Poster: Profiling Edge Resource Demands of Zoom Maneuvers for Autonomous Unmanned Aerial Vehicles

Kevyn Angueira Irizarry and Christopher Stewart
Ohio State University

ABSTRACT

Autonomous Unmanned Aerial Vehicles (AUAVs) use software, rather than human pilots, to decide where and when to fly, when to capture data, and when to land. AUAVs generate detailed and dynamic maps that inform decisions across a wide range of applications from digital agriculture to wildlife conservation to forestry to smart cities. For example, in digital agriculture, farmers use crop health maps to tailor the application of pesticides to the specific needs of each management zone, increasing crop yield and optimizing resources. AUAVs use edge computing to process captured data and determine their flight path. However, edge processing demands differ between competing designs for AUAVs. For example, AUAVs that fly to preset waypoints in an automated fashion consume significant battery resources but fewer computational resources. In contrast, reinforcement learning (RL) designs wherein AUAVs select waypoints to maximize their reward function can save battery but require more computational resources. This poster will discuss our early efforts and research strategies in profiling the trade-offs in battery and computational resources from automated and RL approaches for zoom maneuvers. A pivotal element for low-cost mapping, zoom maneuvers reduce the AUAV's altitude to increase data resolution, capturing details previously obscured at higher altitudes. Zoom maneuvers can qualitatively improve the efficacy of AUAV missions, impacting resource requirements. In practice, a better understanding of zoom approaches will provide additional avenues for optimization of map generation, improving the ability of autonomous AUAVs to scale to large scale missions.

ACM Reference Format:
Kevyn Angueira Irizarry and Christopher Stewart. 2023. Poster: Profiling Edge Resource Demands of Zoom Maneuvers for Autonomous Unmanned Aerial Vehicles. In . ACM, New York, NY, USA, 2 pages.

1 INTRODUCTION

Autonomous unmanned aerial vehicles (AUAVs) conduct complex missions without human piloting [1, 3, 4, 8–10]. Services like Percepto employ AUAVs to inspect large industrial facilities for gas leaks, overheating equipment, and degraded structures, allowing companies to take early actions to address these problems [9].Similarly, AUAVs in digital agriculture capture high-resolution images of crop fields and convert them into maps that characterize crop health [2, 4, 5]. Additionally, AUAVs have further applications in smart cities, forestry, wildlife conservation, and military. Like unmanned aerial vehicles, AUAVs can conduct missions that are too risky for manned aircraft. As a type of unmanned aerial vehicle, AUAVs are not piloted remotely by human operators. Instead, they use data captured by onboard sensors (e.g., cameras and GPS) to decide where to fly next and when to land. AUAVs rely on platforms that allow software to issue commands during flight. Such platforms are provided aircraft manufacturers (e.g., DJI and Parrot), open-source flight-control systems (e.g., Pixhawk and Ardupilot), and AI-driven platforms for navigation (e.g., SoftwarePilot and Aerostack).

Most traditional flight-control platforms only support UAVs through automated waypoint missions or predetermined flight paths. This approach utilized by automated UAVs usually relies on exhaustively scouting all states (e.g. GPS locations) for an entire region. However, if adjacent or similar states convey the same or correlated information, exhaustive automated approaches waste limited battery resources without contributing compensatory benefits. By contrast, reinforcement learning (RL) AUAVs approaches exploit the correlation of adjacent states to maximize the data map accuracy while minimizing the number of states required to be visited. These RL approaches conserve battery by requiring a lower number of states. For this reason, it is worth researching the possible benefits of the RL approach, but its implementation would require novel resources to support it. To address this lack of resources, our group previously developed SoftwarePilot 2.0. This is a middleware designed for rapid implementation of both explicit UAV and UAV swarm autonomy [6].The utility of SoftwarePilot 2.0 is in providing a general RL solution to multiple competing goals and budgets in an autonomous UAV context. SoftwarePilot 2.0's crop health mapping balances crop stressors and health metrics that are competing for budgets, maximizing mapping accuracy within the given budget.

This paper proposes improvements to AUAV RL mapping approaches through an exploration into the zoom maneuver [11]. AUAVs can adjust their altitude during flight to sense their surroundings in greater detail. For example, using a 4K HD camera, an AUAV flying 5 meters will capture images where each pixel represents a 2-millimeter area and the picture spans 2-3 meters. Conversely, flying at 100 meters yields coarse images that can cover a hectare. Incorporating zoom maneuvers can significantly improve the execution of missions by (1) unveiling previously obscured data or (2) reducing the number of visited waypoints and saving crucial battery resources.

2 PROPOSED STUDY

The zoom maneuver includes any policies that adjust the altitude of UAVs or improve data accuracy while localizing data recollection. This paper aims to perform a study on the impacts of combinations of exploration and zoom policies on crop health map accuracy. We

Symposium on Edge Computing, 2023, Delaware, US

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00

will outline a procedure to measure algorithm efficiency in terms of accuracy and resource management. Improvements in health map accuracy from zoom maneuvers will be transferable to other models for autonomous health mapping and will improve the efficiency of UAV use in precision agriculture. The results of the study will inform the utility and implementation of the zoom maneuver into SoftwarePilot 2.0 [6, 7]. In total we defined four key implementations across our study: Auto Exploration (1)Auto Zoom, (2)Auto Exploration RL Zoom, (3)RL Exploration Auto Zoom, and (4)RL Exploration RL Zoom. We define Auto as automated, i.e. a preset route, and RL as a reinforcement learning policy, making live informed decisions. These four implementations were chosen to compare and contrast the improvements of RL strategies for exploration and zoom independently over autonomy, and the improvements of RL strategies for joint exploration and zoom over other methods. Each of the methodologies above must compete for the highest accuracy within the same constraints for battery, time, and number of states. These arrangements allow us to compare different reinforcement learning strategies against a set baseline. We now define the different strategies:

1. **Automated Exploration + Automated Zoom** uses an automated lawnmower pattern to represent exhaustive search for both exploration and zoom. It employs a random policy to decide when to employ the zoom maneuver. This method suffers from the highest flight costs as it must visit at least one state per management zone and all lower states on a zoom. However, as this model does not employ RL, it suffers no hovering costs waiting on the next action from the edge device.
2. **Automated Exploration + RL Zoom** uses an automated lawnmower pattern for exploration and an RL model for zoom. Its RL indicates when to zoom and what lower states to explore. This method must explore at least each state from every management zone; however, it may choose how many states to explore during zoom.
3. **RL Exploration + Automated Zoom** uses RL for exploration and an automated lawnmower pattern for zoom. Its RL strategically decides its next action for exploration while it employs a random policy to decide when to zoom. This method may reduce its total flight costs through strategic location, however it cannot control when to zoom and must explore all four lower states.
4. **RL Exploration + RL Zoom uses RL** for both exploration and zoom. It lumps lower and higher states into as possible waypoints for a single model to learn from, the model may freely choose actions to move to lower states as if they were adjacent states. This method may reduce its total flight cost through strategic locations, however it incurs the highest decision hovering costs as it must wait for a response from the edge for each action.

Experimental Plan: These experiments will allow us to profile the performances of different combinations of exploration and zoom strategies. For each we method we will track the battery consumption and computational costs. Automated searches from method 1 serve as the baseline and represent the traditional exhaustive approaches. By comparing methods 2 and 3 to method 1 we can measure the improvements in mapping accuracy from RL Zoom and RL Exploration respectively. Additionally comparing method 4 to method 1 demonstrates the total improvements from the baseline. This methodology may be then repeated with multiple versions of RL algorithms, different automated strategies to measure the impacts of the Zoom maneuver across different conditions and develop the most optimal solution for the given task. The results of this study will then inform future implementation of the *Zoom maneuver* into SoftwarePilot 2.0 and other AUAV RL strategies.

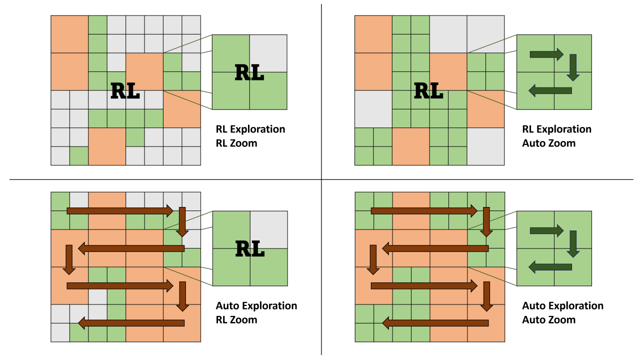

Fig. 1: The four methods being studied: Auto Exploration Auto Zoom, Auto Exploration RL Zoom, RL Exploration Auto Zoom, and RL Exploration RL Zoom

Acknowledgments: This work was funded by NSF Grant OAC-2112606 and the Ohio Soybean Council.

REFERENCES

[1] B. Boroujerdian, H. Genc, S. Krishnan, W. Cui, A. Faust, and V. Reddi. Mavbench: Micro aerial vehicle benchmarking. In *MICRO*, 2018.
[2] J. Boubin, C. Burley, P. Han, B. Li, B. Porter, and C. Stewart. MARbLE: Multi-Agent Reinforcement Learning at the Edge for Digital Agriculture. In *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)*, pages 68–81. IEEE, 12 2022.
[3] J. Boubin, J. Chumley, C. Stewart, and S. Khanal. Autonomic computing challenges in fully autonomous precision agriculture. In *2019 IEEE International Conference on Autonomic Computing (ICAC)*. IEEE, 2019.
[4] J. G. Boubin, N. T. Babu, C. Stewart, J. Chumley, and S. Zhang. Managing edge resources for fully autonomous aerial systems. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pages 74–87. ACM, 2019.
[5] J. del Cerro, C. Cruz Ulloa, A. Barrientos, and J. de León Rivas. Unmanned aerial vehicles in agriculture: A survey. *Agronomy*, 11(2):203, 2021.
[6] K. A. Irizarry, Z. Zhang, C. Stewart, and J. Boubin. Scalable distributed microservices for autonomous uav swarms. In *Proceedings of the 23rd International Middleware Conference Demos and Posters*, pages 1–2, 2022.
[7] T. P. Kelly, A. X. Zhang, and C. C. Stewart. Determining performance of an application based on transactions, May 18 2010. US Patent 7,720,955.
[8] Nothrop-Grumman. Autonomous Systems. https://www.northropgrumman.com/what-we-do/air/autonomous-systems, 2023.
[9] Percepto Inc. The Differences Between UAV, UAS, and Autonomous Drones. https://percepto.co/what-are-the-differences-between-uav-uas-and-autonomous-drones/, 2023.
[10] J. L. Sanchez-Lopez, R. A. S. Fernández, H. Bavle, C. Sampedro, M. Molina, J. Pestana, and P. Campoy. Aerostack: An architecture and open-source software framework for aerial robotics. In *International Conference on Unmanned Aircraft Systems*, 2016.
[11] M.-D. Yang, J. G. Boubin, H. P. Tsai, H.-H. Tseng, Y.-C. Hsu, and C. C. Stewart. Adaptive autonomous uav scouting for rice lodging assessment using edge computing with deep learning edanet. *Computers and Electronics in Agriculture*, 179:105817, 2020.