

DRONE-BASED CORROSION DETECTION ON HIGH-VOLTAGE TRANSMISSION TOWERS USING HYPERSPECTRAL IMAGING

Michiel Vlaminck¹, Zakaria Bnoulkacem², Gonzalo Luzardo¹, Hamed Zivariadab³, Zohreh Zahiri³, Bikram Koirala², Frédéric Mangialetto⁴, Irid Bufi⁴, Ljiljana Platisa¹, Murali Jayapala³, Paul Scheunders², Hiep Luong¹

¹ IPI-URC-imec, Ghent University, Ghent, Belgium;
² VisionLab, University of Antwerp, Antwerp, Belgium;
³ imec, Leuven, Belgium;
⁴ Elia, Brussels, Belgium;

ABSTRACT

High-voltage transmission towers require regular inspections to identify corrosion. Traditionally, these inspections are performed through climbing, involving skilled technicians. This method is both tedious and hazardous, often necessitating the shut-down of sections of the high-voltage grid. In this paper we propose a workflow that relies on drone-based hyperspectral imaging, which enables remote assessment without endangering the technician. Currently, drones are equipped with conventional RGB cameras. However, these cameras have limited spectral resolution and range, which compromises their ability to reliably detect corrosion and often leads to false alarms. Moreover, conventional RGB cameras are unsuitable for accurately assessing the severity of corroded areas. To address these challenges, this study proposes a solution that leverages hyperspectral imaging and a dedicated processing pipeline to robustly detect corrosion and classify it based on severity level. Experiments using drones equipped with imec’s VIS-NIR hyperspectral payload demonstrated the effectiveness of our developed solution.

Index Terms— Hyperspectral imaging, UAV, corrosion detection, high-voltage tower

1. INTRODUCTION

High-voltage electrical pylons are often situated in remote or inaccessible locations, making it perilous for human workers to inspect them for corrosion through climbing. The use of drones offers a safe and efficient solution for conducting inspections from a distance. Conventional RGB cameras are commonly used to perform such microview inspections, but their limited spectral range and resolution hinder their ability to resolve the ambiguity between corrosion and visually similar anomalies like dirt or bird droppings. Moreover, RGB cameras are unsuitable for differentiating between different

levels of corrosion severity, as they tend to present a similar appearance regardless of severity. However, severely corroded areas directly impact the structural integrity as they lead to material loss over time. Therefore, it is crucial to be able to classify the severity of corrosion.

In this study, we investigate the use of imec’s VIS-NIR hyperspectral camera to detect corrosion more reliably and to classify it in three levels of severity, being superficial, moderate and severe corrosion, which is directly correlated to the amount of material loss after cleaning. Our approach eventually allows for a more accurate assessment of the impact on the tower’s strength, facilitating timely maintenance and repair decisions.

2. RELATED WORK

In the current literature, many approaches to detect corrosion on infrastructure rely on conventional RGB cameras. In [1], Ortiz et. al. employ micro-aerial vehicles equipped with regular cameras to detect corrosion on vessels. However, their technique relies heavily on the determination of dominant colours and changes in texture, which limits its scalability for other use cases. This limitation becomes particularly apparent in scenarios where there are similar-looking outliers present. A comprehensive overview of computer vision techniques to detect corrosion using RGB images is given in [2].

Regarding the identification of corrosion with hyperspectral technology, very few studies have been conducted in outdoor scenarios. Despite the challenges associated with accurately reconstructing hyperspectral cubes, the potential of this technology far surpasses that of conventional RGB cameras. In [3], Naik et. al. investigated the use of hyperspectral imaging to detect corroded areas on steel plates under visually ambiguous scenarios and to identify the source of corrosion in such scenarios. In [4], Zahiri et. al. evaluated the potential of hyperspectral data in the Shortwave Infrared (SWIR)



Fig. 1. The dual hyperspectral snapshot UAV payload mounted on a DJI M600 pro.

range for the use case of corrosion detection in steel. However, the steel samples were scanned in a controlled laboratory environment, with no need to correct for outdoor illumination changes or to determine a region of interest.

This paper builds on the work presented in [5], where a few samples were collected from two metallic high-voltage towers with two different coatings. They were scanned and inspected for corrosion in laboratory conditions. In this paper we extend the work by integrating a hyperspectral snapshot camera on a drone to perform outdoor measurements near an operational electrical pylon.

3. MATERIAL AND METHODS

The hyperspectral UAV payload used in this study consists of a stereo set-up with two imec hyperspectral snapshot cameras, together providing a range from 470nm to 860nm, hence covering both the visible and NIR range of the electromagnetic spectrum. The VIS and NIR camera both have a 4x4 mosaic filter array resulting in respectively 16 and 15 bands after calibration. There are thus a total of 31 different bands available. The payload is integrated on a gimbal and mounted on a DJI M600 pro, which is visualized in Figure 1. The methodology is divided into three main components: 1) preprocessing of the RAW data to reconstruct hyperspectral cubes 2) segmentation of the tower in the hyperspectral data and 3) corrosion classification for the different segmented regions.

3.1. Preprocessing

The aim of the preprocessing is to restore hyperspectral cubes from snapshot mosaic imaging data. There are two major challenges involved in this reconstruction process. First, spectral corrections are needed to deal with the inherent effects of the sensors, such as harmonics, cross-talks and leakage. Second, the environment is subject to varying incident solar radiation and the measurements depend on the sun-sensor viewing angle. The measurements are therefore a mixture of reflection and diffusion. Finally, objects are often (self-)shadowed, which further influences the reflectance of the material. Usually, a reference panel is utilized to normalize the images for solar irradiation. However, when conduct-

ing drone measurements, it is often impractical to include reference panels within the field of view. Moreover, ensuring that these panels experience identical lighting conditions as the objects of interest, becomes even more challenging. Therefore, we adopt an alternative approach that utilizes radiance data instead of reflectance data. Our method relies on the inherent consistency of paint spectra across various measurements. This enables us to employ the paint spectrum as a reliable reference for calibrating the images before classification, serving as a substitute for the traditional reference panel. Figure 3 illustrates the various stages involved in the restoration of hyperspectral images while a brief description of each is provided below.

Dark noise correction The images of the data cube are sensor-biased due to the offset voltage of the sensor, thermal noise, etc. As a result they need to be corrected before they can be spectrally calibrated. To address this, a channel-wise subtraction of the dark signal is performed.

Demosaicking The process of demosaicking usually involves splitting the image into different spectral bands, followed by spatial interpolation to fill in the missing data. In the course of this project we developed our own advanced CNN-based super-resolved demosaicking algorithm, which outperforms classical demosaicking algorithms.

Radiometric correction Radiometric correction, also known as radiometric calibration, refers to the essential process of converting the abstract digital numbers assigned to the recorded outputs of each pixel and spectral channel into accurate measurements of radiation intensity.

Angular correction The angular correction deals with the correction of the spectral shift in the measurements, which is caused by the sensitivity of the filters to the angle of incidence. The shift is mathematically corrected by considering the distance to the exit pupil of the lens.

Non-uniformity correction Non-uniformity correction tackles the issue of drift caused by the camera's internal heat, which can potentially interfere with its IR radiation readings. To enhance imaging precision, the camera measures the IR radiation emitted from its optics and utilizes those readings to adjust the gain and offset for each pixel individually, resulting in a more accurate image.

3.2. Segmentation

The purpose of the segmentation is two-fold. First, it aims to extract the pixels of the image that belong to the transmission tower itself (referred to as foreground) and separate them from the background pixels. This step allows to isolate and focus on the tower for further analysis. Second, as various subsections of the tower experience distinct lighting conditions, the segmentation also differentiates between shadowed

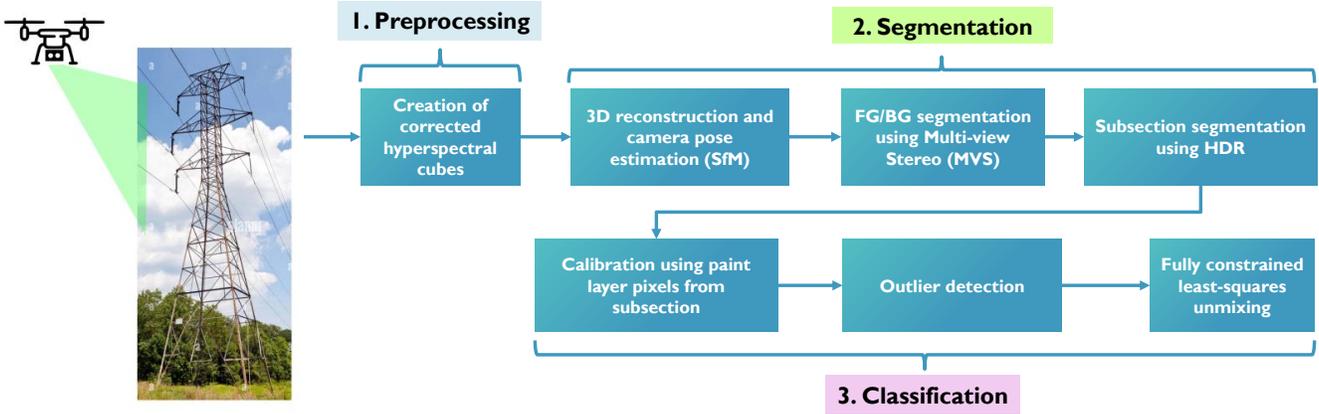


Fig. 2. Overview of the methodology. The primary objective of the segmentation is to isolate and focus on the tower. The classification aims to localize corrosion as well as to classify its severity in three different levels: light, moderate and severe.

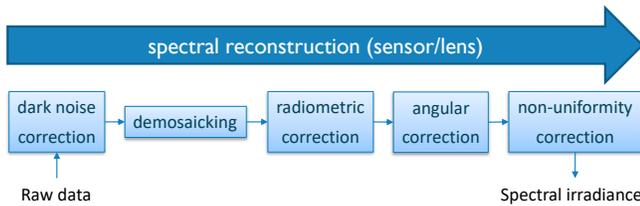


Fig. 3. Overview of the preprocessing pipeline.

areas and areas illuminated by either diffuse or direct sunlight. The latter is done implicitly by adopting varying exposure times similar as to what is done in high dynamic range imaging (HDR).

3.2.1. FG/BG segmentation through SfM and MVS

The foreground/background (FG/BG) segmentation relies on estimating depth for every pixel in the hyperspectral image, which is accomplished by employing multi-view stereopsis (MVS). MVS requires the images to be stereo-rectified, which, in turn, necessitates the estimation of the camera poses. The latter is the primary objective of structure from motion (SfM) algorithms and in this work we use a SfM pipeline that is based on the open-source algorithms from COLMAP [6]. Since these algorithms are not designed to deal with hyperspectral imaging cubes, we first generate a master band and enhance the contrast using contrast limited adaptive histogram equalization (CLAHE).

The multi-view stereo matching domain is still highly active, with significant advances recently driven by the latest developments in deep learning. In order to achieve the highest possible accuracy in generating FG/BG masks, we leverage the state-of-the-art technique developed by Li et. al. [7]. Since the quality of the MVS-based depth estimation is strongly dependent on the accuracy of the camera pose estimation, we have plans to incorporate a LiDAR scanner

on our drone payload in the future. Figure 4 illustrates the workflow for the FG/BG segmentation.

3.2.2. Subsection segmentation using HDR

To differentiate between shadowed regions and areas lit by either diffuse or direct sunlight within various subsections of the tower, we implemented a technique that involves capturing multiple images using different exposure times. This approach simulates high dynamic range (HDR) imaging and avoids overexposure while maintaining a high signal-to-noise ratio for poorly lit regions. By using longer exposure times, we ensure that shadowed regions maintain a sufficiently high signal-to-noise ratio. Conversely, by using shorter exposure times, we prevent the regions directly exposed to sunlight from becoming excessively saturated. In Figure 6, two different images are depicted which were captured using an exposure time of respectively 1ms (a) and 2ms (b). It is clear that the image captured with an exposure time of 2ms shows more supersaturated areas, but in turn the metal bars that are shadowed have more signal left to reliably detect potential corrosion.

3.3. Classification

As illustrated in Figure 2, the classification process comprises three distinct subtasks. First, using the FG/BG masks generated in the segmentation part, the metal bars from the electrical pylon are isolated and only pixels sampled on its surface are selected. We then make the assumption that the paint layer or coating constitutes the majority of the pixels and we use those as a reference to calibrate the classification algorithm. Second, outliers are removed by computing the spectral angle between the reference spectrum and the image spectrum. Finally, we apply the fully constrained least squares unmixing (FCLSU) algorithm to classify corrosion.

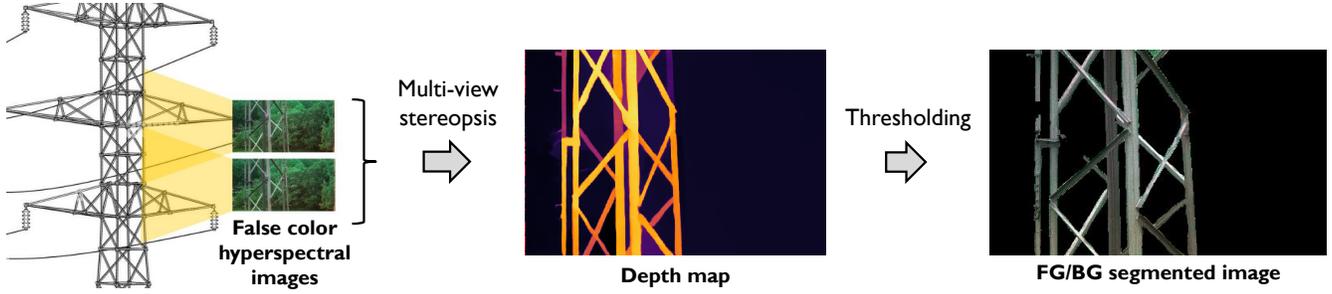


Fig. 4. Workflow for the foreground/background segmentation of the hyperspectral images. Using the known camera poses and applying multi-view stereopsis, depth can be estimated for every pixel of the hyperspectral image. This depth is then thresholded to generate a binary mask, which enables to isolate the tower from the background.

3.3.1. Outlier detection

Obviously, since other anomalies or outliers can be present on the metallic tower, such as dirt or bird poo, we still need to differentiate between those and corrosion. Therefore, using the hyperspectral NIR data, we rely on the spectral angle between the reference spectrum of the corrosion and the image spectrum. We empirically picked a threshold of 7 degrees to minimize the odds for false positives. The result of this detection algorithm is a mask that is used as a pre-classification filter. As a small test, we attached some soil and a dry leaf on a metal bar sample. Both outliers could correctly be distinguished by considering the spectral angle.

3.3.2. Fully constrained least-squares unmixing

After the outlier detection step, we assume that the only remaining pixels contain either paint or corrosion. We then apply the fully constrained least squares unmixing algorithm (FCLSU), that takes as input an end-member library of 3 levels of corrosion and paint spectra. This 4 end-member library was constructed from one sample containing paint and all three corrosion types by manually labelling pixels from all 4 end-members and averaging their spectra. The output of this algorithm is given by abundance maps of the three levels of corrosion and paint, as shown in Figure 5. These abundances obey the positivity and sum to one constraints. Using this type of soft classifier gives us more flexibility than classifiers that provide hard classification labels, as it allows to obtain an in-between classification for more granularity in the labelling. For example, an abundance map containing (0 paint, 0.1 light corrosion, 0.4 medium corrosion, 0.5 severe corrosion) can be directly interpreted as medium to severe corrosion. Another advantage of using FCLSU is that it is less sensitive to the noise commonly associated with infrared hyperspectral imaging.

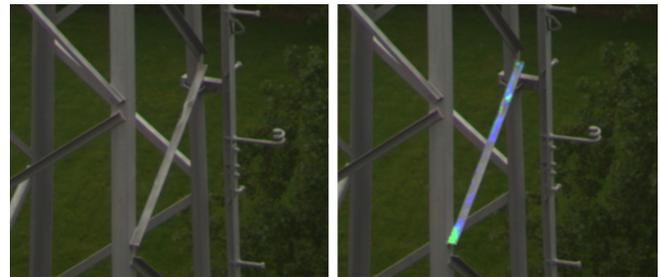


Fig. 5. Left: The RGB view of the suspected corroded segment. Right: Classification results overlaid over the RGB view of the corroded segment (green is medium corrosion and blue is light corrosion).

4. EVALUATION

The evaluation of the classification is conducted on outdoor measurements of dismantled metal bars that were placed on the ground. In total we have 13 bar samples for which ground truth is available. The latter was generated by manually measuring the amount of material loss after cleaning. To compare the classification to this ground truth, we defined light corrosion as corresponding to [0.1 mm - 1.5 mm] loss, medium corrosion to [1.5 mm - 2.5 mm] loss and severe corrosion to a loss of 2.5 mm onward. In total, 11 spots of light corrosion, 8 spots of medium corrosion and 12 spots of severe corrosion could be identified on the samples. If a pixel is assigned a label by the algorithm, but the measured loss is outside the range associated with that class, then it is deemed a false positive (FP). If the corrosion level matches the output of the algorithm, then it is labeled as a true positive (TP). In the case that some class is not detected, and that class is also not present in the sample, then it is labeled as a true negative (TN). If the class is not detected whilst being present in the sample, then it is considered a false negative (FN).

To assess the accuracy of the results, the standard statistical metrics were calculated and listed in Table 1. According to the metrics, the algorithm has a high sensitivity for all three types of corrosion but a relatively higher sensitivity for

	Light	Medium	Severe
N ^o True Positives	9	6	11
N ^o True Negatives	9	13	7
N ^o False Positives	0	0	2
N ^o False Negatives	2	2	1
Sensitivity	0.82	0.67	0.92
Specificity	1	1	0.78
Positive Predictive value	1	1	0.85
Negative Predictive value	0.83	0.8	0.88
Accuracy	0.9	0.86	0.86

Table 1. Statistical metrics for the corrosion classification.

light and severe corrosion and a relatively lower sensitivity for medium corrosion. Thus, all three levels of corrosion are likely to be detected but the medium corrosion is less likely to be detected compared to the other two classes. The algorithm's specificity was high for the light and medium classes, whereas it was lower for the severe class, indicating that the rate of false positives is lower for the first two classes than it is for the severe class. Similar to the specificity, the positive predictive value is generally high for all three classes, but it is slightly lower for the severe corrosion class, compared to the other 2 classes, indicating that the rate of true positives is lower for that class. For all three classes, the negative predictive values are similarly high, which indicates that the true negative rates are also high. The classification accuracy is generally high for all three classes (over 86%), with slightly higher accuracy for light corrosion. The entire workflow has also been tested on data captured during a real drone flight. The results are illustrated in Figure 5 where the abundance map is overlaid on the RGB view of the corroded segment. The effectiveness of using HDR imaging is demonstrated in Figure 6. In image (a) captured using 1 ms exposure, corrosion could be identified in the region that was directly lit by sun light (c), whereas no corrosion could be detected for the side in shadow. Conversely, image (b) captured using 2 ms exposure shows that corrosion could be identified on the shadowed side (d), whereas no corrosion is detected on the directly lit side due to overexposure.

5. CONCLUSION

In this paper, a novel approach was presented to inspect high-voltage transmission towers for corrosion. The use of irradiance-based methods and HDR techniques enabled successful classification in challenging lighting conditions on an actual pylon in the field. The accuracy of our classification pipeline has been validated using provided ground truth and reaches more than 86% for all severity classes. By leveraging our hyperspectral imaging pipeline and adopting our quantitative approach to corrosion severity classification, drone inspections become more effective, enabling timely detection of corrosion-related issues and ensuring the long-term structural integrity of transmission towers.

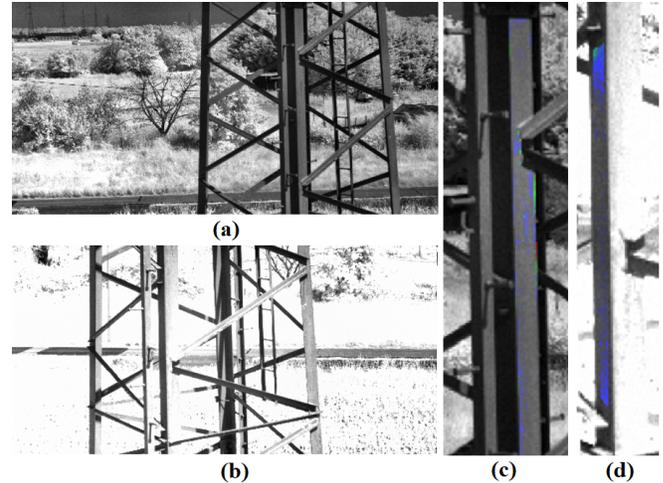


Fig. 6. Images (a) and (b) are captured with an exposure time of respectively 1 ms and 2 ms. In (c) and (d) the classification is overlaid on a zoomed in part of the image (green is medium corrosion and blue is light corrosion).

6. REFERENCES

- [1] Alberto Ortiz, Francisco Bonnin-Pascual, Emilio Garcia-Fidalgo, and Joan P. Company-Corcoles, "Vision-based corrosion detection assisted by a micro-aerial vehicle in a vessel inspection application," *Sensors*, vol. 16, no. 12, 2016.
- [2] Sanjay Kumar Ahuja and Manoj Kumar Shukla, "A survey of computer vision based corrosion detection approaches," in *Information and Communication Technology for Intelligent Systems (ICTIS 2017) - Volume 2*, 2018, pp. 55–63.
- [3] Dayakar Naik, Hizb Ullah Sajid, Ravi Kiran, and Xin Sun, "Hyperspectral imaging for the elimination of visual ambiguity in corrosion detection and identification of corrosion sources," *Structural Health Monitoring*, 07 2021.
- [4] Zohreh Zahiri, Alfredo Lamberti, Jan Wielant, and Paul Scheunders, "Characterization of corrosion products on carbon steel using hyperspectral imaging in short-wave infrared (swir)," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2022, pp. 1–4.
- [5] Zahiri, Zohreh and Luong, Hiep and Platasa, Ljiljana and Vandebriel, Roeland and Jayapala, Murali and Mangialetto, Frédéric and Scheunders, Paul, "Detecting and characterizing corrosion on high voltage metallic towers using hyperspectral imaging," in *2022 12th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS), Abstracts*, 2022, p. 1.
- [6] Johannes Lutz Schönberger and Jan-Michael Frahm, "Structure-from-motion revisited," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] Jiankun Li, Peisen Wang, Pengfei Xiong, Tao Cai, Ziwei Yan, Lei Yang, Jianguo Liu, Haoqiang Fan, and Shuaicheng Liu, "Practical stereo matching via cascaded recurrent network with adaptive correlation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16263–16272.