

Executive Summary

This work presents a UAV-enabled computer vision framework for automated roof damage detection, designed to improve the speed, safety, and consistency of post-storm inspections. The system identifies both hail and wind damage using aerial imagery and addresses a key barrier in the field by combining fragmented roofing datasets into a standardized multi-class dataset of over 10,000 labeled images. To enhance detection of small and subtle damage features, the framework incorporates a SAHI-inspired slicing strategy. The approach delivered a marked increase in detection accuracy, raising mean average precision from 0.35 to 0.65. Overall, the findings demonstrate strong potential for deployment in insurance and property assessment workflows, where faster inspections and more efficient claims handling are critical.

Background

Accurate roof damage assessment is critical for insurance processing, repair planning, and post disaster recovery after severe weather events. Traditional inspections are typically manual, slow, and potentially unsafe, especially when conducted on damaged roofs after hail or windstorms. UAV imagery and computer vision provide a promising alternative by enabling safer, faster, and more consistent inspection workflows. Prior research has demonstrated the feasibility of UAV-based roof inspection [4], but broader progress is limited by fragmented datasets and the difficulty of detecting small, subtle damage such as hail impacts. This makes automated roof damage detection an important problem with practical value for both property assessment and disaster response operations.



Figure 1. UAV-captured image of a residential roof showing missing shingles caused by wind damage. [5]

Experimental Setup

The system was designed to address two main technical challenges: inconsistent data sources and the small, subtle appearance of hail damage in aerial images. To support multi-class detection, several Roboflow datasets [3] were combined into a unified dataset with standardized annotations for hail and wind damage. Because the original datasets were collected under different UAV heights, viewing angles, and lighting conditions, the unified dataset was intended to improve training consistency while preserving real-world variation. The final dataset contained more than 10,000 labeled roof images at 640×640 resolution and was organized for train/test. To improve visibility of small damage regions, especially hail impacts, a SAHI-inspired slicing strategy [2] was applied during data preparation. This generated localized image samples that emphasized fine-scale damage features that may be difficult to learn from full-roof images alone. The detection model was based on RF-DETR Base [1], selected for its transformer-based object detection capability and its ability to be fine-tuned from pretrained weights. The final setup therefore combined three main design elements: a unified two-class dataset, localized sliced training samples, and a fine-tuned RF-DETR architecture.




Figure 2. Pipeline overview for training and evaluation

Experimental Test Plan

The primary evaluation compares two configurations: a baseline RF-DETR model trained without the proposed data preparation strategy, and the proposed RF-DETR model trained using the unified dataset and slicing approach. This comparison is intended to isolate the effect of the dataset design and localized image preparation. Performance will be measured primarily using mean Average Precision at an Intersection over Union threshold of 0.50 (mAP@50), since this metric evaluates both class prediction and bounding box localization quality [6]. Precision, recall, and class-level performance for hail and wind damage will also be reviewed to determine whether improvements are consistent across both damage categories. To ensure fair comparison, all experiments will use the same train, validation, and test splits, while keeping the model architecture and general training conditions consistent across runs. In addition, qualitative review of representative test images will be used to examine false positives, false negatives, and detection behavior under different roof textures and UAV viewing conditions. This test plan is intended to show whether the proposed method provides a measurable improvement over baseline detection and whether it is effective for challenging small-damage cases.

Experimental Results

The baseline model achieved a mAP@50 of 0.35. After applying the unified dataset design and SAHI-inspired slicing strategy, the fine-tuned RF-DETR model achieved a mAP@50 of 0.65. This improvement indicates that the proposed preprocessing approach helped the model better detect roof damage, particularly small and subtle patterns such as hail impacts. These results show that dataset structure and image scale played a significant role in detection performance. Overall, the proposed method improved automated roof damage detection and demonstrated the value of combining dataset unification with localized image slicing for UAV-based inspection.



Figure 3. Top-down UAV image of a residential roof with hail damage detections from our model, shown as red bounding boxes with confidence scores [3]

Conclusions

This project presented an automated UAV-based roof damage detection framework for identifying hail and wind damage. The final design addressed the client's goal of improving the safety and efficiency of roof inspection while working within a major constraint of the field: limited and inconsistent public training data. The results showed that hail damage detection is a scale-sensitive problem, since small and subtle damage patterns are difficult to capture and localize in standard aerial imagery. By combining dataset unification with slicing during training and inference, the proposed approach improved the model's ability to detect fine-scale damage features. This demonstrated that data quality, image scale, and annotation consistency were just as important as model selection in achieving strong detection performance. At the same time, the project revealed important limitations. Variation in UAV height, angle, and lighting across datasets introduced inconsistency, and slicing could only enhance visible detail rather than create missing information. Future work should therefore focus on higher-resolution imagery, more standardized annotations, and broader testing in real-world inspection settings. We acknowledge the client for defining the project goals and the data sources that supported development of this work, including the Roboflow datasets used in training and evaluation.

References

1. Isaac Robinson, undefined., et al, "RF-DETR: Neural Architecture Search for Real-Time Detection Transformers," 2026.
2. F. Akyon, S. Onur Altinuc, A. Temizel, "Slicing Aided Hyper Inference and Fine-Tuning for Small Object Detection," in 2022 IEEE (ICIP)
3. haildamage-devyz, "Hail Damage Test Dataset," , 2024.
4. Zhao, X., & Jebelli, H. (2025). A computational method for real-time roof defect segmentation in robotic inspection.
5. Five Points Roofing. (n.d.). Types of roof storm damage. <https://fivepointsroofing.com/roofing-blog/roof-storm-damage/>
6. Everingham et al., The PASCAL Visual Object Classes (VOC) Challenge