WAGENINGEN
UNIVERSITY & RESEARCH

A drone that learns to efficiently find objects in agricultural fields: from simulation to the real world

Rick van Essen, Gert Kootstra

Introduction

- Drones are promising for data collection in agriculture, but their battery capacity limits the area covered in a single flight.
- Some dominant weed species grow in spatially distinct patterns in the field instead of randomly.
- Reinforcement Learning (RL) can be used to plan paths to find objects, such as weeds, in a field by exploiting these spatial object distributions.
- We propose to train an RL agent in an abstract simulation environment and apply it in the real world by combining it with a detection network (Figure 1).

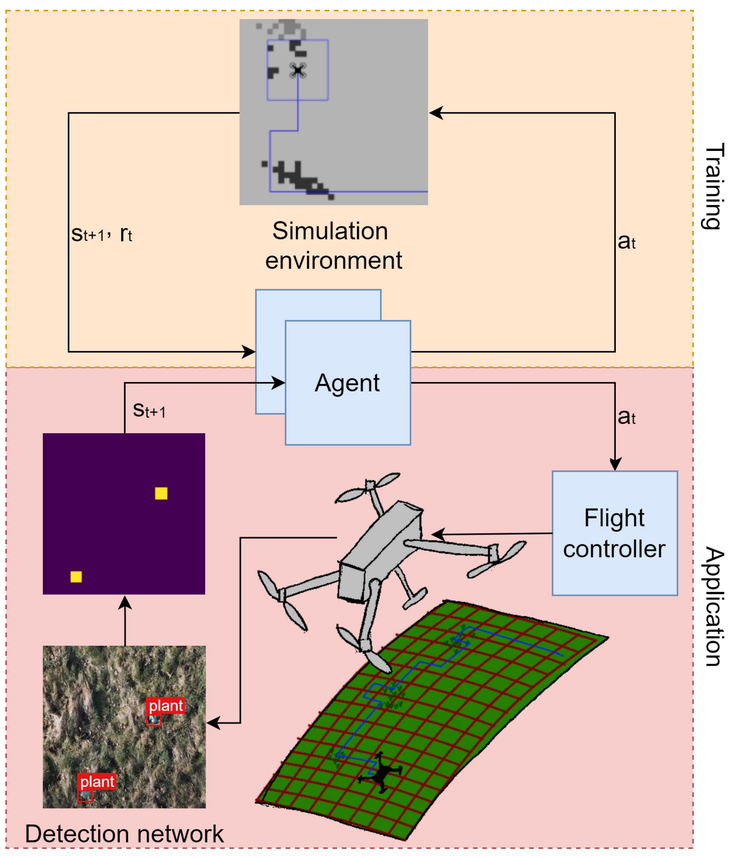


Figure 1. Overview of the learning-based adaptive path planner, which learns a policy in simulation that can be applied in the real world.

Contributions

- Evaluate the simulation-trained RL policy on real-world data.
- Discuss the performance differences between the simulation and the real-world data.

Methods



Figure 2. Example of a drone image containing an artificial plant.

RL agent:

State-space: low-resolution, uncertain prior knowledge about the field, and the output of an object detection network.

Action-space: *north*, *south*, *east*, *west*, and *terminate search.*

Training: 10M timesteps in the simulation environment.

Real-world evaluation dataset:

Data: four high-resolution orthomosaics containing artificial plants (proxy for weeds) (example image in Figure 2).

Repetitions: four rotations per orthomosaic, resulting in 16 evaluations.

Object detection: YOLOv8n.

Experiments:

Realism levels: four realism levels that gradually introduce more real-world data (Table 1).

Baseline: full coverage flight path.

Table 1. Four realism levels that gradually increase the realism from full simulation (level 1) to the real-world dataset (level 4).

Level	Object distribution	Prior knowledge	Detections
1	Simulated	Simulated	Simulated
2	Real-world	Simulated	Simulated
3	Real-world	Real-world	Simulated
4	Real-world	Real-world	Real-world

Results

- Up to around 500 flight steps, the RL agent was faster in discovering the objects (Figure 3).
- Around one-third of the lower performance on real-world data can be attributed to the real-world object distribution, and two-thirds to real-world prior knowledge.
- Most clusters of objects were found, but the agent often terminates too early, resulting in a longer flight path than needed (Figure 4).

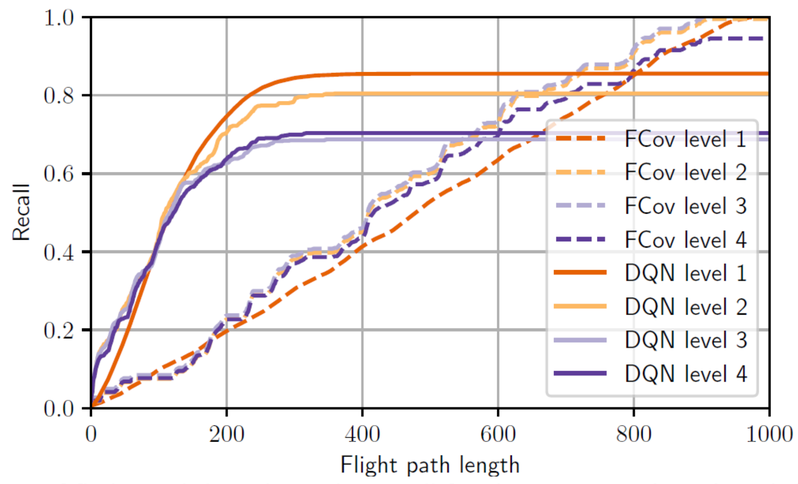


Figure 3. Effect of flight path length on the recall for the RL agent (DQN) and the baseline full coverage planner (FCov) with the different realism levels. Note that the lines for 'FCov level 2' and 'FCov level 3' are equal.

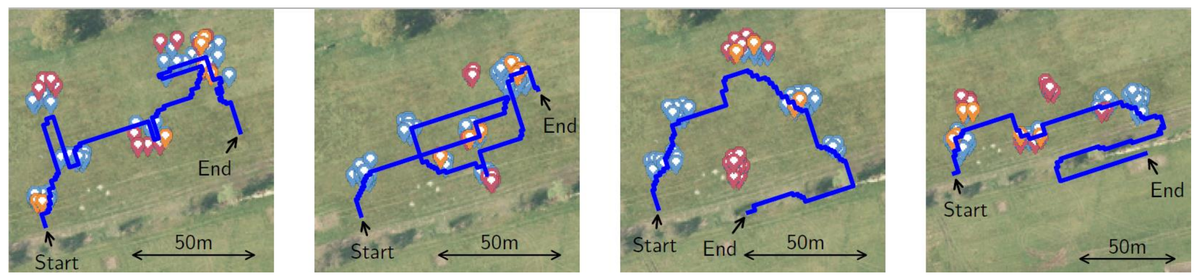


Figure 4. Flight paths for the RL agent on the four real-world datasets. Blue, orange, and red markers indicate true positive, false positive, and false negative markers, respectively. The start and end location of the flight path is indicated with an arrow.

Discussion and conclusion

- **Simulation:** 75% shorter path at the cost of a 14% lower recall.
- **Real-world data:** 72% shorter path at the cost of a 25% lower recall.
- The lower real-world performance was mainly due to prior knowledge errors, which may be solved by increasing its uncertainty in the training simulation.
- When the goal of an application is to find the most objects in a short time window, the learned path planner is suitable and efficient.

Acknowledgements

This research is part of the research program Synergia, funding was obtained from the Dutch Research Council (NWO grant 17626), IMEC-OnePlanet, and other private parties



Rick van Essen
Wageningen University & Research
Agricultural Biosystems Engineering
Contact: rick.vanessen@wur.nl

edu.nl/97bg3

Find our paper here!

Synergia