A Segmentation-Based Evaluation for Aviation Operations and Visual Analysis

Chennakeshava Akhil Pillalamarri
MSDS-692 Data Science Practicum 1
Anderson College of Business and Computing
Regis University, Denver, CO
cpillalamarri@regis.edu

Abstract

The aviation industry generates large amounts of visual data from airport surveillance cameras, aircraft inspection images, and operational monitoring systems. These visual sources are important in terms of maintaining safety, monitoring airport activity, and analysing the aircraft conditions. However, manually analysing those large volumes of aviation images can be time-consuming and difficult to scale. Human inspection may also include any inconsistencies when different observers understand the visual information differently.

Recent developments in computer vision have made it possible to automate many visual analysis tasks. Particularly, deep learning models can detect objects and segment meaningful regions within any kind of images. This project explores the use of a segmentation-based computer vision pipeline for analysing aviation images. The pipeline combines object detection using a pretrained YOLO model with segmentation using the Segment Anything Model (SAM).

In this experimental approach, YOLO is first used to detect objects within an image and generate bounding boxes around possible aviation-related objects. These bounding boxes are then used as prompts for SAM, which produces segmentation masks highlighting those important visual regions. This allows the segmentation model to focus on relevant objects instead of background areas.

Multiple aviation datasets were used in this analysis, including aircraft detection datasets, aircraft damage imagery, and frames extracted from operational airport videos. The results include segmentation outputs as well as different visualization-based analyses which helps understand object detection patterns within the datasets. The overall goal of this project is to see whether a combined detection and segmentation pipeline can support aviation visual analysis tasks using an experimental computer vision pipeline.

I. Introduction / Background

The aviation industry depends heavily on visual monitoring systems to make sure of the operational safety. Airports generate huge amounts of visual data from surveillance cameras, runway monitoring systems, and aircraft inspection imagery. These visual data sources are useful for monitoring airport activity, inspecting aircraft surfaces, and understanding operational conditions.

Traditionally, many of these tasks depend on manual inspection performed by trained personnel. Aircraft surfaces are inspected visually to detect any kind of structural damage, and airport surveillance footage is monitored to observe the aircraft movement and ground activity. While manual inspection can be effective, mostly it becomes difficult to scale when large volumes of visual data needs to be analysed.

Recent advancements in computer vision and deep learning have showed automated approaches for analysing visual data. Research shows that computer vision techniques can support many aviation applications like airport surveillance, aircraft inspection, and operational monitoring [1]. These systems can improve the efficiency by automatically detecting most important objects or regions within the images.

Computer vision models have also been used for aircraft detection and recognition tasks using deep learning architectures [2]. Those works show how convolutional neural networks can identify aircraft structures within complex visual environments. More recent studies have explored object detection models such as YOLO for identifying types of aircraft and other flying objects in real-world images [3], [4].

However, many current computer vision systems are designed for specific tasks. Some models focus on aircraft detection while others are designed for defect detection or airport surveillance. Because of this, separate pipelines are required for different applications. This project explores whether a unified pipeline combining object detection and segmentation can support multiple aviation-related visual analysis tasks.

II. Problem Statement

Aviation visual analysis pose many challenges for modern data systems. One major challenge is relying on manual inspection processes. Aircraft inspection and airport monitoring tasks require human experts to analyse th visual data and gather relevant information. This process can be slow and difficult to scale when large datasets are involved.

Another challenge is the limited availability of labeled aviation datasets. Deep learning models usually require large annotated datasets for training purposes. However, labeling aviation images needs domain expertise, and can be hard and time-consuming.

Previous research has shown the effectiveness of computer vision techniques for specific aviation tasks. For example, automated systems have been proposed for collecting operational data from airport environments using computer vision techniques [5]. Other studies have explored deep learning models for detecting aircraft defects and structural damage in inspection imagery [6].

Computer vision has also been used for air-side surveillance and airport monitoring in low-visibility environments, where automated visual systems can assist operational observation and safety analysis [7]. Even though these approaches perform well for their specific tasks, they often require task-specific models and separate training pipelines.

This project investigates whether a unified segmentation-based framework can support multiple aviation visual analysis tasks by combining object detection and segmentation techniques.

III. Related Work

Several studies have explored the use of computer vision techniques in aviation environments. Survey studies highlight the growing role of visual analytics in airport monitoring, aircraft inspection, and aviation safety systems [1]. These studies show that automated image analysis can help aviation operations by identifying objects and patterns within visual data.

Research has also focused on monitoring airport operations using camera-based systems. Farhadmanesh et al. proposed a computer vision system designed to collect operational data from non-towered airports using visual monitoring methods [5]. Their work shows that visual analytics can help monitor airport activity and aircraft movement.

Deep learning approaches have also been applied to aircraft inspection tasks. Research on aircraft defect detection has shown that convolutional neural networks can identify aircraft surface damage in inspection images [6]. These methods typically require annotated datasets and specialized training procedures.

Object detection research has also been widely explored in computer vision. The YOLO framework introduced a unified real-time object detection architecture capable of detecting multiple objects within an image efficiently [8]. More recent studies have applied YOLO-based models to aircraft recognition and flying object detection tasks [3], [4].

Other research has investigated aircraft detection using convolutional neural networks applied to remote sensing images [2]. Studies on airport object detection have also showed how deep learning models can identify multiple airport-related objects within complex scenes [9].

Recent advances in computer vision have introduced foundation models capable of performing multiple visual tasks. The Segment Anything Model (SAM) is designed to generate segmentation masks across many image domains using prompt-based inputs [10]. Combining object detection with segmentation models provides an opportunity to build flexible computer vision pipelines that can operate across multiple aviation datasets.

IV. Methodology

The methodology used in this project follows a computer vision pipeline that combines object detection and image segmentation techniques. The pipeline was designed to analyse aviation imagery collected from multiple datasets and to generate both visual and analytical outputs.

A. Data Collection

The data used in this project was collected from multiple publicly available aviation sources in order to show different aviation environments. The goal was to include both image datasets and real-world operational images so that the proposed computer vision pipeline could be tested across various visual conditions.

In addition to existing aviation image datasets, video data showing aircraft operations was also used. Short aviation operation videos containing aircraft ground movements, taxiing activities, and takeoff or landing events were selected. Frames were extracted from these videos using the FFmpeg tool, which allows individual frames to be extracted from video sequences.

Frame extraction was performed at regular intervals so that representative frames from each video could be included in the dataset. This helped capture different operational stages of aircraft movement while keeping the dataset size manageable for execution.

Once the frames were extracted, they were used as image inputs and combined with the other aviation datasets used in the study. This approach allowed the dataset to include both structured aviation image datasets and real-world operational scenes.

B. Data Preprocessing

Before applying the computer vision models, several preprocessing steps using baic computer vision techniques were done on the images. Images were resized and formatted to ensure compatibility with the object detection and segmentation models. Preprocessing helps maintain consistent input sizes and improves model stability.

For video-based datasets, frames were extracted at fixed intervals to capture moments within the aviation operations. This approach makes sure that the dataset includes various visual scenes while keeping the total dataset size manageable for execution.

Even though the preprocessing steps are simple, they are important for maintaining consistency on different datasets and image sources.

C. Object Detection using YOLO

Object detection in this project was performed using a pretrained YOLO model. YOLO is widely known and used for real-time object detection tasks because of its efficiency and ability to detect multiple objects within a single image [8].

The YOLO model processes an input image and predicts bounding boxes around the detected objects. Each bounding box shows a region that likely has a relevant object, like an aircraft or another visible structure within the aviation scene.

Using object detection as the first stage of the pipeline helps identify the important regions within an image before applying the segmentation.

D. Segmentation using the Segment Anything Model

After object detection, the bounding boxes generated by YOLO are used as prompts for the Segment Anything Model (SAM). SAM is a foundation segmentation model capable of generating segmentation masks across many types of images [10].

SAM receives the bounding box coordinates and produces segmentation masks that highlight the object regions within the image. This step allows the model to produce more detailed representations of the detected objects compared to simple bounding boxes.

By combining object detection with segmentation, the pipeline is able to identify both the location and the shape of objects in aviation imagery. This approach allows the system to analyse images from multiple datasets without requiring separate models for each specific task.

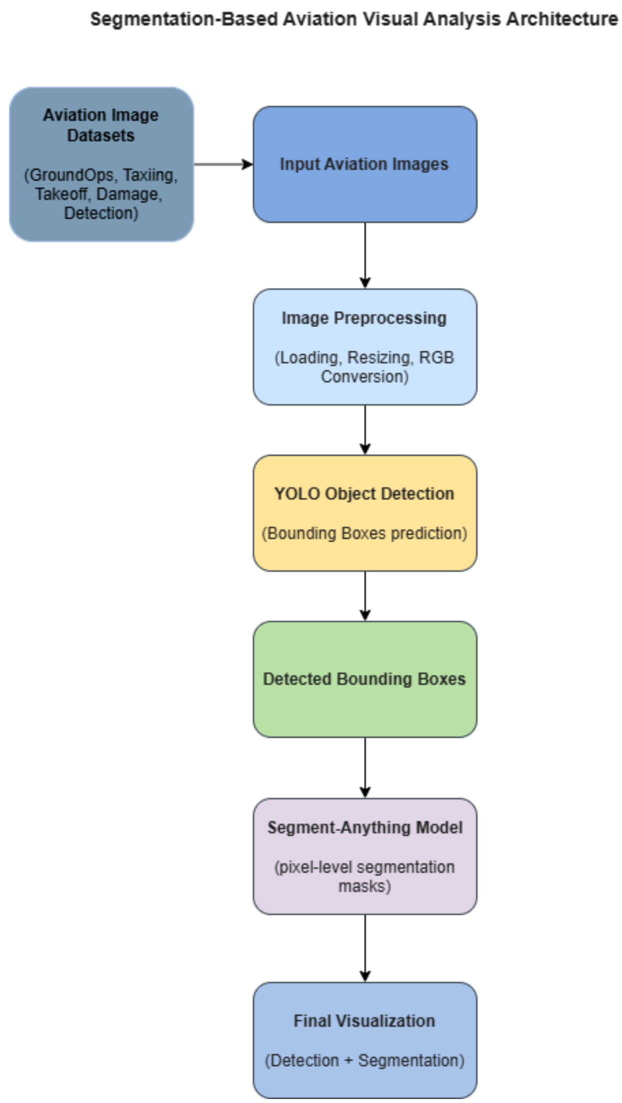

Fig. 1. Overall computer vision pipeline combining YOLO object detection and SAM segmentation.

V. Data Description

The dataset used in this project consists of multiple aviation-related image sources representing different operational scenarios. Using several datasets helps analyse the computer vision pipeline across various aviation environments.

One of the primary datasets used in the study is the FGVC Aircraft dataset. This dataset contains images of aircraft captured from different viewing angles and environments. The images include various aircraft models and provide visual variation in terms of lighting, background scenes, and aircraft orientation. These characteristics make the dataset useful for evaluating object detection performance.

Another dataset used in the project is the Aircraft Damage dataset. This dataset contains images of aircraft surfaces showing defects like dents, cracks, or structural irregularities. These images are particularly useful for examining how segmentation models behave when identifying damaged regions on aircraft surfaces.

In addition to these image datasets, frames extracted from aviation operation videos were included in the dataset. These frames show real-world airport operational scenes such as aircraft taxiing, ground handling activities, takeoff, and landing scenarios. Including these operational scenes helps evaluate how the detection and segmentation pipeline behaves in more complex environments compared to static inspection images.

The combination of aircraft inspection datasets and aviation operational images provides a balanced dataset which allows the pipeline to be evaluated across multiple aviation-related scenarios.

VI. Results and Analysis

After running the detection and segmentation pipeline, some visual outputs and analytical plots were generated to understand how the system behaves across different aviation datasets.

A. Pipeline Output Visualization

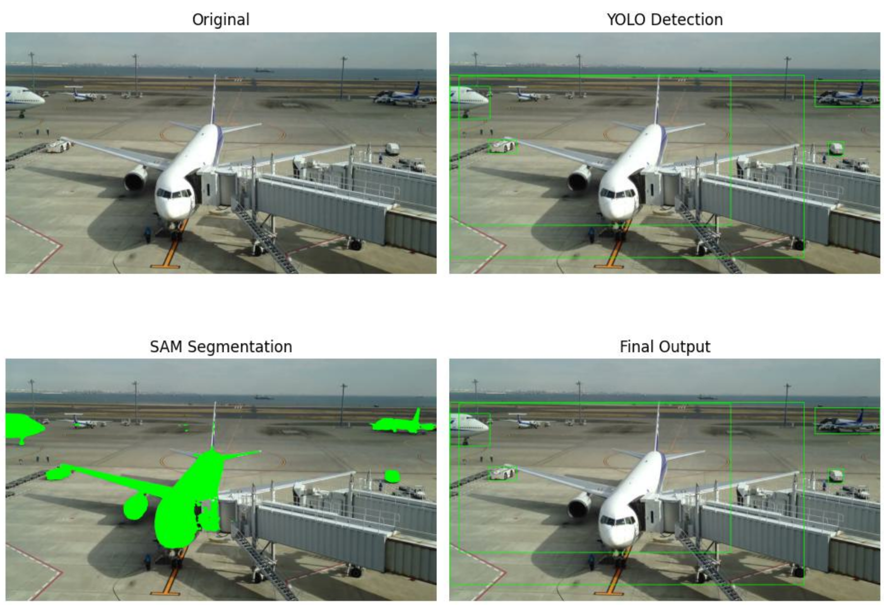

Fig. 2. Example pipeline output showing object detection and segmentation results.

To better understand how the pipeline works, some images were processed through the object detection and segmentation stages. The following figure shows an example of the complete pipeline output. The visualization includes the original image, the YOLO object detection results, the segmentation masks generated by SAM, and the final combined output.

B. Dataset Composition

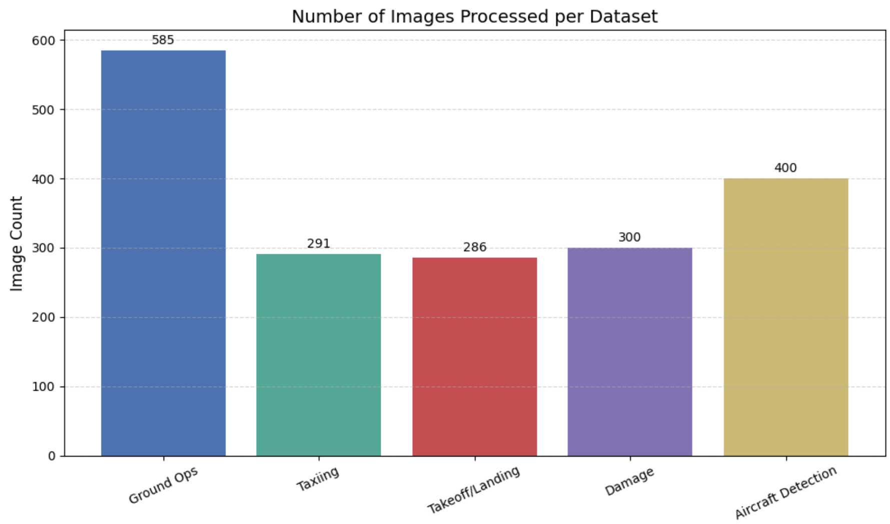

Fig. 3. Number of images processed for each aviation dataset.

Before analysing the detection behavior of the models, it is useful to understand how the images are distributed across the datasets used in this project. The following plot shows the number of images processed from each aviation dataset. This helps provide context for the analysis and ensures that the datasets contribute meaningfully to the overall experiment.

C. Average Objects Detected per Image

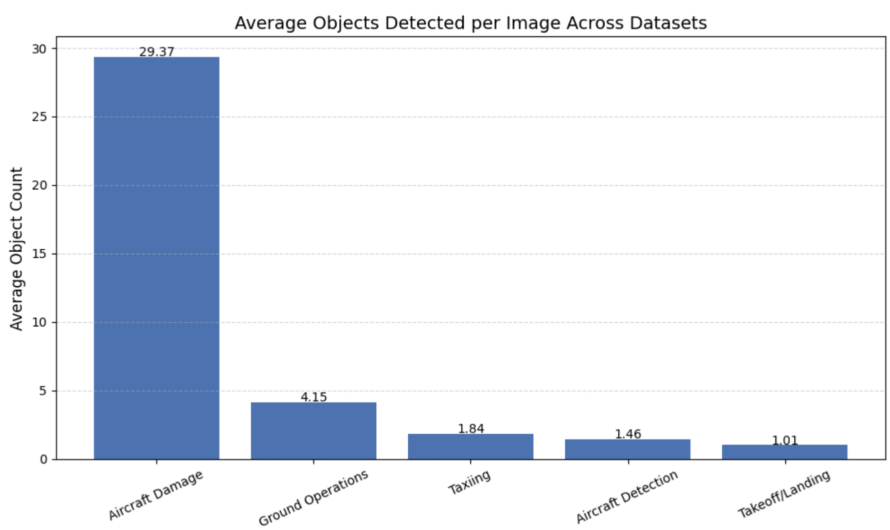

Fig. 4. Average number of detected objects per image across datasets.

To better understand detection activity in each dataset, the average number of objects detected per image was calculated. This provides an overview of how frequently objects appear in the different aviation scenarios. Some datasets contain images with only one or two detectable objects, while others contain scenes with multiple objects.

D. Detection Trends

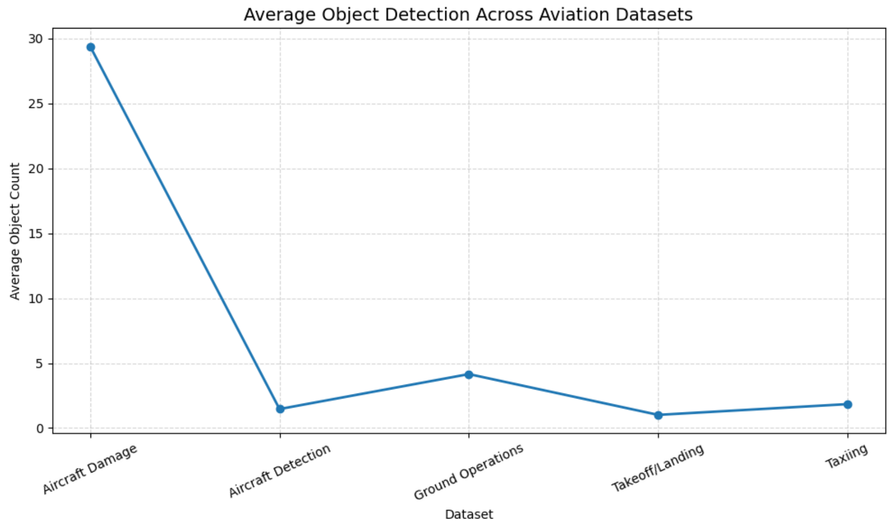

Fig. 5. Average object detection trends across datasets.

The next analysis analyzes the detection trends across the aviation datasets. The following line chart shows the average number of detected objects across the datasets. Differences in detection counts may occur due to variations in image complexity, object visibility, and scene composition.

E. Detection Distribution

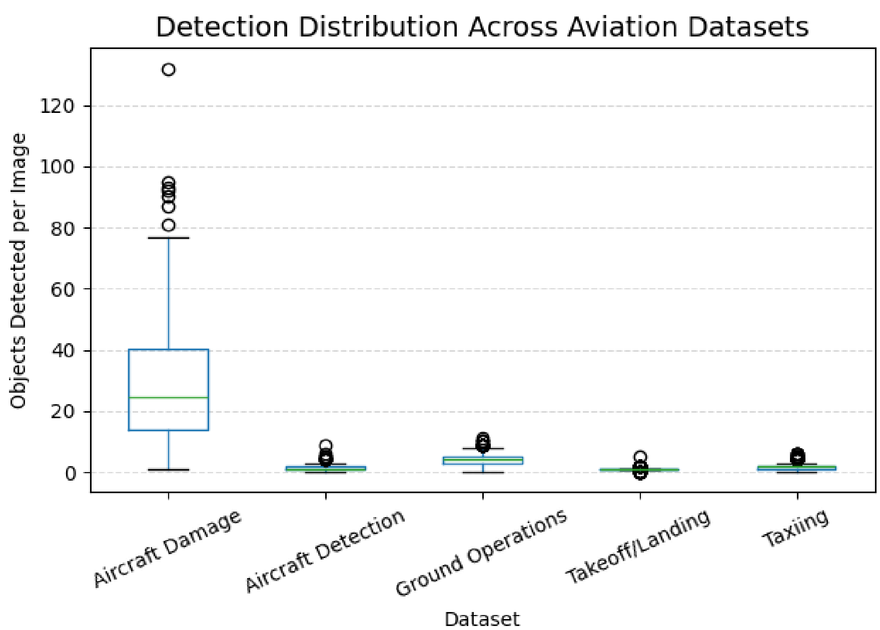

Fig. 6. Distribution of object detections across aviation datasets.

To analyse how detections vary within each dataset, a box plot was created. The box plot shows the distribution of object detections across datasets, including the median and the spread of values. Some datasets show a wider range of detections, indicating that images within those datasets contain different levels of visual complexity.

F. Overall Detection Distribution

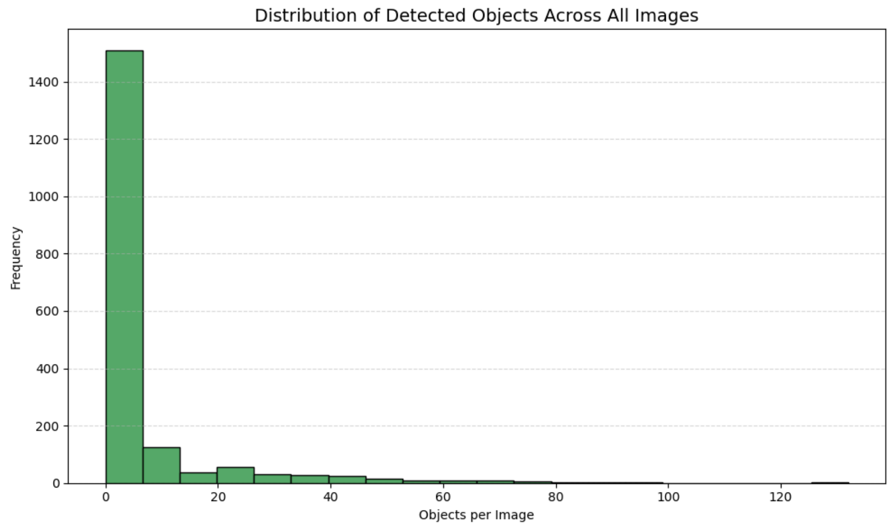

Fig. 7. Histogram showing distribution of detected objects across all images.

Finally, a histogram was generated to observe how object detections are distributed across all processed images. The histogram shows that the distribution of detections is right-skewed. Most images contain a small number of detected objects, while only a smaller number of images contain many detections.

VII. Future Work

While the pipeline shows how object detection and segmentation models can be combined for aviation image analysis, there are many chances for further improvement and expansion.

One possible extension of this work is the using larger and more diverse aviation datasets. In this project, only a limited number of publicly available datasets and video frames were used to show the possibility of the pipeline. Future work could use larger aviation datasets that include different airport environments, weather conditions, and aircraft types to further evaluate the scope of the approach.

Another direction for future work is including quantitative evaluation metrics. In this project, most of the evaluation was based on visual inspection and exploratory analysis. Future research could evaluate segmentation quality using metrics like precision and recall where ground truth annotations or labels are present.

Future work could also explore additional computer vision models for comparison. For example, other object detection architectures or segmentation models could be tested to understand how different models perform on aviation images.

Finally, the pipeline could be expanded into real-time monitoring applications. Integrating the detection and segmentation pipeline into airport surveillance systems may support automated monitoring of aircraft operations, runway conditions, and airport activity in real-world environments.

VIII. Conclusion

This project explored a segmentation-based computer vision pipeline for analysing aviation images. The system combines YOLO object detection with the Segment Anything Model to detect and segment meaningful regions within aviation images.

The results show that combining object detection with segmentation helps guide the segmentation model toward relevant visual areas. This allows the same pipeline to be applied across different aviation datasets without creating separate models for each task.

Through multiple datasets and exploratory analysis, the project shows how modern computer vision techniques can actually support aviation-focused visual analytics tasks.

References

[1] A. Talbi, M. Abdellaoui, J. Pons-Prats, and J. Kuljanin, “An overview on computer vision analysis in airport applications,” *IEEE Access*, vol. 13, pp. 120 446–120 471, 2025.

[2] L. Zhou and H. Yan, “Aircraft detection for remote sensing images based on deep cnn,” *Journal of Electrical and Computer Engineering*, 2021.

[3] C. Lei and J. Zeng, “Aircraft type recognition based on yolov8,” in *Journal of Physics Conference Series*, 2024.

[4] D. Reis and J. Kupec, “Real-time flying object detection with yolov8,” *arXiv preprint arXiv:2305.09972*, 2023.

[5] M. Farhadmanesh, Y. Yu, and B. Smith, “A computer vision-based standalone system for automated operational data collection at non-towered airports,” *IEEE Access*, vol. 12, pp. 152 345–152 360, 2024.

[6] N. Suvittawat and C. Kurniawan, “Advances in aircraft skin defect detection using computer vision,” *Aerospace*, vol. 12, no. 4, p. 356, 2025.

[7] P. Thai and S. Alam, “Airside surveillance by computer vision in low-visibility environments,” in *Air Traffic Management Research Seminar*, 2023.

[8] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified real-time object detection,” in *CVPR*, 2016, pp. 779–788.

[9] Y. Han and S. Ma, “Effective airport object detection using cnn,” *IEEE Access*, vol. 8, pp. 172 652–172 663, 2020.

[10] A. Kirillov *et al.*, “Segment anything,” in *ICCV*, 2023, pp. 4015–4026.