Artificial Intelligence Powered Autonomous and Distributed Multi-Drone Platform

Daniel Miner, William States Lee College of Engineering, UNC Charlotte

Dr. Ran Zhang, Dept. of Electrical and Computer Engineering, UNC Charlotte

UNIVERSITY OF NORTH CAROLINA CHARLOTTE

Introduction

After a natural disaster, search teams need a fast and organized way to scan large areas and locate people, supplies, hazards, or other important objects.

This project supports an Artificial Intelligence Powered Autonomous and Distributed Multi-Drone Platform: a three-drone system intended to sweep a target area, share observations, and respond when a target is found.

Student contributions focused on the vision and communication components needed before the full system can operate in the field. The work included a custom-object-detection training package, water-bottle proof-of-concept testing, and a drone-relay prototype for drone-to-drone communication.

YOLO, which stands for "You Only Look Once," is a real-time object detection model that analyzes an image in a single pass. It identifies what objects are present and also determines where each object is located by drawing a bounding box around it. This makes YOLO especially useful for drone systems because the drone needs more than a simple label. It needs the object's position in the camera view so it can decide where to move, where to aim the camera, or what area to inspect next.

Objectives

Project Objectives

- Make custom object detection easier to repeat for future SmartNet drone tests
- Use a water bottle as a proof-of-concept target before moving to people, supplies, or other search targets
- Measure performance with held-out images instead of examples only seen during training
- Keep camera output simple: target name, confidence, image location, depth/health status, and frame ID
- Start a fallback communication layer that prefers TP-Link but can route small messages through nearby drones when needed

Method

Vision and Training Method

The training package was designed so future targets can follow the same workflow: collect photos, create labels, train YOLO, evaluate results, and prepare the model for OAK camera use.

Positive images showed the target object. Negative images showed similar scenes without the target, so the model could learn what not to detect.

GroundingDINO was used to suggest starting boxes around the object. Those labels still needed human review before training.

Training-only image changes such as glare, blur, lighting shifts, and perspective changes helped test whether the detector could still recognize the same object under harder conditions.

The OAK/DepthAI workflow kept the camera output simple: the drone program received compact detection information instead of making every flight decision from full image streams.

Collected Data

Collected Data and Student-Created Artifacts

- Student-created materials include the YOLO training package, drone-relay prototype, water-bottle evaluation report, and project abstract.
- Training set after augmentation: 4,833 images.
- Validation set: 94 untouched images.
- Test set: 921 images, including 100 untouched positive images, 22 untouched negative images, and 799 augmented stress-test images.
- The current communication work uses TP-Link router communication while the access-point relay prototype is still being developed.

Results

Current Results and Lessons

- Independent test: 117/122 images correct (95.90%).
- Detected 95/100 untouched water-bottle images.
- Rejected 22/22 no-bottle images.
- Stress test: 759/799 transformed target images detected.
- The model stayed nearly as accurate on transformed test images, but the 22-image negative set is still too small to prove field reliability.

What was learned: the workflow can train and test a useful custom detector, but field deployment needs more real outdoor negatives, more missed-case review, and hardware testing on the OAK camera during drone operation.

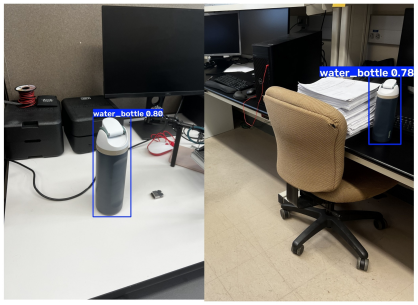

Photos: examples of bounding box detections of water bottles.

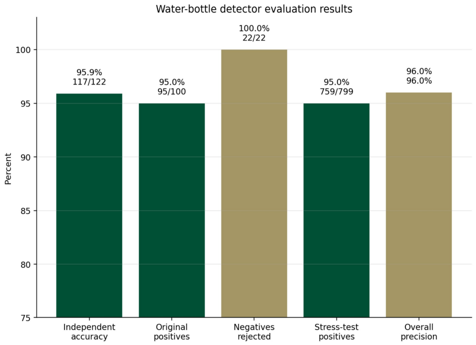

Chart: held-out water-bottle detector results from the YOLO11n evaluation report.

Field Photo: Fly4Future drone being tested outdoors.

Conclusions

Significance and Next Steps

This work is significant because it turns drone vision from a one-time model test into a repeatable pipeline that can be reused for new targets. A water bottle was a safe proof-of-concept; the same workflow could later be adapted for emergency supplies, hazards, or people after proper data collection and safety review.

The results show that the training package can produce a strong first detector, but they also show the next evidence needed: more real no-target field images, more examples of missed cases, and camera-on-drone tests rather than only laptop evaluation.

Next steps are to run the trained model through the OAK camera pipeline, connect detection packets to the three-drone sweep behavior, and measure communication performance such as range, reconnect time, packet loss, and whether a relay drone can keep the group connected when the router signal weakens.

References

Zhang, R. (2026, May 18). SmartNet Lab summer research kickoff meeting [PowerPoint slides]. Laboratory for Smart Wireless Communication Networks, William States Lee College of Engineering, University of North Carolina at Charlotte.

Luxonis. (n.d.). DepthAI software V3 documentation. Retrieved July 13, 2026, from https://docs.luxonis.com/software-v3

Ultralytics. (n.d.). Ultralytics YOLO model documentation. Retrieved July 13, 2026, from https://docs.ultralytics.com/models/

IDEA Research. (n.d.). GroundingDINO [Computer software]. GitHub. Retrieved July 13, 2026, from https://github.com/IDEA-Research/GroundingDINO