

Barely-Visible Surface Crack Detection for Wind Turbine Sustainability

Sourav Agrawal, Isaac Corley, Conor Wallace, Clovis Vaughn, and Jonathan Lwowski



Fig. 1: Can you spot the cracks in these images? Hairline cracks are barely visible and commonly confused with stains and markings from dirt or grease. However, these defects can lead to catastrophic and irreparable damage to wind turbines if left untreated. This figure contains sample tiles from our Zeitview Crack Detection (ZVCD) dataset of annotated barely-visible surface cracks collected from inspections across the globe. This dataset is used to train our turbine crack detection pipeline for the inspection of renewable wind energy sources.

Abstract—The production of wind energy is a crucial part of sustainable development and reducing the reliance on fossil fuels. Maintaining the integrity of wind turbines to produce this energy is a costly and time-consuming task requiring repeated inspection and maintenance. While autonomous drones have proven to make this process more efficient, the algorithms for detecting anomalies to prevent catastrophic damage to turbine blades have fallen behind due to some dangerous defects, such as hairline cracks, being barely-visible. Existing datasets and literature are lacking and tend towards detecting obvious and visible defects in addition to not being geographically diverse. In this paper we introduce a novel and diverse dataset of barely-visible hairline cracks collected from numerous wind turbine inspections. To prove the efficacy of our dataset, we detail our end-to-end deployed turbine crack detection pipeline from the image acquisition stage to the use of predictions in providing automated maintenance recommendations to extend the life and efficiency of wind turbines.

All authors are with the AI Research team at Zeitview {sourav.agrawal, isaac.corley, conor.wallace, clovis.vaughn, jonathan.lwowski}@zeitview.com

I. INTRODUCTION

Wind energy stands as a pivotal component in the pursuit of sustainable development, playing a crucial role in mitigating the impacts of climate change and reducing reliance on fossil fuels. However, the efficient and long-term operation of wind turbines, which are at the heart of wind energy production, is encumbered by the challenges of maintaining their structural integrity. Traditional methods of turbine inspection are not only time-consuming and costly but also fall short in detecting minute yet critical anomalies, such as barely-visible surface cracks, that can lead to irreparable damage and catastrophic failures costing upwards of \$5-8M to replace a modern 3.3MW onshore wind turbine [1], [2]. The advent of autonomous drones has revolutionized this process, offering a more efficient means of monitoring the condition of turbine blades. Nevertheless, the true potential of this technology is hindered by the limitations of current machine learning models, such as object detection and segmentation, often requiring extensive post-processing to eliminate erroneous predictions when applied in real-world scenarios.

A. Background

Surface Crack Types Among the different classifications of damages wind turbine blades experience, structural cracks are one of the most critical yet also generally undetectable even with routine external inspections performed by humans [3]. Blade cracks can originate in many ways and their impact to the blade structure may vary. Coating cracks may occur due to impact from foreign objects, fatigue, or secondary to an underlying structural defect such as delamination. Laminate cracks can also occur from prolonged fatigue and overload events, or they can form over time due to manufacturing defects such as wrinkles in the laminate layers or gaps in the core materials [4].

Once a crack forms, it is difficult to predict how quickly it will grow. Failure rates can occur within hours of detection to years depending on size and location. The structural significance to the blade is difficult to determine without referencing the unique specifications of each particular blade model, which is not always possible [5]. As a result, it is important to identify and classify cracks as soon as possible so that the appropriate repairs can be made and further damage can be mitigated.

In most cases, cracks will appear as black lines on a blade that is painted white, but the length and width of the line can vary significantly. These cracks may only be a few centimeters long and less than a millimeter wide in some cases, making identification difficult even on a clean blade that is fully painted white [6]. Furthermore, detection can suffer when the blade is dirty or when image quality is poor. One of the greatest sources of a false identification is a dirt or grease mark that has streaked across the blade. There are subtle visual differences between these cases and true cracks, resulting in automated methods which struggle with overproducing false positives.

Severity Categorization There is no universal industry standard for blade defect categorization, however, there is a commonly accepted *best practice* which classifies damages into one of five ordinal severity levels [7], ranging from low to high risk of existing or potential future damage to the blade structure. Cosmetic defects and surface-level damage generally fall within the *Severity* [1-2] range. *Severity* 3 can encompass larger surface-level defects, or early, low risk structural damage. *Severity* [4-5] defects are generally considered to be moderate to high severity structural damage. Cracks typically fall within the *Severity* [3-5] range, with emerging cracks, referred to as *hairline cracks*, often falling under the lower end of that range [8].

Related Work While there are numerous prior works in the field of external anomaly and defect detection in wind turbine blades [9], [10], [11], few of them focus specifically on crack detection [12], [13], [14], [15]. Other works utilize infrared cameras for detecting cracks which can provide insights not apparent in the visible spectrum [16], [17], [18]. However the hardware to effectively make use of these thermography-based algorithms are more costly than RGB cameras, hence,

their slow adoption in industry [19].

However, all of these works focus specifically on easily visible and detectable surface cracks. Our work focuses specifically on barely visible and high severity hairline cracks which can easily go undetected until it is too late and catastrophic damage occurs to the wind turbine between inspections that is beyond repair.

B. Contributions

In this work, we address the aforementioned challenges by introducing a comprehensive approach that combines the capabilities of autonomous drones with machine learning specifically for detecting barely-visible cracks in wind turbine blades. Our approach involves all aspects of machine learning from dataset curation to model training, resulting in an accurate and end-to-end intelligent system for automated inspection of wind turbine blades. We detail a fully-deployed pipeline for more reliable and efficient maintenance of wind turbines, ultimately contributing to the sustainable production of wind energy. Our contributions are as follows:

Zeitview Crack Detection (ZVCD) Dataset

Throughout several years of inspections and numerous turbines analyzed for defects, we have compiled a novel, large, high severity surface crack detection dataset. We find a dataset of this size, geographical diversity, and quality, key in producing machine learning models which can accurately detect barely-visible turbine blade cracks.

Deployed Turbine Crack Detection Pipeline We detail Zeitview's end-to-end crack detection pipeline, providing industry insights for how to deploy a model trained on our ZVCD dataset at scale in real-world wind turbine inspection scenarios.

II. METHODS

A. Dataset

There is a scarcity of existing turbine anomaly detection datasets. Many of the recently proposed datasets are proprietary making them difficult to benchmark against. Additionally, several prior works suffer from issues such as only containing simulated imagery or using intensive image augmentations to artificially inflate the dataset size. Furthermore, of the few datasets that are available, crack detection is only a subset of the anomaly categories making it a relatively rare defect, further limiting the amount of available data to train and evaluate against. Rarer still are hairline cracks which are barely visible and only occur on a small subset of turbines. This makes it incredibly difficult to collect a comprehensive dataset. For example, after inspecting roughly 35,000 turbines, we expect to observe approximately 3,000 turbines with hairline cracks.

ZVCD dataset Hairline cracks have a wide variety of appearances ranging from visible and great in length to nearly invisible and short in length. Wind turbines can also vary in model type and manufacturer as well as location across the globe. Machine learning models trained to detect

Dataset	# Cracks	Barely Visible	High Severity	Real	# Locations	# Turbines	# Manufacturers	# Models
Blade30 [20]	9	×	×	✓	6	-	-	-
IDR [21]	12	×	×	✓	1	-	-	-
UAD [22]	30	×	×	✓	1	-	-	-
ADO [12]	50	×	×	✓	1	33	-	-
IEM [23]	55	×	×	✓	1	-	-	-
IRO [24]	<200	×	×	✓	1	-	-	-
CAI-SWTB [25]	258	×	×	×	1	-	-	-
DOC [26]	2,189	×	×	✓	1	-	-	-
ZVCD (Ours)	9,107	✓	✓	✓	988	4,684	34	195

TABLE I: **Comparison of turbine blade crack detection datasets.** Existing datasets suffer from several issues such as being simulated or using image augmentations to artificially inflate the dataset size. Further, these datasets typically contain visibly obvious and/or low severity defects by industry standards, creating them infeasible for detecting barely-visible surface cracks. Our proposed ZVCD dataset consists of real inspection imagery from a diverse set of geographic sites and turbine models and manufacturers. Note that some literature does not disclose how many unique turbines, model, and manufacturers are captured in their datasets. Some datasets use image augmentations to artificially inflate number of images in the dataset, therefore we report number of unique cracks in each dataset.

defects need to generalize to these variations and be robust enough to learn to extract features only on the blade. Therefore, we create a comprehensive, geographically diverse dataset of barely-visible cracks from wind turbine inspections which we refer to as the Zeitview Crack Detection (ZVCD) dataset. Figure 2 displays a map of the turbine farm locations captured in our dataset. Table I describes the attributes of our dataset compared to pre-existing datasets found in the literature. We observe a much greater diversity of wind turbines and locations, thus allowing our models to generalize better when deployed in a real-world inspection scenario.

B. Models

Crack Detection Modeling in Practice While segmentation and object detection methods are popular for detecting turbine defects and potentially provide more accurate localization, we instead choose to utilize classification for this task for the following reasons:

- Turbine inspection imagery are commonly acquired at high resolutions resulting in images too large to use with modern machine learning models, therefore images are patched into sub-tiles making localization of detections not as critical because of the known location of the sub-tiles.
- Segmentation and Object Detection are sensitive and commonly require significant post-processing to remove erroneous predictions and false positives.
- Segmentation and object detection labels are more costly to annotate in comparison to classification.

While classification isn't as appealing as the alternatives, we find in practice that crack detection classifiers operating on sub-tiles of an image dramatically reduce the amount of false positive detections while also maintaining efficiency for use with on-board hardware or deployed on a server.

Model Architectures In our experiments we train the following classifier models on our ZVCD dataset. We select several models ranging in parameter sizes while targeting

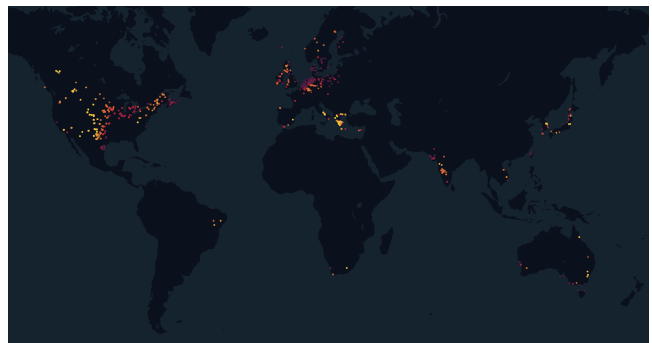


Fig. 2: **Geographic locations of the inspections in our ZVCD dataset.** Our dataset consists of a diverse set of imagery of barely-visible surface cracks acquired from 4,684 turbine inspections from nearly 1,000 unique sites across the world.

efficiency for use on-board drone hardware. All models are initialized with ImageNet pretrained weights [27].

- **ResNet-18**: [28] is a canonical ConvNet classifier model which uses skip connections after each block to improve flow of features and gradients.
- **EfficientNet** [29] is a model designed specifically to maintain similar performance to large classifier models while reducing the number of parameters to improve efficiency and usage in hardware constrained settings. In our experiments we utilize the EfficientNet-B3 variant.
- **MobileNetV3** [30] is a model designed using neural architecture search (NAS) methods to both improve performance on mobile hardware while maintaining performance on ImageNet. In our experiments we utilize the MobileNetV3 Large and Small variants.

III. EXPERIMENTS

A. Preprocessing

To generate our datasets for training and evaluation we preprocess the ZVCD dataset described in Section II-A. Our dataset contains roughly annotated polygons and a severity

Backbone	# Params (M)	FLOPs (G)	Runtime (ms)	ImageNet Acc.	Accuracy	F1	Precision	Recall
ResNet-18 [28]	11.2	38.00	15	73.16	70.34	66.94	63.85	70.34
EfficientNet-B3 [29]	10.7	20.59	47	80.87	82.20	74.33	67.83	82.20
MobileNetV3-S [30]	1.5	1.20	11	67.92	81.35	74.42	68.57	81.36
MobileNetV3-L [30]	4.2	4.63	16	75.512	74.58	72.73	70.97	74.58

TABLE II: **Test set performance of classifier models trained on our ZVCD dataset.** We report accuracy, f1-score, precision, and recall. Because missing cracks can be costly, we place more emphasis on precision when determining the best performing model. We further report metrics for model size in millions (M) of parameters and runtime efficiency in milliseconds (ms).

category for each crack labeled by our inspection analysts. We then tile each image into patches with a 25% overlap, keeping all patches that contain a crack and 1% of patches which do not contain a crack. We then filter by severity, removing all *Severity* [1-2] cracks, and further filter easily visible cracks.

B. Dataset Splits

Given our preprocessed dataset, we then split our dataset by turbine, not by image, using a 90/10 train/val split criterion, such that no tile from the same image will be in both sets. We further manually create a very hard subset of our val set as test set which contains the least visible surface cracks. Samples of this set can be viewed in Figure 1.

C. Metrics

For each model variant, we report the original ImageNet accuracy from pretraining for clarity. Additionally we report the accuracy, F1 score, precision, and recall on the ZVCD very hard test set. We place particular importance on precision as our objective is to minimize missed crack detections.

D. Training Details

To keep comparisons fair, we retain the same training details across all model variants. We use an initial learning rate of 10^{-3} paired with the AdamW optimizer [31] and a cosine learning rate scheduler. The dataset is tiled with an image size of 1024×1024 with an overlap of 25%. We train for 50 epochs with no early stopping condition. For data loading, we use a batch size of 8 images and we apply a balanced weighting to the two classes to eliminate imbalance bias towards the negative class.

We use the model architectures implemented in the PyTorch Image Models (timm) library [32] and trained using the PyTorch Lightning library [33]. All models were trained with the same train-time image augmentations: 1) random horizontal flip with $p = 0.5$, 2) random vertical flip with $p = 0.5$, 3) random affine transform with angle between $\alpha = [0, 90]$ and $p = 0.5$, and 4) color jitter with $brightness = 0.1$, $contrast = 0.1$, $saturation = 0.1$, $hue = 0.1$ and $p = 0.25$. We use implementation from the Kornia library [34]. All images are preprocessed by normalizing with ImageNet statistics of $mean = [0.485, 0.456, 0.406]$ and a $std = [0.229, 0.224, 0.225]$. All experiments were run on an AWS compute instance with 8 vCPUs and a Tesla V100-SXM2-16GB GPU.

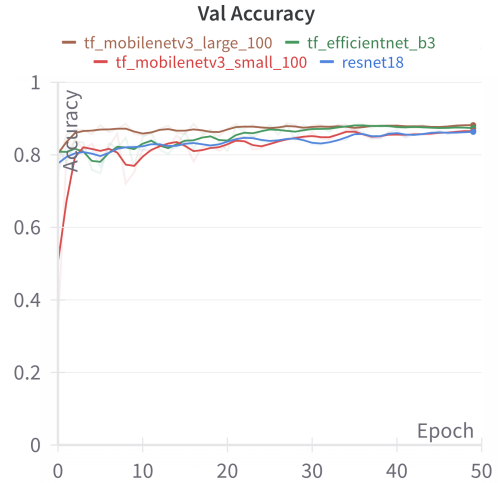


Fig. 3: **Validation accuracy performance during training of different models.** We observe that all the models have converged well and show similar performance.

E. Training Results

We report the metrics for each model on the ZVCD test set in Table II. For additional insight we display the validation accuracy performance throughout training in Figure 3. We further compute efficiency metrics – total number of parameters, floating-point operations per second (FLOPs), and runtime (ms) – for each model using the fvcare [35] library with a batch size of 1 and an image size of $3 \times 1024 \times 1024$.

While these models are by no means bleeding edge, we found that popular state-of-the-art transformer-based architectures like ConvNextV2 [36] and SwinV2 [37] would not converge to the same performance during training, likely requiring significant hyperparameter tuning to reach similar performance to the other models we used.

IV. PIPELINE

In this section, we discuss Zeitview’s currently deployed turbine crack detection pipeline. The pipeline is split into 3 components: image acquisition, inference, and post-processing. The pipeline is entirely automated until the last stage which requires a human-in-the-loop for final review. A visualization of the end-to-end pipeline is provided in Figure 5.



Fig. 4: **Sample Crack Detection Region Proposals using the GradCAM method [38]**. Our pipeline provides region proposals to a human analyst who then reviews and further diagnoses barely-visible surface hairline cracks for severity categorization and repair recommendations. We first compute the GradCAM attributions w.r.t the crack detection class model output. We then post-process the attribution heatmap using the `normalize` and `polygonize` operations.

A. Image Acquisition

For each turbine, we utilize DJI drones to fly a pre-planned flight path which performs a pass of each of a turbine’s three blades on all four sides: suction side, pressure side, trailing edge, and leading edge. We utilize a gimbal to orient the camera in such a way that it is orthogonal to the surface of the blade along the entire pass. For each pass, the drone stops to capture an image of the blade at waypoints spaced approximately three meters apart. This means that for a typical 50 meter blade, we capture 18 images per pass, generating a total set of 216 images per turbine.

B. Inference

After image acquisition, the pilot then uploads the imagery to our cloud infrastructure. Images are then tiled into 1024×1024 patches.

Once the images have been tiled, we run our classifier in batched processing form and filter out low confidence predictions with a threshold, $\tau = 0.5$. Because the images are tiled we have an approximate location on the blade which contains potential cracks. We further localize the cracks into region proposals by computing the gradient attributions w.r.t the crack output class using the GradCAM method [38].

C. Post-Processing

GradCAM attribution heatmaps are typically noisy and require additional post-processing to create an accurate localization of detected cracks. Furthermore, storage of heatmaps in the cloud is costly, therefore post-processing to convert region proposals into polygons is necessary. We perform the following steps to convert the attribution heatmaps to localized polygons and provide visual examples in Figure 4.

Normalizing We normalize the attributions within a tile using the min and max statistics.

Clipping We remove attribution outliers by clipping the normalized heatmap to the percentile range $[0, 98]$.

Contour Detection We utilize contour detection to segment regions with high attribution. We find best results by retrieving only the extreme outer contours and compressing to only obtain the endpoints. We find this as a simple, yet effective method for localizing the cracks within a tile’s heatmap.

Polygonization Given the raster-level heatmaps outlines, we then convert them into N-point polygons.

Segmentation Finally, we utilize our U-Net[39] based turbine outline segmentation model masks to filter polygons with less than 50% overlap in pixels with the blade.

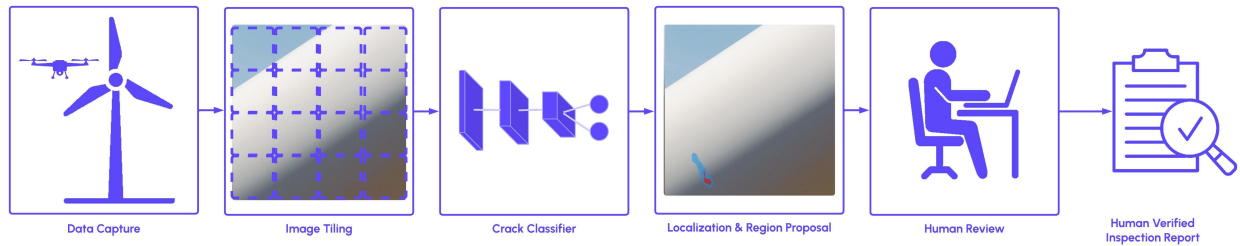


Fig. 5: **Our turbine crack detection pipeline diagram.** Our pipeline consists of automated acquisition of overlapping imagery of each turbine blade using DJI UAVs and our flight planning algorithms. Each image is then decomposed into 1024 x 1024 tiles which we then run our classifier on. Next, we compute and postprocess the gradient attribution heatmap on the high confidence predicted tiles. Our analysts are then shown these proposal regions containing cracks for review and in-depth analysis for inspection report generation.

This model is pretrained for multiclass turbine part segmentation, however, we convert the predictions to binary foreground/background masks for this use case.

D. Human Analysis and Review

We treat the attribution polygons as region proposals where high severity cracks may exist. From here we incorporate a human-in-the-loop system to refine proposed defects into further severity types and tighter localizations. This is also where any potential false positives are removed. We believe that with further research, we can remove the human from the loop entirely, making the system entirely self-sufficient and scale the number of inspections further.

V. CONCLUSION

In this paper we have proposed a comprehensive, barely-visible, high-severity hairline crack detection dataset for wind turbine inspections, which is significantly more diverse and rich in quality than prior literature. We have also detailed our tiled classification-based approach to hairline crack detection with pseudo region proposals using GradCAM attributions and showed how our system is designed in an end-to-end fashion. We have developed our system with the ability to scale to the vastly increasing number of wind turbines being constructed to generate enough renewable energy to support the world population and reduce reliance on fossil fuels.

We have chosen model architectures which have the ability to be easily ported to drone hardware, allowing inspections to be performed in real time. We hope to move towards this process in future iterations of this research.

Furthermore, we note there has been little research into comparing the performance of infrared vs RGB crack detection. While we have collected a sample dataset for this task, we leave the deeper technical evaluation of this for future work.

Barely-visible surface cracks prove to be a challenging and unsolved problem and can be some of the most severely

damaging defects as the turbine is typically only inspected on a yearly basis, giving more than enough time for these minimal cracks to grow and cause catastrophic damage. We hope to further refine our models to improve performance on even the most minute cracks to prevent such disasters from happening. Lastly, we hope this research sheds light to the importance of wind turbine inspection and maintenance to wind energy sustainability.

REFERENCES

- [1] Jun 2023. [Online]. Available: <https://weatherguardwind.com/how-much-does-wind-turbine-cost-worth-it/>
- [2] L. Hartman, "Wind turbines: The bigger, the better," Aug 2023. [Online]. Available: <https://www.energy.gov/eere/articles/wind-turbines-bigger-better>
- [3] L. Mishnaevsky Jr, "Root causes and mechanisms of failure of wind turbine blades: Overview," *Materials*, vol. 15, no. 9, p. 2959, 2022.
- [4] A. Dashtkar, H. Hadavinia, M. N. Sahinkaya, N. A. Williams, S. Vahid, F. Ismail, and M. Turner, "Rain erosion-resistant coatings for wind turbine blades: A review," *Polymers and Polymer Composites*, vol. 27, no. 8, pp. 443–475, 2019.
- [5] D. A. Katsaprakakis, N. Papadakis, and I. Ntintakis, "A comprehensive analysis of wind turbine blade damage," *Energies*, vol. 14, no. 18, p. 5974, 2021.
- [6] W. Wang, Y. Xue, C. He, and Y. Zhao, "Review of the typical damage and damage-detection methods of large wind turbine blades," *Energies*, vol. 15, no. 15, p. 5672, 2022.
- [7] M. Malkin, "A white paper on wind turbine blade defect and damage categorization: Current state of the industry," Oct 2020. [Online]. Available: <https://www.epri.com/research/products/00000003002019669>
- [8] EPRI, "Wind turbine blade maintenance: With a focus on damage and defect categorization," Oct 2020. [Online]. Available: <https://www.epri.com/research/products/00000003002017731>
- [9] I. Gohar, J. See, A. Halimi, and W. K. Yew, "Automatic defect detection in wind turbine blade images: Model benchmarks and re-annotations," in *2023 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*. IEEE, 2023, pp. 290–295.
- [10] I. Gohar, A. Halimi, J. See, W. K. Yew, and C. Yang, "Slice-aided defect detection in ultra high-resolution wind turbine blade images," *Machines*, vol. 11, no. 10, p. 953, 2023.
- [11] C. Zhang, T. Yang, and J. Yang, "Image recognition of wind turbine blade defects using attention-based mobilenetv1-yolov4 and transfer learning," *Sensors*, vol. 22, no. 16, p. 6009, 2022.
- [12] L. Wang and Z. Zhang, "Automatic detection of wind turbine blade surface cracks based on uav-taken images," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 9, pp. 7293–7303, 2017.

- [13] L. Wang, Z. Zhang, and X. Luo, "A two-stage data-driven approach for image-based wind turbine blade crack inspections," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 3, pp. 1271–1281, 2019.
- [14] Z. Xiaoxun, H. Xinyu, G. Xiaoxia, Y. Xing, X. Zixu, W. Yu, and L. Huaxin, "Research on crack detection method of wind turbine blade based on a deep learning method," *Applied Energy*, vol. 328, p. 120241, 2022.
- [15] A. Joshua and V. Sugumaran, "Crack detection and localization on wind turbine blade using machine learning algorithms: A data mining approach," *Structural Durability & Health Monitoring*, vol. 13, no. 2, p. 181, 2019.
- [16] Y. Jia, R.-m. Zhang, W. Zhang, G.-f. Jin, J.-f. Liu, G. Tian, and Z.-w. Yang, "Rapid detection of cracks in turbine blades using ultrasonic infrared thermography," in *IOP Conference Series: Materials Science and Engineering*, vol. 389, no. 1. IOP Publishing, 2018, p. 012023.
- [17] B. E. Jaeger, S. Schmid, C. U. Grosse, A. Gögelein, and F. Elischerberger, "Infrared thermal imaging-based turbine blade crack classification using deep learning," *Journal of Nondestructive Evaluation*, vol. 41, no. 4, p. 74, 2022.
- [18] C. Wang and Y. Gu, "Research on infrared nondestructive detection of small wind turbine blades," *Results in Engineering*, vol. 15, p. 100570, 2022.
- [19] T. X. B. Nguyen, K. Rosser, and J. Chahl, "A review of modern thermal imaging sensor technology and applications for autonomous aerial navigation," *Journal of Imaging*, vol. 7, no. 10, p. 217, 2021.
- [20] C. Yang, X. Liu, H. Zhou, Y. Ke, and J. See, "Towards accurate image stitching for drone-based wind turbine blade inspection," *Renewable Energy*, vol. 203, pp. 267–279, 2023.
- [21] Y. Yu, H. Cao, S. Liu, S. Yang, and R. Bai, "Image-based damage recognition of wind turbine blades," in *2017 2nd International Conference on Advanced Robotics and Mechatronics (ICARM)*. IEEE, 2017, pp. 161–166.
- [22] Y. Wang, R. Yoshihashi, R. Kawakami, S. You, T. Harano, M. Ito, K. Komagome, M. Iida, and T. Naemura, "Unsupervised anomaly detection with compact deep features for wind turbine blade images taken by a drone," *IPSJ Transactions on Computer Vision and Applications*, vol. 11, pp. 1–7, 2019.
- [23] J. Zhang, G. Cosma, and J. Watkins, "Image enhanced mask r-cnn: A deep learning pipeline with new evaluation measures for wind turbine blade defect detection and classification," *Journal of Imaging*, vol. 7, no. 3, p. 46, 2021.
- [24] X. Yang, Y. Zhang, W. Lv, and D. Wang, "Image recognition of wind turbine blade damage based on a deep learning model with transfer learning and an ensemble learning classifier," *Renewable Energy*, vol. 163, pp. 386–397, 2021.
- [25] B. Altice, E. Nazario, M. Davis, M. Shekaramiz, T. K. Moon, and M. A. Masoum, "Anomaly detection on small wind turbine blades using deep learning algorithms," *Energies*, vol. 17, no. 5, p. 982, 2024.
- [26] A. Reddy, V. Indragandhi, L. Ravi, and V. Subramaniaswamy, "Detection of cracks and damage in wind turbine blades using artificial intelligence-based image analytics," *Measurement*, vol. 147, p. 106823, 2019.
- [27] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [29] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," 2020.
- [30] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for mobilenetv3," 2019.
- [31] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [32] R. Wightman, "Pytorch image models," <https://github.com/rwightman/pytorch-image-models>, 2019.
- [33] L. AI, "Pytorch lightning," <https://github.com/Lightning-AI/pytorch-lightning>, 2019.
- [34] E. Riba, D. Mishkin, D. Ponsa, E. Rublee, and G. Bradski, "Kornia: an open source differentiable computer vision library for pytorch," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 3674–3683.
- [35] M. Research, "fvcore," <https://github.com/facebookresearch/fvcore>, 2019.
- [36] S. Woo, S. Debnath, R. Hu, X. Chen, Z. Liu, I. S. Kweon, and S. Xie, "Convnext v2: Co-designing and scaling convnets with masked autoencoders," 2023.
- [37] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, F. Wei, and B. Guo, "Swin transformer v2: Scaling up capacity and resolution," 2022.
- [38] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [39] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.