AutoWindLoc: Precise Localization of Wind Turbines in High-Resolution Orthophotos for Enhanced Registers

Justus Middendorf [1], André Kelm [2], Simone Frintrop [3]

[1] University of Hamburg, 22527, Hamburg, Germany - middendorf.justus@studium.uni-hamburg.de
[2] Department of Computer Vision, University of Hamburg, 22527, Hamburg, Germany - andre.kelm@uni-hamburg.de
[3] Department of Computer Vision, University of Hamburg, 22527, Hamburg, Germany - simone.frintrop@uni-hamburg.de

Keywords: Wind Turbine Detection, High-Resolution Orthophoto Analysis, Automated Localization, Geospatial Data Processing, Negative Sampling.

Abstract

The paper proposes a novel framework for automatically detecting wind turbines in orthophotos, transferring this information to a database, and linking detected turbines to an existing registry to minimize location inaccuracy. This inaccuracy has a significant impact on planning and identifying new potential locations for wind turbines, as existing turbines must be considered in these processes. The existing public data frequently exhibit discrepancies from the actual location, and existing work also exhibits relatively large discrepancies from the actual location, even though the exact location of a wind turbine is so important for these processes. Moreover, existing work has not produced a new or improved database that could be used in the long term for processes in the wind energy sector. The development of the AutoWindLoc framework creates a fully automated data basis from which locations and possible further information can be retrieved. The recognition process utilizes a two-stage approach, incorporating a You Only Look Once model with negative sampling and a binary classification Convolutional Neural Network, which attains an average deviation of 0.85 m from the actual location.

1. Introduction

Global heating of land and ocean is progressing, with dramatic consequences for humans, animals, and our environment (Minière et al., 2023). At the same time, high energy prices are leading to competitive problems and thus to layoffs in some energy-intensive industries (Olk, 2023). Wind energy, which is both carbon-neutral and cost-effective, is a key component in the solution to both problems. However, modern wind energy companies face bureaucratic hurdles, such as inaccurate public data or long waiting times for official data requests, which make it difficult to identify lucrative sites. This is where our approach comes in: a framework that **Auto**matically searches for **Wind** turbines to **Loc**ate (**AutoWindLoc**) them, so that, in turn, unused areas can be identified more easily and faster.

In detail, the expansion of wind energy is confronted with significant challenges. A significant challenge pertains to the identification of suitable locations, necessitating the collection and consideration of extensive data (Sliz-Szkliniarz et al., 2019). For instance, the identification of existing turbines or the utilization of wind maps is necessary. This data is needed to determine the preload of a site or the efficiency of a new wind turbine. The process of obtaining the locations of existing turbines is time-consuming. In numerous countries, there are public registers for wind turbines, predominantly for general green energy or general power generation. However, accessing these is often complicated, or it is not possible to query individual sites from the public register, as is the case in the UK (Department for Business, Energy & Industrial Strategy, 2013). Furthermore, a discrepancy was observed between the locations shown and their actual geographical location, a phenomenon we found to be particularly common in Germany.

In Germany there is the Markstammdatenregister (MaStR) (Markstammdatenregister, 2019), in which all energy producers, including wind turbines, are registered and publicly accessible. However, the accuracy of this register is often questionable, as it contains significant deviations in its location data, as evidenced in Figure 1, or as substantiated by the research of Maximilian Kleebauer (Kleebauer et al., 2024), which consequently hinders effective planning processes.

Therefore, all locations of wind turbines have to be queried, and as the authorities have at least one month to provide the information (Umweltinformationsgesetz (UIG), 2005), this usually takes a long time. This delay is often attributed to the substantial workload of these authorities, which includes the administration of wind turbine permits. This has a negative impact on the green energy business, as customers need explanations of the process and must be patient. In many cases, the business relies on real orders and contracts, making this an unnecessary stress test, especially for smaller companies. Each subsequent request for new information continues the cycle, consuming more time and testing the patience of both parties.

We hereby propose a system, designated AutoWindLoc, that automatically searches for existing wind energy sites on the basis of orthophotos. Orthophotos are images created by flying over an area. AutoWindLoc then localizes the sites, stores the wind turbines found, and searches for the real wind energy asset in the MaStR based on the location found to provide further information, such as the manufacturer or the power output. We expect that the integration of AutoWindLoc will result in an enhanced and sustainable database, a feat that will be achieved through the presentation of two key contributions:

– The developed framework automates the downloading, processing, and saving of images, as well as the extraction of additional information from existing registers.

– We present a two stage detection process that has been developed to identify wind turbines with small deviations in

(a) On the orthophoto, all entries that are included in the area in MaStR have been marked with green markers. (wind-turbine.com, 2025)

(b) An image in which all existing wind turbines have been correctly recognized by AutoWindLoc. The detected wind turbines are outlined in blue and have a label indicating what was found, in this example wind turbines and the safety as a percentage number, with a detection highlighted for further analysis.

Figure 1. Two orthophotos which both show the same section, where in (a) the entries of the MaStR are marked and in (b) the recognitions of AutoWindLoc

their positioning. This data can then be utilized in various processes within the wind energy sector.

While existing methods employ orthophotos and satellite images to localize wind turbines, these approaches do not generate a novel data foundation or address the issue of discrepancies between the actual and perceived locations (Zhai et al., 2024). The deviations of the locations are of particular relevance as they are still in the meter range, which rapidly distorts the results, as numerous calculations, such as those for turbulence, are highly sensitive to displacements.

The solution employed in this study utilizes a two-stage process to enhance the efficacy of the results. Initially, a You Only Look Once (YOLO) model (Jocher et al., 2023) is employed, and subsequently, the outcomes are validated using a self-developed classification Convolutional Neural Network (classification CNN) as shown in Table 1. The incorporation of negative examples into the data sets serves to enhance the precision and stability of the system. The Django framework (Django Development Team, 2025), in conjunction with a PostgreSQL (PostgreSQL Development Team, 2000) database, is utilized to optimize the processing and storage of data obtained because they support geodata and are very stable when used together.

Section 2 is dedicated to a review of related work. Thereafter, the methodology of AutoWindLoc and the two-stage detection process are explained. This is complemented by a detailed description of the data sets. Finally, the results achieved are shown and discussed.

2. Related Work

The detection of offshore wind turbines, unlike onshore wind turbines, has been the subject of extensive research. The results of research in this area have consistently achieved good results with accuracies of up to 99.93%, as shown in the framework of (Zhang et al., 2024). In addition, the work creates a sustainable database of identified wind turbines. Furthermore, the framework of Yichen Zhai et al. (Zhang et al., 2024) has broad applicability, as demonstrated by its successful use in both Chinese waters and the North Sea. However, it should be noted that the applicability of this system is limited to offshore installations and does not extend to onshore applications.

In a more recent work, a separate model was developed (Zhai et al., 2024). This model represents a modification of the YOLO framework with version 5 and is called Wind Turbine YOLO. However, it should be noted that this version 5 is an older version from 2020, after which several other versions have already been released (Jocher and Qiu, 2024). In this model, the shadow of the turbine is utilized as a distinctive feature to enhance the efficacy of the system. To this end, a range of image resolutions are employed to ensure broad applicability, even in cases where the orthophotos are of suboptimal resolution. The approach yielded average precisions ranging from 92.53 % to 98.84 %, contingent on the resolution. However, this work is marred by a problem in the accuracy of location detection. While this is already adequate for values between 4.76 m and 8.13 m for the mean distance error, it is too high for planning processes or similar, as these processes are highly sensitive to shifts in wind turbines. In addition, direct applicability is hindered by the lack of processing or storage of the data in geospatial data storages.

Another study also deals with the detection of wind turbines in Germany (Kleebauer et al., 2024). The study utilizes a RetinaNet (Lin et al., 2017), which is naturally good at dealing with unevenly distributed features. The network's two-stage architecture enables it to operate more efficiently with large, high-resolution data sets without requiring exhaustive scrutiny of every detail. Utilizing this approach, an average accuracy of 96 % is achieved, yielding commendable outcomes. However, it should be noted that no meaningful value for a deviation of the found position from the original position is considered in this work. Instead, the study introduces a categorization system that classifies all detected instances into certain deviation classes from the MaStR. This categorization system hinders the determination of the mean deviation from the actual location, thereby preventing comparisons. The authors acknowledge the potential enhancement of the existing dataset, yet this proposal remains unimplemented. This is problematic because it does not create a database that could accelerate the expansion of renewable energy.

Parallel to this work, a further framework was developed (He et al., 2025). This framework creates a map of all found wind turbines, which can then be used as a database and achieves an F-Score of 0.963. However, the comparability of results is hindered in this case. The limitation of the study to China and the absence of a published data set for comparison are significant obstacles. Additionally, the paper does not contain any part on the methodology or further results regarding the comparability of both precision and recall as well as the deviation from the actual location.

3. Methodology

In this section, the methodology is explained. In section 3.1, the structural composition of AutoWindLoc is considered and elaborated. This is followed by section 3.2, which describes a detailed description of the two stage recognition process, and finally section 3.3, which describes the data sets required to train the two stage recognition models.

3.1 Framework Architecture

This section will address the three primary components of AutoWindLoc, as illustrated in Figure 2. These components are as follows: data acquisition and preprocessing (highlighted in blue), object detection and output generation (highlighted in green), and data extraction (highlighted in purple).

The primary function of data capture and preprocessing is to facilitate the seamless integration of orthophotos and their metadata into object detection and the output generation. The preparatory actions encompass the organization of queries for orthophotos and their metadata and the subsequent cutting of this data in preparation for its detection. The resultant cut data, comprising images and metadata, is then passed on to the next subprocess. The processing of each image is initiated individually, and this process is then repeated sequentially for each image. This process is further delineated in section 3.1.1.

In the subsequent stage of the process, the resulting image is processed with a YOLO model, the purpose of which is to recognize wind turbines. In cases where the recognition is uncertain, the image section of the object in question passes through the second stage, which involves a classification CNN model that performs a second check. If the classification is confirmed, the coordinates of all wind turbines found are then determined from the metadata and saved. A comprehensive explanation can be found in section 3.1.2.

The final step of the Data Extraction section 3.1.3 follows only after all the data has been retrieved and processed. This subsequent step involves the acquisition of additional information related to the location of the wind turbine in question. The location found is used to try to find a wind turbine from the MaStR, which stores information such as rated power and manufacturer. A breadth-first search is employed for this purpose; however, it should be noted that this approach does not guarantee correctness due to the potential for significant variations in the data provided by the MaStR.

3.1.1 Data Acquisition and Preprocessing A GeoJSON file is employed for the download process, which is provided by the State Office for Geoinformation and Surveying of Lower Saxony (LGLN). However, it is imperative to note that this process requires further refinement to ensure that only the most recent images of a specific area are downloaded. To this end, a check is made for each specific area to see if multiple images are available. If multiple images are present, the most recent one is identified and utilized based on its publication date. All other entries for older images will be deleted. For each image, an XML file containing the metadata is downloaded. This metadata encompasses not only the conventional attributes of metadata but also the coordinates encompassing the area. It is noteworthy that the coordinates are expressed in the UTM32 projection (Bundesamt für Kartographie und Geodäsie, 2018). This coordinate system is frequently utilized and employs a conforming representation of the earth's surface. Each GeoJSON entry is then processed as described below.

In the initial processing step, as previously referenced, an image is downloaded in loss-free TIFF format, accompanied by the corresponding XML. In the subsequent step, the image is divided into sixteen parts. This division is imperative in order to satisfy the requirements of the YOLO model. The necessity for this division arises from the alternative option of an enormous computational capacity, which is further analyzed in section 3.2.

The XML file must then undergo a similar splitting process. This is the best way to get the given coordinates for each part of the image after the splitting process. Furthermore, it is necessary that both the images and the XML files are numbered consecutively so that they can be correctly assigned in the subsequent steps of the process.

3.1.2 Object Detection and Output Generation The subsequent step of recognition and output generation is based on a two-stage recognition process, which will be considered in more detail in section 3.2. In this stage, the image section undergoes the recognition process. This process returns coordinates in the image of bounding boxes for each wind turbine found. These bounding box coordinates are then utilized to calculate the center point of the bounding box, which is the exact base point of a wind turbine.

These coordinates can then be translated into real-world coordinates in conjunction with the metadata XML. Given the property of UTM32 as a conformal mapping of the Earth's surface, the conversion into real-world coordinates can be straightforwardly accomplished. The transformation is linear, whereby the coordinate value is incremented by a single unit for each meter moved in a given direction. Subsequent to this, an entry is created in the database for each of these points, containing the coordinates and a generated name.

The database has been constructed using the Django framework (Django Development Team, 2025). The selection of Django was primarily influenced by its Python foundation, which enables the integration of machine learning capabilities. Additionally, Django's abstraction of database accesses facilitates faster

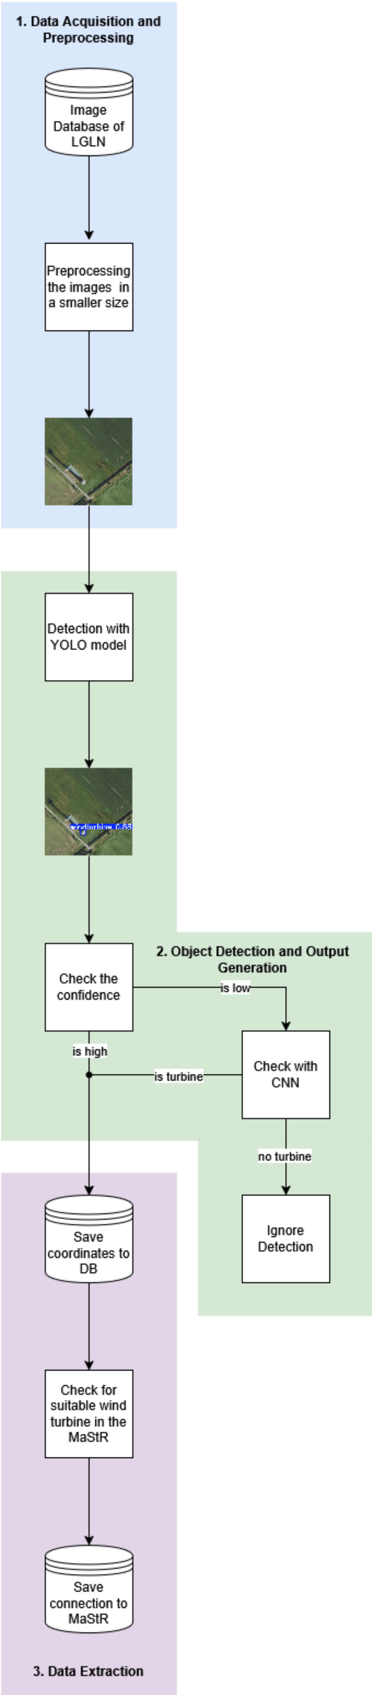

Figure 2. Flowchart of the automated process for detecting and validating wind turbines using satellite images. The process includes data pre-processing (with image data from the LGLN - Landesamt für Geoinformation und Landesvermessung Niedersachsen), object recognition with YOLO and binary CNN and comparison with the MaStR.

and more structured development. PostgreSQL (PostgreSQL Development Team, 2000) was selected as the underlying database. A pivotal consideration in this decision was the robust support for geodata, facilitated by the PostGIS extension. The selection was further influenced by the open-source nature of PostgreSQL, its capacity for scalable functionality, and the active community of users and developers. The combination of PostgreSQL and Django has been demonstrated to be highly stable and powerful. This stability and power are primarily attributable to the support of geodata, but also to the comprehensive functionality of PostgreSQL supported by Django.

After being entered into the database, the next image part goes through this process. Once this process has been completed for all sub-parts of an image, the subsequent steps of splitting 3.1.1 and the above-mentioned process start again from the beginning until the processing of all images has been completed.

3.1.3 Data Extraction The federal state, district, and municipality are determined and saved for each entry in the database. To this end, polygons for each federal state, district, and municipality are stored in the Geographic Coordinate System. The determination of the federal state is initiated by ascertaining the location of the wind turbine point within the polygon, a process facilitated by the Django function that converts the coordinates. This is then repeated for all federal states until the polygon is found. Subsequently, the same procedure is followed, first for the district and then for the municipalities.

The subsequent step involves searching for a matching entry in the MaStR for each wind turbine identified. The analysis of all wind turbines is conducted through the implementation of a breadth-first search. The initial step involves ascertaining whether there are any wind turbines in the MaStR that are within a radius of one meter of the turbines found. If such a turbine is found, the MaStR turbine is assigned to the turbine found, and the size of the radius is determined. This process is repeated for all wind turbines, and once completed, the radius is increased. The radius is increased adaptively, as most wind turbines are assigned at the beginning. The radius is increased in increments of 1 m until it reaches 50 m, then by 2 m up to 100 m, and from there by 5 m up to 500 m, and finally by 10 m up to 3000 m. At this juncture, the pursuit of the MaStR wind turbine has been discontinued due to the improbability of achieving a meaningful classification.

In the event that a wind turbine that has been found is assigned to a wind turbine from the MaStR list by the process previously described, that turbine is then removed from the set of wind turbines found that still need to be assigned. This same principle applies to the wind turbine of the MaStR. Consequently, if these are assigned to a wind turbine found, they are no longer considered as possible candidates. This ensures that each wind turbine found can only be assigned to one MaStR wind turbine, and vice versa.

The procedure in effect creates a database that contains much more information than just locations. In addition to the data of the MaStR, the radius in which the wind turbine was found is also stored, as well as the real federal state, district, and municipality. The radius value can serve as an indicator of the relationship between the paired wind turbines found and in the MaStR.

3.2 Recognition process

The recognition process is a two-stage process. We chose this two-stage process, as it has been demonstrated to enhance precision, leading to the elimination of numerous false positives in the subsequent stage. Achieving a high level of precision is imperative, as the absence of wind turbines in the database is of paramount importance. Consequently, the benefit derived from the enhanced precision of the locations becomes less significant as the results become contaminated by false positives.

The initial stage of the process employs a YOLO model of version 8 (Jocher et al., 2023). This YOLO model is capable of identifying wind turbines from a provided orthophoto and subsequently detecting bounding boxes. The decision to employ a YOLO model was predicated on its proven capacity to deliver expeditious results of a commendable caliber. This is of particular importance given that over 20,000 images are analyzed for Lower Saxony a federal state of Germany alone.

As previously referenced in 3.3.1, the model was meticulously trained with a resolution of 2500 x 2500 pixels, in order to circumvent an excessively substantial computational burden. During this training process, the YOLO model internally downscales the images, and later rescales them, as is customary for the YOLO model, to a resolution of 2500 x 2500 pixels. The model was trained over 100 epochs. This number was determined through a series of trials, which identified this particular number as the optimal value for the number of repetitions. For the base model for YOLO, the model size m was chosen. This decision was guided by the finding from preliminary experimentation with alternative model sizes, which indicated that m consistently exhibited minimal overfitting and underfitting tendencies.

In the second part of the recognition process, a simple, self-developed classification CNN was utilized, as illustrated in Table 1, to predict the probability of a wind turbine being visible in the input image. The architecture of this model comprises two convolutional layers, the first with 32 filters and max-pooling, and the second with 64 filters and max-pooling. The final stage of the network involves a flattening and output layer with sigmoid activation.

Id	Layer	Kernel Size / Operation	Output Size
1	Input Layer	-	256×256×3
2	Conv2D + ReLU	3×3×3×32, padding=same	256×256×32
3	MaxPooling2D	2×2	128×128×32
4	Conv2D + ReLU	3×3×32×64, padding=same	128×128×64
5	MaxPooling2D	2×2	64×64×64
6	Flatten	-	262144
7	Dense + Sigmoid	1	1

Table 1. CNN architecture

The selection of the classification CNN was made on the basis that it is relatively simple and that it carries out its function in an efficient manner. The classification CNN has been found to rapidly and reliably detect the presence of a wind turbine.

The training was initiated with 100 epochs. However, the training was configured in such a way that it is automatically cancelled if the validation loss does not improve for ten epochs in a row. In such a scenario, the model with the most recent optimal validation loss is employed.

The combination of these two models involves the classification CNN model utilizing the outputs from the described YOLO model. Initially, the probability of a bounding box is analyzed to determine whether it corresponds to a wind turbine. The probability and the bounding box of a detection can be read directly from the result. In instances where this value falls below 75 %, the subsequent phase of detection is initiated, which is otherwise bypassed. The bounding box is then processed. In the second part, the center point of a bounding box is determined, and a square image is cut out around this center point, which is 200 x 200 pixels in size. This image then covers the entire base of a wind turbine and is used as input for the CNN after it has been scaled appropriately.

If the classification CNN subsequently asserts with a probability exceeding 50 % that a wind energy plant is visible in the image, the subsequent process is initiated. Conversely, if the classification CNN's prediction falls below this threshold, the result is disregarded without further consideration.

3.3 Dataset

It is acknowledged that the recognition process is comprised of two distinct models, thus necessitating two separate data sets. These data sets are addressed in the following sections.

3.3.1 Object Detection Dataset The data set for the YOLO model is based on publicly accessible orthophotos. The images originate exclusively from Lower Saxony and were published by the State Office for Geoinformation and Surveying of Lower Saxony (LGLN Open Geodata , 2023). The choice of Lower Saxony as a training dataset is based on the fact that the region is characterized by a diverse range of terrains. The topographical variety is notable, encompassing a coastline in the northernmost region, extensive agrarian terrain, urban areas such as Hanover, Osnabrück, and Oldenburg, and the mountainous landscape of the Harz. In addition, the state of Lower Saxony provides a complete, openly accessible interface for orthophotos, with the images available in high resolution. This is a particularly salient point, as the availability of freely accessible data is otherwise limited. Furthermore, the expense associated with procuring orthophotos or satellite images is a significant deterrent, thereby further reinforcing the appeal of Lower Saxony as a resource.

The original images are characterized by a resolution of 20 cm per pixel, with each image encompassing an area of 2 km by 2 km. However, due to constraints in computing capacity, it is impractical to employ a YOLO model trained with images of 10,000 x 10,000 pixels. It is evident that, in order to resolve this issue, a reduction in image size was necessary. Consequently, the area covered by each image is now 500 x 500 m, equivalent to 2500 x 2500 pixels. Each image has been subdivided into sixteen parts. In the event that a wind turbine is situated on and proximate to the section boundary due to the splitting of the images, the wind turbine in question has been labeled in the image in which the largest part is shown. Consequently, this approach results in the inclusion of even incomplete wind turbines within the dataset.

The labeling of these images was conducted manually using the LabelImg tool (TzuTa Lin, 2015). It is noteworthy that the base of the wind turbine was the sole component placed within the bounding box. That means that no rotors or shadows were labeled as shown in Figure 3. Two arguments support this approach. The initial argument pertains to TrueDOP (Arbeitsgemeinschaft der Vermessungsverwaltungen der

Figure 3. Image on which a wind turbine is labeled.

Länder der Bundesrepublik Deutschland (AdV), 2019), an approach to orthophoto processing that involves adapting orthophotos to produce a digitally decertified aerial image. This approach ensures that buildings no longer overlap roads in images, thereby enhancing the clarity and interpretability of the visual data. However, this approach also results in the exclusion of wind turbine towers and rotors, which typically render both as poorly identifiable features. Secondly, this approach facilitates a more precise and reliable calculation of the base point.

The data set consisted of 2536 images. Of these, approximately 20 % did not depict wind turbines and were manually added to the data set. The introduction of the negative sample resulted in a marked enhancement in the outcomes across all parameters, a topic that will be further explored in section 4.

Of the 2,536 images, 2,012 were allocated for the training set, constituting 80 % of the total, while 419 images were designated for the validation set, accounting for approximately 16 % of the entire dataset. The remaining 105 images (4%) were reserved for testing at the culmination of the development process.

3.3.2 Image Classification Dataset The second data set, i.e., the one for the classification CNN, is based on the YOLO data set. The YOLO model was configured to recognize all images once, thus enabling the generation of bounding boxes for wind turbine bases and false detections. Sections were extracted from the images comprising the YOLO dataset. These images were sized at 200 x 200 pixels, with the midpoint of each section corresponding to the center of the bounding box determined. The images, when categorized into two-state groupings and sliced, are shown on 4.

In this particular instance, a dimension of 200 x 200 pixels proved to be adequate, as it encompassed an area of 400 m^2, which is sufficient to cover the base of a wind turbine. This was followed by the extraction of the bounding boxes from the images and their subsequent saving as a separate file. These new images were then manually binary classified, which indicates whether they show a wind turbine or not.

The data set under consideration comprises a total of 218 images. Of these, 50 % correspond to 109 images, which are correctly identified by the YOLO model, while the remaining

(a) Example detections of the YOLO model showing three different wind turbines. These detections were cut out and used for training the classification CNN.

(b) Negative detections of the YOLO model that show no relevant objects for the classification. You can see a silo in the first image, a power pole in the second image and a wave in the third image.

Figure 4. Example detections of the YOLO model, in which (a) the detections of wind turbines and (b) negative detections are shown. These were used for training the classification CNN.

50 % correspond to images that are misidentified by the YOLO model. For the purposes of training, the data set was divided into three sections. The initial segment constitutes the training dataset, encompassing 70 % of the images (153). The subsequent segment is designated as the validation dataset, comprising 20 % of the images (44). The final segment, constituting the test dataset, encompasses the remaining 10 % of the images (21).

4. Results

The approach employed is predicated on sustainable data preparation. The platform employs a robust storage system, founded on Django and PostgreSQL, to ensure efficient data management and retrieval. The comprehensive data set provides additional insights into the characteristics of each individual wind turbine, thus increasing the information value. However, it is important to mention that the dataset contains errors. These errors pertain to both the geographical location and the mapping of the found wind turbines and MaStR, so sometimes incorrect information is output.

With regard to the discrepancy between the predicted and actual locations, the software displays an average deviation of 0.85 m. This average deviation is the sum of all deviations from the true location, divided by the number of locations. Caution should be taken when interpreting this value, however, as the lack of data sets on the exact location of each turbine introduces a degree of uncertainty. Consequently, manual delineation of each center point is imperative, though this approach can introduce additional deviations.

The software achieved a precision of 97.1 % on the test data set. Precision is calculated by summing true positives and false positives and dividing by the number of true positives. This value indicates the number of objects found and, in this case, included in the database, despite not being wind turbines.

The software reported a recall of 96.94 % on the test dataset. For this parameter, the true positives are added to the false negatives and divided by the true positives. This calculation enables the estimation of the proportion of wind turbines that were not identified.

It has been shown that the negative sampling technique used in the YOLO dataset 3.3.1 improves the effectiveness of the YOLO model. This enhancement was observed when training was performed with and without negative sampling. Specifically, the precision of AutoWindLoc was enhanced by 14 percentage points. Additionally, a notable enhancement in the recall score was observed, with an increase of 10 percentage points, suggesting a substantial improvement in the model's performance.

5. Discussion

The mean deviation of the framework is 0.85 m, indicating that it possesses a low deviation from the original location. This outcome is notably superior to the mean distance error, which ranges from 4.76 m to 8.13 m. However, it is crucial to acknowledge that this is not the same parameter, as it is not feasible to calculate the mean distance error. Nevertheless, a statement can be made about the deviation from the real location, which is significantly smaller in this work. This substantial enhancement has the potential to facilitate the planning and approval of wind turbines, as it could lead to a reduction in the number of hurdles encountered during the process.

It is imperative to acknowledge that the mean deviation is not exact, as previously shown in 4. This inaccuracy arises from the unavailability of a publicly accessible dataset that contains the precise base points of wind turbines. Consequently, manual delineation of all base points of the wind turbines was necessary to calculate the deviation, which cannot be absolutely accurate.

The values of approximately 97 % for the precision are indicative of a highly favorable outcome. This means that very few wind turbines are added to the database that are not wind turbines at all, resulting in a usable database with few errors. It is also worth noting that both models are characterized by their simplicity and lack of complexity. The CNN model consists of only three layers, and the YOLO model is a model that can be trained quickly. This simplicity speeds up the recognition process and successfully avoids the overfitting of AutoWindLoc.

Moreover, AutoWindLoc has an F-score of 0.9702, which is an improvement over the work described in (He et al., 2025). However, this is not the primary advantage. The most significant advantage of these models lies in their database and the fully automated process of recognizing wind energy sites. This development establishes a sustainable and valuable data foundation for the expansion of wind energy in Lower Saxony and has the potential to be extended beyond the region. The uniqueness of this fully automated process is underscored by the fact that no other database has been developed with such a sophisticated and comprehensive approach to recognizing wind energy sites.

However, AutoWindLoc is not without its limitations. Its efficacy is dependent on the data formats available in terms of resolution per pixel, the overall size of the image, and the need for TrueDop images. This limitation is further compounded by the fact that many orthophotos or satellite images are inaccessible due to restrictions on free availability or are financially prohibitive to purchase, thereby limiting their overall utility. Consequently, its practical application is constrained to regions where the necessary data is accessible, which encompasses approximately 50 % of Germany.

6. Conclusion

The present paper proposes a framework capable of recognizing and processing wind turbines from orthophotos. The data is meticulously prepared and processed using the Django framework and the PostgreSQL database. This approach ensures that the data obtained can be readily retrieved at a future date.

The recognition process, which is two-stage and based on a YOLO model and a binary classification classification CNN, was developed for this study. This approach has been demonstrated to yield precision and recall values of approximately 97 %, thereby ensuring the integrity and stability of the database with few errors. For instance, the framework exhibits a deviation of 0.85 m compared to a mean distance error of 4.76 m to 8.13 m, signifying a substantial enhancement. It is important to note, however, that the actual value may vary due to the manual definition of the locations.

This work has the potential to significantly improve the data basis for wind turbines in Lower Saxony by using this framework. The federal government wants to increase the wind energy output in Germany in the next 5 years (Sander, 2025), which will be facilitated by the framework.

It is imperative to note that orthophotos intended for processing must be of superior quality and in TrueDOP format to ensure optimal results. Failure to meet these criteria can substantially compromise the deviation of the recognition process.

The potential exists for the project to be expanded even further in the future. The processing of orthophotos from all regions of Germany is a potential avenue for expansion, given that all newly generated orthophotos are required to adhere to the same standards as the established framework. A subsequent expansion throughout Europe would also be a plausible option.

Acknowledgements

This research was conducted as part of a base.camp student project at the University of Hamburg.

References

Arbeitsgemeinschaft der Vermessungsverwaltungen der Länder der Bundesrepublik Deutschland (AdV), 2019. Leitfaden zur Qualitätssicherung von True Orthophotos (TrueDOP). Arbeitsgemeinschaft der Vermessungsverwaltungen der Länder der Bundesrepublik Deutschland (AdV), adv-online.de/AdV-Produkte/Standards-und-Produktblaetter/Standards-der-Geotopographie/binarywriterservlet?imgUid=1e220307-0b71-ee71-7657-80b6a757628auBasVariant=11111111-1111-1111-1111-111111111111 (10 April 2025).

Bundesamt für Kartographie und Geodäsie, 2018. Geografische Gitter für Deutschland in UTM-Projektion (GeoGitter national). Bundesamt für Kartographie und Geodäsie. gdz.bkg.bund.de/index.php/default/geographische-gitter-fur-deutschland-in-utm-projektion-geogitter-national.html (10 April 2025).

Department for Business, Energy & Industrial Strategy, 2013. Renewable Energy Planning Database. Department for Business, Energy Industrial Strategy. data.barbour-abi.com/smart-map/repd/desnz (9 April 2025).

Django Development Team, 2025. Django Documentation, Version 5.1. Django Software Foundation. docs.djangoproject.com/en/5.1 (31 March 2025).

He, T., Hu, Y., Li, F., Chen, Y., Zhang, M., Zheng, Q., Jin, Y., Ren, H., 2025. Mapping land- and offshore-based wind turbines in China in 2023 with Sentinel-2 satellite data. *Renewable and Sustainable Energy Reviews*, 214, 115566.

Jocher, G., Chaurasia, A., Qiu, J., 2023. Ultralytics yolov8, version 8. Ultralytics, github.com/ultralytics/ultralytics, (11 April 2025).

Jocher, G., Qiu, J., 2024. Ultralytics yolo11, version 11.0.0. Ultralytics. github.com/ultralytics/ultralytics (11 April 2025).

Kleebauer, M., Braun, A., Horst, D., Pape, C., 2024. 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*, 7863–7867.

LGLN Open Geodata , 2023. Digitales orthophoto dop. Landesamt für Geoinformation und Landesvermessung Niedersachsen. ni-lgln-opengeodata.hub.arcgis.com/apps/lgln-opengeodata::digitales-orthophoto-dop/explore (31 March 2025).

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollar, P., 2017. *Focal Loss for Dense Object Detection*. IEEE, Venice, 2980–2988.

Markstammdatenregister, 2019. Bundesnetzagentur, markstammdatenregister.de (19 March 2025).

Minière, A., von Schuckmann, K., Sallée, J.-B., Vogt, L., 2023. Robust acceleration of Earth system heating observed over the past six decades. *Scientific Reports*, 13(1), 22975.

Olk, J., 2023. Hohe Preise bedrohen den Standort Deutschland. Handelsblatt, handelsblatt.com/politik/deutschland/energiekosten-hohe-preise-bedrohen-den-standort-deutschland/29361760.html (11 April 2025).

PostgreSQL Development Team, 2000. PostgreSQL Documentation. PostgreSQL Global Development Group. postgresql.org/docs/ (31 March 2025).

Sander, L., 2025. Wie läuft der ausbau erneuerbarer energien in deutschland? NDR. ndr.de/nachrichten/ndrdata/Wie-laeuft-der-Ausbau-von-Solar-Windkraft-Batteriespeicher-Erneuerbare-Energien-in-Deutschland,erneuerbare104.html (4 April 2025).

Sliz-Szkliniarz, B., Eberbach, J., Hoffmann, B., Fortin, M., 2019. Assessing the cost of onshore wind development scenarios: Modelling of spatial and temporal distribution of wind power for the case of Poland. *Renewable and Sustainable Energy Reviews*, 109, 514-531.

TzuTa Lin, 2015. labelimg. HumanSignal. github.com/HumanSignal/labelImg (31 March 2025).

Umweltinformationsgesetz (UIG), 2005. Bundesministerium der Justiz, gesetze-im-internet.de/uig/_2005/__3.html (10 April 2025).

wind-turbine.com, 2025. Karte aller Windkraftanlagen in Deutschland. wind-turbine.com, wind-turbine.com/tools/wkamap (7 April 2025).

Zhai, Y., Chen, X., Cao, X., Cui, X., 2024. Identifying wind turbines from multiresolution and multibackground remote sensing imagery. *International Journal of Applied Earth Observation and Geoinformation*, 126, 103613.

Zhang, S., Wang, F., Hou, Y., Wang, J., and, J. G., 2024. Global offshore wind turbine detection: a combined application of deep learning and Google earth engine. *International Journal of Remote Sensing*, 45(18), 6601–6623.