Revista ALCONPAT
www.revistaalconpat.org
eISSN 2007-6835
Revista de la Asociación Latinoamericana de Control de Calidad, Patología y Recuperación de la Construcción

YOLOv8-based model for automatic detection of residential roof damage.

A. S. Silva[1]*, A. R. Azevedo[2], F. H. A. M. Neto[2], P. H. Ferreira[2]
*Contact author: so_alisson@hotmail.com
DOI: https://doi.org/10.21041/ra.v15i1.783

Received: 17/09/2024 | Received in revised form: 19/11/2024 | Accepted: 10/12/2024 | Published: 01/01/2025

ABSTRACT

This study developed an automated image recognition model for inspecting residential roofs using the YOLOv8 architecture to identify three types of damage. The methodology involved images from 167 buildings captured by drones and annotated in CVAT, which were used to train and test the model. YOLOv8 was applied for anomaly detection and classification, achieving 79% precision. The limitations were the small dataset and the limited variety of capture angles. The originality of the work lies in the innovative use of YOLOv8 for roof inspection. Future research will focus on developing the YOLOv9 and YOLOv10 architectures and expanding the dataset and damage classes.

Keywords: maintenance management; machine learning; You Only Look Once (YOLO); roof inspection; rooftop assessment.

Cite as: Silva, A. S., Azevedo, A. R., Neto, F. H. A. M., Ferreira, P. H. (2025), *"YOLOv8-based model for automatic detection of residential roof damage."*, Revista ALCONPAT, 15 (1), pp. 50 – 63, DOI: https://doi.org/10.21041/ra.v15i1.783

[1] Departamento de Ingeniería Estructural y de la Construcción, Universidad Federal de Bahía, Salvador, Brasil.
[2] Departamento de Matemáticas y Estadística, Universidad Federal de Bahía, Salvador, Brasil.

Contribution of each author
A.S.S: conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, project administration, supervision, validation, visualization, writing-original draft, writing-review & editing. A.R.A: data curation, formal analysis, investigation, methodology, validation, visualization. F.H.A.M.N: data curation, formal analysis, investigation, methodology, validation, programs, visualization, writing-original draft, writing-review & editing. P.H.F: data curation, formal analysis, investigation, methodology, programs, validation, visualization, writing-original draft, writing-review & editing.

Creative Commons License
Copyright 2025 by the authors. This work is an Open-Access article published under the terms and conditions of an International Creative Commons Attribution 4.0 International License (CC BY 4.0).

Discussions and subsequent corrections to the publication
Any dispute, including the replies of the authors, will be published in the third issue of 2025 provided that the information is received before the closing of the second issue of 2025.

© 2025 ALCONPAT Internacional

Modelo basado en YOLOv8 para la detección automática de daños en tejados residenciales.

RESUMEN

Este estudio desarrolló un modelo automatizado de reconocimiento de imágenes para la inspección de techos residenciales utilizando la arquitectura YOLOv8, con el objetivo de identificar tres tipos de daños. La metodología incluyó imágenes de 167 edificios, capturadas por drones y anotadas en CVAT, que se utilizaron para entrenar y probar el modelo. YOLOv8 se aplicó para la detección y clasificación de anomalías, logrando un 79% de precisión. Las principales limitaciones fueron el pequeño tamaño del conjunto de datos y la limitada variedad de ángulos de captura. La originalidad del trabajo radica en el uso innovador de YOLOv8 para la inspección de techos. Investigaciones futuras se centrarán en el desarrollo de las arquitecturas YOLOv9 y YOLOv10, además de ampliar el conjunto de datos y las clases de daños.

Palabras clave: gestión de mantenimiento; aprendizaje automático; You Only Look Once (YOLO); inspección de techos; evaluación de techos.

Modelo baseado em YOLOv8 para detecção automática de danos em telhados residenciais.

RESUMO

Este estudo desenvolveu um modelo automatizado de reconhecimento de imagens para inspeção de telhados residenciais, utilizando a arquitetura YOLOv8, com o objetivo de identificar três tipos de danos. A metodologia contou com imagens de 167 edifícios, capturadas por drones e anotadas no CVAT, utilizadas para treinar e testar o modelo. A YOLOv8 foi aplicada na detecção e classificação de anomalias, alcançando 79% de precisão. As principais limitações foram o tamanho reduzido do conjunto de dados e a pouca variedade de ângulos de captura. A originalidade do trabalho está no uso inovador da YOLOv8 para a inspeção de telhados. Pesquisas futuras incluirão o desenvolvimento das arquiteturas YOLOv9 e YOLOv10, bem como a ampliação do conjunto de dados e das classes de problemas.

Keywords: gestão da manutenção; aprendizado de máquina; You Only Look Once (YOLO); inspeção de telhado; avaliação de telhado.

Legal Information

Revista ALCONPAT is a quarterly publication by the Asociación Latinoamericana de Control de Calidad, Patología y Recuperación de la Construcción, Internacional, A.C., Km. 6 antigua carretera a Progreso, Mérida, Yucatán, 97310, Tel.5219997385893, alconpat.int@gmail.com, Website: www.alconpat.org

Reservation of rights for exclusive use No.04-2013-011717330300-203, and ISSN 2007-6835, both granted by the Instituto Nacional de Derecho de Autor. Responsible editor: Pedro Castro Borges, Ph.D. Responsible for the last update of this issue, ALCONPAT Informatics Unit, Elizabeth Sabido Maldonado.

The views of the authors do not necessarily reflect the position of the editor.

The total or partial reproduction of the contents and images of the publication is carried out in accordance with the COPE code and the CC BY 4.0 license of the Revista ALCONPAT.

1. INTRODUCTION

The building envelope plays a critical role in the overall performance of a structure, acting as a protective layer that shields the building from external environmental degradation (Silva et al., 2015). Roofs and coverings primarily protect buildings, creating indoor environments shielded from external agents such as temperature, humidity, rain, wind, and noise (Garcez et al., 2012; Conceição et al., 2017). The maintenance of this system is a crucial component in building maintenance programs, as the roof system is the most vulnerable part of buildings due to its direct exposure to the environment (Rosenberger, 2012). When exposed to severe weather conditions, these systems often experience accelerated deterioration, leading to higher energy consumption, compromising the structural integrity of the roof and other building components, and reducing the building's lifespan (Alashari et al., 2022). Therefore, regular assessments are essential, as they can lead to substantial cost reductions (Liu et al., 2020; Alzarrad et al., 2022).

Manual inspection is the primary method for periodically assessing the structural and functional requirements to ensure compliance with basic performance standards (Dang et al., 2021). Traditional inspection methods rely on human labor to identify issues and defects. However, these tasks are often tedious and exhausting, resulting in errors due to human factors like inattention, optical illusions, subjective assessments, and inspector fatigue (Radopoulou and Brilakis, 2017; Hoang et al., 2018). Moreover, human inspectors require training and experience, which is time-consuming. Employee turnover is another significant drawback of manual labor (Olimov et al., 2022). Beyond these challenges, two additional factors complicate periodic assessments: the risk of falling from the roof and the extra costs arising from a lack of precision and human error (Brown et al., 2021).

Based on these findings, we can classify the limitations of roof inspections into two main areas: data acquisition and data analysis. To address the first limitation related to data acquisition, drones, also known as Unmanned Aerial Vehicles (UAVs), stand out. Drones are highly flexible, with versatile movements, allowing for rapid exploration of built environments and requiring minimal labor costs, as they only need battery changes and pre-planned routes that eliminate the need for manual operations during flight (Chu et al., 2024). Many studies have utilized this technology for roof inspections (Gajjar & Burgett, 2020; Silveira et al., 2020; Staffa et al., 2023; Santos et al., 2023).

Regarding data analysis, automated Artificial Intelligence (AI) algorithms have shown promising results in overcoming the limitations of subjective assessments by individual inspectors (Alizadeh & Ma, 2021; Avola et al., 2022). Notable works in this area include studies by Alzarrad et al. (2022), Hou et al. (2023), and Wang et al. (2023). Combining these two strategies can significantly enhance inspection results, enabling regular inspections and addressing previously mentioned limitations, as evidenced by research from Mostafa et al. (2023), Santos et al. (2023), and Staffa et al. (2023).

However, despite the relevance of these studies, some limitations persist. Mostafa et al. (2023) developed a two-stage Convolutional Neural Network (CNN) model for defect detection and classification in roofs using 2D images from over 5,500 images of 21 roofs. The model achieved 95% accuracy in detection and 97% in classification. However, the absence of defects such as cracks and blisters in the training images limits its effectiveness, highlighting the need for further research and improvements.

In the study by Santos et al. (2023), the authors optimized the maintenance of flat roofs using 330 images captured by drones and tested six deep learning algorithms for object detection, including Region-based CNN (Faster R-CNN), RetinaNet, and Adaptive Training Sample Selection (ATSS). Faster R-CNN achieved the best results with 93% recall and F1-score, while ATSS reached 95% accuracy. Error analysis highlighted the complexity of flat roofs, suggesting the need to improve

algorithms to handle their variations. Although the study effectively detected condenser units, there is a need to expand the focus to other issues, such as extra tiles and oxidation in gutters, which also require inspection.

To address these research gaps, this work proposes an image recognition model for residential roofs. A comprehensive dataset of images from 167 residential buildings collected by UAVs was used to detect and classify three roof damage types: extra tiles across the roof, oxidation in gutters, and open hatch covers. Annotations were made using CVAT, and the study employed the eighth generation of the You Only Look Once (YOLO) architecture, YOLOv8.

2. APPLICATIONS OF YOLO IN CONSTRUCTION

In computer vision, the challenging task of object detection has undergone significant evolution, driven by the revolutionary introduction of the YOLO architecture in 2016 (Redmon et al., 2016). YOLO's innovative approach differed from traditional two-stage object detection architectures by proposing a unified architecture capable of simultaneously predicting bounding boxes and class probabilities, meeting real-time processing demands (Koubaa et al., 2021). YOLOv8 is one of the latest versions in the YOLO series, utilizing a network architecture that incorporates a Feature Pyramid Network (FPN) and layer connections to integrate feature information across multiple scales efficiently. This version combines attention mechanisms and optimization strategies to enhance accuracy and performance in object detection. The core structure includes a backbone network for feature extraction from images, typically using deep CNN such as Darknet or ResNet, and a detection head composed of convolutional and fully connected layers responsible for predicting the bounding box and class probability of objects (Cao et al., 2024).

In this context, significant studies in construction have explored the advanced capabilities of YOLOv8. Sohaib et al. (2024) developed a model to detect and segment cracks in concrete, while Han et al. (2024) automated the detection and classification of small tools on construction sites. Zhao et al. (2024) created a system for automated monitoring of defects in 3D-printed concrete. Jiang et al. (2024) proposed a method to optimize concrete vibration timing with robots, using YOLOv8 to identify the best vibration positions and remove visual interference from reinforcement bars. Ribeiro et al. (2024) proposed a YOLOv8-based method for detecting building fissures. The images used for training were collected from the internet, and the model achieved 80% accuracy in validation with images not included in the training. However, the model focused solely on fissures, limiting its applicability to other use cases.

In a more specific context, for evaluating roofs and building coverings, Hou et al. (2023) developed an enhanced detection algorithm based on YOLOv3 to identify decorative elements known as "ridge beasts" on roofs of ancient Chinese buildings. This technology is crucial for the 3D reconstruction of historic buildings, historical dating, and cultural and tourist services. They improved the YOLOv3 network by incorporating local features and deep separable convolutions to expand the receptive field and capture more diverse characteristics. Additionally, they optimized the residual structure of the feature extraction network and redesigned the model's prediction head by introducing extrusion and excitation modules and using convolutions of various sizes. Experiments showed that the enhanced algorithm achieved a mean Average Precision (mAP) of 86.48%, 3.05% higher than the original YOLOv3, while reducing model parameters by 70% and improving detection performance. Although this research focused on detecting decorative elements, the model could also be applied to identify common issues in old buildings, such as broken or missing tiles, as well as moss and oxidation on roofing elements, which are essential for the proper maintenance of these historic structures.

In addition to this study, Alzarrad et al. (2022) developed a method using deep learning for the automated evaluation of residential roofs, focusing on identifying damaged areas, such as missing tiles. They employed transfer learning methodology using YOLOv5 to process images and identify these damages. The study found that despite the limited data, the model achieved approximately 81% precision and 86% accuracy. However, while the study showed good results in detecting issues with residential roof coverings, it focused only on one type of problem. In practice, many other issues with residential roofs need to be identified during inspections for proper resolution, such as oxidation on gutters and flashings. Neglecting these issues can lead to structural deterioration, significantly increasing repair costs.

Qiu et al. (2024) proposed a two-stage method based on YOLOv8. In the first stage, the YOLOv8 model is used to identify historically significant buildings in the area, achieving an mAP rate of 97.2%. The same model is employed in the second stage to segment the selected images, detecting potential roof defects, such as collapses, missing tiles, inappropriate architectural additions, and vegetation encroachment. Although this research represents a significant advancement and utilizes one of the latest versions of YOLO, it still focuses on the roofs of culturally historic buildings, such as Buddhist temples, which generally have specific roofing systems different from those used in modern residential buildings.

Given this, it is clear that immediate detection and classification of defects and problems in residential building roofs are crucial. While the studies mentioned have significantly advanced the automation of detecting certain roof issues, further research is still needed to address other types of problems. Considering the advancements in YOLO architectures, there is a clear opportunity to apply these networks in this context. In this study, we evaluate the use of YOLOv8 to detect and classify three issues in residential building roofs. As we found no prior studies in the literature addressing this topic, we highlight the application of one of the latest YOLO versions to automate the inspection process. Furthermore, we provide the codes used in this research to allow other researchers to enhance the model and add more problem classes, as different types of defects also need to be identified and addressed with equal importance.

3. RESEARCH METHOD

This study adopted an experimental simulation research strategy, developed in four stages (Figure 1): (1) refining the database of residential roof images with surface damage; (2) creating the automated damage recognition model using the YOLOv8 architecture; (3) training the model using 70% of the database; and (4) testing and validating the trained model with new images (30%). Precision, Recall, and mAP indicators evaluated the model's performance.

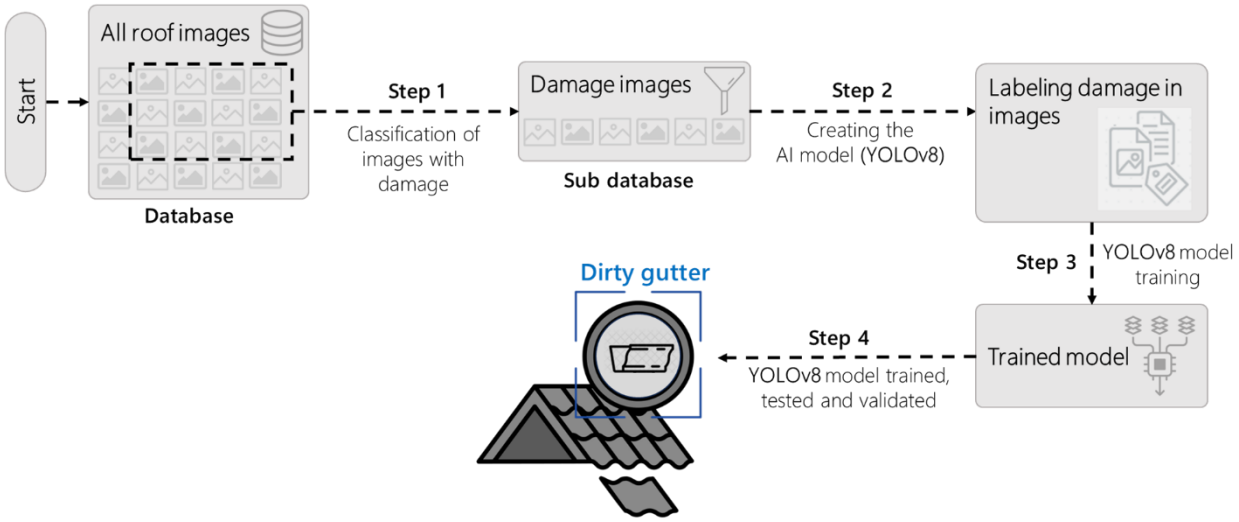

Figure 1. Research design.

In this work, we utilized YOLOv8. According to Lalinia and Sahafi (2024), one of the main advantages of incorporating YOLOv8 into computer vision projects is its greater accuracy compared to previous YOLO models. YOLOv8 supports multiple tasks, such as object detection, segmentation, and image classification, enhancing its versatility for various applications. The YOLOv8 architecture is shown in Figure 2; further details about the architecture can be found in Terven et al. (2023).

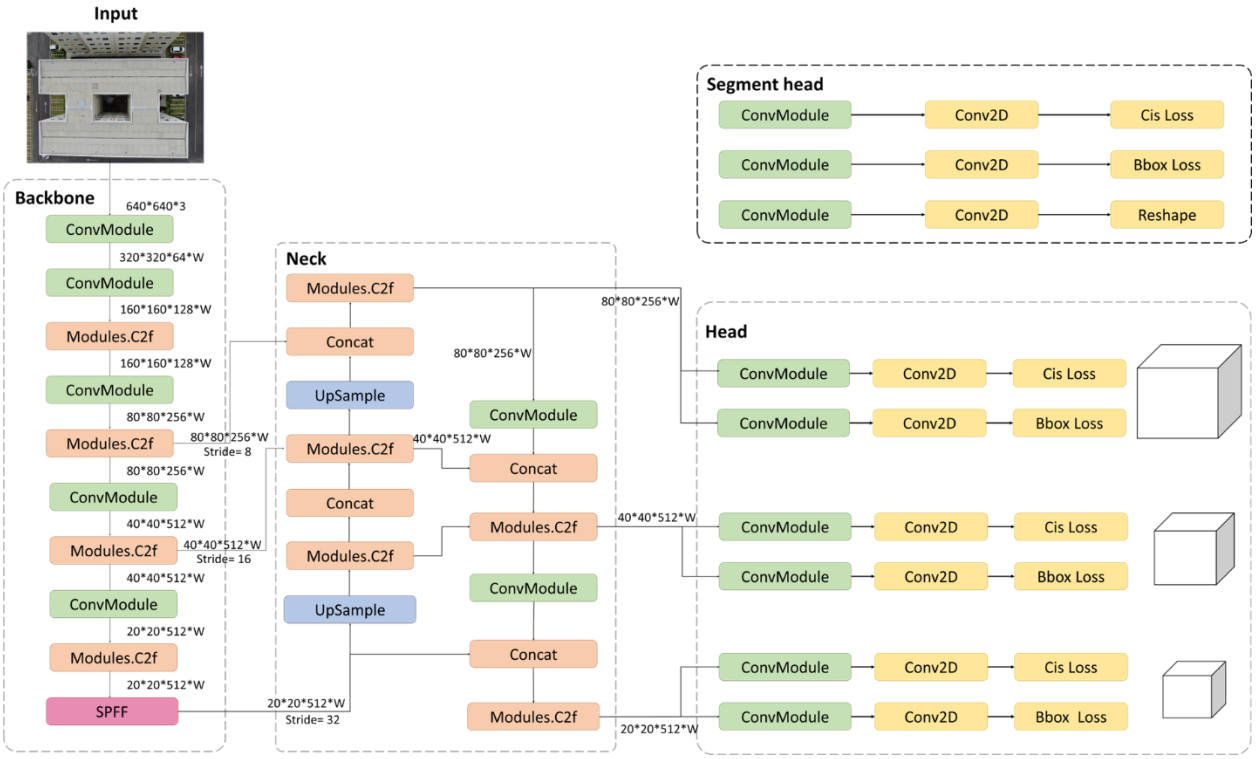

Figure 2. YOLOv8 architecture used.

Figure 2 shows that the architecture is organized into three main components: Backbone, Neck, and Head, each playing a specific role in image processing and prediction generation.

The input image is the starting point of the process, where the network begins to extract relevant features to identify potential defects. The Backbone extracts basic features from the image through several convolutional layers that detect initial patterns, such as edges and textures. As the image progresses through these layers, more complex modules, known as Modules.C2f, refine the extracted features, allowing for a more detailed image analysis. The Backbone also includes the Spatial Pyramid Pooling Fast (SPPF) module, aggregating information at different scales and enhancing detection robustness.

The Neck is an intermediary between the Backbone and the Head, refining and combining the extracted features through Concatenation (Concat), Upsampling, and additional convolutional layers. The Neck's goal is to prepare these features for more precise multi-scale detection, ensuring that both fine details and larger structures in the image are considered in the final prediction.

Finally, the Head is responsible for generating the final predictions. The features refined by the Neck are used to identify and classify objects or defects in the image. The Head consists of two main types: the Segment Head, which handles segmentation tasks by identifying exact areas where defects are present, and the Detection Head, which focuses on object detection and classification by generating bounding boxes and object classes. Additionally, losses (Cis Loss and Bbox Loss) associated with these predictions are calculated to optimize the model's performance.

This architecture allows YOLOv8 to process images efficiently, from initial feature extraction to final prediction generation, ensuring accurate detection of defects in residential roofs.

3.1 Database

GETEC-UFBA ("Grupo de Pesquisa e Extensão em Gestão e Tecnologia das Construções da Escola Politécnica da Universidade Federal da Bahia", in Portuguese) provided the dataset used in this research. These images were collected from residential buildings aged between one and five years. Over 167 residential roofs were inspected from manual drone flights, with an approximate altitude of 5 meters above the building being analyzed. The data acquisition followed inspection protocols proposed by Silveira et al. (2021). After analyzing the images in this dataset, 394 images were selected, distributed among 425 annotations for the class "Presence of extra tile in the transverse direction of the roof", 473 annotations for the class "Gutter integrity (oxidation)", and 29 annotations for the class "Open trapdoor". The dataset was split into approximately 70% for training, 20% for validation, and 10% for testing. After data augmentation in the training set, 498 images were obtained from the training set, with examples of these classes shown in Figure 3.

Figure 3. Examples of roof damage addressed in this research: a) Presence of extra tile in the transverse direction of the roof; b) Gutter integrity (oxidation); c) Open trapdoor.

3.2 Creating the YOLOv8 model

After defining the anomalies, the training images were labeled using Roboflow. An example of this labeling is shown in Figure 4a, where some extra tiles on the roof, oxidation, and an open skylight are labeled. Figure 4b illustrates the complete workflow from data entry to model training and

subsequent inference process, where annotations are made; then preprocessing, data augmentation, and finally, the data is exported in YOLOv8 format for analysis in Google Colab.

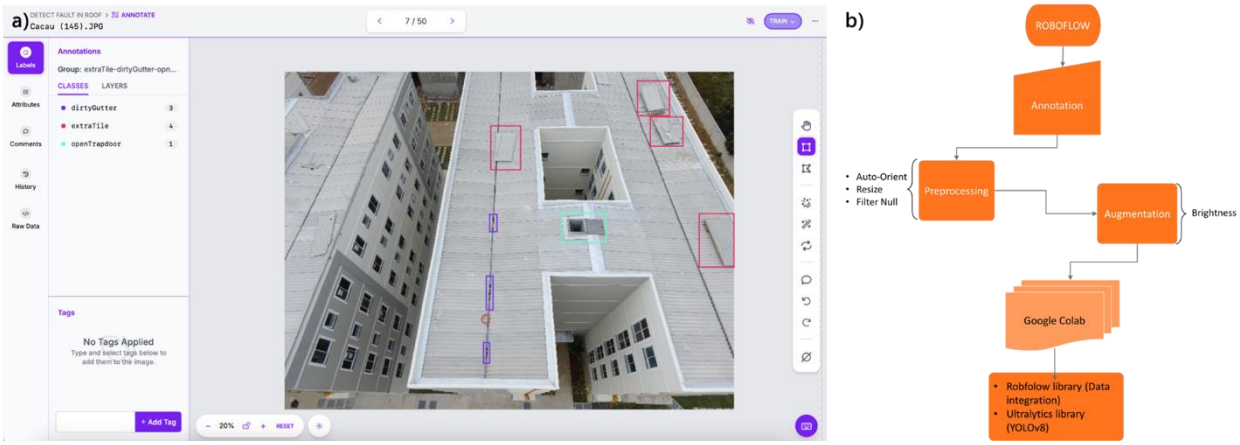

Figure 4. a) Damage notes; b) Workflow.

3.3 Model training, testing and validation

Precision, Recall, and mAP are the evaluation methods used to assess the models created on the Custom Vision and Roboflow platforms. Precision represents the number of correct cases out of all available cases. High precision indicates that anomalies are being identified accurately. Recall, consequently, is the number of correct instances out of all positive cases. A model with high recall effectively identifies the greatest number of correct anomalies. Similarly, mAP is the mean of the average precision, which is the precision obtained at various levels of recall across all classes (anomalies). The following formulas represent the Key Performance Indicators (KPIs):

(1) $$\boldsymbol{Precision} = \frac{TP}{TP + FP}$$

(2) $$\boldsymbol{Recall} = \frac{TP}{TP + FN}$$

(3) $$\boldsymbol{mAP} = \frac{\mathbf{1}}{\boldsymbol{n}} \sum_{\boldsymbol{k=1}}^{\boldsymbol{n}} \boldsymbol{APk}$$

Where TP (True Positives) are hits in positive instances, FN (False Negatives) are errors in positive instances, TN (True Negatives) are hits in negative instances, and FP (False Positives) are errors in negative instances and APk is the Average Precision for each class k. These metrics provide insights into the model's classification and quality.

4. RESULTS AND DISCUSSION

During the model training, the YOLOv8 architecture performs feature activations at different neural network stages. These activations demonstrate how the network processes visual information to detect patterns and important features in the input images (Figure 5). These visualizations illustrate how YOLOv8 performs object detection, from basic feature extraction to combining information across different scales and resolutions, enabling precise identification of roof defects.

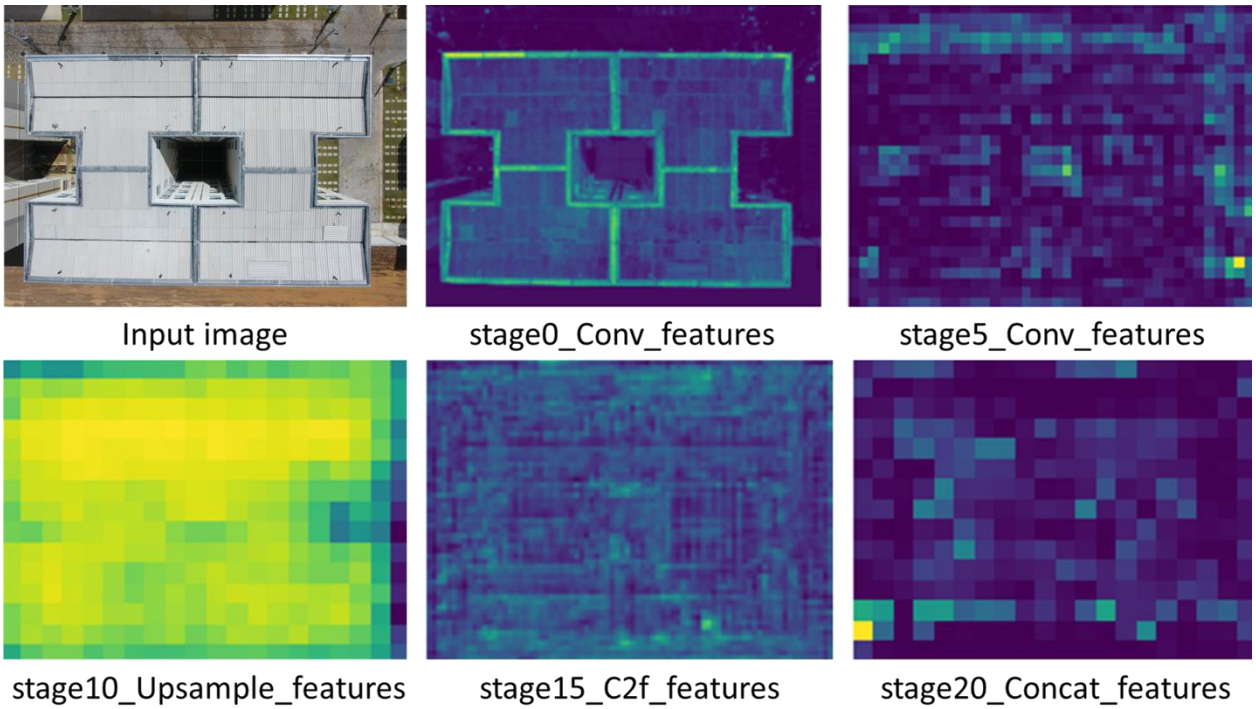

Figure 5: Structure of the trained YOLOv8 layers.

The "Conv_features" are the outputs of the convolutional layers, responsible for identifying local patterns such as edges and textures. As we progress through the different stages, from "stage0_Conv_features" to "stage20_Concat_features," the activations reflect increasingly complex and abstract representations of the image features. The "Upsample_features" are related to the upsampling process, where the spatial resolution of the activations is increased, allowing the network to detect objects at different scales. This process is crucial for detecting smaller details on roofs.

The "C2f_features" refer to combinations of convolutions followed by activation functions designed to capture more complex interactions between features extracted in previous layers, enhancing the network's ability to identify patterns. Finally, the "Concat_features" result from concatenation operations, which combine activations from different layers or stages of the network. This technique allows the network to integrate information from other levels of abstraction, merging features extracted at various stages to produce the final prediction.

The YOLOv8 model was trained using Google Colab. The model training and Streamlit application development code is available at: (OMITTED FOR REVIEW). The trained model has 168 layers and approximately 3 million parameters. The input images for the model were 640 × 640 pixels. The performance metrics of YOLOv8 on the test set are presented in Table 3, where mAP50 represents the mAP calculated with an Intersection over Union (IoU) threshold of 0.50.

Table 2. Evaluation metrics of the YOLOv8 model for the test set.

Class	Images	Precision (%)	Recall (%)	mAP50 (%)
All	73	0.79	0.59	0.66
dirtyGutter	73	0.69	0.44	0.51
extraTile	73	0.78	0.66	0.68
openTrapdoor	73	0.90	0.67	0.78

For extra tiles across the roof, the model achieved a Precision of 78%, Recall of 66%, and mAP50 of 68%. The model effectively identifies extra tiles, although it is slightly less accurate than detecting open skylights. Finally, for the integrity of the gutters, the model performed the worst, with a Precision of 69%, Recall of 44%, and mAP50 of 51%, indicating greater difficulty in detecting and accurately classifying this issue. Figure 6 shows the training and validation loss graphs, along with the Precision and Recall metrics; these indicate that the model is learning robustly and steadily (with decreasing losses and increasing metrics).

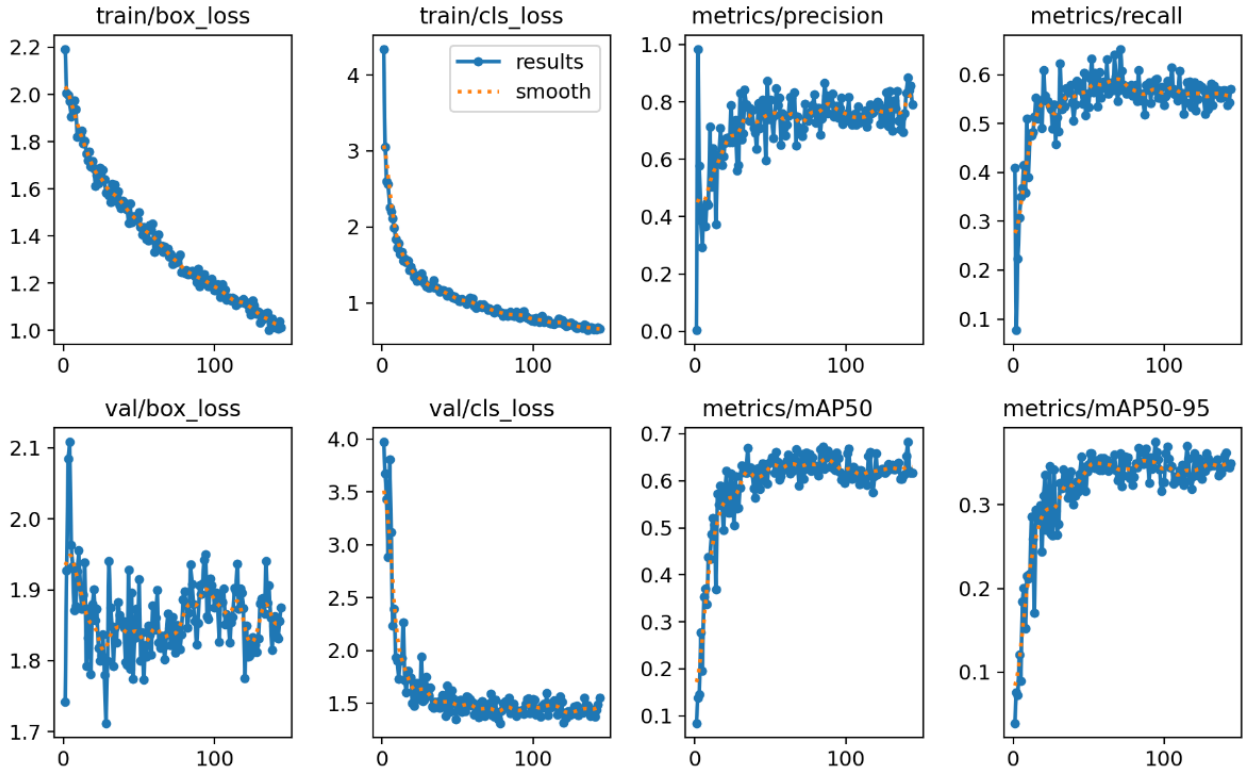

Figure 6. Training and validation losses, Precision and Recall metrics of YOLOv8.

Figures 7 and 8 display the classification results from the test set. From these images, it is clear that the model can identify extra tiles arranged in different orientations (Figure 7a shows extra tiles; Figure 7b shows an extra tile and oxidation on the gutter) and open skylights in various positions (Figures 8a and 8b show an open skylight and some extra tiles). These classifications are vital for preventive maintenance, ensuring the roofs' structural integrity and proper functionality. Despite the model's successes, some elements were not identified in Figures 7 and 8.

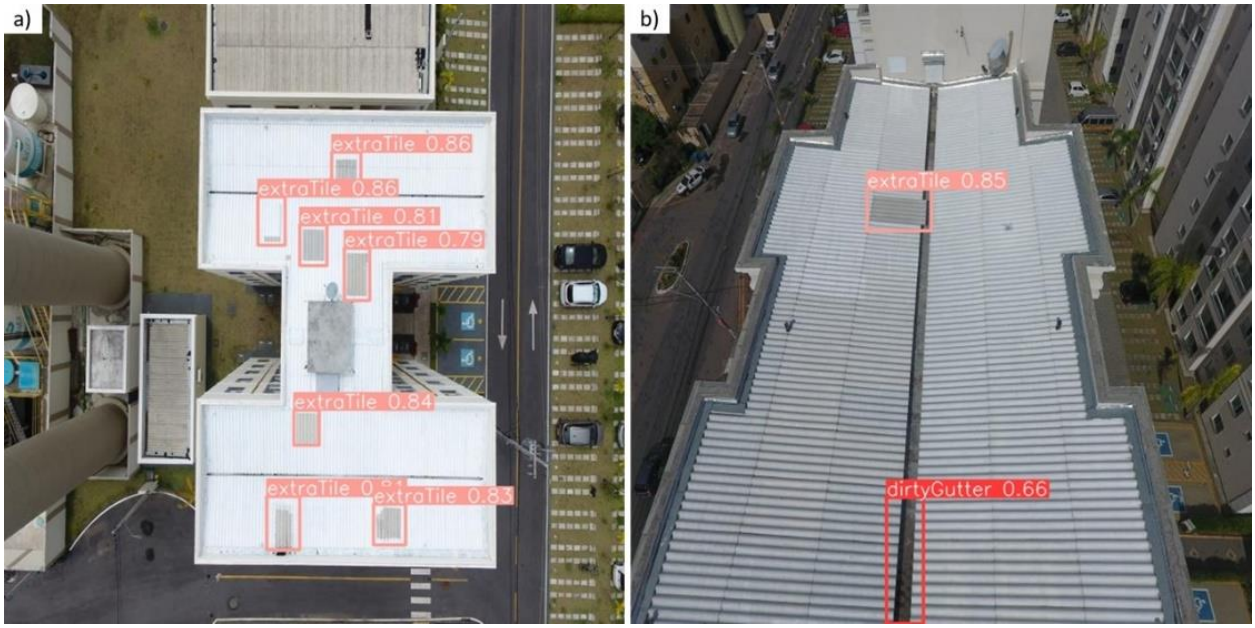

Figure 7: Test set classification for extra tiles and oxidation.

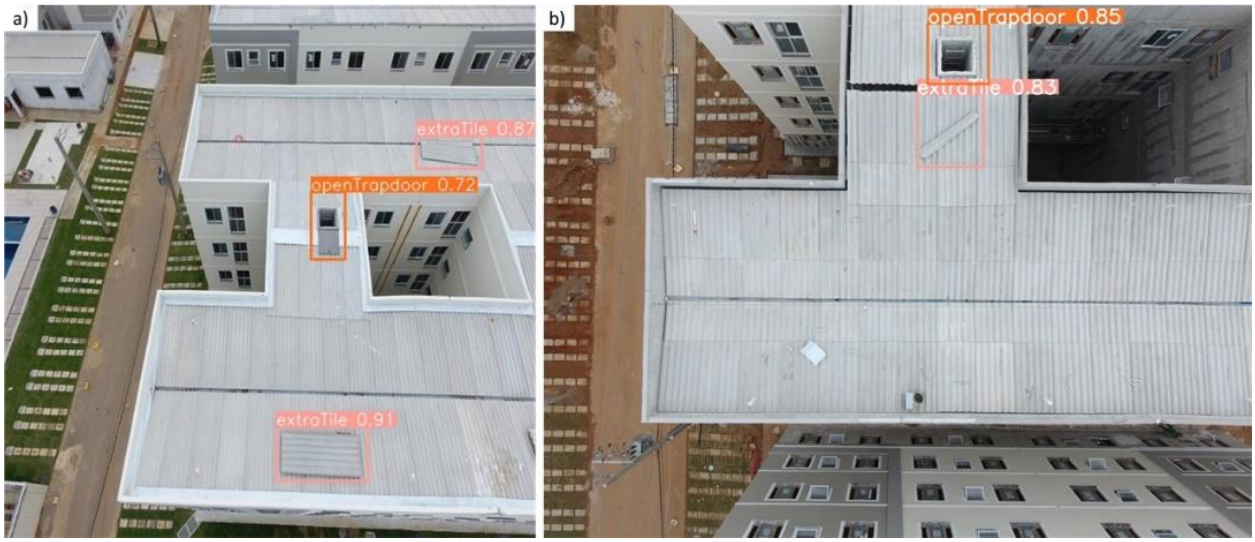

Figure 8: Test set classification for open trapdoor and extra tile.

The literature review on the detection and classification of defects in pitched roofs of residential buildings highlighted the need for studies focusing on the use of YOLOv8 algorithms, given the promising results of earlier versions for similar purposes (Alzarrad et al., 2022; Hou et al., 2023; Wang et al., 2023; Fan, 2024). This study addressed this gap by developing an automated recognition model for three defects in residential roofs using images collected by drones from 167 roofs. The model enabled the identification of extra tiles, oxidation in gutters, and open skylights, which are crucial for maintaining roof integrity. The results achieved 90% accuracy in detecting these defects.

This study's results were positive compared to similar studies (Alzarrad et al., 2022; Hou et al., 2023; Qiu et al., 2024). Additionally, while this research focuses on defects in residential roofs, previous studies concentrated on detecting decorative elements on historical building roofs. Nevertheless, all studies underscore the importance of early problem identification and the need for routine inspections to ensure the physical integrity of buildings. Furthermore, other studies, such as those by Santos et al. (2023), use YOLO for detecting air conditioning condensers, noting

that these elements are more distinct on roofs, making them easier to visualize than detecting oxidized gutters. The results from this study showed that the model recognized this damage with up to 78% precision during testing.
Moreover, comparing this study with other literature that did not use YOLO architectures reveals some common points. Firstly, the results of this study align with the findings of Staffa et al. (2023) regarding the need for inspections on residential roofs, considering the maintainability of these structures. Another important point is that although Staffa et al. (2023) addressed a larger number of classes to be detected on roofs, they used algorithms from an Application Programming Interface (API) that do not allow for hyperparameter adjustments, complicating model improvements.

5. CONCLUSION

This study proposed an image recognition model for residential roofs using the eighth generation of the YOLO architecture, YOLOv8. A rich dataset consisting of images from 167 buildings collected by Unmanned Aerial Systems (UAS) was used to evaluate the actual conditions of residential building roofs comprehensively. These images were annotated in CVAT to detect and classify three types of roof damage: extra tiles in the transversal direction, oxidation in gutters, and open skylights. The model was built using 394 images, with 70% used for training, 20% for validation, and 10% for testing. Despite the limited number of images used for training, the results were promising.
The YOLOv8 architecture demonstrates significant advantages over its previous versions. YOLOv8 allows for rapid detection and classification, offering easy implementation and efficiency in image analysis. Additionally, YOLOv8 can be scaled for real-time classification. One limitation in this study was the small dataset of only 167 buildings, which may have impacted the model's final performance. Moreover, image capture through drones may limit the variety of capture angles. Future work could consider a more diverse data collection approach.
Future studies will focus on developing YOLOv9 and YOLOv10 architectures for roof inspection. Additionally, object recognition models specific to roofing will be developed to enhance maintenance management. These models will integrate the generated data into the maintenance management system and evaluate potential improvements. Opportunities will also be explored to incorporate other technologies into the process, optimize automation, and include a wider variety of roof types and materials. To strengthen the model, the authors will add more images and expand the set of problem classes in residential roofs, aiming to increase robustness and accuracy in the analysis.

6. ACKNOWLEDGMENTS

The authors would like to thank the funding agencies CAPES and CNPq.

7. REFERENCES

Alashari, M., El-Rayes, K., Attalla, M., Al-Ghzawi, M. (2022). *Multivariate time series and regression models for forecasting annual maintenance costs of EPDM roofing systems.* Journal of Building Engineering, 54, 104618.
Alizadeh, M., Ma, J. (2021). *A comparative study of series hybrid approaches to model and predict the vehicle operating states.* Computers & Industrial Engineering, 162, 107770. https://doi.org/10.1016/j.cie.2021.107770

Alzarrad, A., Awolusi, I., Hatamleh, M. T., Terreno, S. (2022). *Automatic assessment of roofs conditions using artificial intelligence (AI) and unmanned aerial vehicles (UAVs).* Frontiers in Built Environment, 8, 1026225.

Avola, D., Cascio, M., Cinque, L., Fagioli, A., Foresti, G. L., Marini, M. R., Rossi, F. (2022). *Real-time deep learning method for automated detection and localization of structural defects in manufactured products.* Computers & Industrial Engineering, 172, 108512. https://doi.org/10.1016/j.cie.2022.108512

Brown, S., Harris, W., Brooks, R. D., Dong, X. S. (2021). *Fatal injury trends in the construction industry*.

Cao, Y., Pang, D., Zhao, Q., Yan, Y., Jiang, Y., Tian, C., Li, J. (2024). *Improved yolov8-gd deep learning model for defect detection in electroluminescence images of solar photovoltaic modules.* Engineering Applications of Artificial Intelligence, 131, 107866.

Chu, J. C., Shui, C. S., & Lin, K. H. (2024). *Optimization of trucks and drones in tandem delivery network with drone trajectory planning.* Computers & Industrial Engineering, 189, 110000. https://doi.org/10.1016/j.cie.2024.110000

Conceição, J., Poça, B., De Brito, J., Flores-Colen, I., Castelo, A. (2017). *Inspection, diagnosis, and rehabilitation system for flat roofs.* Journal of Performance of Constructed Facilities, 31(6), 04017100.

Dang, L. M., Kyeong, S., Li, Y., Wang, H., Nguyen, T. N., Moon, H. (2021). *Deep learning-based sewer defect classification for highly imbalanced dataset.* Computers & Industrial Engineering, 161, 107630. https://doi.org/10.1016/j.cie.2021.107630

Fan, C. L. (2024). *Using convolutional neural networks to identify illegal roofs from unmanned aerial vehicle images.* Architectural Engineering and Design Management, 20(2), 390-410.

Gajjar, D., Burgett, J. (2020). *Evaluating the Use of Unmanned Aerial Systems (UAS) to Perform Low-Slope Roof Inspections.* EPiC Series in Built Environment, 1, 214-222.

Garcez, N., Lopes, N., de Brito, J., Silvestre, J. (2012). *System of inspection, diagnosis and repair of external claddings of pitched roofs.* Construction and Building Materials, 35, 1034-1044.

Han, S., Park, W., Jeong, K., Hong, T., Koo, C. (2024). *Utilizing synthetic images to enhance the automated recognition of small-sized construction tools.* Automation in Construction, 163, 105415. https://doi.org/10.1016/j.autcon.2024.105415

Hou, M., Hao, W., Dong, Y., Ji, Y. (2023). *A detection method for the ridge beast based on improved YOLOv3 algorithm.* Heritage Science, 11(1), 167. https://doi.org/10.1186/s40494-023-00995-4

Jiang, D., Kong, L., Wang, H., Pan, D., Li, T., Tan, J. (2024). *Precise control mode for concrete vibration time based on attention-enhanced machine vision.* Automation in Construction, 158, 105232. https://doi-org/10.1016/j.autcon.2023.105232

Koubaa, A., Ammar, A., Kanhouch, A., AlHabashi, Y. (2021). *Cloud versus edge deployment strategies of real-time face recognition inference.* IEEE Transactions on Network Science and Engineering, 9(1), 143-160. https://doi.org/10.1109/TNSE.2021.3055835

Li, C., Li, L., Jiang, H., Weng, K., Geng, Y., Li, L., Wei, X. (2022). *YOLOv6: A single-stage object detection framework for industrial applications.* arXiv preprint arXiv:2209.02976. https://doi.org/10.48550/arXiv.2209.02976

Liu, Y. F., Nie, X., Fan, J. S., Liu, X. G. (2020). *Image-based crack assessment of bridge piers using unmanned aerial vehicles and three-dimensional scene reconstruction.* Computer-Aided Civil and Infrastructure Engineering, 35(5), 511-529.

Ma, X., Li, Y., Yang, Z., Li, S. (2024). *Lightweight network for millimeter-level concrete crack detection with dense feature connection and dual attention.* Journal of Building Engineering, 109821. https://doi.org/10.1016/j.jobe.2024.109821

Mostafa, K., Hegazy, T., Hunsperger, R. D., Elias, S. (2023). *Using image analysis to quantify defects and prioritize repairs in built-up roofs.* Facilities, 41(7/8), 498-509.
Ni, Y. H., Wang, H., Mao, J. X., Xi, Z., Chen, Z. Y. (2024). *Quantitative detection of typical bridge surface damages based on global attention mechanism and YOLOv7 network.* Structural Health Monitoring, 14759217241246953. https://doi.org/10.1177/14759217241246953
Olimov, B. A., Veluvolu, K. C., Paul, A., Kim, J. (2022). *UzADL: Anomaly detection and localization using graph Laplacian matrix-based unsupervised learning method.* Computers & Industrial Engineering, 171, 108313. https://doi.org10.1016/j.cie.2022.108313
Qiu, H., Zhang, J., Zhuo, L., Xiao, Q., Chen, Z., Tian, H. (2024). *Research on intelligent monitoring technology for roof damage of traditional Chinese residential buildings based on improved YOLOv8: taking ancient villages in southern Fujian as an example.* Heritage Science, 12(1), 231.
Redmon, J., Divvala, S., Girshick, R., Farhadi, A. (2016). *You only look once: Unified, real-time object detection.* In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 779-788).
Ribeiro, W. S., Zanetti, J., Totola, L. B., Junqueira, S. Â. C., Lauff, P. H. P. (2024). *Crack detection in buildings using the YOLO v8 network.* Revista ALCONPAT, 14(3), 288-298. https://doi.org/10.21041/ra.v14i3.76.
Rosenberger, T. (2012), *Roofing Systems in Costa Rica Need Careful Consideration*, The Costa Rica Star, https://news.co.cr/roofing-systems-in-costa-rica-need-careful-consideration/7258/. Acessado em outubro de 2023.
Santos, L. M. A., Zanoni, V. A. G., Bedin, E., Pistori, H. (2023). *Deep learning applied to equipment detection on flat roofs in images captured by UAV.* Case Studies in Construction Materials, 18, e01917.
Silva, C. M., Flores-Colen, I., Coelho, A. (2015). *Green roofs in Mediterranean areas–Survey and maintenance planning.* Building and Environment, 94, 131-143.
Silveira, B., Melo, R., Costa, D. B. (2020). *Using UAS for roofs structure inspections at post-occupational residential buildings.* International Conference on Computing in Civil and Building Engineering (pp. 1055-1068). Cham: Springer International Publishing.
Sohaib, M., Jamil, S., Kim, J. M. (2024). *An ensemble approach for robust automated crack detection and segmentation in concrete structures.* Sensors, 24(1), 257. https://doi.org/10.3390/s24010257
Solawetz, J. (2023). *What is yolov8? the ultimate guide.* Roboflow Blog, 25.
Staffa, L. B. J., Costa, D. B., Nogueira, J. L. T., and Silva, A. S. (2023). *Web platform for building roof maintenance inspection using UAS and artificial intelligence*. International Journal of Building Pathology and Adaptation. https://doi.org/10.1108/IJBPA-12-2022-0186.
Wang, C. Y., Bochkovskiy, A., Liao, H. Y. M. (2023). *YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors.* In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 7464-7475).
Wang, W., Chen, J., Huang, Z., Yuan, H., Li, P., Jiang, X., Lin, Q. (2023). *Improved YOLOv7-based algorithm for detecting foreign objects on the roof of a subway vehicle.* Sensors, 23(23), 9440. https://doi.org/10.3390/s23239440
Zhao, H., Wang, X., Sun, J., Wang, Y., Chen, Z., Wang, J., Xu, X. (2024). *Artificial intelligence powered real-time quality monitoring for additive manufacturing in construction.* Construction and Building Materials, 429, 135894. https://doi-org/10.1016/j.conbuildmat.2024.135894