

Power Plant Classification from Remote Imaging with Deep Learning

Michael Mommert, Linus Scheibenreif, Joëlle Hanna, Damian Borth
University of St. Gallen, Institute of Computer Science

Presented at the 2021 IEEE International Geoscience and Remote Sensing Symposium (IGARSS).

Abstract

Satellite remote imaging enables the detailed study of land use patterns on a global scale. We investigate the possibility to improve the information content of traditional land use classification by identifying the nature of industrial sites from medium-resolution remote sensing images. In this work, we focus on classifying different types of power plants from Sentinel-2 imaging data. Using a ResNet-50 deep learning model, we are able to achieve a mean accuracy of 90.0% in distinguishing 10 different power plant types and a background class. Furthermore, we are able to identify the cooling mechanisms utilized in thermal power plants with a mean accuracy of 87.5%. Our results enable us to qualitatively investigate the energy mix from Sentinel-2 imaging data, and prove the feasibility to classify industrial sites on a global scale from freely available satellite imagery.

1 Introduction

The availability of satellite remote imaging data with global coverage and high imaging cadence enable the detailed investigation of land use patterns and changes therein. Land use classification creates a mapping of how land is being utilized by humans; such information is highly valuable to better understand socio-economic, environmental, and ecological problems and to inform policy-makers. However, the value of land use information is limited by both the spatial and categorical resolution of the data, the latter referring to the level of detail utilized in the classification or segmentation.

We investigate the use of deep learning applied to satellite remote imaging to identify the nature of industrial sites in an attempt to improve the level of detail of land use classification maps. Previous works have successfully combined both approaches for land use classification [1], to detect solar [2] and wind power facilities [3], as well as oil and gas processing facilities [4]. However, all these works rely on high-resolution remote imaging data, which are typically associated with high cost and have a low imaging cadence. In this work, we investigate the use of freely available medium-resolution and high-cadence satellite imaging data to classify different types of power plants as a test bed for a future application to industrial sites. This application also addresses the topic of climate change mitigation by enabling the qualitative analysis of the energy mix from satellite imagery.

2 Data Set

For this study, we select a total of 450 different power plant sites from the JRC Open Power Plants Database (JRC-PPDB-OPEN, [5]), which in addition to general power plant type information also provides approximate geolocation data for each site. For each power plant class, we extract up to 50 different sites for which we verify and improve geolocation data based on high-resolution remote imaging data provided by the Google Maps service. Our resulting sample contains the following power plant classes (nomenclature as used in JRC-PPDB-OPEN, number in parentheses refer to the number of different sites in the sample): “Fossil Brown coal/Lignite” (50), “Fossil Gas” (50), “Fossil Hard Coal” (50), “Fossil Oil” (30), “Hydro Pumped Storage” (50), “Hydro Run-of-river and poundage” (50), “Hydro Water Reservoir” (50), “Nuclear” (50), “Solar” (20), “Wind Onshore” (50).

Images utilized in this work were acquired with the Multi-Spectral Imagers (MSI) onboard the two Sentinel-2 Earth-observing satellites operated by the European Space Agency. Both satellites observe the land surface of

the Earth in 13 bands covering the optical to short-wave infrared wavelengths at spatial resolutions of 10 m to 60 m per pixel and with an imaging cadence of ≥ 5 days for most locations on Earth.

For each site, we download a total of 10 different and mostly cloud-free Sentinel-2 rasters taken during 2020 to ensure proper coverage of seasonal and other temporal changes in each site. Around each location, we crop a square image region with an edge length of 1 km. We utilize only those MSI imaging channels that do not focus on atmospheric features (bands 2, 3, 4, 5, 6, 7, 8, 8A, 11, and 12) and upsample all bands to the highest image resolution (10 m/pixel). We perform the same extraction procedure on 4 randomly selected locations on each raster, resulting in our background images, which are used as a control group to quantify the detectability of power plants against random settings.

Our data set consists of 4,154 images of 450 different power plant sites across Europe and 16,064 random background images. Example images are shown in Figure 2.

3 Methods

Deep Neural Network approaches are well-suited to deal with the complexity inherent to remote imaging data. We approach the multi-class classification problems in this work with a ResNet-50 deep residual neural network [6] that was slightly modified: the first convolutional layer takes in 10-channel input images (see Section 2) and the final fully connected layer returns a number of logits that is equal to the number of classes of the respective problem. Furthermore, the first convolutional and maxpool layers have been modified to use a 3×3 kernel size and a stride of unity to enforce sensitivity to small-scale structure in our images.

For all experiments we utilize stochastic gradient descent with momentum as our optimizer. The learning is based on a cross entropy loss function. Our deep learning methods are implemented utilizing PyTorch [7].

To boost the learning and results of our classifiers, we pretrain our models on the EuroSAT dataset [8], which contains a sample of Sentinel-2 image patches with corresponding land use classification labels.

We randomly split our data set on a per-image basis into a training, validation, and test data set using a 80%/10%/10% split. Before provided to the model for either training or evaluation, each image is normalized based on the mean and standard deviation of each Sentinel-2/MSI band across all images. Furthermore, random image flipping, mirroring, and rotations by integer multiples of 90 degrees are applied to provide data augmentation. For each epoch during training, validation, and testing, balanced samples are drawn from the corresponding data sets, using either subsampling or oversampling depending on the number of available images in each class.

4 Results

4.1 Power Plant Classification

We train our modified ResNet-50 architecture on the classification task involving 10 different power plant classes (see Section 2) and a background class. After tuning the hyperparameters of the model on the validation data set, we evaluate the model performance on the test data set. We find an overall accuracy over all classes of 90.0%; the accuracy of individual classes varies between 77.4% for the background class and 97.5% for solar power plants. Figure 1 shows the per-class accuracies and corresponding uncertainties in detail. The accuracy variations between power plant types correlate well with the presence or absence of unique feature that would enable a clear identification.

4.2 Power Plant Cooling Classification

Thermal power plants utilize different methods of cooling necessary for the mechanical generation of power from steam. As provided by JRC-PPDB-OPEN, our full data set contains 4 different cooling types among the

Ground Truth	Background	77.4±3.7	0.5±0.7	1.7±0.9	0.1±0.3	0.3±0.5	0.1±0.3	3.2±1.2	4.4±2.1	2.1±1.2	1.7±1.7	8.5±2.5
	Brown Coal	0.2±0.4	97.4±1.1	0.3±0.5	2.0±1.1	0.1±0.3	0.0	0.0	0.0	0.0	0.0	0.0
	Gas	0.3±0.5	0.5±0.5	90.1±2.4	5.7±1.9	1.8±0.6	1.1±0.7	0.1±0.3	0.0	0.0	0.2±0.4	0.2±0.4
	Hard Coal	0.4±0.7	5.3±1.2	4.9±2.0	88.7±2.8	0.6±0.9	0.0	0.1±0.3	0.0	0.0	0.0	0.0
	Oil	1.5±0.8	0.1±0.3	0.9±0.8	0.9±0.8	95.9±1.9	0.3±0.5	0.0	0.0	0.0	0.0	0.4±0.5
	Nuclear	0.1±0.3	0.3±0.5	1.5±1.2	1.7±1.4	0.3±0.5	96.1±2.0	0.0	0.0	0.0	0.0	0.0
	Run-of-River	6.5±1.4	0.1±0.3	0.0	0.3±0.5	0.1±0.3	0.4±0.7	84.4±2.3	5.4±1.8	1.2±0.6	0.0	1.6±0.9
	Reservoir	3.2±1.9	0.1±0.3	0.0	0.1±0.3	0.0	0.0	6.9±2.1	83.6±3.6	4.3±1.6	0.1±0.3	1.7±0.8
	Pumped Storage	3.0±1.7	0.0	0.0	0.0	0.0	0.0	2.7±1.6	7.6±1.8	86.7±2.8	0.0	0.0
	Solar	0.8±0.6	0.0	0.0	0.0	0.0	0.6±0.5	0.5±0.5	0.0	0.0	97.5±0.8	0.6±0.5
	Wind Onshore	5.9±2.3	0.0	0.1±0.3	0.0	0.6±0.5	0.0	0.7±0.8	0.5±0.7	0.0	0.0	92.2±2.4
Prediction Results		Background	Brown Coal	Gas	Hard Coal	Oil	Nuclear	Run-of-River	Reservoir	Pumped Storage	Solar	Wind Onshore



Figure 1: Confusion matrix for the classification of 10 different power plant classes (classes “Brown Coal”, “Gas”, “Hard Coal”, and “Oil” refer to fossil fuel power plants, whereas “Run-of-River”, “Reservoir”, and “Pumped Storage” refer to hydroelectric plants) and one background class. Values are accuracies expressed as percentages. Uncertainties, derived as standard deviations across ten different randomly drawn balanced samples from the test data set, are quoted where significant. While most classes have high individual detection accuracies, there are chances of confusion among the different hydroelectric classes (see example images on the right) and the fossil fuel plant classes, as well as between the renewable energy classes and the background class.

thermal power plants (all fossil fuel plants and nuclear plants): air cooling (29 examples), mechanical draft tower (22 examples), natural draft tower (84 examples), once-through cooling (90 examples). We ignore 5 examples of thermal power plants that do not utilize active or passive cooling. The full data set of thermal power plants includes 2,112 images.

We investigate the possibility to also classify the type of cooling used with the same ResNet-50 architecture as outlined in Section 3. Utilizing the same approach as introduced above, we train the model on the reduced data set of thermal power plants and train against the 4 different target classes (see above). We utilize the trained model from Section 4.1 as a pretrained model for this task.

Our trained model achieves a mean accuracy of 87.5% over all cooling classes. The individual class accuracies range from 70.0% for mechanical draft towers to 99% for once-through cooling. This range is easily explainable as once-through cooling systems require access to a water body that is readily detectable, whereas mechanic draft towers can resemble other buildings and structures.

5 Discussion and Conclusions

In order to verify that predictions made by our architecture are based on salient features relevant to the classification task, we utilize Class Activation Maps (CAM, [9]). CAMs highlight the most important image areas in the model’s classification process. Examples, shown in Figure 2, show that CAMs indeed focus on relevant image regions. For instance, coal-powered plants are identified based on the presence of coal piles, oil and gas power plants based on storage tanks, solar plants based on their unique shape and spectral behavior,

and hydroelectric plants based on the presence of gorges, water bodies, and dam structures. In the case of the background class, activation maps are typically more uniform than for power plant images due to the lack of characteristic features.

We would like to point out that our results are based on a spatially limited sample of power plants (all located in Europe), which is likely to have a beneficial effect on our classification accuracies. For instance, brown coal power plants in our sample are likely to be located in Eastern Europe, whereas hydroelectric plants are mainly located in Scandinavia or the Alps. This localization and its effect on site surroundings, but also similarities in the design of plants, are likely to be learned by our model and are thus reflected by its results.

Nevertheless, the results presented in this work underline the suitability of medium-resolution remote imaging data to classify different power plant types, providing the ability to estimate the energy mix in a given region. We conclude that extending this classification method to other industrial sites would be worthwhile and strongly improve the value of land use classification information.

References

- [1] P. Zhang, Y. Ke, Z. Zhang, M. Wang, P. Li, and S. Zhang, “Urban land use and land cover classification using novel deep learning models based on high spatial resolution satellite imagery,” *Sensors (Basel, Switzerland)*, vol. 18, 2018.
- [2] J. Yu, Z. Wang, A. Majumdar, and R. Rajagopal, “Deepsolar: A machine learning framework to efficiently construct a solar deployment database in the united states,” *Joule*, vol. 2, pp. 2605–2617, 2018.
- [3] S. Zhou, J. A. Irvin, Z. Wang, E. Zhang, J. Aljubran, W. Deadrick, and R. Rajagopal, “Deepwind: Weakly supervised localization of wind turbines in satellite imagery,” in *Tackling Climate Change with Machine Learning, NeurIPS 2019*, 2019.
- [4] H. Sheng, J. Irvin, S. Munukutla, S. Zhang, C. T. Cross, K. Story, R. Rustowicz, C. W. Elsworth, Z. Yang, M. Omara, R. Gautam, R. B. Jackson, and A. Ng, “Ognet: Towards a global oil and gas infrastructure database using deep learning on remotely sensed imagery,” *ArXiv*, vol. abs/2011.07227, 2020.
- [5] K. Kanellopoulos, M. D. Felice, I. H. Gonzalez, and A. Bocin, “Jrc open power plants database (jrc-ppdb-open),” Dec. 2019.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [7] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *NeurIPS*, 2019.
- [8] P. Helber, B. Bischke, A. Dengel, and D. Borth, “Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, pp. 2217–2226, 2019.
- [9] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, “Learning deep features for discriminative localization,” 2015.




Figure 2: Image examples for all power plant types and the background class. For each image pair, the left image shows an RGB image of the site, while the right image shows the Class Activation Map (CAM, see Section 5) of the corresponding class. All examples shown here have been identified correctly by our model (see Section 4.1). In most cases, the most important image regions for the classification (brightest region in the CAM) correlate well with characteristic image features of the corresponding power plant type.