ENHANCING WIND TURBINE LOCATION ACCURACY: A DEEP LEARNING-BASED OBJECT REGRESSION APPROACH FOR VALIDATING WIND TURBINE GEO-COORDINATES

Maximilian Kleebauer[*] [1,2], *Axel Braun*[2], *Daniel Horst*[1,2], *Carsten Pape*[2]

Department of Energy Management and Power System Operation, University of Kassel, Germany[1]
Energy Meteorology and Geo Information System, Fraunhofer IEE, Kassel, Germany[2]

ABSTRACT

Remote sensing and deep learning-based methods can be combined to obtain location information automatically on a large scale. This paper introduces an approach for enhancing the geo-coordinate accuracy of existing wind turbines. By employing a RetinaNet-based method for regressive object localization, turbines can be precisely located in images in addition to being identified. Utilizing semi-automatically processed and manually filtered high-resolution image data, a model is trained with an average precision of 96 %. Subsequently, the model is applied to Germany's MaStR wind turbine dataset. The application illustrates the advantageous implementation of the method and emphasizes its considerable potential for improving the accuracy of geo-coordinates. While 73.72 % of existing coordinates can be confirmed as correct with a deviation of less than 10 meter, for more than 15 % of the turbine locations coordinates between 10 and 100 meters can be corrected, and for 5.6 % locations a deviation of more than 100 meter can be determined. This showcases the real-world application of the proposed methodology and underscores its significant potential for enhancing the quality of geo-coordinates.

***Index Terms*—** wind turbines, renewable energy systems, object regression, geo-coordinate validation

1. INTRODUCTION

The future of energy supply faces a pivotal challenge, necessitating a substantial expansion of renewable energy sources. This shift towards renewable energy often occurs at the local level, characterized by a strong decentralization trend. This development is evident in the spatial fragmentation of renewable energy expansion. Accurate representation of existing producers and consumers is crucial, particularly in grid operations but also for forecasting wind potentials. Effective planning, considering priority areas and land availability for further wind turbine installations, becomes feasible through this detailed depiction. Ensuring the reliability of wind turbine facilities requires precise predictions.

[*]Corresponding author: maximilian.kleebauer@iee.fraunhofer.de

Looking at datasets of existing renewable energy systems on a global scale, different approaches to providing such datasets can be found. For example, datasets derived from aggregated and digitally processed information on the locations and performance of wind and solar farms have already been published by Dunnet et al. [1]. In contrast, Zhang et al. [2] created a global dataset of offshore wind turbines using Sentinel-1 Synthetic Aperture Radar time series images. Hoeser et al. [3] also used Sentinel-1 Synthetic Aperture Radar data for the automatic derivation of offshore wind turbine locations. Other published studies perform segmentation of onshore wind turbines in high-resolution aerial imagery [4, 5] and wind turbine detection using Sentinel-2 RGB imagery [6]. An improvement in detection accuracy was achieved through multiple acquisition times with Sentinel-2, as shown in a multimodal approach [7]. According to our research, it has not yet been achieved to derive a precise data set of existing energy facilities such as wind turbines, especially onshore, on a large scale based solely on remote sensing data and machine learning methods.

In Germany, the Core Energy Market Data Register provides an overview of existing energy facilities. Location information, performance values or specific plant characteristics are provided for wind turbines, solar photovoltaic systems, and other plant types [8]. However, certain coordinates given here are very inaccurate. Studies have consequently made several revisions to improve data quality [9, 10].

Our approach aims to enhance the accuracy of existing wind turbine location data through an deep learning-based object detection approach using aerial images. There are a variety of DL techniques and methods for object detection, starting with early approaches of the Single Shot MultiBox Detector [11], a variety of real-time object detection variants based on YOLO [12], or RetinaNet [13], which combines the ResNet [14] with an FPN [15]. After studies have already shown high accuracy [16, 17], RetinaNet will be trained with semi-automatically processed and manually filtered high-resolution imagery and then applied to all wind turbines in Germany. This case study is designed to demonstrate the potential for improving the location accuracy of wind turbines.

2. MATERIALS

The Core Energy Market Data Register (German: Marktstammdatenregister, MaStR), administered by the Federal Network Agency for the German electricity and gas market, serves as a comprehensive database for energy market. Commencing operations in 2019, the register undergoes daily updates and encompasses detailed information on entities and facilities within the grid-bound energy supply market. Alongside registered electricity generation units, the register also includes extensive listings of large-scale consumers. On the producer side, the MaStR provides location information, performance values, and specific plant characteristics for various energy sources, such as wind turbines, solar photovoltaic systems, biomass plants, hydro power plants, and conventional plants. The register offers a wealth of information specifically tailored to wind turbines like registration date, commissioning date, rated power, remote controllability, current operating status, manufacturer, type designation, hub height, and rotor diameter. Publicly accessible address data is available at the zip code level, while detailed coordinates are provided for the majority of the turbines. As of November 21, 2023, out of the 32,788 listed turbines in operation, 31,892 include coordinate information [8].

The Digital Orthophotos (DOP) of Germany, as documented by the Federal Agency for Cartography and Geodesy [18], constitute georeferenced and differentially rectified aerial imagery, sourced from the surveying administrations of Germany's federal states. These images faithfully represent the Earth's surface, within the confines of the Federal Republic of Germany, employing a ground resolution of 0.2 meter for the purposes of this investigation. The dataset encompasses both color images in the RGB spectrum. The images maintain a positional accuracy of 0.4 meter standard deviation. The entire dataset covers Germany and is presented in tiles measuring 1,000 $\times$ 1,000 meters, equivalent to 5,000 $\times$ 5,000 pixels each. Each tile is accompanied by a file containing meta-information, notably the timestamp indicating when the respective image was captured. The image data is updated cyclically, usually available every 3 years in site-specific overflight intervals.

3. METHODS

3.1. Data preprocessing

The preparation of training data comprises multiple steps. Initially, all wind turbines registered in the MaStR are loaded. A pre-filtering process is then applied, focusing exclusively on wind turbines with operational status labeled as "in operation". Additionally, the turbines must be categorized as "onshore", and only systems with available geo-coordinates are considered. Finally, turbines situated outside the German federal border are excluded under the assumption that their location data is inaccurate. Existing coordinates of the turbines are provided with a static buffer of 30 meter radius in order to obtain an area-like imprint of the point coordinates. These are required for the subsequent regressive localization method. To generate training and application image data, the wind turbine location data is integrated with DOPs. To align with the requirements of RetinaNet, DOP tiles, each measuring 5,000 $\times$ 5,000 pixels, are further divided into 1,000 $\times$ 1,000 pixel tiles. A methodical differentiation is applied to training and application images. For training, the cut edges are statically selected to generate 15 tiles from each original tile, as illustrated on the left side in Figure 1.

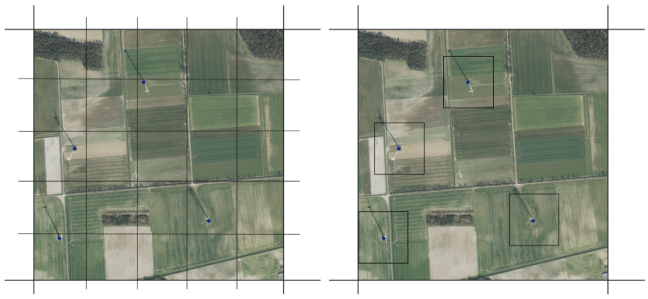

Fig. 1. The two methods of processing the DOPs are shown in comparison. The left-hand side shows an example of static cutting of the training images, while the right-hand side shows cutting using the coordinates of the wind turbines as a centroid. The black lines represent the cutting edges, the blue dots the coordinates of the wind turbines.

This approach ensures that wind turbines are not consistently positioned at the center of the image sections. Conversely, for application images, the wind turbine location is designated as the centroid of the image, as depicted on the right side in Figure 1. Approximately 12,000 images are produced, each containing at least one wind turbine.

Fig. 2. Samples based on their suitability for training. The images marked in red are unsuitable due to incorrect position or poor image resolution, the images marked in yellow contain wind turbines that are clearly visible but were rejected for fine-tuning due to their inaccurate position. The images marked in green contain turbines whose tower base is located directly in the center of the respective boxes.

This dataset is utilized for the initial training. To en-

sure the utilization of highly suitable image data in the subsequent, second training, samples generated automatically undergo manual inspection. This process identifies and removes instances with incorrect coordinates stored in MaStR, imprecise coordinates, and image scenes with insufficient resolution. The re-selection leads to the reduction of a further 5,000 unsuitable images, yielding a dataset of 7,000 images for the second training. The primary emphasis is on the precise localization of wind turbines, ensuring that the center of the regression boxes accurately represents the tower's exact ground location. As illustrated in Figure 2, several samples are depicted to exemplify their suitability. The training is divided into two parts. First, all 12,000 samples automatically derived from the data preprocessing are used, whereas in the second training, the number of samples is reduced to 7,000 highly suitable samples by manual filtering. All other parameters remained the same for both the first and second training: 100 epochs, 100 steps, 80 % training and 10 % independent validation and test data set each. Further parameter configurations are the default settings defined in the package [19].

3.2. Deep learning approach

The machine learning method used in the work is called RetinaNet. The network combines several common variants of deep learning into a classification and regression network. One key feature is the ResNet architecture, including Convolutional Neural Network (CNN) [13]. Furthermore, adding skip connections between the layers enables residual learning as a a widely used basis for deep learning [14]. A Feature Pyramid Net is also used. It consists of a top-down architecture with horizontal connections to create high-level semantic feature maps at all scales [15]. The classification task is performed at the output of the backbone network using Focal Loss, which was developed to train extremely unevenly distributed foreground and background classes [13]. The regression task, on the other hand, is implemented for the regressive delineation of the objects. The regression targets are output as rectangles that are located entirely within the images shown and have predefined aspect ratios of 1:2, 1:1, 2:1 [13]. The regression loss uses the Smooth L1 Loss, which was originally developed as part of the Fast R-CNN [20]. Average Precision (AP) is used to measure the performance of the model. The sum of correctly detected objects in relation to all detected objects is considered. RetinaNet uses the COCO detection evaluation metric under the AP, whereby the overlap of the areas from the regression and the ground truths must be greater than 50 %, in order to be considered a correctly detected object. The keras-retinanet [19] package is used, which was specially developed for the use of RetinaNet. To quickly achieve good generalization of the network, pre-trained weights on 500 classes from the Open Image Dataset where used [21].

4. RESULTS

As Training Progress Summary, the progression of the two losses from the classification and regression networks, as well as the AP, were validated to determine the networks' performance. Figure 3 shows the losses on the left-hand and the AP on the right-hand side.

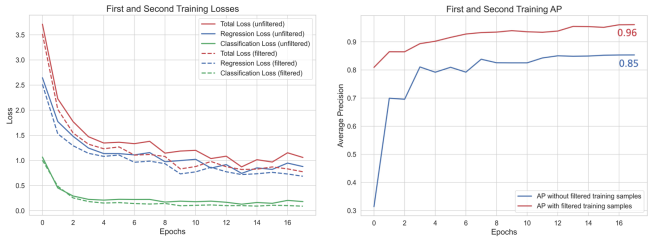

Fig. 3. The illustrations depict the losses and the AP.

The training is terminated by early stopping after 17 epochs in each case, indicating no further progress in training. The trend in losses exhibits a nearly constant decrease for both the initial and subsequent training. This is observed for both the focal loss in classification and the Smooth L1 loss in regression. The total loss represents the cumulative sum of the individual losses. A consistent upward trend can be observed in the AP. Finally, the AP is 85 % for the first training and 96 % for the second training with post-filtered samples. Overall, the loss and AP's curves clearly show the strong generalization of the network based on the training examples. Incorrect recognition are shown in Figure 4.

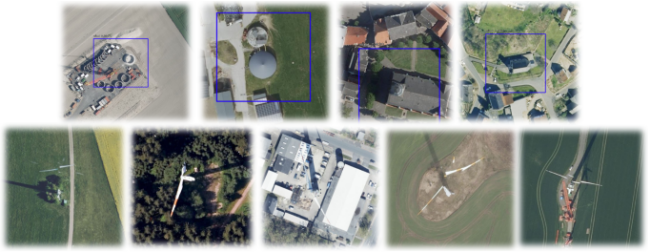

Fig. 4. False positive and false negative examples from the application are summarized in the following. The top line represents incorrectly identified wind turbine, false positives. The bottom line shows turbines that have not been detected.

This includes a construction site, a biogas plant and two churches. Secondly, some of the poorly represented turbines are not recognized by the network. This applies to different backgrounds, so that turbines in open fields, in the forest and also in the settlement are not recognized. However, they are also hard to identify during a visual inspection. Examples of correctly recognized wind turbines, conversely, are shown in Figure 5. In addition to turbines with good resolution, poorly resolved turbines can also be identified in the images. All

images show that the regression locates the towers of the turbines exactly in the centers of the bounding boxes. In other words, the centers of the regression boxes can be interpreted as exact geo-coordinates of the wind turbines.

Fig. 5. True positive examples from the application are presented as follows. The upper row displays instances featuring clearly visible and accurately identified wind turbines, where the centroid of the regression boxes serves as the base of the tower. The bottom row, shows correctly detected wind turbines, with less clear representation in the images.

The results of the geo-location site correction for onshore wind turbines are described below. The dataset comprised a total of 30,326 operational onshore wind turbines. Through geographical selection, it was revealed that 55 wind turbine (0.18 %) coordinates were situated outside the country's borders. Additionally, 1,397 turbines (4.61 %) were flagged for having outdated aerial images captured before their commissioning. The validation methodology was applied to identify 1,699 instances (5.60 %) of incorrect location data, where a deviation of at least 100 meter from the coordinates in the MaStR was observed. Consequently, 27,130 wind turbine location data points (89.46 %) were identified within 100 meter accuracy and are treated with sensitivity in subsequent analyses due to their deviation characteristics. Table 1 visually depicts the dispersion of wind turbines across various distance ranges, offering crucial insights into the efficacy of the applied coordinate corrections. The categorization of turbines based on deviation ranges in meter provides a detailed view. Notably, a substantial percentage (73.72 %) of turbines exhibits a deviation of less than 10 meter, underscoring the high precision of their coordinates. In contrast, 7.66 % of turbines deviate between 10 and 20 meter, while more than 8 % show deviations exceeding 20 meter.

5. DISCUSSION

Comparing the results of the first and second training reveals a significant increase in recognition accuracy due to the manual post-filtering of the training samples. The high accuracy of the second model suggests a strong generalization of the network based on the validation data. This can also be illustrated by Figure 5. It should be noted, however, that the data used in the validation is very homogeneous due to the experimental setup. A large-scale application would result in a reduction with regard to the AP. Nevertheless, the misidentified objects observed in the test application, as shown in Figure 4, provide insights into two important considerations. Several objects are incorrectly recognized as wind turbines. Facing this issue, an improvement in recognition accuracy could be achieved by including misidentified objects as part of the negative class during training. However, the unrecognized wind turbines also illustrate the limitations of the application based on DOPs. Moreover, Figure 4 demonstrates that turbines that lack consistent representation in the images may go unnoticed by the network during the application. Supplementing the image data with additional acquisition times could address this issue. The application shows a strong variance in the accuracy of the geo-coordinates. The locations of almost three quarters of the wind turbines can be confirmed. However, nearly 8 % of the locations deviate between 10 and 20 meter, and more than 8 % deviate by more than 20 meter. With 5.6 % of the completely incorrectly stored coordinate locations, a considerable proportion of the wind turbines are incorrectly located by more than 100 meter. Due to outdated aerial images, the 4.61 % of turbines with outdated aerial images could not be validated. Expanding the methodology to regularly updated image datasets, such as the utilization of Sentinel-2 imagery, is conceivable. Nevertheless, the adoption of Sentinel-2 imagery, characterized by considerably lower resolution, may result in a substantial decline in accuracy during facility identification.

Range in meter	wind turbines	wind turbines in %
< 10	22,357	73.72
10 - 20	2,323	7.66
20 - 30	872	2.88
30 - 40	514	1.70
40 - 50	319	1.05
> 50	745	2.46

Table 1. Distances between MaStR and validated coordinates.

6. CONCLUSIONS

This paper introduces an approach for enhancing the geo-coordinate accuracy of existing wind turbines. Employing a RetinaNet-based approach for regressive object localization, wind turbines can be precisely located in images in addition to being identified. The validation results demonstrate the network's high generalization capability. The exemplified application, focusing on enhancing wind turbine location accuracy through MaStR data, illustrates the practical implementation of the method and emphasizes its considerable potential for improving the accuracy of geo-coordinates.

7. REFERENCES

[1] Sebastian Dunnett, Alessandro Sorichetta, Gail Taylor, and Felix Eigenbrod, "Harmonised global datasets of wind and solar farm locations and power," *Scientific Data*, vol. 7, no. 1, pp. 1–12, 2020.

[2] Ting Zhang, Bo Tian, Dhritiraj Sengupta, Lei Zhang, and Yali Si, "Global offshore wind turbine dataset," *Scientific Data*, vol. 8, no. 1, pp. 191, 2021.

[3] Thorsten Hoeser, Stefanie Feuerstein, and Claudia Kuenzer, "Deepowt: A global offshore wind turbine data set derived with deep learning from sentinel-1 data," *Earth System Science Data*, vol. 14, no. 9, pp. 4251–4270, 2022.

[4] Min Han, Huabin Wang, Guanghui Wang, and Yu Liu, "Targets mask u-net for wind turbines detection in remote sensing images," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, pp. 475–480, 2018.

[5] Narayana Darapaneni, A Jagannathan, Vigneshwaran Natarajan, Guruprasadh Vadakkupattu Swaminathan, S Subramanian, and Anwesh Reddy Paduri, "Semantic segmentation of solar pv panels and wind turbines in satellite images using u-net," in *2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS)*. IEEE, 2020, pp. 7–12.

[6] Michael Mommert, Linus Scheibenreif, Joëlle Hanna, and Damian Borth, "Power plant classification from remote imaging with deep learning," in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 2021, pp. 6391–6394.

[7] N Mandroux, S Drouyer, and R Grompone von Gioi, "Multi-date wind turbine detection on optical satellite images," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 2, pp. 383–390, 2022.

[8] Federal Network Agency (BNetzA), "Core energy market data register (mastr)," https://www.marktstammdatenregister.de/MaStR/, Accessed on 02 November 2023.

[9] David Manske, Lukas Grosch, Julius Schmiedt, Nora Mittelstädt, and Daniela Thrän, "Geo-locations and system data of renewable energy installations in germany," *Data*, vol. 7, no. 9, pp. 128, 2022.

[10] Deniz Tepe, Florian Kotthoff, Christoph Muschner, Esther Vogt, and Ludwig Hülk, "Improving data reliability in germany's energy system: A validation of unit locations of the marktstammdatenregister," *arXiv preprint arXiv:2304.10581*, 2023.

[11] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg, "Ssd: Single shot multibox detector," in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. Springer, 2016, pp. 21–37.

[12] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi, "You only look once: Unified, real-time object detection," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.

[13] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar, "Focal loss for dense object detection," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 22.10.2017 - 29.10.2017, pp. 2999–3007.

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[15] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie, "Feature pyramid networks for object detection," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 21.07.2017 - 26.07.2017, pp. 936–944, IEEE.

[16] Maximilian Kleebauer, Daniel Horst, and Christoph Reudenbach, "Semi-automatic generation of training samples for detecting renewable energy plants in high-resolution aerial images," *Remote Sensing*, vol. 13, no. 23, pp. 4793, 2021.

[17] Wenwen Qi, "Object detection in high resolution optical image based on deep learning technique," *Natural Hazards Research*, vol. 2, no. 4, pp. 384–392, 2022.

[18] Bundesamt für Kartographie und Geodäsie, "Dokumentation digitale orthophotos," 2023.

[19] Hans Gaiser, Maarten de Vries, Valeriu Lacatusu, vcarpani, Ashley Williamson, Enrico Liscio, András, Yann Henon, jjiun, Cristian Gratie, Mihai Morariu, Charles Ye, Martin Zlocha, Ben Weinstein, Rodrigo Meira de Andrade, Pedro Conceição, Alexander Pacha, hannesedvartsen, Daniyal Shahrokhian, Wudi Fang, Mike Clark, meagerYak, Iver Jordal, Max Van Sande, Jin, Etienne-Meunier, Andrew Grigorev, Guillaume Erhard, Eduardo Ramos, and Denis Dowling, *fizyr/keras-retinanet 0.5.1*, 2019.

[20] Ross Girshick, "Fast r-cnn," in *2015 IEEE International Conference on Computer Vision (ICCV)*. 07.12.2015 - 13.12.2015, pp. 1440–1448, IEEE.

[21] GitHub, "Github - zfturbo/keras-retinanet-for-open-images-challenge-2018: Code for 15th place in kaggle google ai open images - object detection track," 2021.

Autorenversion von: M. Kleebauer, A. Braun, D. Horst and C. Pape, "Enhancing Wind Turbine Location Accuracy: a Deep Learning-Based Object Regression Approach for Validating Wind Turbine Geo-Coordinates," IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium, Athens, Greece, 2024, pp. 7863-7867, doi: 10.1109/IGARSS53475.2024.10641018.