Multi-Agent Reinforcement Learning for Heterogeneous UAV Swarm Enabling Detailed Crop Health Assessment

Rugved Katole[1], Kevyn Angueira [2], Arpita Sinha [1] and Christopher Stewart [2]

Abstract— **Over the last few years, precision agriculture has advanced significantly with the aid of unmanned aerial vehicles (UAVs) and multi-agent systems (MAS). Traditionally, UAVs exhaustively scout the field and predict crop health, but this practice levies higher costs in terms of energy and execution time. In this paper, we propose an alternative approach where UAVs sample only a part of the field to predict the overall crop health. The selection of areas in the field to be sampled is based on different indices such as NDVI (Normalized Difference Vegetation Index), GLI (Green Leaf Index), and NDWI (Water Index). These vegetation indices indicate various factors of plant health. By correlating and quantifying these indices, we can assess the overall health of the crop field. Moreover, the individual indices provide a finer level of detail in precision agriculture, allowing for targeted measures to enhance yield. Our approach employs reinforcement learning and deep learning techniques to autonomously scout and predict the crop health map. Preliminary results show that by sampling only 40% of the field, we can generate a health map with 90% accuracy. This approach reduces labor costs by 4.8 times and increases profits by 36% compared to traditional methods.**

I. Introduction

The rapid growth of the global population and the increasing demand for food necessitates swift advancements in agricultural practices to meet the pressing needs of humanity. Projections indicate that the combination of population expansion and heightened per capita food consumption will require a substantial 70% increase in agricultural yields by the year 2050 [1], [2]. However, the looming challenge of climate change is exacerbating the complexity of farming, as it contributes to stressors on crop health, such as drought, diseases, and pest infestations [3]. These adverse effects are projected to lead to a considerable reduction in crop yields by up to 11% by 2050 [4].

Precision agriculture occupies a pivotal role in meeting the escalating global food demand, as it centers upon the optimized utilization of resources and the maximization of yields through the strategic integration of technology. This, in turn, contributes significantly to the stabilization and reliability of the global food supply chain. Precision agriculture is a promising step toward improving efficiency and reducing adverse impacts of agriculture production [5]. It assesses the variation across the crop fields and divides the field into multiple management zones. So they can be treated efficiently and effectively [6], [7]. The advent of digital agriculture, or data-driven precision agriculture, employs a suite of tools including remote sensors (e.g., satellites and UAVs), in-field sensors (such as embedded soil sensors), and data processing techniques (e.g., machine learning)[8]. This combination informs decisions related to planting, harvesting, and crop treatment, all aimed at maximizing yield and minimizing the environmental impact of agricultural activities. Frequent sensing using these technologies enables the detection of crop health stress due to factors like drought and heat, the identification of diseases, pests, and other harmful phenomena [9], [10], [11], [12]. A pivotal task in digital agriculture involves transforming the collected data into health maps, providing valuable geospatial insights into crop health, and guiding effective crop treatment strategies, commonly referred to as crop scouting.

Traditionally, data acquisition is approached through two main methods: human piloting of UAVs to capture high-resolution images or UAVs autonomously scouting the entire fields. Human pilots, while capable of capturing accurate data, tend to escalate operating costs due to the need for frequent field mapping. Conversely, the autonomous UAV scouting method is cost-effective, but it often leads to redundant data due to a 60-70% side overlap in captured images. Moreover, both approaches necessitate frequent battery replacements due to limited flight times [13], which subsequently extend execution times and have an impact on profit margins.

The UAVs are equipped with various imaging sensors, including RGB, multi-spectral, thermal, and hyper-spectral. RGB cameras are well-suited for tasks such as growth prediction, biomass estimation, and canopy height measurement. On the other hand, multi-spectral cameras excel in early detection of drought stress, identification of pests, yield prediction, and their combination with thermal and hyper-spectral data enables estimation of nutrient status, pathogen presence, and weed detection [14]. Unlike RGB cameras, multi-spectral cameras capture both visible and invisible light spectra, enhancing the assessment of crop conditions and thereby enabling more informed agricultural decisions [15], [16]. Hyper-spectral and thermal cameras capture distinct bands of the invisible light spectrum. Hyper-spectral sensors are particularly effective in the early detection of pathogens and diseases [17], while thermal cameras are effective in identifying drought stress in crops [18]. RGB and multi-spectral cameras are commonly used whereas hyper-spectral and thermal are less common due to relatively higher costs [19].

Given the limited payload capacity of UAVs, they are constrained to carrying only one imaging sensor at a time

*This work was supported by any ICICLE AI Institute, TIH Foundation, IIT bombay

[1] Systems and Control Engineering, IIT Bombay

[2] Department of Computer Science, Ohio State University

[20]. Consequently, achieving a comprehensive analysis of a crop field necessitates the deployment of multiple drones. However, employing multiple drones for scouting an entire field introduces additional operational and maintenance expenses that may outweigh the potential gains. The increased costs associated with utilizing multiple drones can present a challenge in terms of maintaining profitability.

Contributions: In this paper, we present an efficient method for detailed analysis of whole-field without exhaustively scouting entire field. We employ a swarm of heterogeneous UAVs with distinct capabilities. We utilize multi-agent reinforcement learning to scout crucial areas through competing rewards, as a result battery replacements and payload requirements are minimized. The collected data is combined and extrapolated to provide deeper insights on crop health eliminating the need of exhaustive scouting of the whole field.

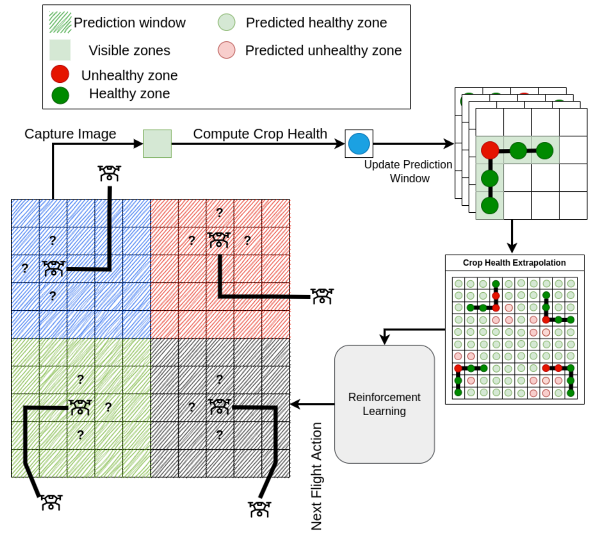

Fig. 1: An illustration of crop scouting using 4 UAVs with distinct capabilities through CNN extrapolation and Reinforcement learning

II. Methodology

This approach can be divided into two major alternating components: 1) RL algorithm for exploration and sensing, 2) extrapolation algorithm for creating a health map from sensed data. Each UAV will continually cycle between selecting their next location based on the estimated health map and updating the estimated health map by extrapolating from sensed data. Together all agents will pool their results into a combined extrapolated health map.

A. Reinforcement Learning Algorithm

We use a modified version of Q-learning and a MARbLE architecture to select the path to be sensed through multi-agent reinforcement learning [21]. As shown in figure 1 the UAVs explore during the initial phase and then try to maximize the utility of visiting a particular management zone. The states are the x and y coordinates of the management zones and the utility or reward of each action is the error between the predicted values from the CNN and the observed values. Ultimately, the Q-table is updated with the reward from the combined goal and budget preferences from the MARbLE algorithm. These rewards are updated using Bellman's Equation, as shown below.

$$Q(s_i, a_i) = (1-\alpha) * Q(s_i, a_i) + \alpha * [R(s_i, a_i, s_{i+1}) * \gamma \max(Q(s_{i+1}, a_{i+1}))] \quad (1)$$

Equation 1 calculates the maximum reward with learning rate α and a discount factor γ taking into account the immediate and long-term rewards.

To create a generalized and transferable model we use filters while populating the Q-table. The observed rewards are quantified based on their variance such that the observed pattern can be transferred to different fields as well.

B. Extrapolation

While the UAV explores the field, the health map is continuously extrapolated using the newly gathered data at each step. This extrapolation is crucial, as it provides an accurate foundation for decision-making within the RL algorithm. The RL algorithm aims to maximize the percentage error gain between predicted and ground truth values, thereby systematically reducing the error associated with the projected health map. To extrapolate the health maps, we employ Convolutional Neural Networks (CNN). This extrapolation is built on the premise that distinct sensor readings correspond to various aspects of plant health. This concept is akin to human health diagnostics, where different tests reveal diverse health issues that may also be interconnected. For instance, in humans the identification of calcium deficiency through the Total Calcium Test (TCT) can potentially indicate the presence of osteoporosis [22].

We quantify comprehensive crop health by combining data from different health indicators such as NDVI, NDWI, and GLI. This combination allows us to present an overarching picture of the crop field's health. Moreover, the extrapolated individual health maps associated with these indicators offer more intricate insights, aiding in the identification of precise measures necessary to enhance the crop's well-being. The quantification of overall health based on different health indicators (NDVI, NDWI, GLI). Furthermore, the overall health map complements the decision making in multi-agent reinforcement learning.

The design of the CNN is based on U-net Architecture [23]. The input to the CNN is observed health maps, and the output is fully predicted health maps.

III. Preliminary Results

Our previous work using RL for crop scouting and CNN for extrapolation with single agent employing RGB camera provide promising results [24]. A health map is predicted with 90% accuracy by scouting only 40% field and hence reducing labor costs 4.8 times and boosting profit by 36%.

REFERENCES

[1] H. Charles J. Godfray, John R. Beddington, Ian R. Crute, Lawrence Haddad, David Lawrence, James F. Muir, Jules Pretty, Sherman Robinson, Sandy M. Thomas, and Camilla Toulmin. Food security: The challenge of feeding 9 billion people. *Science*, 327(5967):812–818, 2 2010.

[2] Lee T. Hickey, Amber N. Hafeez, Hannah Robinson, Scott A. Jackson, Soraya C. M. Leal-Bertioli, Mark Tester, Caixia Gao, Ian D. Godwin, Ben J. Hayes, and Brande B. H. Wulff. Breeding crops to feed 10 billion. *Nature Biotechnology*, 37(7):744–754, 7 2019.

[3] Meetpal S. Kukal and Suat Irmak. Climate-Driven Crop Yield and Yield Variability and Climate Change Impacts on the U.S. Great Plains Agricultural Production. *Scientific Reports*, 8(1):3450, 2 2018.

[4] Gerald C. Nelson, Hugo Valin, Ronald D. Sands, Petr Havlík, Helal Ahammad, Delphine Deryng, Joshua Elliott, Shinichiro Fujimori, Tomoko Hasegawa, Edwina Heyhoe, Page Kyle, Martin Von Lampe, Hermann Lotze-Campen, Daniel Mason d'Croz, Hans van Meijl, Dominique van der Mensbrugghe, Christoph Müller, Alexander Popp, Richard Robertson, Sherman Robinson, Erwin Schmid, Christoph Schmitz, Andrzej Tabeau, and Dirk Willenbockel. Climate change effects on agriculture: Economic responses to biophysical shocks. *Proceedings of the National Academy of Sciences*, 111(9):3274–3279, 3 2014.

[5] Naiqian Zhang, Maohua Wang, and Ning Wang. Precision agriculture—a worldwide overview. *Computers and Electronics in Agriculture*, 36(2-3):113–132, 11 2002.

[6] Robin Gebbers and Viacheslav I. Adamchuk. Precision Agriculture and Food Security. *Science*, 327(5967):828–831, 2 2010.

[7] Francis J. Pierce and Peter Nowak. Aspects of Precision Agriculture. pages 1–85. 1999.

[8] Laurens Klerkx, Emma Jakku, and Pierre Labarthe. A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda. *NJAS: Wageningen Journal of Life Sciences*, 90-91(1):1–16, 12 2019.

[9] Sumanta Das, Scott Chapman, Jack Christopher, Malini Roy Choudhury, Neal W. Menzies, Armando Apan, and Yash P. Dang. UAV-thermal imaging: A technological breakthrough for monitoring and quantifying crop abiotic stress to help sustain productivity on sodic soils – A case review on wheat. *Remote Sensing Applications: Society and Environment*, 23:100583, 8 2021.

[10] Xin Zhang, Liangxiu Han, Yingying Dong, Yue Shi, Wenjiang Huang, Lianghao Han, Pablo González-Moreno, Huiqin Ma, Huichun Ye, and Tam Sobeih. A Deep Learning-Based Approach for Automated Yellow Rust Disease Detection from High-Resolution Hyperspectral UAV Images. *Remote Sensing*, 11(13):1554, 6 2019.

[11] Everton Castelão Tetila, Bruno Brandoli Machado, Gilberto Astolfi, Nícolas Alessandro de Souza Belete, Willian Paraguassu Amorim, Antonia Railda Roel, and Hemerson Pistori. Detection and classification of soybean pests using deep learning with UAV images. *Computers and Electronics in Agriculture*, 179:105836, 12 2020.

[12] Zichen Zhang, Sami Khanal, Amy Raudenbush, Kelley Tilmon, and Christopher Stewart. Assessing the efficacy of machine learning techniques to characterize soybean defoliation from unmanned aerial vehicles. *Computers and Electronics in Agriculture*, 193:106682, 2 2022.

[13] Perry J. Hardin and Ryan R. Jensen. Small-Scale Unmanned Aerial Vehicles in Environmental Remote Sensing: Challenges and Opportunities. *GIScience & Remote Sensing*, 48(1):99–111, 1 2011.

[14] Wouter H Maes and Kathy Steppe. Perspectives for Remote Sensing with Unmanned Aerial Vehicles in Precision Agriculture. 2018.

[15] Jorge Rodríguez, Iván Lizarazo, Flavio Prieto, and Victor Angulo-Morales. Assessment of potato late blight from UAV-based multispectral imagery. *Computers and Electronics in Agriculture*, 184:106061, 2021.

[16] Weiguang Yang, Weicheng Xu, Changshen Wu, Bingyu Zhu, Pengchao Chen, Lei Zhang, and Yubin Lan. Cotton hail disaster classification based on drone multispectral images at the flowering and boll stage. *Computers and Electronics in Agriculture*, 180:105866, 1 2021.

[17] Peyman Moghadam, Daniel Ward, Ethan Goan, Srimal Jayawardena, Pavan Sikka, and Emili Hernandez. Plant Disease Detection Using Hyperspectral Imaging. In *2017 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8. IEEE, 11 2017.

[18] Zheng Zhou, Yaqoob Majeed, Geraldine Diverres Naranjo, and Elena M.T. Gambacorta. Assessment for crop water stress with infrared thermal imagery in precision agriculture: A review and future prospects for deep learning applications. *Computers and Electronics in Agriculture*, 182:106019, 3 2021.

[19] Narmilan Amarasingam, Arachchige Surantha Ashan Salgadoe, Kevin Powell, Luis Felipe Gonzalez, and Sijesh Natarajan. A review of UAV platforms, sensors, and applications for monitoring of sugarcane crops. *Remote Sensing Applications: Society and Environment*, 26:100712, 4 2022.

[20] Dimosthenis C. Tsouros, Stamatia Bibi, and Panagiotis G. Sarigiannidis. A Review on UAV-Based Applications for Precision Agriculture. *Information*, 10(11):349, 11 2019.

[21] Jayson Boubin, Codi Burley, Peida Han, Bowen Li, Barry Porter, and Christopher 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.

[22] B.E. Christopher Nordin and Howard A. Morris. The Calcium Deficiency Model for Osteoporosis. *Nutrition Reviews*, 47(3):65–72, 4 2009.

[23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. pages 234–241. 2015.

[24] Zichen Zhang, Jayson Boubin, Christopher Stewart, and Sami Khanal. Whole-Field Reinforcement Learning: A Fully Autonomous Aerial Scouting Method for Precision Agriculture. *Sensors*, 20(22):6585, 11 2020.