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]

[1]Indian Institute of Technology, Bombay, [2]Ohio State University

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.

Motivation

- The burgeoning global population necessitates rapid advancements in agricultural practices to meet the increasing demand for food.
- Climate change and crop health stressor including drought, diseases, and pest infestations lead to 11% reduction in crop yields.
- Precision agriculture emerges as a savior in meeting the escalating global food demand. Strategic integration of technology helps in optimizes resource utilization and maximizes yield.

- Digital Agriculture uses data collected through sensors, UAVs to provide data analytics for crop management
- Traditional data collection methods using UAVs have higher operating costs and require frequent battery replacements due to limited flight time

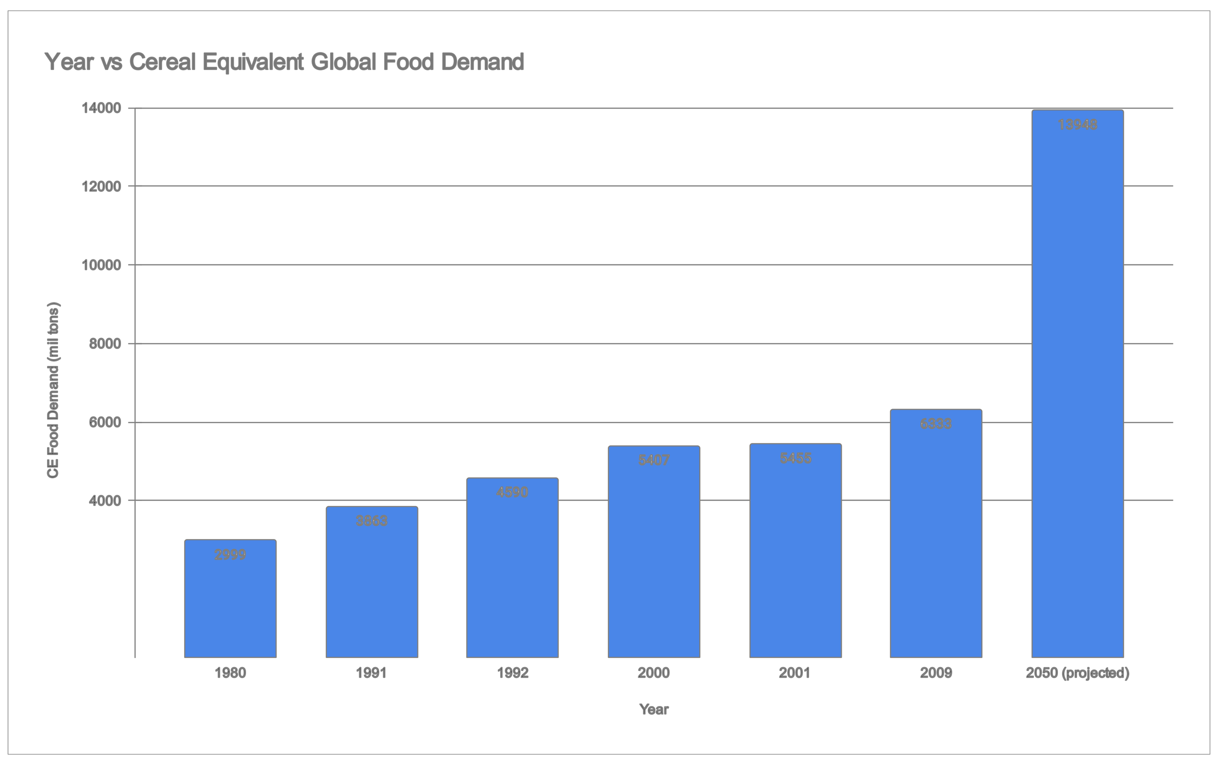

- Collecting multi-modal data for detailed crop analysis tend to escalate these costs
- We use Reinforcement learning to scout crucial areas of the field, resulting in lesser battery replacements.
- Data collected through RL is extrapolated using Convolution Neural Networks.

Reinforcement Learning

- We use a modified version of Q-learning and a MARbLE architecture to select the path to be sensed through multi-agent reinforcement learning.
- Utility of visiting a management zone is the error between predicted values and observed values.
- Q-values are obtained by combining goals and budgets from MARble Architecture. Then Q-table gets populated using bellman equation.

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

- Model Generalization through use of filters, and quantifying rewards based on variance in the field.

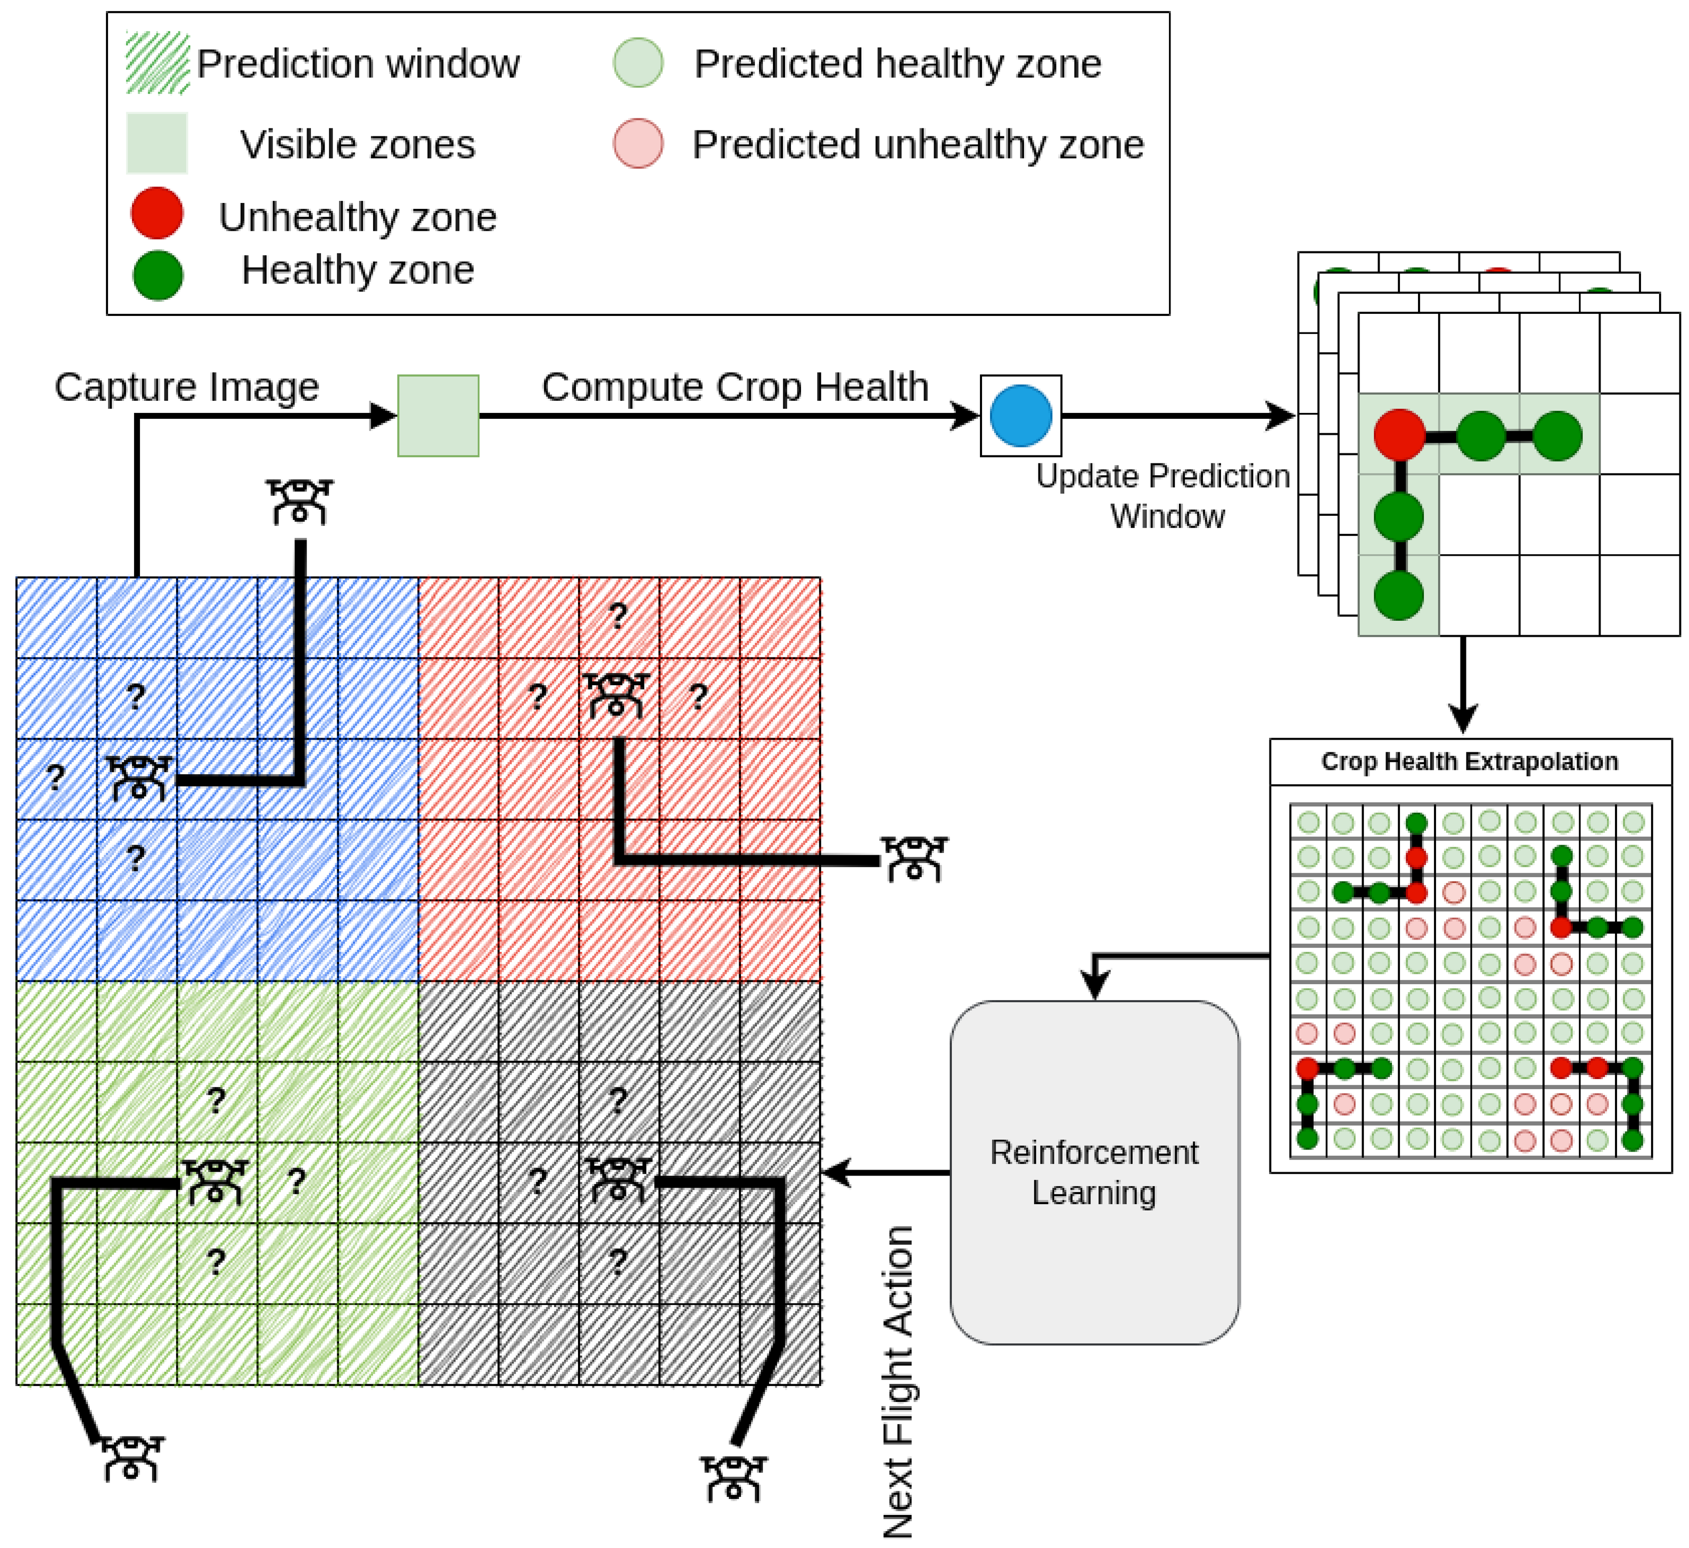

Extrapolation

- The UAV continuously updates a health map as it explores the field.
- This health map is essential for decision-making within the RL algorithm.
- The RL algorithm aims to minimize the error between predicted and actual values in the health map.
- Convolutional Neural Networks (CNN) are used for extrapolating the health maps.
- The concept is analogous to medical diagnostics, where different tests reveal interconnected health issues.
- Comprehensive crop health is quantified by combining data from various indicators like NDVI, NDWI, and GLI.
- This combination provides an overall view of the crop field's health.
- Individual health maps based on these indicators offer detailed insights.
- These insights help identify precise measures to improve crop well-being.
- The overall health map aids decision-making in multi-agent reinforcement learning.
- The CNN design is based on the U-net Architecture.
- Input to the CNN consists of observed health maps, and the output is fully predicted health maps.