Autonomous Inspection and Reliability Assessment of Renewable Energy Assets

James J. Slade Honors Thesis

Dev J. Patel
Rutgers Industrial Engineering
Rutgers University, New Brunswick

Under the direction of
Dr. Ahmed Aziz Ezzat

I. Abstract

Despite the promising potential, the large-scale integration of wind energy into modern-day electricity systems is still hampered by its relatively high operations and maintenance (O&M) expenditures. In recent years, there have been numerous efforts to reduce these costs with emerging technologies. The most promising method seems to be the combination of drone-based inspection/maintenance with deep learning-based approaches. The goal of this research is to develop a fully autonomous deep learning-based approach that can accurately classify the type of damage on a wind turbine blade, given a set of drone images. The fundamental piece in the deep learning method proposed herein is the YOLOv4 model, a powerful deep learning object detection framework. Tested on a real-world drone inspection dataset, the proposed approach achieved up to 90% accuracy in correctly classifying the inspection dataset, which was found to be 22% better than an existing deep learning benchmark model.

II. Introduction

A. Background and Motivation

In their quest towards mitigating climate change, several countries have been heavily investing in renewable energy assets such as wind and solar farms. In particular, wind energy has shown to be promising in that regard, as evident by the rapid growth in its adoption worldwide. Despite this, one major obstacle facing this transition is the Levelized Cost of Energy (LCoE), one of the most commonly used metrics for assessing energy projects' financial viability. In its simplest form, LCoE represents the average revenue per unit of electricity generated that would be required to recover the costs of developing and operating a plant during an assumed financial life" [1]. Today, the Levelized Cost of Energy (LCoE) poses the overall driver of the wind energy sector, encompassing operation and maintenance (O&M) costs, which account for up to 20-25% of the total LCoE cost [2]. Inevitably, these costs pose a limitation to the scale and dependability of wind energy turbines.

A fundamental task in O&M is wind turbine insepction, which is typically conducted manually by wind farm technicians, who have to climb the wind turbine, typically sitting at 262 feet high. Not to mention to inspect the blades, which are 120 feet long, is labor-intensive and risky. Offshore wind farms are even more difficult to access, and thus operation & maintenance (O&M) costs of offshore wind farms tend to be higher than onshore wind farms. An emerging alternative is to leverage the advancements in drone technologies to perform manual inspection tasks. Currently, drone inspections can cover up to 10 or 12 turbines daily, reviewing each blade between four to nine minutes. This compares to a manual inspection rate of two to five turbines a day [7]. Drones can deliver clear, precise imagery of the entire blade, with up to 40 megapixels for ultra-high resolution [6]. Moreover, drones today are equipped with enhanced sensors, grabbing arms and probes, and lightweight batteries, allowing them to fly longer, act independently and replace dangerous tasks [4]. When using technologies like LiDAR, detailed measurements of any defects can be obtained with high resolution. These images can then be analyzed to provide actionable inspection results.

Towards that, there have been numerous efforts in recent years to reduce O&M inspection costs with emerging technologies such as robotic inspection, Unmanned Aerial Vehicle images, coupled with either manual classifications and/or machine learning approaches. One method that proves promising within machine learning is a deep learning-based approach. The goal of this research is to develop a fully autonomous deep learning-based approach that can accurately classify the type of damage on a wind turbine, given a set of drone images. Considering offshore wind farm locations, which can be challenging to access due to unfavorable weather and marine conditions, using drones equipped with machine vision can potentially allow farm operators to efficiently detect wind turbine damages at minimal inspection costs. This can help prevent major failures that can result in significant downtime and lost revenues. This is evident by a recent study by Bloomberg New Energy Finance (BNEF) which estimates that the use of drones on offshore wind farms in Europe could shave off more than $1,000 per turbine per year in inspection costs (reducing the cost of producing electricity by 1 percent)" [4]. These cost-savings show the importance of developing a fully autonomous deep learning-based approach, which may help alleviate the cost limitations of developing more renewable energy assets.

B. *Overview and Paper Breakdown*

The deep learning method proposed herein is based on the YOLOv4 model, a deep learning object detection framework. YOLO stands for You Only Look Once and is based on a single Convolutional Neural Network (CNN). CNN divides an image into regions, and then, predicts the boundary boxes and probabilities for each region. YOLO sees the entire image during training and testing time, so it implicitly encodes contextual information about classes as well as their appearance [5]. In conjunction with this CNN-based architecture, Roboflow was employed to pre-process and add augmentations to the images. Augmentations help generalize the learned model and enhances its performance for unseen test images.

Thus, our method can be summarized in the following set of steps: First, drone images are imported and processed. Then, the following step is to use annotation software to draw bounding boxes on the images manually. After that, the images were uploaded into Roboflow, where the images were resized and augmented. Then, the YOLOv4 model is trained on the annotated and augmented images. Within the YOLOv4 framework, the CNN layers are deep, converting images into high-level spatial features called feature maps. YOLOv4 uses a single-stage deep learning object detector. After the training, the deep learning model is run against the test images to assess its performance. When the deep learning model was run against the test images, the metrics used to assess the model's performance are the mean average precision (mAP) and accuracy. The mAP for leading edge erosion, a specific type of blade defects considered herein, was 84.13%, and for no damage, the mAP was 88.32%, while the accuracy metric was 90% in classifying drone images into defective versus healthy.

III. Experimental

A. *Data*

The dataset that was used in this research is the DTU Drone Inspection Dataset. The dataset is hosted at [3] and has 559 high-resolution images. This dataset contains temporal inspection images of 2017 and 2018 covering the "Nordtank" wind turbine located at DTU Wind Energy's test site at Roskilde, Denmark. For the current research, a subset of the images was used. The DTU Drone Inspection Dataset contained images with a variety of object classes. These classes include vortex generator panel missing teeth, vortex generator panel, leading edge erosion, and lightning receptor. Vortex generator panel missing teeth and leading edge erosion were two different defects on wind turbines. However, for this research, the focus was specifically on leading edge erosion as it was a defect that was most easily recognizable to the human eye.

B. Methods

To begin working with the YOLOv4 model, the first step was to define the object detection classes. Therefore, the two classes that were defined: leading-edge erosion (LE Erosion), and no damage. Leading edge erosion is often caused by raindrops, hailstones or other particles impacting the leading edge of the blade. In turn, this causes the material to be removed off the surface of the blade, leaving a rough profile that degrades the performance of the blade over time. Figure 1 below provides examples of each of these classes, the left image in Figure 1 shows a wind turbine blade with leading edge erosion and the image on the right is of a wind turbine blade with no damage.

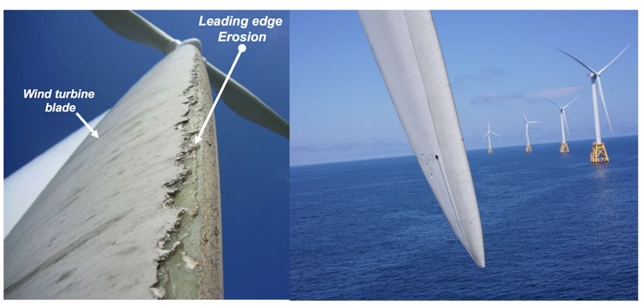

Figure 1: Example of leading edge erosion and no damage

After defining the different classes, the images were uploaded into Labelbox; this is where the bounding boxes and labels were annotated onto each image. Figure 2 below shows what the images look like with the bounding boxes annotated onto each image.

Figure 2: Examples of manually annotated damages related to wind turbine blades

Figure 3 below summarizes the number of annotations for the two classes. After the annotations were complete, the images were then uploaded to Roboflow. Within Roboflow, the images were split; 70% were used for training, 20% for validation, and 10% for testing, and preprocessing steps such as auto orientation and resizing were applied.

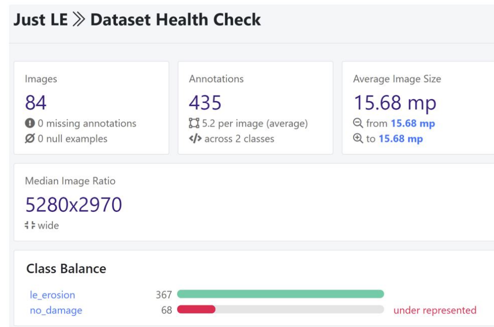

Figure 3: Number of annotations for each class in subset of data

Since the dataset had a limited number of images, some damage types had less representation. However, to effectively train a deep learning model, the image dataset needs to be adequately balanced, so that the model can recognize the distinctive damage features, and generalizes well to the testing images. Therefore, using Roboflow, the images went through augmentations, such as a horizontal flip, 90-degree rotation, crop, shear, saturation, and brightness. Along with augmenting the image, the bounding box annotations were also augmented with the same augmentation steps.

The model architecture used for this research is based on the YOLOv4 model. YOLOv4's architecture comprises of a CSPDarknet53 backbone, spatial pyramid pooling additional module, PANet path-aggregation neck, and YOLOv3 head. The spatial pyramid pooling block is added over CSPDarknet53 to increase the receptive field and separate out the most significant context features. The PANet is used as the method for parameter aggregation for different detector levels [5]. The process by which CNN learns about the features is during the training process. A CNN contains multiple layers, where the information is passed from one layer to the next. When the image reaches the final layer, a feature map is created for that image.

The mean average precision (mAP) metric was used for the validation and testing of images. YOLOv4 models work by predicting a bounding box. Therefore, the prediction is considered correct if the overlap between the predicted bounding box and the ground truth exceeds a certain threshold. Intersection over Union (IoU) is the metric used to measure the overlap. For this research, the threshold used is 0.2; thus, the prediction is considered true if the value is more than 0.2. After IoU is calculated, the next step is to determine the precision, which is the number of true positives over the total number of objects. Then the average precision is calculated by taking the sum of the precision over the total number of images.

Another metric that was used to judge the model's performance was an accuracy metric. This metric measured whether the model correctly predicted if there was damage (or no damage) in an image when the ground truth contained damage (or no damage), respectively. The accuracy metric does not consider the spatial location of the damage but only the classification task.

Regarding the software aspect of the research, PyTorch deep learning API was used in Google Collaboratory, which is like Jupyter notebook but with the added benefit of running the code on the cloud and having access to GPU computing. Deep learning models require heavy computational abilities; therefore, using Google Collaboratory helps bridge the computational power gap.

C. Results

Below Figure 4 one can see the training loss for bounding boxes predicted versus the bounding boxes from the ground truth with regards to the IoU threshold during the training process. From this, we can see that if we were to run the model for a longer period, then the loss would further decrease. However, due to computational limits within Google Collaboratory, we were not able to increase training time. The study also posed a limitation to the number of images one can run, as that would also increase the training time.

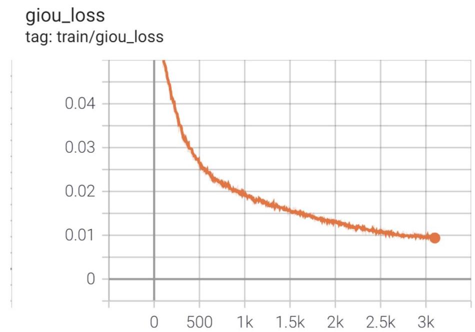

Figure 4: GIoU Loss

When the data was fed into the custom YOLOv4 model, the mAP was as follows: no damage 88.32% and LE erosion 84.13%. The accuracy metric was 90%. In addition, to further assess the model's performance, the data was also fed into a YOLOv4 that was not tuned to the data set and a faster R-CNN model that was tuned to the data set. Faster R-CNN uses a two-stage deep learning object detector: first, it identifies regions of interest and then passes these regions to a convolutional neural network. These results are shown in Table 1. Figure 5 also shows the proposed YOLOv4 model's output on two test images containing an LE defect (left) and no damage (right), respectively.

Table 1: Comparison of Results
Bold-Face Value Indicate Best Performance

	mAP		Accuracy
	No Damage	LE Erosion	
YOLOv4	**88.32%**	**84.13%**	**90%**
YOLOv4 Benchmark	45.28%	34.87%	70%
Faster R-CNN	37.05%	26.62%	50%

Figure 5: YOLOv4 Model output

D. Discussion

After analyzing the results, it is evident that the YOLOv4 model performs better than the Faster R-

CNN model. One reason for this success could be because the YOLOv4 model is a single-stage detection model, and it simultaneously predicts multiple bounding boxes and probabilities for those classes. Similarly, unlike the Faster RCNN, YOLO models see the entire image during training, while in Faster R-CNN, only certain parts of the image are seen. Not only that, but being that the damage we are trying to detect is relatively small compared to the rest of the image, YOLOv4's spatial pyramid pooling block is extremely beneficial because it helps separate out the most significant context features.

The high accuracy value shows that even if the model may not be correct with regards to exactly pinpointing the spatial location of the damage, it is generally accurate when it comes to the classification of whether there is damage or no damage. Future research can look into increasing the mAP at even higher IoU thresholds, which may necessitate more training images, naturally increasing the computational cost. Nevertheless, this research shows that, even with computational limitations, the model is still relatively accurate and brings drone-based inspection of renewable energy assets one step closer into full autonomy.

IV. Conclusion

O&M costs constitute a major determinant of wind's LCoE. One promising route to reduce those expenditures is by enabling autonomous inspection of wind turbine assets. In this research, we demonstrated the promise of a drone-based deep learning approach to autonomously inspect wind turbine blades. Starting from image data processing, augmentation, labeling, we adopted a deep learning architecture which is fundamentally based on YOLOv4 model, a deep object detection framework. Our numerical experiments show that this approach can detect leading edge erosion defects with an accuracy of up to 90%, significantly better than existing benchmarks in the related literature.

Although there have been multiple attempts to reach autonomous inspection of renewable energy assets, one of the major limitations is the lack of data available. When building a deep learning model, one of the key determinants of model performance is whether there is enough data and good quality data. We must also consider training cost as well; as the amount of data increases in size, the cost to train a model increases. As research into deep learning applications with drone inspection of wind turbines increases, we hope to see more data available. As more data is available, the more represented the damage types are. With increased representation, we anticipate that the model's performance will further improve. Also, as drone-based inspection is adopted to inspect wind turbines, there is room to research the implications of adapting drone-based inspection to solar farms. Especially as drones are becoming equipped with advanced sensors, like infrared cameras, they can collect various types of data, which can be leveraged in a data fusion framework to further boost the economic outlook of renewable energy assets.

V. Acknowledgment

The guidance and support provided by Dr. Ezzat were crucial for the completion of this research.

References

[1] U.S. Energy Information Administration. (2020). Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2020, 1–22.

[2] Shihavuddin, A., Chen, X., Fedorov, V., Nymark Christensen, A., Andre Brogaard Riis, N., Branner, K., Bjorholm Dahl, A., Reinhold Paulsen, R. (2019). Wind Turbine Surface Damage Detection by Deep Learning Aided Drone Inspection Analysis. Energies, 12, 676.

[3] Shihavuddin, A.; Chen, X. DTU—Drone Inspection Images of Wind Turbine; This Dataset Set Has Temporal Inspection Images for the Years of 2017 and 2018 of the Nordtank Wind Turbine at DTU Wind Energy's Test Site in in Roskilde, Denmark.

[4] Bullard, N., & Curry, C. (2018, September 7). Why Power Companies Love Drones. Retrieved from https://www.bloomberg.com/opinion/articles/2018-09-07/drones-poised-to-save-big-bucks-for-power-companies

[5] Bochkovskiy, Alexey & Wang, Chien-Yao & Liao, Hong-yuan. (2020). YOLOv4: Optimal Speed and Accuracy of Object Detection. Available at https://arxiv.org/abs/2004.10934

[6] Measure,"Drones for Wind Turbine Inspections," Measure White Paper, Measure.com, 2021.

[7] J. Deign, "Fully automated drones could double wind turbine inspection rates," Reuters Events Renewables, September 27, 2016. Available at https://www.reutersevents.com/renewables/wind-energy-update/fully-automated-drones-could-double-wind-turbine-inspection-rates