Low-cost System for Visual Inspection of Corrosion: An Industrial Case Study

Karina Hernández Oliver, Giovanna Martínez-Arellano, Joel Segal

Advanced Manufacturing Technology Research Group, Faculty of Engineering, University of Nottingham, Advanced Manufacturing Building, Jubilee Campus, Nottingham, NG7 2GX, UK

Abstract: The use of digital technologies around the world has increased considerably, modifying the way in which daily activities are conducted. The manufacturing sector is no exception. Over the last decade, digital technologies have become a key element for manufacturing companies to deliver high quality products, which is a critical factor in their success and competitiveness. However, for most companies, cost and lack of understanding of the benefits of digital solutions are still a main barrier to digitalisation of their operations. In this paper, a low-cost visual system is proposed and developed for in-process detection of defective parts in a manufacturing company that produces hoses with fittings and connectors. The company often faces problems with the angle, length, and condition of parts, particularly with brazing residues and corroded surfaces, causing rework and rejection of products by clients. The proposed low-cost system is based on a deep learning approach and uses an off-the-shelf camera to identify corrosion in parts before assembly. Using a low-cost solution contributes to optimizing costs and operations in the production line, reducing waste and rework due to the assembly of defective parts, and minimizing human errors due to differing expertise or diverse criteria. Also, the solution can potentially be implemented in other production lines, reducing the learning curve as regards implementation and extending the lessons learned to the solution of other problems.

1. Introduction

Clients and customers are essential stakeholders, prompting companies to continuously explore alternatives for more effectively fulfilling their needs and demands. This pursuit is driven by the recognition that clients and customers exclusively only will pay for products that are 'good' for them, that is, quality products. Therefore, ensuring product quality has emerged as a substantial concern for companies. The shipment of defective products creates quality-related issues, while abstaining from shipping defective items translates into financial losses for the manufacturer [1]. Hence, several quality methodologies, such as Six Sigma, have received significant attention over the past decades, aiming to mitigate process variability during the manufacturing of products.

In modern manufacturing and industrial processes, the early detection and accurate assessment of defects within in-process products or during final inspection are critical to ensuring product quality, safety, and reliability. To be more specific, within the metal-mechanic industry, corrosion is a type of defect that companies strive to identify promptly. Corrosion can compromise structural integrity, functionality, and overall performance. Traditional methods of corrosion identification, often relying on manual visual inspection, are limited by subjectivity, time consumption, and inconsistent results.

Recent advances in computer vision and deep learning techniques have revolutionized the field of image analysis and pattern recognition. Among these techniques, the You Only Look Once (YOLO) algorithm has garnered attention due to its real-time object detection capabilities. YOLO's ability to rapidly process images and identify objects within them, coupled with its accuracy, makes it a promising tool for developing robust and efficient visual inspection systems [2].

This paper presents an approach to corrosion detection through the integration of a low-cost visual inspection system with deep learning methodologies, specially by using the YOLO algorithm. The proposed system aims to address the limitations of traditional visual inspection methods, which are often time-consuming, subjective, and dependent on the expertise of inspectors. By harnessing the power of deep learning, the system enhances the accuracy and efficiency of corrosion identification, contributing to improved decision-making processes, a reduction of in-process rework, and more effective resource allocation.

2. Background

2.1. Testing Techniques and Corrosion

In the context of manufacturing processes, the thorough inspection of components assumes paramount importance to ensure structural integrity, performance, and conformance with engineering specifications and standards. A range of techniques are conventionally employed to meticulously evaluate the quality and integrity of these components, in particular those involving metallic parts. Nondestructive testing techniques are frequently used to assess the integrity of materials by examining surface or internal defects, as well as evaluating metallurgical conditions, without adversely impacting the material's structural composition or its viability for its designated function. Among the most common techniques include visual inspection, microscopy, radiography, dye penetrate, ultrasonic, magnetic particle, eddy current for metals, and acoustic emission [3].

In one hand, the Visual inspection technique is useful for identifying macroscopic flaws. It involves direct observation either with the naked eye or aided by magnification to identify surface irregularities such as cracks, corrosion, and fractures [3]. Nevertheless, this technique also requires that the person conducting the inspection has the necessary training and skills, in addition to the fact that the final decision of whether the inspected part is accepted becomes subjective and falls entirely on the experience of the inspector. All of this leads to the inspection process not being consistent and, although doing so requires repetitive tasks, the result will not necessarily be the same, therefore repeatability is not ensured [4]. In-process inspection becomes complicated and frequently presents a challenge when companies lack the resources to invest in tools and equipment that mitigate uncertainty in outcomes. Conducting exhaustive inspections to complete batches using conventional methods not only escalates production costs but also is very time consuming.

On the other hand, corrosion denotes a material's response upon interaction with its environment, leading to alterations encompassing changes such as consumption, dissolution, or deterioration of the material [5]. Corrosion monitoring involves a systematic observation and evaluation of material changes, caused by inherent degradation mechanisms like atmospheric corrosion, chemical dissolution, or oxidation, among others. Typically, this is accomplished by employing specialized corrosion assessment methods or instruments within structures, parts or products on a small and big scale. The purpose of this monitoring is to assess the integrity of the structure of assets (e.g. infrastructure, facility, machinery) or parts (e.g. raw materials, in-process products or final products) to ensure that the element inspected is not affected by corrosion damage. The primary aim of corrosion monitoring is to assess the structural soundness of infrastructure, ensuring that the longevity of the asset remains unaffected by corrosion-related impairments. Corrosion monitoring also may encompass quality assurance by minimizing contamination arising from corrosion, as well as averting safety lapses and potential incidents. Currently, initiatives related to corrosion monitoring have experienced a notable surge in adoption, spanning a wide array of applications. These applications encompass various domains, including but not limited to pipelines, refinery complexes, architectural structures, aircraft, maritime vessels, automobiles, electronic devices, computing systems, and even biomedical implants [6].

2.2. *Machine Learning Approaches and Low-cost Solutions*

Considering the importance for companies to ensure the quality of their products and compliance with standards, in recent years the realm of inspection and quality control has witnessed a transformative shift with the incorporation of new techniques using technology. In this context, two particularly noteworthy approaches, especially deep learning and image processing, have emerged as powerful tools in recent times. These approaches can be categorized as 'thinking/reasoning' technologies due to their inherent capabilities that strive to emulate, and in certain instances, exceed human cognitive faculties. These methodologies align with the broader pursuit of artificial intelligence, seeking to replicate and augment human-like cognitive processes in machines [7]. Both deep learning and image processing methodologies hold the potential to transcend the limitations of traditional inspection methods, offering an unprecedented level of accuracy, speed, and adaptability.

Deep learning, a subset of artificial intelligence introduced in the 2000s, exhibits a remarkable capacity to automatically extract intricate patterns and features from complex data, rendering it particularly adept at discerning subtle anomalies in diverse applications such as manufacturing, healthcare, and infrastructure assessment. Deep learning, situated within the purview of machine learning, is characterized by its ability to autonomously acquire intricate knowledge from large datasets, thereby enabling the identification of intricate patterns and representations. In this sense, it strives to replicate the cognitive processes associated with human pattern recognition and abstraction, albeit in a more streamlined and expedited manner. Similarly, image processing, underpinned by computational algorithms, seeks to emulate human visual perception by processing and extracting relevant information from visual data. As such, these thinking/reasoning technologies exhibit the potential to not only mimic but also potentially surpass certain aspects of human cognitive functioning, thereby reshaping conventional paradigms in inspection and quality control procedures [7-9].

Image processing techniques have gained prominence by leveraging computational algorithms to manipulate, analyse, and interpret visual data, enabling accurate defect detection and classification, and enhancing the accuracy, efficiency, and objectivity of inspection processes across various domains. Image processing encompasses a range of fundamental tasks, including seemingly straightforward operations such as image resizing. In the context of deep learning, uniformity of images is essential for the development of such models, making resizing a common image processing task.. This requires resizing all images to a consistent size, a preprocessing step that facilitates their compatibility with the network's architecture. Beyond resizing, an array of additional processing tasks can be undertaken to optimize image inputs for subsequent analysis. Geometric transformations, for instance, enable the augmentation of the dataset by applying rotations, translations, or reflections to images, these transformations are essential for achieving good performance when training deep learning models. Color transformations offer a possibility to standardise and manipulate the color distribution of images, while conversion to grayscale reduces computational complexity and eliminates color-based features. The amalgamation of these preprocessing techniques underscores the role of image processing in priming raw visual data for effective interpretation by deep learning techniques, exemplifying its significance in the broader landscape of machine vision and artificial intelligence applications [8].

In this context, object detection is a computer vision task that pursues a dual-faceted objective. Primarily, it involves the localization of one or more objects present within a given image, followed by the secondary task of ascertaining the classification of each individual object contained therein. This process is executed through the delineation of a bounding box encompassing the identified object, concomitant with the attribution of its anticipated class label. Thus, it diverges from the conventional purview of image classification, wherein the predictive scope is limited solely to the classification of the image's entirety. In contrast, object detection entails an extended prognostic ambit, encompassing not only the categorisation of the object but also the prediction of the spatial coordinates demarcating the bounding box that optimally encapsulates the detected object. This computational pursuit assumes a remarkable behaviour, as it necessitates the successful achievement of both accurate object localization, thereby enabling the delineation of precise bounding boxes around distinct objects present within an image, and thoughtful object classification, ensuring the accurate anticipation of the specific object class corresponding to the localised entity [8, 9].

As mentioned in the Introduction section, YOLO ("You Only Look Once") is a family of object detection architectures that have exhibited continuous refinement subsequent to its first launch in 2015, commencing with the inception of YOLOv1. The compendium of YOLO models constitutes a succession of end-to-end deep learning constructs meticulously tailored to expedite the process of object detection. Underpinned by an open-source framework, these models were conceptualized and realized by Joseph Redmon and Ali Farhadi, heralding a pioneering venture in the domain of swift real-time object localization. Distinguished by its expeditious computational performance, YOLO occupies a relevant niche within the spectrum of object detection algorithms. Central to the YOLO methodology is its distinctive operational paradigm, wherein predictions are exclusively formulated for a delimited set of bounding boxes. This hinges upon the partitioning of the input image into an array of discrete cells, with each cell imbued with the capability to directly infer both bounding box specifications and the classification label of the encompassed object. This predictive computations contributes to the rapidity and efficiency that distinguishes YOLO within the set of object detection methodologies [2, 10]. YOLOv8 is the most recent version, a cutting-edge, state-of-the-art (SOTA), which not only introduces novel attributes but also encompasses a spectrum of enhancements to augment overall performance, flexibility, and efficacy [10].

By leveraging advanced algorithms and computational methods, image processing empowers practitioners to transcend the limitations of human perception and provides a quantitative foundation for decision-making. Segmentation, in addition to finding the bounding boxes, it adds a mask, delimiting the object within the box. This introduction highlights the growing significance of image processing as a contemporary methodology in inspection practices, underscoring its potential to revolutionize conventional approaches and foster unprecedented levels of precision and reliability.

3. Research Methodology

This section outlines the methodology adopted for the development of the visual inspection system aimed at identifying brazing residue or corrosion in fitting parts, utilizing the YOLOv8 image segmentation model and a preloaded dataset.

Danfoss, a multinational corporation headquartered in Denmark, demonstrated willingness to engage in research collaboration providing hypothetical samples of brazing residue or corrosion manifested in their parts. With a global presence spanning over twenty countries, Danfoss operates within three principal business segments: power solutions, climate solutions, and power electronics and drives [11]. The participating site is situated in the United Kingdom and specializes in offering power solutions.

The initial step involved identifying the specific fitting parts subject to inspection. The characteristics of these parts, including various types and degrees of corrosion, were defined to establish a comprehensive framework for subsequent dataset collection and model training. To enable effective training of the visual inspection system, a suitable dataset was sought, obtained from samples provided by Danfoss and images readily accessible from the internet. This dataset was selected to align with the identified fitting parts and their corresponding corrosion characteristics. The dataset encompassed images showcasing varying levels of corrosion, ensuring relevance and applicability to the inspection task.

The experimental environment was meticulously configured to facilitate both image acquisition and programming. For image acquisition, a physical setup was established, encompassing the positioning of fitting parts, appropriate lighting, and camera placement. Additionally, hardware and software components were selected and configured, ensuring compatibility with the YOLOv8 model. An emphasis was placed on utilizing low-cost infrastructure and off-the-shelf components. The pre-trained YOLOv8 object detection model was employed as the foundation for the inspection system. Using transfer learning, the model was further trained using the acquired dataset containing images of fitting parts, including instances of general corrosion.

In one hand, the dataset employed for the pre-training of the model comprises 2580 images showcasing various corrosion issues evident on commonplace objects. This collection shows different real world objects that present some level of corrosion, such as automobiles, vessels, and pipelines, among other examples [12, 13]. On the other hand, the data set used for testing includes 280 images from discontinued fitting parts, provided by Danfoss to simulate actual parts with corrosion problems. The photographs were captured from varying perspectives, with the components positioned at different angles, and illuminated using distinct light intensities, achieved through the utilization of LED lighting, in-camera flash, and ambient lighting conditions.

By initializing the model with pre-existing weights, training efficiency was optimised, enhancing convergence. Upon successful training, the trained YOLOv8 model was used to analyse real photos of fitting parts. These photos represented actual instances of corrosion and allowed for the assessment of the inspection system's performance in a practical context. The model's ability to accurately detect and delineate corroded regions was evaluated against the real parts.

The performance of the visual inspection system was evaluated through quantitative and qualitative measures. Metrics such as train/segmentation loss and validation/segmentation loss were computed to assess the system's accuracy and reliability in identifying corrosion. Additionally, visualizations of the model's predictions were examined to gain qualitative insights into its behaviour. Feedback obtained from the evaluation stage was utilized to iteratively refine the visual inspection system. Misclassifications and areas of improvement were identified and addressed through adjustments to the dataset, model parameters, or post-processing techniques. This iterative process aimed to enhance the system's overall performance and effectiveness. Finally, the deployment potential of the developed visual inspection system was explored. Considerations were made regarding its integration into existing quality control workflows, compatibility with camera systems, and real-world implementation scenarios. Challenges associated with deployment and potential solutions were assessed.

4. Results and Discussion

The subsequent section presents the outcomes and corresponding discussion stemming from the development of a low-cost system tailored for the visual inspection of corrosion in fitting parts. This comprehensive exploration delves into the findings derived from the implementation, emphasising its efficacy in detecting and analysing instances of corrosion within the specified context. The results and subsequent discussion offer valuable insights into the performance of the inspection system, its strengths, limitations, and potential future work for refinement, thereby contributing to the broader discourse on low-cost solutions for corrosion assessment.

The equipment employed for the execution of the experimental procedures includes:

a) Laptop, processor Intel(R) Core(TM) i5-8265U CPU @ 1.60GHz 1.80 GHz, 8.0 GB RAM.
b) Comercial camera in mobile phone for tests and Raspberry Pi, HQ, Camera Module for the future pilot in Danfoss.
c) All photos were taken with 328 lux average.

Starting with the identification of parts and their characteristics, the products considered as part of this study were crimp fittings. Danfoss offers a broad line of fittings that are described as high-quality and high-performing products. They are designed for extremely high-pressure applications, and their superior resistance to corrosion contributes to enhanced performance of the equipment in use. An example of these fittings is presented in Fig. 1.

Fig. 1. *Crimp Fittings [14]*

Specific requirements and inspection criteria have been delineated for these parts within their quality management system documentation. The standards addressing corrosion and discoloration explain —via illustrative examples as presented in the Fig. 2 and 3— the conditions under which the parts are considered acceptable and satisfactory. Furthermore, it outlines instances when seeking the approval of the team leader becomes imperative, as well as scenarios where the parts manifest defects or damages, warranting their rejection.

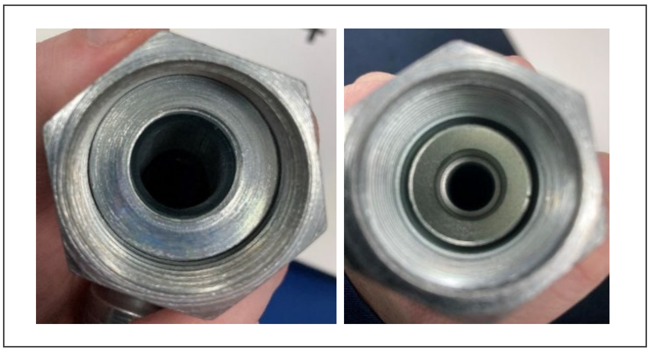

Fig. 3. *Parts acceptable for use [15]*

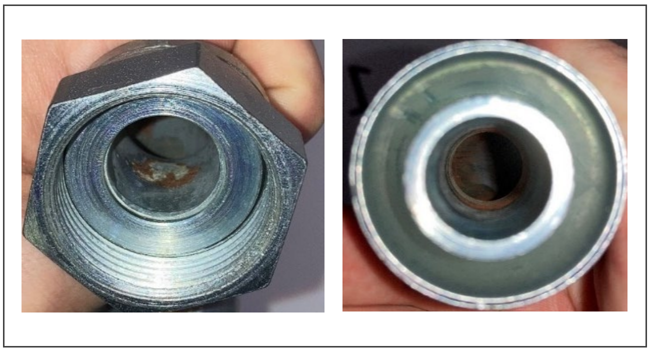

Fig. 4. *Parts with rust identified within end fitting [15]*

In order to obtain the dataset of images, an initial classification process was undertaken to categorise the images as either "CORROSION" or "NO CORROSION." The compilation of images was facilitated through the extraction of visual data from Google using a scraping technique [13]. Afterwards, the data set also had to be annotated with the bounding boxes. Subsequent to the programming of the code, the training model underwent testing using the collection of images sourced from Google.

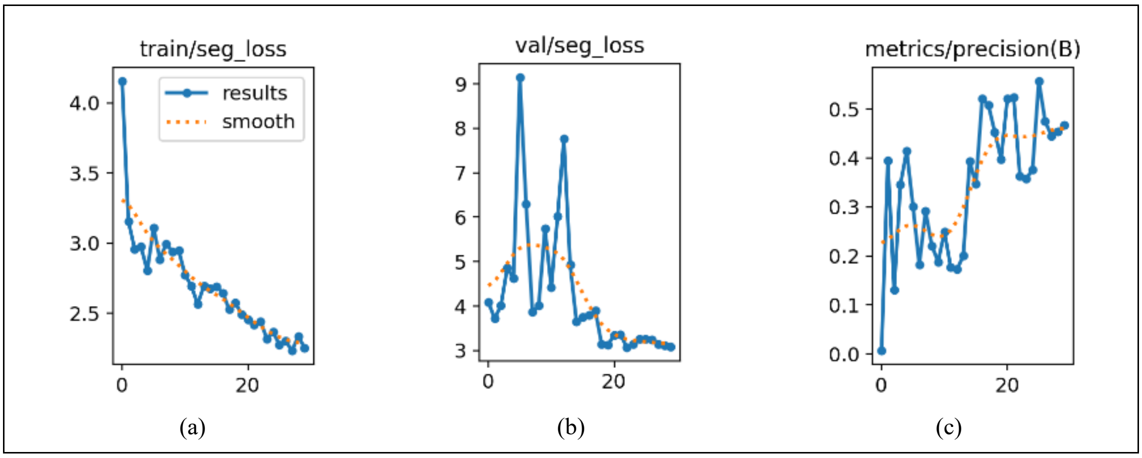

Fig. 2. *Performance of the deep learning model on the training set*

The performance of the deep learning model during the training process is illustrated in Figure 4. As evident from the visualization, both the training (Fig. 4(a)) and validation (Fig. 4 (b)) graphs depict a discernible downward trend in Loss. This trend represents a progressive enhancement in image segmentation performance with each successive iteration of the model. Opposite to the training and validation loss graphics, the precision graphic (Fig. 4 (c)) exhibits a positive trend. Precision, a critical metric in evaluating the model's performance, measures the ratio of true positive predictions to the total number of predicted positive instances. In this context, the ascending trend of the precision curve denotes an improvement in the model's ability to accurately identify and classify instances of interest among the positively predicted cases. The positive trend underscores the model's increasing capability to minimize false positives and enhance its precision as the training iterations progress. This observation signifies a commendable advancement in the model's capacity to precisely delineate areas of interest, contributing to the overall efficacy of the image segmentation task.

Upon completing the model training process, the subsequent step involves the execution of the prediction code, which serves the purpose of discerning distinct segments within the provided images. This discrimination is centered upon the identification of regions exhibiting the presence or absence of the specific attribute under consideration, namely "corrosion" or "no corrosion." Through the utilization of the trained model, this predictive procedure enables the accurate classification and localization of the aforementioned segments within the visual data.

Fig 5. comprises two unprocessed photographs, each depicting distinct sections of interest from two different parts. In these images, notable instances of corrosion manifest as discernible irregularities and discolored areas, indicative of material degradation. The presence of these corroded regions within the images is apparent, although visual identification alone might not suffice for accurate quantification and comprehensive analysis of the corrosion extent. Subsequently, to enhance the precision of corrosion localization and identification, the YOLOv8 model is deployed for predictive analysis.

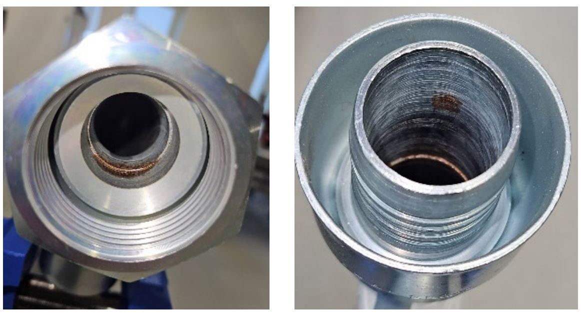

***Fig. 5.** Crimp Fittings before the image segmentation*

Upon conducting a series of seven tests, each characterized by distinct image configurations, results and corresponding metrics are presented in Table 1. These tests encompassed a diverse array of image setups for one fitting, ranging from differing lighting conditions to varying angles of object placement. For the task of detecting corrosion in components, precision is considering as the metric of interest. Considering that the main objective is to identify the maximum number of corroded parts relative to the overall count of the actual corroded items.

As discernible from the data, test number 6 stands out as the best iteration. The test yields a precision value of 0.96, signifying that the model's predictions regarding corrosion presence in parts are correct approximately 96 percent of the times. It is pertinent to highlight that this particular test exhibits the highest values not only in precision, but also across other relevant metrics. Given this noteworthy performance, the specific imaging conditions applied in this test were determinant for the success of the test. Consequently, an additional test was conducted, incorporating 60 images from different fitting parts.

# Test	1	2	3	4	5	6	7
True Positive	7	6	3	8	22	24	13
True Negative	1	0	0	3	2	0	0
False Positive	5	4	8	11	31	1	1
False Negative	17	10	4	7	23	5	6
Sample	**30**	**20**	**15**	**29**	**78**	30	**20**
Precision	0.58	0.60	0.27	0.42	0.42	**0.96**	0.93
Recall	0.29	0.38	0.43	0.53	0.49	**0.83**	0.68
Accuracy	0.27	0.30	0.20	0.38	0.31	**0.80**	0.65
F1-score	0.39	0.46	0.33	0.47	0.45	**0.89**	0.79

Table 1. Results from tests for corrosion prediction in fitting

After performing the last test, an observable reduction in precision becomes apparent, shifting from 0.96 to 0.69., results are presented in Table 2. Notably, despite this decrease, the precision value remains notably substantial. This occurrence can be attributed to the fact that, while the imaging conditions were replicated for this trial, the sections within which potential corrosion was discerned exhibited less distinct clarity or were situated at greater depths within the fittings, relative to the previous trials.

True Positive	True Negative	False Positive	False Negative	**Sample**	Precision	Recall	Accuracy	F1-score
29	4	13	14	**60**	0.69	0.67	0.55	0.68

Table 2. Results from test, best conditions, different parts

In the second set of figures (Fig. 6), which corresponds to the same original photographs as depicted in Fig. 5 but post YOLOv8 prediction, a significant advancement in corrosion detection and annotation becomes evident. Employing the YOLOv8 model, the predicted outcome effectively demarcates and highlights, via a red coloration overlay, the areas within the images that have been successfully identified as exhibiting corrosion. This computer-assisted prediction is notably accurate in localizing and categorizing corroded segments within the visual data. The contrasting red annotations, juxtaposed against the original images, unequivocally delineate the spatial distribution of corrosion occurrences, thus providing a comprehensive and insightful visualization of the extent and locations of corrosion.

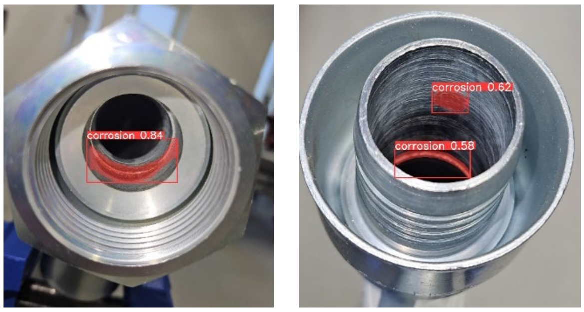

Fig. 6. *Crimp Fittings after the image segmentation*

The prediction code's execution thus contributes to the overarching objective of automated image analysis for the purpose of corrosion detection, thereby facilitating efficient and informed decision-making in domains reliant upon the assessment of material degradation and structural integrity.

5. Conclusions

The development and successful validation of a low-cost inspection system for identifying corrosion using deep learning and YOLOv8 is presented in this paper. A set of eight tests was conducted to systematically evaluate optimal imaging conditions for future applications. Among the initial seven tests, the iteration that garnered the highest precision, at 96%, involved the use of a flash or integrated camera lighting. However, upon replicating the environmental parameters to generate an alternate image dataset featuring diverse components, the precision score experienced a decline to 69%. This reduction can be attributed to the intricate accessibility of camera capture to sections housing potential corrosion, which were positioned in challenging-to-reach locations within the fittings. In light of these observations, a future experiment is envisioned, entailing the use of two cameras deployed at varying angles. This approach aims to maximize surface coverage for part examination, with the intent of addressing the limitations posed by the obscured sections encountered in previous trials.

The conclusive results obtained through experimentation at Danfoss underscore the system's efficacy and practicality. This study has yielded several noteworthy outcomes.

Firstly, the inspection system's construction using off-the-shelf components and open-source software has a paramount implication for industries. This approach mitigates the financial risk associated with investing in new technology, making it accessible for a wide range of companies, regardless of their financial capacity. For this reason, enterprises consistently endeavour to integrate components into their operations that involve minimal investments or implementation efforts. This drive for innovation in product manufacturing methods underscores their pursuit of novel approaches.

Secondly, the implementation of the proposed inspection system contributes to a substantial enhancement in accuracy. The system's reliance on deep learning and YOLOv8 eliminates the inherent subjectivity stemming from human interpretation, experience, or individual viewpoints. This objectivity ensures consistent and reliable inspection outcomes, reducing the likelihood of errors and discrepancies in corrosion detection.

Furthermore, the principal objective of this research, which was to overcome the limitations of traditional visual inspection methods, has been met. The conventional methods often suffer from time-intensive procedures, subjectivity influenced by inspectors' expertise, and potential variability in results. The developed system provides a solution that is not only efficient but also standardised, mitigating these limitations and advancing the state of corrosion inspection.

The next phase of the research involves on-site testing, where the system will be piloted in one of the production lines. This real-time inspection will provide insights into the system's adaptability and performance in a live industrial environment. Subsequent adjustments can be made to tailor the system to the specific requirements of the production line. The successful integration of the system into one production line will pave the way for its implementation across the remaining three lines, signifying a comprehensive adoption of the technology.

6. Acknowledgments

The authors would like to acknowledge the support of the National Council of Science and Technology (Conacyt) for providing funding to Karina Hernandez Oliver for this research work. The acknowledgement is extended to Danfoss Power Solutions 9125 for their support and cooperation to conduct the case study.

7. References

1. Taguchi, G., *Taguchi's quality engineering handbook [elektronische middelen] / Genichi Taguchi, Subir Chowdhury, Yuin Wu associate editors, Shin Taguchi and Hiroshi Yano*, ed. S. Chowdhury, et al. 2005, Livonia, Mich: Livonia, Mich : ASI Consulting Group.
2. Jiang, P., et al., *A Review of Yolo Algorithm Developments.* Procedia Computer Science, 2022. **199**: p. 1066-1073.
3. Dwivedi, S.K., M. Vishwakarma, and P.A. Soni, *Advances and Researches on Non Destructive Testing: A Review.* Materials Today: Proceedings, 2018. **5**(2, Part 1): p. 3690-3698.
4. Andriosopoulou, G., et al., *Defect Recognition in High-Pressure Die-Casting Parts Using Neural Networks and Transfer Learning.* Metals, 2023. **13**(6): p. 1104.
5. Tian, Z., et al. *Corrosion Identification of Fittings Based on Computer Vision*. in *2019 International Conference on Artificial Intelligence and Advanced Manufacturing (AIAM)*. 2019.
6. He, Y., *Corrosion Monitoring*, in *Reference Module in Materials Science and Materials Engineering*. 2016, Elsevier.

7. Slack, N., *Operations management / Nigel Slack, Alistair Brandon-Jones*. Ninth edition ed, ed. A. Brandon-Jones. 2019: Harlow : Pearson.
8. Elgendy, M., *Deep Learning for Vision Systems*. 2020, Manning Publications Co. : Manning. p. 480.
9. Diwan, T., G. Anirudh, and J.V. Tembhurne, *Object detection using YOLO: challenges, architectural successors, datasets and applications.* Multimedia Tools and Applications, 2023. **82**(6): p. 9243-9275.
10. Ultralytics. *Ultralytics YOLOv8*. 2023 [cited 2023 17 July 2023]; Available from: https://docs.ultralytics.com/.
11. A/S, D., *Annual Report 2022. Investing to build a better future*. 2023, Danfoss A/S: https://www.danfoss.com/en/. p. 150.
12. Yin, B., et al., *Corrosion Image Data Set for Automating Scientific Assessment of Materials*, in *British Machine Vision Conference (BMVC)*. 2021: Online.
13. Sun, P. *Deep Learning for Automated Corrosion Detection*. 2021.
14. A/S, D. *Engineering Tomorrow*. Crimp fittings 2023 [cited 2023 12 June 2023]; Available from: https://www.danfoss.com/en/.
15. A/S, D., *Corrosion and Discolouration Standards*. 2023, Danfoss A/S: Internal Quality Management System.