Electrical grid automated visual asset inspection - an application to power line insulators

Pedro Rocha[1]
a21200219@isec.pt

Fernando Lopes[1,2]
flopes@isec.pt, flopes@co.it.pt

[1] Polytechnic Institute of Coimbra,
Coimbra Institute of Engineering,
Coimbra, PT

[2] Instituto de Telecomunicações,
Coimbra, PT

Abstract

Electric grid asset inspections are critical to ensure that modern society can continue to depend on a stable and uninterrupted supply of electricity. This paper represents an initial investigation to test the efficiency of a deep learning model to address the problem of identifying defects on visible light images of power line insulators. We tested a publicly available dataset with 1688 images containing over 6000 shell insulators, using a Faster R-CNN architecture on the Detectron2 framework. We achieved a mean average recall (mAR@50:95) of 82,6% with a 79,9% mean average precision (mAP@50:95) for the three considered classes.

1 Introduction

Utility companies employ a range of strategies to maintain the availability and reliability of electric power distribution systems [1]. Common practices include scheduled maintenance routines involving checks, servicing and replacements. These routines proactively identify and address deterioration before it leads to significant issues. Additionally, emergency response plans ensure prompt reactions to unexpected outages, swiftly mobilizing repair crews and efficiently allocating resources. These tasks are facilitated by Digital Asset Management systems, that track equipment conditions and locations, maintenance history and performance data, prioritizing maintenance tasks and allocating resources effectively.

In recent decades, real-time remote monitoring of critical equipment has been implemented, enabling operators to identify anomalies without requiring a physical presence. Skilled technicians and engineers play a vital role in interpreting data, and making informed decisions based on their expertise. Critical equipment often utilizes sensors for real-time assessment of operational conditions. When deviations from normal parameters are detected, it triggers alerts for immediate investigation. By combining sensor information and advanced data analysis, the aim is to predict equipment failures based on performance trends and operating conditions, enabling timely interventions, and subsequently reducing downtime and costs [2]. Some examples are given next.

Infrared cameras are utilized to detect abnormal temperature variations in equipment. Ultraviolet inspection is employed to identify contaminants on the surface of insulators and, in other cases, to detect the ionization of the air surrounding a conductor, a phenomenon known as corona discharge. Ultrasound Technology is employed to detect sound waves emitted by malfunctioning equipment, revealing issues like electrical arcing, leaks, and mechanical faults. LIDAR can detect and analyze vegetation growth around power lines and substations. Drones equipped with sensors such as cameras can inspect power lines, towers, insulators and substations from above, while ground-based robots access confined spaces and hazardous environments for inspections. AI and big data analytics can analyze extensive data to predict equipment failures by identifying patterns, enhancing maintenance decisions. However, the increasing volume of collected data such as images, during infrastructure inspections, creates a bottleneck in the human interpretation task resulting in a expensive and slow processe. Minor inspection responsibilities can be automated using computer vision and deep learning techniques.

2 Literature review

Haiyan Cheng *et al*. [3] proposed a method for self-shattering defect detection of glass insulators based on spatial features using classical computer vision methods. The approach begins by applying a double-limit threshold in the RGB color space to achieve insulator segmentation and simple morphological operations to eliminate noise. The class evaluation involves analyzing the number of pixels in each Region of Interest (ROI), along with region length and the distances between regions, considering specified tolerance levels. Despite the fact that this and similar classical approaches achieve an accuracy exceeding 90%, these techniques do not demonstrate generalization across various insulator materials with differing colors. The primary constraint lies in their capacity to solely identify the absence of a skirt in the insulator string, while overlooking the detection of other prevalent issues like surface contamination.

High-resolution aerial images captured by UAVs with intricate backgrounds and small target detection present a critical challenge for insulator defect detection. Qiaodi Wen *et al*. [1] overcome these complexities with two deep learning methodologies, Exact R-CNN and CME-CNN, both based on Faster R-CNN. Exact R-CNN incorporates advanced techniques such as FPN, cascade regression, and GIoU, along with innovative approaches like RoI Align and depthwise separable convolution to enhance accuracy while reducing computational demands. Additionally, CME-CNN introduces an insulator mask extraction network to eliminate background interference and employs Exact R-CNN for defect detection. Experimental results demonstrate the superior effectiveness of these methods, with CME-CNN-ResNet50 achieving an average precision of 88.7%, outperforming various established detection algorithms.

Van Nhan Nguyen, Robert Jenssen, and Davide Roverso [4] identified some of the main challenges in DL vision-based UAV inspection and proposed a set of solutions. Being the first challenge the lack of training data, they collected a total of 28,674 images with a resolution of 6048x4032. By cropping the bounding boxes, they managed to create three additional datasets, resulting in a total of 94,477 256x256 images. Recognizing the time-consuming nature of labeling a large dataset, they used data augmentation to address class imbalance. This included making several crops around the bounding boxes and applying operations such as flipping, mirroring, blurring, adding noise and zooming during training, increasing the dataset size by 12 times. To address the challenge of detecting small power components and small faults, they proposed a multi-stage component detection and classification pipeline. It allows for a 'zoom-in' operation during inspection, enabling the detection of small faults on power line components, such as cracks on poles. Their best results in terms of mAP reached 81.3%, outperforming the other two methods in 7 out of 10 classes and achieving an inference speed of 300 images per minute.

3 Methodology and dataset

Our work focus on the central core of an automated inspection pipeline. The implementation receives high resolution color images, obtained in real life scenarios of electric power transmission and distribution, containing different types and amounts of electric insulators, and generates a localization bounding box with a respective class of the predicted defect.

We used a publicly available dataset [5] that contains over 1600 high-resolution images of transmission line ceramic insulator strings, in a diversified range of colors, captured from various angles and lighting conditions. Each image is annotated using COCO-style annotations, which include the bounding box top-left coordinates, width, height and a class label designation corresponding to the defects zoomed-in in Figure 1: Flashover damage insulator shell, Broken insulator shell or Good insulator shell. We used 5780 single insulator shell assets labeled for training and 300 insulator shell assets spread across 88 images used for testing.

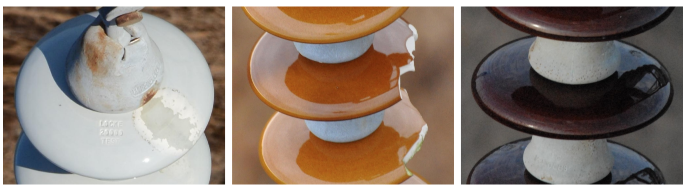

Figure 1: Flashover damage, Broken and Good class examples

4 Model architecture

Our defect detection system relies on a robust computer vision framework developed by Facebook AI Research (FAIR), Detectron2 [6], that stands out for its adaptability and effectiveness across multiple computer vision tasks, including object detection, instance segmentation, and keypoint detection. Its modular design simplifies customization for precise specifications and streamlines data handling with utility functions for tasks such as data augmentation, resizing, training, evaluation, and even architecture selection, making it a comprehensive choice for computer vision projects.

Our architecture selection was a Faster R-CNN model [7] that is well-known for addressing the computational bottleneck of its predecessor, Fast R-CNN, replacing a traditional region proposal algorithm for a Region Proposal Network (RPN) that shares the convolutional features for better inference performance. The images feeding the CNN originate features, such as edges, shapes and object parts, that are encoded in feature maps and processed by the RPN, which tests different box shapes through these maps using larger sliding steps. This process results in a set of anchor boxes, each with a respective score related to the probability of an object's existence. Subsequently, the model applies Non-Maximum Suppression (NMS) to eliminate redundant or highly overlapping proposals. Finally, the classifier block utilizes the proposal regions to efficiently reduce the feature map search. It achieves this through an ROI pooling operation, which converts irregularly sized regions into a fixed-sized format. This approach allows Faster R-CNN to focus its object identification efforts on regions with high probability, rather than inefficiently conducting a sequential or random search.

We obtained similar results by mirroring the approach of [8], that was tested in the same dataset, characterized by three output classes to suit the defect detection task. Preprocessing involved image resizing and pixel value normalization. Transfer learning was used by loading the pre-trained Faster R-CNN weights from the model-zoo community. For hyperparameter configuration, the base learning rate was set at 0.001, a batch size of 4 images per batch and a maximum of 5000 iterations. This learning process was further refined by implementing a schedule with 200 warm-up iterations, dynamically adjusting the learning rate based on the performance. To continually assess the model's progress, periodic evaluations every 1000 iterations were made, allowing to closely monitor its performance trends. Test-time augmentation (TTA) was used during inference to improve prediction accuracy, reducing overfitting and ensuring more robust results by making multiple predictions of multiple input transformations and averaging the predictions.

5 Results and evaluation

We conducted quantitative evaluations using standard metrics, including the mean Average Precision (mAP) and the mean Average Recall (mAR), across typical Intersection over Union (IoU) thresholds within the MS COCO benchmark. As shown in Table 1, the average precision is relatively high, and higher for IoU@50 and IoU@75 threshold limits in all classes. These metrics can not be blindly trusted because of the overlapping objects in this contex. This is why we see a noticeable drop in performance for mAP@50:95, especially in the 'Broken' class. This drop was expected, as higher threshold values only consider more accurately localized decisions. Looking at recall, the 'Broken' class exhibits the highest recall in the initial detection. This suggests that approximately half of the top-scored region proposals are related to broken shell insulators. This result is not surprising, as broken insulators often exhibit more distinct features such as sharp white edges. In the defective classes, within the first 10 detections (in the same image), the model identifies a significant portion of all the objects it will eventually detect. An exception to this trend is observed in the 'Good' or 'No-issues' class, which still experiences a significant increase in recall after the 10th detection. This is primarily because some images contain more than 10 good insulator shells.

Table 1: Scores for metrics on the test dataset

Metric	IoU	max Dets	Class			
			Good	Broken	Flash	All
mAP	@50:5:95	100	0.819	0.726	0.822	0.799
	@50	100	0.966	0.955	0.969	0.951
	@75	100	0.954	0.860	0.965	0.914
mAR	@50:5:95	1	0.108	0.521	0.290	0.305
	@50:5:95	10	0.723	0.781	0.856	0.776
	@50:5:95	100	0.866	0.785	0.870	0.826

Observing examples of wrong detections, like the first 3 images in Figure 2, we conclude that the model fails to detect a broken shell insulator if it doesn't have a brighter, sharp edge visible. The second group of 2 images fails to identify good insulators near broken ones, what is not as critical as failing to identify defective ones. Illustrated by the 6th image is the confirmation that allowing a high number of maximum detections can lead to less accurate predictions, with a wrong class overlapping a correct prediction. Next, the misclassification of a good insulator for a broken one suggests that the elliptical shape of a shell insulator doesn't carry as much weight in the final decision as needed. The model is clearly not well-fitted yet for evening light conditions and prominent shadows due to a lack of training examples like the 8th image. Lastly, translucent glass insulators cannot be correctly detected because of the lack of training examples.

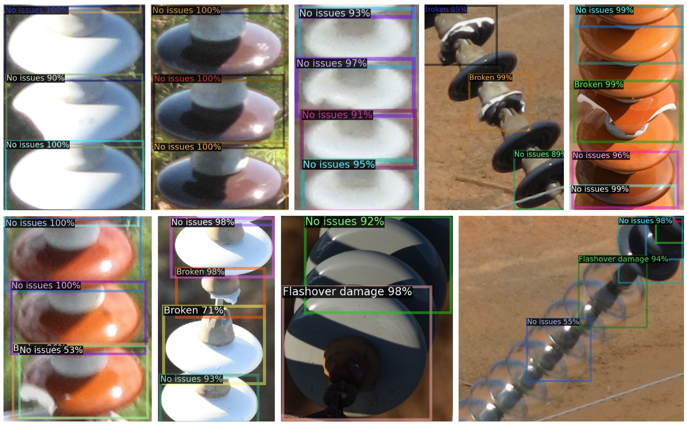

Figure 2: Examples of miss predictions during test inference

6 Conclusions and future work

In this work we tested the use of a deep learning model to detect defects in power line insulators through the analysis of visible light images. The obtained results demonstrate the high practical value of this approach in the context of automated electrical grid asset inspection. After this successful initial investigation we intend to evolve to a broader project in collaboration with a Distribution System Operator (DSO). For the future work, we propose to adapt the model architecture in order to enable it to classify different environmental conditions such as lighting, focus, and object distance. This approach will allow for precise optimization aiming to enhance performance and robustness under varying conditions. We also aim to integrate inspection results with maintenance management.

References

[1] Qiaodi Wen, Ziqi Luo, Ruitao Chen, Yifan Yang, and Guofa Li. Deep Learning Approaches on Defect Detection in High Resolution Aerial Images of Insulators. *Sensors*, 21(4):1033, February 2021.

[2] Fangrong Zhou, Gang Wen, Yi Ma, Hao Geng, Ran Huang, Ling Pei, Wenxian Yu, Lei Chu, and Robert Qiu. A Comprehensive Survey for Deep-Learning-Based Abnormality Detection in Smart Grids with Multimodal Image Data. *Applied Sciences*, 12(11):5336, May 2022.

[3] Haiyan Cheng, Yongjie Zhai, Rui Chen, Di Wang, Ze Dong, and Yutao Wang. Self-Shattering Defect Detection of Glass Insulators Based on Spatial Features. *Energies*, 12(3):543, February 2019.

[4] Van Nhan Nguyen, Robert Jenssen, and Davide Roverso. Intelligent Monitoring and Inspection of Power Line Components Powered by UAVs and Deep Learning. *IEEE Power and Energy Technology Systems Journal*, 6(1):11–21, March 2019.

[5] Dexter Lewis and Pratik Kulkarni. Insulator defect detection. Electric Power Research Institute, IEEE Dataport, 2021. doi:10.21227/vkdw-x769. URL https://ieee-dataport.org/competitions/insulator-defect-detection.

[6] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2, 2019. URL https://github.com/facebookresearch/detectron2.

[7] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. January 2016. arXiv:1506.01497 [cs].

[8] Monkira. Faster r-cnn using detectron2. Kaggle, 2023. URL https://www.kaggle.com/code/monkira/faster-rcnn.