.*, 2019). In the past, studies of these sources have been limited to some hundreds of “events,” such as a ship or heavy industry site producing a visible change in clouds. Research is ongoing to enlarge such data sets, for example, by automatically detecting such cloud changes caused by ships with machine learning (Watson-Parris *et al.*, 2022). Another approach is to use more information about the sources themselves: Manshausen *et al.* (2022) use millions of ship paths and reconstruct where their pollution has changed cloud properties, uncovering previously undetected changes in cloud water. A similar approach seems promising for large industrial sites.

However, existing data such as that collected by the International Energy Agency or by private companies is not publicly available. The databases that are available such as the EDGAR data do not include point sources outside Europe and North America (Janssens-Maenhout *et al.*, 2015). Here, we build on two openly accessible databases of point sources to fill in such gaps (see Section 2).

We also note that there is a growing interest in independent estimates of greenhouse gas emissions using satellite observations, and therefore in heavy industry sites as large emitters. Among others, [www.climate TRACE.org](http://www.climate TRACE.org) and its partners are accounting for global emissions. Recently, large point sources of methane have also been discovered from satellite imagery (Lauvaux *et al.*, 2022).

This project is informed by recent advances in detecting large industrial sites and land cover classification: Sheng *et al.* (2020) use a deep neural network approach to identify oil refineries in the United States. However, their work uses satellite data at relatively high resolution of 2.5 m, which requires handling large amounts of data for continental-scale detection. For land cover classification, Sumbul *et al.* (2021) propose “BigEarthNet,” a deep neural network which classifies 10–60 m resolution Sentinel2 images into 19 land use classes. However, their “industrial and commercial units”-class is the one that performs second to worst. If the locations of power plants are known, Mommert *et al.* (2021) have shown that ResNet architectures are skillful at identifying their type, for example, distinguishing nuclear, gas, and coal. With known locations, Hanna *et al.* (2021) propose a method to then estimate power plant CO<sub>2</sub> emissions.

Our work addresses the gap of detecting heavy industry in Sentinel2 data, with the aim of developing a method that can, in principle, be deployed globally. To this end, we propose a two-step approach in which we first perform classification on 10 m resolution data and then additional filtering of false positives at 1.2 m resolution, using Bing Maps (Microsoft Corporation, 2023).

## 2. Data

This work aims to extend some of the existing and freely available data sets of heavy industry. Namely, we base the project on the location data from Global Energy Monitor (2022a, 2022b) and from McCarten *et al.* (2021) for coal and steel plants worldwide. These data sets include over 4,000 samples of heavy industry sites. To include nonindustry sites, we add the same number of random land cover patches by choosing the patch at a fixed distance of 0.1° longitude to the east of each industry site center.

We download Sentinel2 data for the 10 m resolution bands 02, 03, 04, and 08 (i.e., RGB and near-infrared), in patches of 240 × 240 pixels centered around the locations from the above databases. Querying the database for cloud-free Sentinel2 overpasses, we use the date range from March to October 2021 and take the output with the lowest cloud cover in the scene and in the patch. The data is normalized to have

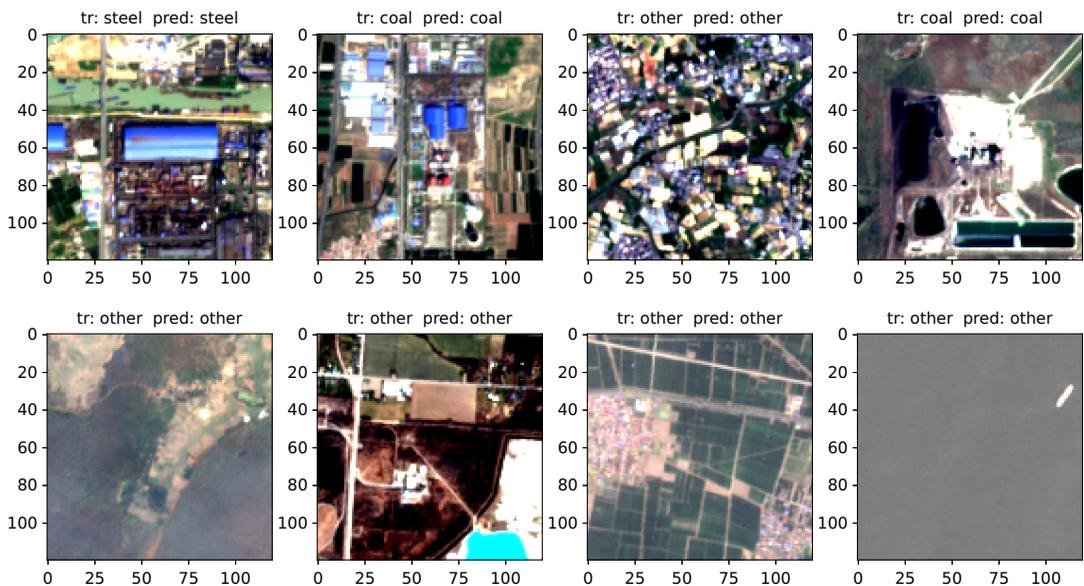
zero mean and unit standard deviation in each channel. Leaving out the NIR band slightly reduces the performance during training, so all models include it.

As the Sentinel cloud mask is not completely accurate, we iterate over the data set twice, training only on half the data in each case. The models obtained that way are used to predict the labels of the held-out data. We review the most mismatched prediction/ground truth pairs by eye, which in some cases are cloud-covered, or do not show any industry (errors in the site data). This leaves a data set of 7,829 patches, with 28% of coal, 22% of steel, and 50% of nonindustry samples. For training, we use 80% of the data with 10% held out for testing and 10% for validation. We employ data augmentation using random rotation, random crop, and random flip. We find that we obtain the best results with a center crop to  $140 \times 140$  pixels and then a random crop to the target  $120 \times 120$  pixel patches (the same size as those used by Sumbul et al., 2021). This cropping is in order to vary where the industry plant is with respect to the scene. It should always be at the center of the  $240 \times 240$  patch, then center and random crops allow it to be offset from the center in the final  $120 \times 120$  training patch. Each patch is assigned to one of three classes coal, steel, other/no industry. Examples of what this data looks like in the three visible channels can be seen in Figure 1.

For the second step, we download Bing Maps (Microsoft Corporation, 2023) aerial RGB imagery in the same locations as discussed above, and filter in the same way. The resolution here is 1.2 m, and we limit the image size to 1,400 pixels, giving an on-the-ground size of just below 1,200 m, the same as the size of the Sentinel2 images. Splitting and augmentation are as above.

### 3. Models and Training

For step one, following Sumbul et al. (2021), we initially try a number of CNN architectures, trained with categorical crossentropy loss. These are: ResNet50, VGG16, EfficientNet, InceptionNet, and ResNet50v2. For each of the architectures, we add a dense layer of three neurons at the end of the network for the three classes. We find that among the models, ResNet50 (the original, and v2, He et al.,



**Figure 1.** Some example visualizations the RGB channels of the Sentinel2 patches in the data set. For a human, characteristic features of steel plants are, for example, the dark surface together with industrial structures like in the top first panel. For coal plants, cooling towers are common, note the bright crescent shape in the top second panel, bottom right corner. Labels show truth (tr.) and prediction (pred.) of our best model.

2016) converge the quickest. Compared to ResNet50, ResNet50v2 reaches slightly higher categorical accuracy. We train ResNet50v2 for up to 200 epochs, using early stopping after 50 epochs of no decrease in validation loss. We start training with a learning rate of  $10^{-3}$  and decay with  $e^{-0.03}$  each epoch after 30 epochs.

For step two on high-resolution data, we train the same ResNet50v2 architecture, using early stopping and a decaying learning rate as above. In both cases, training from scratch performs better than using pretrained weights from imagenet (Deng et al., 2009).

## 4. Results

### 4.1. Step 1: Lower resolution

The Sentinel2 model performs well on the held-out data, with precision values of up to 0.96 for coal plants. Table 1 shows performance for all classes in terms of precision, recall, and F1-score. It seems that generally steel plants are harder to detect (lower recall), as they are often not as visually distinctive as coal plants. Figure 1 shows a random selection of predictions on the test data together with the ground truth. Supplementary Figure A3 shows a selection of patches that were misclassified by the model.

### 4.2. Step 2: Higher resolution

In the higher resolution case, the model does not perform as well on the test data, as shown in Table 1. This is surprising given the larger level of detail in the data. The model performs worst on steel plants, with the confusion matrix, shown in Table 2 indicating that it mistakes many of them for either coal plants or no-industry sites—of the 23% steel plants, just over half are correctly identified (compare also the recall of

**Table 1.** Performance evaluation of the best model (ResNet50v2) on the held-out test data for both the lower and higher resolution cases

Category	Lower resolution			Higher resolution		
	Precision	Recall	F1-score	Precision	Recall	F1-score
Coal	0.96	0.95	0.95	0.72	0.82	0.76
Steel	0.88	0.86	0.87	0.84	0.53	0.65
Other	0.95	0.97	0.96	0.85	0.93	0.89
Weight Avg.	0.94	0.94	0.94	0.76	0.77	0.76

*Note.* Overall accuracy of the top-1 prediction is 0.94 for the lower and 0.77 for the higher resolution case.

**Table 2.** High resolution model confusion matrix, normalized by the number of samples

Truth	Prediction		
	Coal	Steel	Other
Coal	0.20	0.02	0.03
Steel	0.06	0.12	0.05
Other	0.03	0.01	0.48

*Note.* Each line shows the fraction of samples belonging to a category that was predicted as each of the categories, that is, the first line means out of all patches, 20% are coal and were identified as coal, 2% are coal and were identified as steel, and 3% are coal and were identified as other. This is evaluated on the held-out test data.

53% in Table 1). The overall lower skill may be due to insufficient training samples to learn high resolution properties, different hyperparameters being needed for the high resolution case, or the model making more cautious predictions with a higher spread of probabilities. A promising result is that the model has high recall on the “other” category, which is crucial for using it as a second-stage filter on the positive predictions of the low resolution stage.

### 4.3. Large scale deployment

In a deployment test, we deploy our first stage model to a  $4^\circ \times 4^\circ$  area centered around ( $21^\circ\text{N}$ ,  $82^\circ\text{E}$ ) in north-eastern India, focusing on coal plants. Checking the results manually, we find that the model finds some coal plants that are present in the region, but also returns many false positives even at high probability thresholds. This is most likely due to the more imbalanced real-world data, with a large fraction of no-industry tiles. A similar behavior was also observed by Sheng et al. (2020).

Supplementary Figure A4 shows some example positive predictions for coal. For some false positives, we can guess why the model confuses them for coal plants. For instance, they are often situated near rivers (a), produce small clouds (b) near the cooling towers, show built structures that may look similar to urban fabric (c), or are close to open-air mines (d) or coal piles. The figure also shows two true positives (e), with the characteristic dark surface, industrial structures, cooling tower, and river (f).

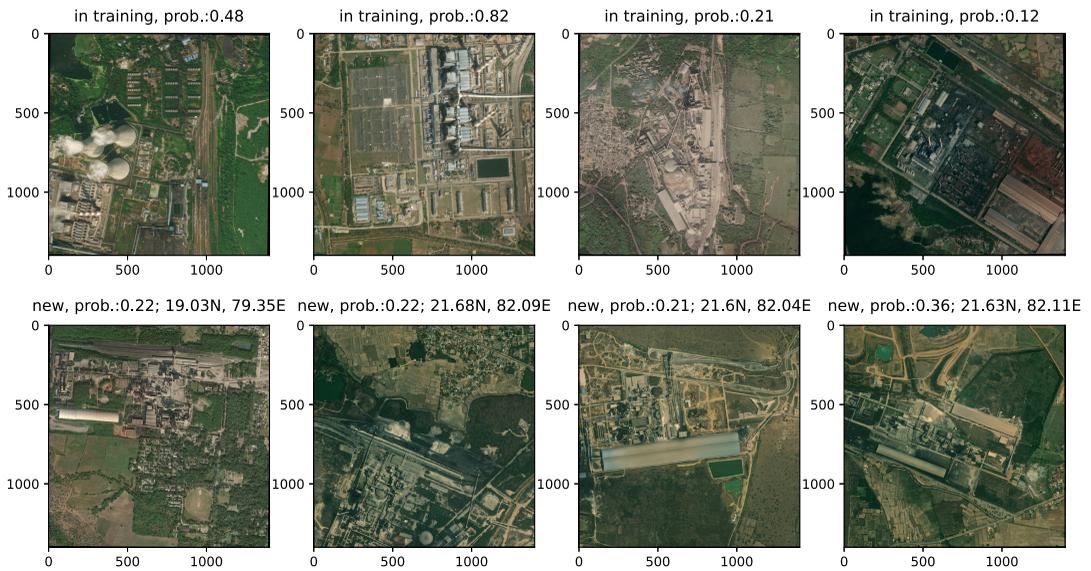
More quantitatively, of the 51 coal power plants in the study area which are present in the training data set, the model correctly identifies 24 (probability of  $>50\%$  for coal). This seems low (recall of 0.47) given the recall of 0.95 on the testing data (cf. Table 1). This may be linked to poorer data quality in the deployment set, with sites off-centered, cut in half, or covered by cloud. The model makes 890 positive predictions in total, most of which are false positives (giving a precision of only 0.03, much lower than on the testing data, because of the imbalanced deployment data with many more “other” patches). This is where the second stage, higher resolution model comes in: We download the high resolution Bing Maps scenes in the locations of the 890 positives from the first-stage and make predictions on them with the second stage model. Even though the accuracy of the high resolution model is lower on the test data, it performs well at this confirmation/rejection task: Setting a threshold of 10% probability for coal confirms 23 out of the 24 known coal plants (recall of 0.96), while reducing the 890 potential positives to 55 (precision of 0.42).

Reviewing the remaining positives, we find five additional coal plants, which were not present in the training, validation, or test set. Four of these, together with four of the 23 that were in the data sets, are shown in Figure 2. The five detected sites are cement production facilities (confirmed by an internet search of the location), which form a complex with coal power plants for their large energy consumption. (The fifth site, not shown in the figure is at  $21.37\text{N}$ ,  $83.61\text{E}$ .) Two additional metal processing sites are detected which likely also use coal power. These new sites were probably identified from the architectural or surface elements that are there because they are coal-powered, such as coal piles, smokestacks, and polluted, dark surfaces. The above values for precision and recall after each step were given without taking into account the new sites, and are therefore a conservative estimate of the actual values.

This demonstrates that our approach can be used to fill in gaps in existing data sets. At the same time, it raises the question of how to treat complex industrial facilities which combine coal power with manufacturing of products like steel or cement.

## 5. Discussion and Future Work

The above results depend strongly on the choice of threshold values for detection in both stages of the algorithm. To illustrate the problem, true and false positive rates of the first and second-stage coal plant detections are shown in Supplementary Figure A5. In the first stage, the priority is a low false positive rate, as we want to download as little high-resolution data as possible for the second stage. In order not to miss too many true positives, we chose a threshold of 50% here.



**Figure 2.** RGB images of the second stage higher resolution predictions in the deployment region. Scenes that the first stage classed as  $>50\%$  probability of coal plant are downloaded at higher resolution from Bing Maps and fed into the model. The second stage detects 23 out of 24 coal plants that were present in the training set (top row, labeled “in training,” and with probability output). It also detects five “new” coal plants not in the data set from Global Energy Monitor, four of which are shown in the bottom row with their respective locations.

For the second stage, the priority is a high true positive rate, so that industry sites present in the high-resolution data are not missed. This is why we chose a (relatively low) threshold of 10% in this study, which still allows to filter out many false positives, as shown by the steep decline in false positives in [Supplementary Figure A5](#). Ultimately, these thresholds should be chosen depending on computational resources (imposing a lower limit on the first stage threshold, as a low threshold will mean many positives that need to be verified in the second stage), as well as the resources for human review (imposing a lower limit on the second stage threshold, as a low threshold will mean many positives that need to be verified by eye).

We would like to further improve step one and lower its false positive rate. For this, we propose: (a) Deploying on two or three Sentinel2 overpass dates for each tile, which would presumably eliminate the “small cloud” false positives. (b) Using human-in-the-loop learning, that is, retraining the model with the correctly labeled false positives from the deployment stage after human review, or just increasing the number of nonindustry cases in the training data. (c) Rerunning on shifted tiles to exclude the case where a site is cut in half and therefore not recognized.

This last point is connected to the cropping hyperparameter choices: We start with  $240 \times 240$  pixel scenes, then center crop to  $140 \times 140$ , and then random crop to  $120 \times 120$ . While the initial and final size are modeling choices, the center crop size should be seen as a hyperparameter.  $140 \times 140$  pixels is the best choice for training and then testing on the data. It balances variation in where the site is in the final scene with ensuring enough of it is present for the model to learn important features. However, for the deployment, a less conservative choice, that is, a larger value than 140, which will train the network with more off-center scenes, could do better on the random positions that it will be deployed to. This could potentially improve the model’s false negative rate.

To improve step two, in order to achieve accuracy comparable to the low-resolution first step, independent hyperparameter tuning seems promising. Ultimately, the goal is to roll out this approach globally, and to fill the gaps in existing data sets. To do this, careful evaluation of detection thresholds at both stages is needed, balancing the threshold’s effect on false positive and false negative detections.

## 6. Conclusions

There is a need for publicly available databases of locations of heavy industry, which act as large point sources of aerosol and greenhouse gas pollution. We show that heavy industry complexes can be detected in globally available 10 m-resolution satellite data. ResNet architectures are particularly well adapted to this task and reach high accuracy for both steel and coal plants. However, this approach has many false positive detections at deployment. This can be remedied by a second step, which makes use of the available high-resolution data, but only in the locations where detection seems likely based on low-resolution data. Our approach combines the advantages of being not too data-intensive, not yielding too many false positives (high precision), and requiring minimal human review of detections. Furthermore, the method succeeds at finding new heavy industry sites not present in the existing data sets it was trained on. Finally, these efforts may contribute to the larger effort of independently and bottom-up quantifying not only air pollution, but also greenhouse gas emissions.

**Acknowledgments.** We would like to acknowledge the important work done by Global Energy Monitor which provided training data for this project. We thank Lilli Freischem for helpful comments on the manuscript.

**Author contribution.** Conceptualization: P. Man, D.W.-P., G.R., P. Mai., P.S.; Data curation: P. Man., L.W., S.J.M.; Data visualization: P. Man.; Methodology: All authors; Model training and analysis: P. Man.; Writing—original draft: P. Man. All authors approved the final submitted draft.

**Competing interest.** The authors declare no competing interests exist.

**Data availability statement.** Replication code can be found on GitHub: [https://github.com/ManshaP/sentinel\\_industry](https://github.com/ManshaP/sentinel_industry). It has also been archived in Zenodo at DOI: [10.5281/zenodo.7991163](https://doi.org/10.5281/zenodo.7991163). Sentinel2 data is freely available and can be downloaded with the sentinelhub API. Bing Maps is available noncommercially and was downloaded using the code referenced above.

**Ethics statement.** The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

**Funding statement.** This work was supported by the European Union's Horizon 2020 research and innovation program under Marie Skłodowska-Curie grant agreement No. 860100 (iMIRACL). D.W.-P. and P.S. were supported by the UK Natural Environment Research Council project ACRUISE (NE/S005099/1). P.S. additionally acknowledges support from the European Research Council Project RECAP under the European Union's Horizon 2020 research and innovation program grant 724602 and from the FORCeS project under the European Union's Horizon 2020 research program with grant agreement 821205. P. Man acknowledges GAF AG for hosting a three-month research visit to their Munich office.

**Provenance statement.** This article is part of the Climate Informatics 2023 proceedings and was accepted in *Environmental Data Science* on the basis of the Climate Informatics peer review process.

**Supplementary material.** The supplementary material for this article can be found at <https://doi.org/10.1017/eds.2023.20>.

## References

- Christensen M, Gettelman A, Cermak J, Dagan G, Diamond M, Douglas A, Feingold G, Glassmeier F, Goren T and Grosvenor DP (2021) Opportunistic experiments to constrain aerosol effective radiative forcing. *Atmospheric Chemistry and Physics* 22, 641–674.
- Coakley JA and Walsh CD (2002) Limits to the aerosol indirect radiative effect derived from observations of ship tracks. *Journal of the Atmospheric Sciences* 59(3 PT 2), 668–680. [http://doi.org/10.1175/1520-0469\(2002\)059%3c0668:ltair%3e2.0.co;2](http://doi.org/10.1175/1520-0469(2002)059%3c0668:ltair%3e2.0.co;2)
- Deng J, Wei Dong RS, Li L-J, Li K and Fei-Fei L (2009) ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. Miami, FL: IEEE, pp. 248–255. <http://doi.org/10.1109/cvprw.2009.5206848>
- Global Energy Monitor (2022a) Global Coal Plant Tracker. Available at <https://globalenergymonitor.org/projects/global-coal-plant-tracker/> (accessed 1 July 2022).
- Global Energy Monitor (2022b) Global Steel Plant Tracker. Available at <https://globalenergymonitor.org/projects/global-steel-plant-tracker/> (accessed 1 July 2022).
- Hanna J, Mommert M, Scheibenreif L and Borth D (2021) Multitask learning for estimating power plant greenhouse gas emissions from satellite imagery. In *NeurIPS 2021 Workshop on Tackling Climate Change with Machine Learning*.
- He K, Zhang X, Ren S and Sun J (2016) Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2016. Las Vegas, NV: IEEE, pp. 770–778. <http://doi.org/10.1109/CVPR.2016.90>

- Janssens-Maenhout G, Crippa M, Guizzardi D, Dentener F, Muntean M, Pouliot G, Keating T, Zhang Q, Kurokawa J, Wankmüller R, Van Der Gon HD, Kuenen JJ, Klimont Z, Frost G, Darras S, Koffi B and Li M (2015) HTAP-v2.2: A mosaic of regional and global emission grid maps for 2008 and 2010 to study hemispheric transport of air pollution. *Atmospheric Chemistry and Physics* 15(19), 11411–11432. <http://doi.org/10.5194/acp-15-11411-2015>
- Lauvaux T, Giron C, Mazzolini M, D’Aspremont A, Duren R, Cusworth D, Shindell D and Ciais P (2022) Global assessment of oil and gas methane ultra-emitters. *Science* 375(6580), 557–561. <http://doi.org/10.1126/science.abj4351>
- Malavelle FF, Haywood JM, Jones A, Gettelman A, Clarisse L, Bauduin S, Allan RP, Karset IHH, Kristjánsson JE, Oreopoulos L, Cho N, Lee D, Bellouin N, Boucher O, Grosvenor DP, Carslaw KS, Dhomse S, Mann GW, Schmidt A, Coe H, Hartley ME, Dalvi M, Hill AA, Johnson BT, Johnson CE, Knight JR, O’Connor FM, Stier P, Myhre G, Platnick S, Stephens GL, Takahashi H and Thordarson T (2017) Strong constraints on aerosol-cloud interactions from volcanic eruptions. *Nature* 546(7659), 485–491. <http://doi.org/10.1038/nature22974>
- Manshausen P, Watson-Parris D, Christensen M and Stier P (2022) Invisible ship tracks show large cloud sensitivity to aerosol. *Nature* 610, 101–106. <http://doi.org/10.1038/s41586-022-05122-0>
- Masson-Delmotte V, Zhai P, Pirani A, Connors SL, Péan C, Berger S, Caud N, Chen Y, Goldfarb L, Gomis MI, Huang M, Leitzell K, Lonnoy E, Matthews JBR, Maycock TK, Waterfield T, Yelekçi O, Yu R and Zhou B. (eds) (2021) *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Technical Report, IPCC.
- McCarten M, Bayaraa M, Caldecott B, Christiaen C, Foster P, Hickey C, Kampmann D, Layman C, Rossi C, Scott K, Tang K, Tkachenko N and Yoken D (2021) Global Database of Iron and Steel Production Assets. Spatial Finance Initiative.
- Microsoft Corporation (2023) Bing Maps. Available at <https://learn.microsoft.com/en-us/bingmaps/> (accessed 30 January 2023).
- Mommert M, Scheibenreif L, Hanna J and Borth D (2021) Power plant classification from remote imaging with deep learning. In *International Geoscience and Remote Sensing Symposium (IGARSS)*. Brussels, Belgium: IEEE, pp. 6391–6394. <http://doi.org/10.1109/IGARSS47720.2021.9553219>
- Sheng H, Irvin J, Munukutla S, Zhang S, Cross C, Story K, Rustowicz R, Elsworth C, Yang Z, Omara M, Gautam R, Jackson RB and Ng AY (2020) OGNNet: Towards a Global Oil and Gas Infrastructure Database using Deep Learning on Remotely Sensed Imagery (NeurIPS), pp. 1–8. Available at <http://arxiv.org/abs/2011.07227> (accessed 30 January 2023).
- Sumbul G, De Wall A, Kreuziger T, Marcelino F, Costa H, Benevides P, Caetano M, Demir B and Markl V (2021) BigEarthNet-MM: A large-scale, multimodal, multilabel benchmark archive for remote sensing image classification and retrieval [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine* 9(3), 174–180. <http://doi.org/10.1109/MGRS.2021.3089174>
- Toll V, Christensen M, Quaas J and Bellouin N (2019) Weak average liquid-cloud-water response to anthropogenic aerosols. *Nature* 572(7767), 51–55. <http://doi.org/10.1038/s41586-019-1423-9>
- Watson-Parris D, Christensen MW, Laurensen A, Clewley D, Gryspeerd E and Stier P (2022) Shipping regulations lead to large reduction in cloud perturbations. *Proceedings of the National Academy of Sciences of the United States of America* 119(41), e2206885119. <http://doi.org/10.1073/pnas.2206885119>