Article

Fully Autonomous Aerial Scouting for Precision Agriculture

Zichen Zhang [1]*, Jayson Boubin [1], Christopher Stewart[1] and Sami Khanal[2]

[1] Department of Computer Science and Engineering, The Ohio State University, Columbus OH 43210, USA
[2] Department of Food, Agricultural and Biological Engineering, The Ohio State University, Columbus OH 43210, USA
* Correspondence: zhang.9325@osu.edu; Tel.: +1-614-302-0528 (Z.Z.)

Version October 20, 2020 submitted to Sensors

Abstract: Precision agriculture uses tools and technologies to monitor and manage in-field soil and crop variability by dividing a crop field into multiple management zones. This approach helps to implement site-specific management practices which can improve crop productivity and profitability of farming operations while minimizing the environmental impacts from agricultural fields. However, large crop fields may comprise many zones, and thus the cost of collecting data and modeling crop health for each zone may not result in higher return on technology investments. While the use of unmanned aerial systems (UAS) can facilitate timely and cost-effective collection of crop and soil health data compared to human-based scouting approach, UAS performance can be limited due to battery life. To collect exhaustive data from a large crop field, human operators are required to exchange depleted batteries many times, which can be costly and time consuming. In this study, we developed a novel fully autonomous aerial scouting approach, the whole-field based reinforcement learning algorithm, that uses reinforcement learning (RL) and convolutional neural networks to choose sample sections of a field for sensing to predict crop health for an entire crop field. This approach minimizes data collection while maximizing the accuracy for predictions of an entire field. The performance of this approach is compared with prior work focused on a local-field based RL algorithm, and conventional aerial scouting approaches in terms of accuracy, modeling cost, and potential overall cost saving. To develop and test the approach, we ran flight simulations on an aerial image data set. The aerial images were collected from an 80-acre corn field and divided into 40,320 management zones (each zone is around 4.3 square meters), and used the Excess Green Vegetation Index as proxy for crop health condition. The novel scouting approach modeled crop health with 89.8% accuracy, reduced labor cost by 4.8X and increased agricultural profits by 1.36X compared to the conventional scouting approach which exhaustively scouts the entire field with a higher redundancy data collection scheme.

Keywords: convolutional neural networks; reinforcement learning; unmanned aerial systems (UAS); autonomous aerial scouting

1. Introduction

There is currently an unprecedented demand to increase food and energy production, and a desire for sustainability. This demand is accompanied by increased water scarcity and weather variability. It is predicted that the global population will increase to 9.7 billion in 2050 [1,2], and that agricultural production must double to meet the needs of this growing population and shift in dietary preference while balancing against energy and water constraints [3,4]. This goal cannot be reached by simply doubling the agricultural inputs because of constrained resources, already developed agricultural land limits, and environmental concerns [5]. The future efficiency gains of agricultural production systems must be radically improved and adaptable to be able to increase yields with respect to large variances expected in weather locally and growing globally.

Precision agriculture (aka site-specific management practice) is a promising step towards improving efficiency and reducing adverse impacts of agricultural production [6]. It focuses on assessing variation across and within crop fields to divide a field into multiple management zones and treats each management zone accordingly [7,8]. Thus, it is critical to have spatial and temporal maps of crop and soil health on a timely fashion [9]. Accurate crop and soil health maps are critical to support site-specific management practices in a cost-effective manner. Management decisions based on inaccurate crop and soil health maps can result in unwanted crop yield loss, excessive fertilizer application, and increased nutrient loads to waterbodies [10]. For instance, let's assume that a farmer applies fertilizers to only those crops that fall within the unhealthy zones, based on crop health maps. If crop health maps inaccurately label unhealthy areas as healthy, those misrepresented unhealthy sections of a field would not receive treatment, and thus crop yield of those areas could be poor. Alternatively, if healthy zones are mislabeled as unhealthy, they would receive unwanted fertilizer application, which mean loss of farm resources as well as increased risk of nutrient runoffs and leaching without much increase in crop yield.

Accurate representation of field conditions via maps depends on temporal and spatial resolutions of data, which vary across sensors and platforms (e.g., satellite, weather stations, and aircraft) used for data collection, which in turn influence data collection and processing costs [11]. UAS have emerged as a cost-effective approach for aerial scouting [12]. Compared to satellites, UAS can fly to waypoints, hover, and collect high resolution data (millimeters per pixel) from large areas quickly with no or little risk. Compared to human piloted aircraft, UAS are 3X less expensive, achieve better spatial resolution, and pose fewer safety risks [11]. Traditional UAS-based approach for scouting of a field involves a grid mission, which captures images from multiple areas (hereinafter defined as zones) [12–14]. To be more specific, for scouting a whole field, a UAS is given a set of waypoints (i.e., GPS coordinates) to follow and it takes one picture at each waypoint. Various vegetation indices, such as Excess Green Vegetation Index (ExG) [15], are computed to indicate crop health conditions for each zone [16]. In order to provide accurate crop health information of a field, traditional exhaustive scouting approach involves redundant data collection (i.e., 65-80% front and side overlap between images), which results in massive computation costs. In the meantime, batteries on commodity UAS allow just 15–25 minutes of flight. UAS must land and recharge repeatedly to cover large fields. Human operators must monitor flights and battery capacity, swap and recharge batteries and possibly fly aircraft manually by remote control. These activities also delay missions. It can take a full 8-hours workday to exhaustively collect high definition images from every zone in an 80-acre crop field [11,17]. Thus, for UAS with onboard IoT systems, it is crucial to collect as much information as possible within a time frame. Autonomous systems sense and potentially alter their environment without human intervention. Instead, they manage IoT actuators (e.g., UAS flight controls) to achieve high utility (e.g., low prediction error). Fully or partially autonomous tractors, planters and monitoring equipment already perform complex tasks in critical settings today. While autonomy can reduce labor costs, standardize and improve tasks, it also loses robust human problem-solving ability, incurs engineering costs and makes it hard to model compute needs (closed-loop systems). Lin et al. relied on narrowly defined tasks to trace and model compute needs for autonomous cars [18]. Boubin et al. broadened Lin's compute modeling by capturing environmental factors for UAS [17]. In-situ AI [19] and Boroujerdian et al. [13] generalized these approaches via environmental simulation.

In this project, we show that with the help of RL [20] and spatial ensembles of convolutional neural network (CNN) [21], UAS can get accurate crop health maps and reduce flight time and costs by scouting only a fraction of a field. Specifically, we design a whole-field based fully autonomous aerial scouting system for UAS, an alternative to exhaustive scouting where the UAS (1) are piloted by software we designed, (2) can generate an accurate crop health map with only partial coverage of the field and (3) can autonomously set their flight paths to maximize the accuracy of the crop health map. The latter feature distinguishes our approach from random sampling. In this work, we attempt to answer the following questions: *(1) can UAS autonomously select and fly over 20% to 40% of a field's management zones and create accurate crop health maps? (2) Can precision agriculture translate such cost savings in data collection as profit?*

The rest of the paper is arranged as follows: Section 2 states the design steps of the whole-field RL approach including implemented CNN and RL algorithm as well as the experiment environment. Section 3 provides the results of whole-field RL and compares it with results of local-field RL and some traditional scouting methods. Section 4 discusses the limitations and future work followed by a conclusion of the research in section 5.

2. Methods

2.1. Design

To reduce data collection and computation costs, we present a new RL approach, whole-field RL, to guide UAS in aerial crop scouting, and compare the differences in crop health maps generated by these approaches with traditional methods as well as our previous work, local-field RL [11]. Since the traditional approach involves exhaustive scouting of a field, it is assumed that the crop health information based on this approach is 100% accurate, and thus used as ground truth data to evaluate findings based on two RL approaches (Figure 1).

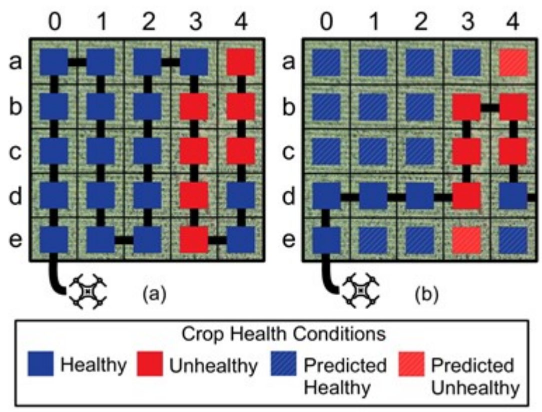

Figure 1. (a) Exhaustive scouting of a field wherein UAS visits all zones in a grid fashion and crop conditions are classified as healthy and unhealthy, and (b) RL based fully autonomous aerial scouting wherein UAS visits over a fraction of a field (e.g., 8 areas) and predict crop conditions for unvisited areas.

Whole-field RL uses a full history of images captured by a UAS during a scouting mission and implements complex CNN models and a RL algorithm to extrapolate a whole-field crop health map from sensed data. That is, during a flight mission, whole-field RL method uses all images (one image per zone) of previous UAS-visited zones as inputs to CNN models to construct crop health prediction maps which serve as inputs to the RL algorithm to decide the next zone to fly over. It is, however, a computationally intensive approach (discussed in detail in section 2.2). In contrast, local-field RL only uses one image from the last UAS-visited zone to predict its next path. After sampling enough points, whole-field RL then extrapolates crop health information for a whole field. Compared to the traditional approach, which uses a predefined path for exhaustive data collection, both RL based scouting approaches visit less areas within a field and thus reduce data collection costs. For example, as shown in Table 1 and Figure 1, since UAS batteries drain at the same rate, exhaustive scouting incurs the cost of landing, recharging and flying back to its most recent zone. Exhaustive scouting takes two flights to scout all 25 zones while fully autonomous aerial scouting only takes one flight. However, these approaches can introduce error as health conditions of the zones that are not directly observed are predicted. Thus, our work also focused on minimizing prediction errors as well as finding a balance between prediction error and UAS coverage rate of the field.

Table 1. The benefits of the fully autonomous aerial scouting captured by empirical traces of battery drain.

Exhaustive Scouting	Step	0	1	5	9	13	17	21	25
	Battery %	100	95	75	55	35	85	65	45
	Mission %	0	4	20	36	52	68	84	100
	Current Zone		[e,0]	[a,0]	[d,1]	[c,2]	[b,3]	[e,4]	[a,4]
Fully Autonomous Aerial Scouting	Step	0	1	5	9	10			
	Battery %	100	95	75	55	50			
	Mission %	0	10	50	90	100			
	Current Zone		[e,0]	[d,3]	[c,4]	[d,4]			

Figure 2 outlines fully autonomous aerial scouting approach. First, UAS fly over management zones and collect images. Crop health is computed for each visited zone. Prior observations of crop health data for all visited zones, associated flight actions and their outcomes are stored as training data. Then, this RL algorithm computes the next flight action, wherein prior action and observation pairs predict future utility. That is, given extant crop health data, this RL algorithm computes mean utility of similar prior observations and chooses the best action. Here, utility gained after taking an action is defined as the improvement in crop health map accuracy. For autonomous aerial scouting, the utility function seeks to maximize accuracy of the final extrapolated crop health map.

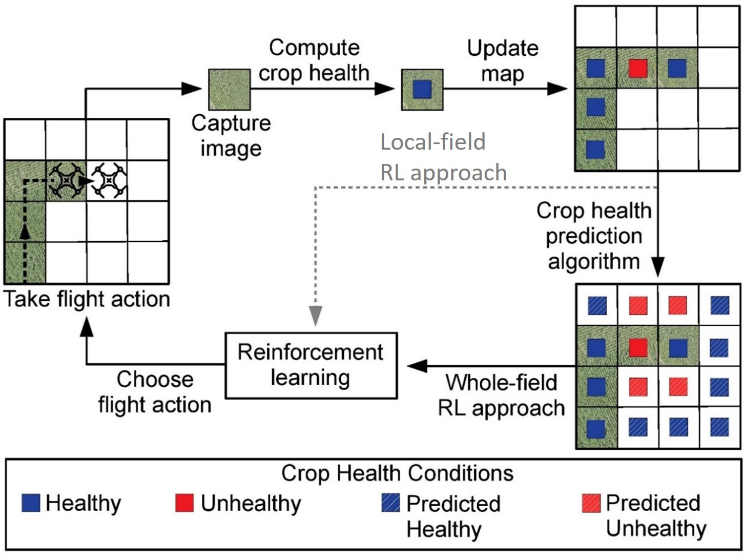

Figure 2. Fully autonomous aerial crop scouting uses RL to decide flight actions and covers only a fraction of the field.

2.2. The Whole-Field RL Algorithm

The whole-field RL approach has three components - 1) a CNN to model crop health, 2) an algorithm to extrapolate crop health predictions over a whole field, and 3) a RL approach to improve future outputs. As an overview, CNN used spatial information of ExG to improve crop health prediction accuracy in areas scouted by UAS. The algorithm then expands predictions beyond spatial neighbors. Then, RL chooses flight paths that wisely sample zones to maximize the accuracy of predicted crop health maps (details provided later in this section). Finally, when desired coverage is reached, CNN-based crop health predictions are used to extrapolate all data sensed by the UAS to create a whole-field crop health map. It was assumed that UAS have access to edge computing systems powerful enough for RL and CNN inference. Edge servers or laptops could sufficiently augment compute available on UAS. Wireless networks could allow data transfer between UAS and compute devices, but this is outside the scope of this paper.

2.2.1. Convolutional Neural Network to Model Crop Health

Based on the first law of geography, things that are closer together tend to be more related than things that are far apart, and this is often evident while monitoring agricultural fields [16,22]. For instance, the root causes of poor crop health, such as diseases and pests, often spread to nearby crops. We leveraged this property to extrapolate crop health given nearby ground truth. This can be accomplished by providing surrounding zones as the input to a CNN that predicts crop health for a targeted zone.

We chose and modified vgg16 [21] neural network as our CNN model to predict crop health conditions. Our design trained CNN models, one for each of the eight neighbors adjacent to the center management zone in a 3x3 grid. Crop health is computed directly in zones visited by the UAS by using vegetation indices [16]. Given the location of a nearby observation, the corresponding one of these CNN models predicts crop health given the observed image of that zone. These models are designed to leverage the spatial crop health distribution property to predict the condition of the management zones not visited by the UAS. This approach makes a few key assumptions. First, each image captured by UAS represents one management zone. Second, since the UAS captures images in flight, management zones must be connected in the field. This assumption can be problematic in urban settings but reflects common practice in rural environments. Also, as a corollary, the UAS visit a connected subset of all the management zones in the whole field, which we define as the visible (or observed) area.

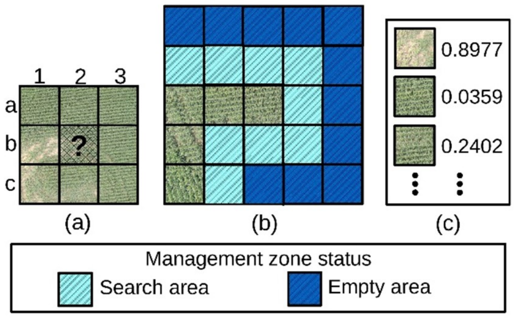

Figure 3. Crop health map prediction using CNN.

As an illustration, eight observed surrounding management zones represent eight positions regarding the unobserved middle one, which serves as inputs to the eight spatial CNN models (Figure 3a). Models were trained by using the feature of the observed surrounding management zone with the label (health condition) of the unobserved middle one. In this case, each of the eight images on the surrounding has a corresponding CNN model that can be used to predict the crop health at position [b,2]. Thus, by using CNN, crop health for all the zones adjacent to the visible area are modeled. These zones are called the search area (as shown in green in Figure 3b). We defined a prediction window as a UAS-centered square area in a field (Figure 3b, a 5x5 prediction window). The prediction window is conducted with three parts of the area. The rest of the prediction window forms the empty area. Management zones in the search area may have multiple adjacent visible zones. In this event, predictions from multiple adjacent CNN models will be used to improve accuracy, an approach akin to ensemble models (illustrated in Figure 3c for zone [b,2]). The ensemble leads to an accurate prediction because one CNN model predicted poor health (0.89), where a label of 1.0 represents an unhealthy zone, and the others predict good health (0.03 and 0.24). The average of these values reflects the crop health for [b,2].

2.2.2. Crop Health Prediction Algorithm

A two-stage crop health prediction algorithm was designed to extrapolate crop health information beyond the search area to reach out to the empty area and create crop health maps on a whole-field scale (Figure 4). For every flight step, the algorithm predicts the health conditions of all management zones within the prediction window given ground truth data sensed along the UAS flight path in it. Only zones within this region are predicted in lieu of all zones in the field, which saves on compute time and improves accuracy. This map serves as the input for RL to make the best decision for the next flight direction. After the mission completes, the crop health prediction algorithm generates a crop health map for the whole crop field.

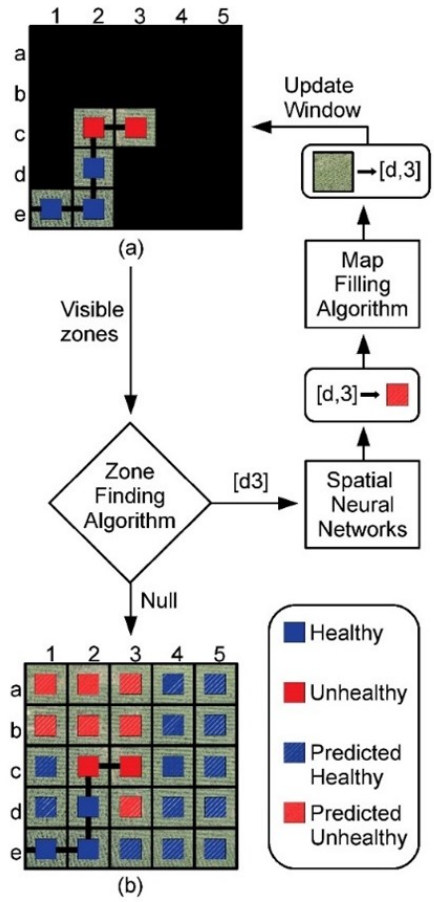

Figure 4. Crop health map prediction using CNN.

First, the algorithm initiates a prediction window with a chosen size (Figure 4a). The initial window only contains the sensed data of all zones on the flight path (visible area). Next, the CNN models are used to predict the health conditions of zones in the search area. To improve prediction accuracy, zones in search area are prioritized based on the number of visible zones they have in their neighborhood. After the prediction for a given zone are obtained, the crop health map will be updated by not only adding the health prediction, but also adding a replacement image at that position to expand the visible area. That visible image is found using a reference data set based on the prediction of that zone from CNN models.

In this study, 2% of all UAS images were used to build a reference data set. The reference data set leverages the feature that the texture of a crop field is similar throughout, which means you can potentially find two zones that are quite similar given enough samples. Every zone image in this data set was already associated with a health prediction from the CNN, which saved time by not needing to load CNN models to make predictions during a flight mission. After we get the prediction of a zone in the search area, we compare it to all the predictions in the reference data set. The reference image with the closest prediction is chosen and placed in that location in the map. In this way the process keeps filling the crop health map with visible images which makes it possible to predict all zones that are not originally adjacent to the visible area. After we update the crop health map with the prediction and the visible image for that management zone, the visible area and the search area are updated as well. The visible area expands as reference data set images replace blank spaces in the prediction window. The search area then removes resolved management zones and adds their neighbors from the area that is not adjacent to visible area. After this round of updates completes, the process continues to find the management zone in the search area with the highest priority. The algorithm completes when the prediction window contains no blank management zones.

2.2.3. Reinforcement Learning Algorithm

Once crop health conditions are evaluated, RL algorithms, modified Q-learning, are used to determine the next UAS movement direction. UAS are designed to maintain one map for every flight step while exploring a crop field. Each map is one state of the field that uses two kinds of information, collected crop health conditions (aerial images on the flight path) as the ground truth data and predicted crop health generated by the crop health prediction algorithm. This data is used to develop a state model which provides a series of possible flight directions from the current zone. The final flight decision will be made by the state model using a list of prior observations of similar field states. In order to build a dataset with plenty of field states, we developed an algorithm to randomly simulate flight path and collect map data (detailed in section 2.4.3).

The model was trained with a data set of 73,000 unique field state combinations. The K-Nearest Neighbors (KNN) [23] algorithm was used to determine which prior examples are most relevant for a given combination of ground truth and prediction maps. KNN determines the 11 most similar prior states to the current state. To determine the final movement direction, RL compares predicted labels for each flight action from KNN to CNN predictions. The flight action with the largest distance between its KNN and CNN prediction (i.e., highest error) is chosen to ensure that UAS explore the locations that they least understand.

2.3. Local-Field RL Algorithm

Local-Field RL is similar in many ways to whole-field RL. Both algorithms use RL to navigate the field and generate a final crop health map. There are, however, a series of important differences between them. Unlike whole-field RL which uses a prediction map generated from the crop health prediction algorithm as an input to RL, local-field RL extracts image properties (i.e., ExG, RGB saturation) from visible zones and feeds such data as an input to RL. Once the UAS has covered a certain amount of the field, local-field RL extrapolates the crop health map of a field using a KNN-based recursive dilation procedure instead of relying on a CNN. The dilation procedure finds every management zone in the map that has not been predicted or observed and assigns that zone the consensus of its directly adjacent neighbors, if it has any. If it has no neighbors, the position remains unassigned. This process is performed recursively until the entire map is full.

2.4. Implementing Autonomous Aerial Scouting

Autonomous scouting algorithms were implemented using the SoftwarePilot simulation environment [24], which has been used in prior work to implement and simulate local-field RL simulation [11,17]. In this study, SoftwarePilot was modified to use CNN-based model outputs for crop health map prediction as well as for RL-based path finding algorithms. This simulator performs both the local-field and whole-field RL algorithms until a user-specified amount of the crop field is explored, then uses the extrapolation algorithms to predict any areas of the final map that are still empty.

2.4.1. Data Set

CNN training and flight simulation were performed on a data set collected in the Molly Caren Research farm near London, Ohio in August 2017. This data set includes UAS-collected images of corn fields covering 950,000 individual management zones in 684 aerial images. Images were collected at 200 feet from the ground using an eBee UAS from senseFly with a ground sample distance of 1.9 cm/px. For this paper, we used 30 out of 684 aerial images conducting of 40,320 zones as our training data set.

2.4.2. Whole-Field RL Implementation

CNN-based crop health modeling is used in two parts of the experiment. First, we need to use the crop health prediction models to build data sets for the RL algorithms. Second, during simulation, the crop health prediction models provide near real-time predicted crop health map for each management zone the UAS captures and for the whole field based on the final flight path. In this subsection we mainly discuss RL data set construction. The extrapolation procedure is discussed in the following subsection.

In order to build a large RL data set and perform thorough analysis, data sets were prepared using 6 coverage rates - 10%, 20%, 30%, 40%, 50%, and 60%. For each coverage rate, five whole-field RL prediction window sizes - 7x7, 11x11, 15x15, 19x19, 23x23 were used. For each combination, 1000 flight paths were randomly generated. The process begins with choosing a random start point for the UAS on the edge of the field. Each subsequent flight step is chosen randomly from the neighboring management zones of the current position. The simulated flight path keeps growing until it meets the specified coverage rate. If the UAS has sampled all its directly adjacent neighbors, but has not reached the coverage threshold, it flies to the nearest unsampled management zone. For each zone image the UAS collects, the crop health map for the prediction window is calculated using the crop health prediction algorithm. Crop health is estimated based on ExG [16], a vegetation index derived using visible aerial image; this data is also used as ground truth. ExG of each management zone is compared with the average ExG of the entire field. If ExG is less than 80% of average field ExG, the management zones are classified as unhealthy, and if ExG is at least 80% of average ExG, zones are defined as healthy.

2.4.3. Simulation Environment

Our simulation environment was modified from the SoftwarePilot, the local-field RL simulator. To test our whole-field RL approach, we used 30 images from all collected aerial images to crop into individual management zones. Each image is broken into a set of 1344 management zones in a 42x32 management zone grid for a total of 40,320 management zones. For the purpose of simulation, we consider one of these images to represent the flight area of the fully autonomous aerial scouting system, and for each management zone to represent sensed data from the simulated UAS. The ground truth ExG of each zone is calculated and provided to the simulated UAS's visible set and used to calculate the accuracy of the resultant crop health maps.

2.5. Non-Autonomous Scouting Approach

Other than the exhaustive aerial scouting approach as shown in Figure 1a, we also compared our fully autonomous aerial scouting with two naive approaches: random scouting and non-scouting. Random scouting entails UAS randomly choosing flight directions until the provided coverage rate is reached. Non-scouting refers to the naive approach of applying fertilizers uniformly without considering internal field variability.

2.6. Comparison Between Scouting Approach

We compared whole-field RL scouting approach with local-field RL approach and some traditional methods (the exhaustive, random walk, and non-scouting approaches) based on metrics such as accuracy, positive precision, positive recall, negative precision, and negative recall (discussed below) relative to the ground truth health determined using ExG. Positive and negative indicate healthy and unhealthy crop conditions, respectively. Positive recall represents the ratio between all correctly classified true positives, and all true positives (both true positives and false negatives), where a higher ratio represents efficient avoidance of false negatives. Negative precision similarly represents the ratio of correctly classified true negatives to all true negatives (classified true negatives and classified false positives), where a high negative precision represents an efficient avoidance of false positives. While false positives refer to unhealthy management zones that are misclassified as healthy, false negatives refer to healthy management zones that are misclassified as unhealthy. True positives and negatives indicate management zones that are correctly classified as healthy and unhealthy management zones. Management decisions based on false positives can result in untreated crops, which may result in low crop yield. Similarly, if unhealthy management zones are predicted as healthy (i.e., false negative), they could lead to excessive use of resources such as fertilizer thereby increasing the likelihood of higher nutrient load to air or water.

For the fully autonomous scouting approaches, we used two coverage rates, 20% and 40%. And for each experiment, we executed both algorithms across the 40,320 management zone crop data set. These performance metrics were calculated using the SoftwarePilot energy models for the DJI Mavic Pro which has a 3830mAh, 11.4v battery. We profiled the execution time of local-field and whole-field RL using a Lenovo ThinkPad T470 as the edge system. This system has an i7-7500u processor, and 24GB of RAM, and runs Ubuntu 18.04.

2.6.1. Energy and labor costs estimation

In order to compare the performance of all approaches from energy and cost perspectives, a simple cost-benefit model was developed by adding up revenue based on crop yield from all the management zones and subtracting the cost of treating misclassified zones (healthy classified as unhealthy and vice versa) as well as UAS deployment costs. The labor costs were considered to be $10 and $20/hour for unskilled and skilled workers, respectively [25]. It was also assumed that autonomous scouting approaches require only one unskilled worker to complete the entire survey, whereas exhaustive scouting requires an additional skilled worker (i.e., two in total) to plan and complete UAS surveys including setting up the system, planning the routes, and swapping UAS batteries. In non-scouting approach, it was assumed that farmers classify every zone as unhealthy and thus treat the field equally.

2.6.2. Nutrient Runoff Risk

We estimated potential risk of nutrient runoffs under various scouting approaches with two assumptions - 1) farmers tend to apply fertilizer uniformly throughout a field if they don't have site specific information from scouting (i.e., non-scouting), and thus the nutrient runoff risk of a field is 100%, and 2) if they have site specific information (i.e., various types of scouting), they apply treatments only to poor (i.e., unhealthy) sections of a field, which reduces the nutrient runoff risk. Thus, the potential of autonomous scouting approaches to reduce nutrient runoff risks is dependent directly on the false negative rates from the classification as unnecessary nutrient runoffs can occur when healthy zones are unnecessarily fertilized. To determine how these two autonomous scouting approaches help minimize nutrient runoff risk, we estimated the percentage of healthy zones that are unnecessarily fertilized.

3. Results and Discussion

In this study, the accuracy, scouting cost, revenue, and energy consumption of the proposed fully autonomous scouting techniques were assessed and compared with state of the practice automated scouting and non-scouting approaches used in both precision agriculture and general agriculture.

3.1. Comparing fully autonomous aerial scouting and conventional methods

Accuracy differences between whole-field and local-field RL at both 20% and 40% coverage settings showed that whole-field RL at 20% coverage rate provides 2.3% better accuracy than local-field RL at 40% coverage rate (Figure 5). Local-field RL provided 74.5% and 80.3% accuracy at 20% and 40% coverage, respectively. This is compared to 82.6% and 87.3% accuracy at 20% and 40% coverage, respectively, for whole-field RL.

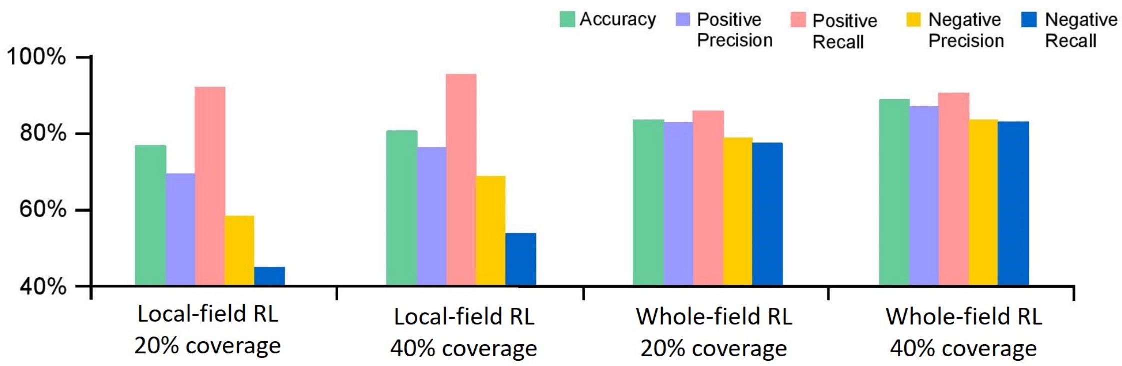

Figure 5. Accuracy of maps generated by autonomous scouting at different coverage rates.

Local-field RL outperformed whole-field RL considerably at avoiding false negatives. It experienced 8.3% and 8% higher positive recall than that of whole-field RL at 20% and 40% coverage, respectively. However, there was 67% increase in negative recall for whole-field RL over local-field RL at 20% coverage, and 58% at 40% coverage.

Table 2. Accuracy, precision, and recall for maps generated using whole-field and local-field RL at different coverage rates.

Coverage Rate		10%	20%	30%	40%	50%	60%
Local Field RL	Accuracy	0.73	0.75	0.77	0.80	0.84	0.88
	Positive Precision	0.69	0.69	0.72	0.75	0.79	0.83
	Positive Recall	0.89	0.91	0.93	0.94	0.95	0.97
	Negative Precision	0.49	0.57	0.64	0.72	0.78	0.85
	Negative Recall	0.48	0.46	0.46	0.51	0.59	0.61
Whole Field RL	Accuracy	0.70	0.83	0.85	0.87	0.89	0.90
	Positive Precision	0.71	0.83	0.85	0.87	0.89	0.89
	Positive Recall	0.71	0.84	0.87	0.89	0.93	0.94
	Negative Precision	0.64	0.78	0.81	0.82	0.85	0.87
	Negative Recall	0.65	0.77	0.80	0.81	0.83	0.83

When comparing two autonomous scouting methods for coverage rates of 10% to 60%, whole-field RL outperformed local-field RL considerably between 20% and 50% coverage while local-field RL outperformed whole-field RL at 10% coverage with a higher positive precision rate and less false positives (Table 2). At 10% coverage, local-field RL classified nearly the entire field as healthy compared to whole-field RL. At 60% coverage, there was a small difference in overall accuracy between whole-field and local-field RL; however, negative recall was significantly higher in whole-field RL. While local-field RL experienced consistent accuracy gains as coverage improves, whole-field RL experienced gains at lower coverage, with a considerable drop-off at 50%; its performance is restricted by our CNNs, whose average accuracy is around 90%. This suggests that the marginal benefit of increasing whole-field RL coverage past 60% is likely not worth the marginal cost of labor and equipment. For over 60% coverage rate, the extra cost is more than the extra revenue comparing to a 40% coverage rate.

The main difference between the two autonomous scouting approaches is how they consider surrounding zones for prediction of crop health and pathfinding. Whole-field RL uses a CNN-based prediction window for path finding, in which a few poor management zones can reinforce the scouting of UAS for poor management zones in their surroundings. Local-field RL simply uses KNN-based linear approach, which predicts health condition of a management zone based on its adjacent neighbors, and thus, can achieve high positive recall. While it's important to achieve high positive recall, negative recall can have significant cost implications when using this data for implementing site-specific management practices. Treating an unhealthy zone as healthy (negative recall) is estimated to cost 8 times higher than treating a healthy zone as unhealthy (positive recall) (discussed in section 2.6).

3.2. Autonomous Pathfinding and Extrapolation Comparison

Both autonomous scouting approaches appear to alternate between two natural behaviors, exploration and scouting. Exploration involves the UAS traversing the field, covering large swaths in search of a region that contradicts current map conditions. Scouting involves the UAS moving in an exhaustive fashion across a region that the RL algorithm perceives as important. Whole-field RL was found to take better advantage of these two pathfinding behaviors by quickly finding areas that it perceives to be problematic, and more thoroughly scouting those areas. This contrasts with local-field RL approach where the first cluster is traversed but barely explored and a large chunk of the second cluster is ignored (Figure 6a). These discrepancies are likely due to the quality of the inputs provided to the RL algorithm in each approach. Whole-field RL provides its entire prediction and ground truth windows which more accurately locate relevant prior examples in the data set collected beforehand than the local features used in local-field RL.

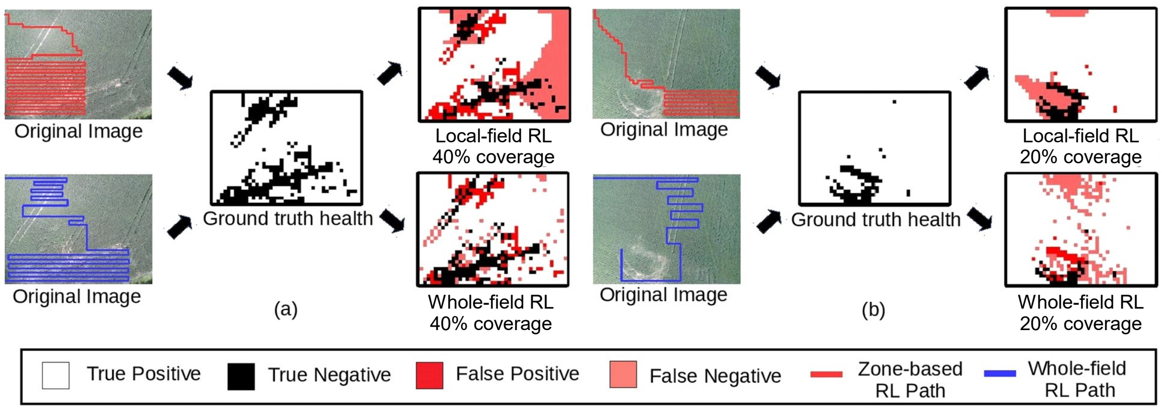

Figure 6. Whole-field and local-field RL beget different paths. [Note: A subset of UAS imagery was used to better illustrate the difference in flight paths taken by local-field and whole-field RL approaches along with resultant crop health maps at 20% and 40% coverage. The final predicted crop health maps were evaluated in four colors, representing true positive, true negative, false positive, and false negative, respectively.]

However, local-field RL outperformed whole-field RL in some cases at a low coverage rate (Figure 6b). Distracted by traces of tractors, whole-field RL spends a considerable amount of time scouting a questionable narrow area in the top left of the field, while local-field RL rushes to the bottom to explore a region that is partially negative or unhealthy. While whole-field RL eventually finds the bounds of the large negative cluster, the quick decision by local-field RL that leads to finding the bad cluster earlier leads to improved accuracy at this low coverage setting.

Differences between the local-field and whole-field RL extrapolation methods are also apparent when examining the four output maps. Both local-field RL maps show significant clusters of false negatives (pink areas). This is due to the underlying KNN-based extrapolation algorithm. When local-field RL's extrapolation algorithm encounters a cluster of similarly classified points, it tends to reinforce that classification across nearby unpredicted zones. One cluster of negative zones would be easily extrapolated such that a huge area of a field could be falsely predicted as negative (details shown in predicted crop health map by local-field RL in Figure 6a). This behavior is not as apparent in whole-field RL which uses online predictions to fill zones instead of binary extrapolation.

3.3. Autonomous Scouting on Fields with Various Crop Health Conditions

The performance of whole-field RL was explored on two very different regions, one which is primarily comprised of healthy zones, and another primarily comprised of unhealthy zones (Figure 7). In the first image, crop density was lower in portions of the field where areas appeared to be compacted by agricultural machinery. Other portions appeared to be low elevation areas, where crops emergence was impacted by prolonged saturation of water (i.e., ponding). In contrast, the second image showed a healthy corn field where lush green corn rows can be seen, separated only slightly by intermittent gaps.

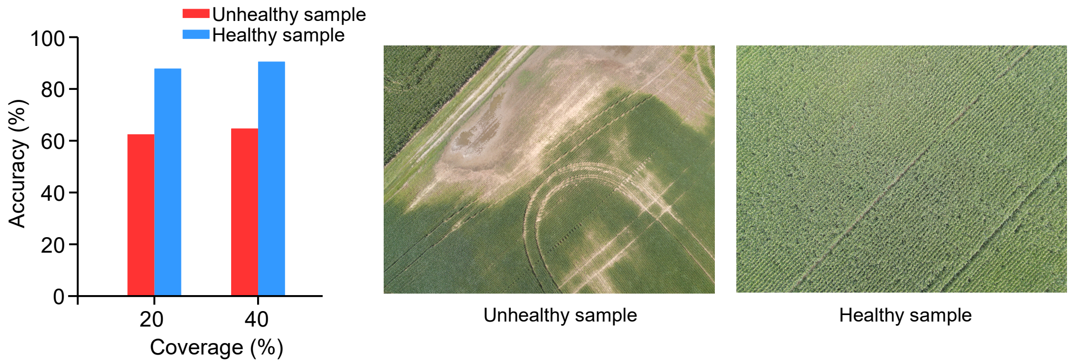

Figure 7. Accuracy of whole-field RL at 20% and 40% coverage rates for largely healthy and largely unhealthy samples.

Whole-field RL achieved over 90% accuracy for both coverage rates in the healthy sample. The accuracy increased over the average accuracy of the total data set is due largely to the uniformity of the healthy image, which indicates that whole-field RL can accurately extrapolate from a set of entirely healthy management zones. However, the size of the reference set impacts accuracy. Sometimes unknown healthy images can be replaced with unhealthy images due to their similarity in comparison to the rest of the reference set, which keeps whole-field RL from achieving 100% accuracy on sample healthy image at both coverage rates. Doubling coverage only improved accuracy by 3.2% for the healthy image. In contrast, doubling coverage for the unhealthy image improved accuracy by 10.3% (Figure 7).

The uniformity of the healthy image limits the accuracy gains from increasing coverage, which simply decreases the prevalence of false negatives generated erroneously by the reference set. Increased coverage in the non-uniform unhealthy image allows the system to get a better understanding of the topology of negative clusters to improve extrapolation. The importance of high coverage is apparent for predominantly unhealthy fields, but the low net accuracy must also be explained. Whole-field RL is largely negative-biased as discussed previously. When confronted with a majority negative image, the extrapolation procedure will reinforce negative regions, resulting in a larger number of false negatives than we see in, for instance, the health image in Figure 7. While whole-field RL experiences decreased overall accuracy on predominantly negative fields due to a high false negative rate, the decreased cost of treating false negatives as compared to false positives discussed earlier in this section implies only modest losses from treatment costs as compared to crop loss from such incorrectly predicted zones.

3.4. Effects of Prediction Window Size

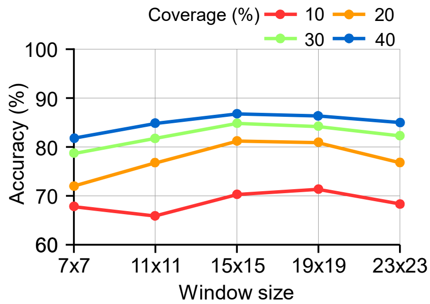

(a)

	Pathfinding	Inference	Extrapolation
Whole-field RL	1.75	0.6	81.77
Local-field RL	0.5	N/A	0.25

Component execution time (s)

(b)

Figure 8. (a) the effects of prediction window size on final crop health map accuracy at different coverage rates for whole-field RL, (b) the execution times of software components for local-field RL and whole-field RL with a window size of 15. [Note: Crop health map was generated offline.]

While comparing different combinations of prediction window sizes and coverage rates on accuracy, it was found that increasing prediction window size is not always beneficial (Figure 8a). For 20% to 40% coverage rate, accuracy is highest for the 15x15 window size, with lower accuracy proportional to both increases and decreases of the window size. 10% coverage rate had the highest accuracy at a window size of 19. Accuracy decreases at smaller window sizes can be attributed to a lack of iteratively updated prediction information with which to generate a final map. As the UAS moves around the field, it iteratively updates unseen but nearby areas to the flight path. If this window does not extend out far enough from the flight path, the only update that some areas will receive is the final extrapolation. At smaller window sizes, it is clear that some areas could have benefited from iterative updates, which would in turn increase accuracy. The opposite can be said for accuracy decreases with increased window size. If the prediction window is too large, the CNN approach may not be able to accurately predict their health due to their distance from ground truth. This result is critical to performance. Given the quadratic increase in latency as the prediction window increases, it is imperative that a whole-field RL system balances accuracy against increased costs due to latency. Given that we have found accuracy's inflection point as a function of window size, a simple solution could be to use the most accurate window size, which we have done for these experiments.

It is worth noting that as prediction window size increases, so does the size of the RL data set required for pathfinding. Both the increased number of predictions and larger pathfinding overhead increase system latency, so larger windows should be avoided to increase throughput unless accuracy returns justify them. The process to compute the final crop health map at the edge system offline by extrapolation after a mission is complete took on average 320X longer for whole-field RL than for local-field RL (Figure 7b). Whole-field RL's CNN models required considerably more time due to computational complexities than local-field RL's KNN-based approach. This process is, however, performed offline. Latencies simply determine how long the farmer must wait after mission execution to receive a crop health map. While whole-field RL experiences much higher latency in crop health map generation, both approaches return a map to the farmer in reasonable time.

3.5. Energy, Labor Costs and Nutrient Runoff Risk

The amount of charges required to map a hectare of crops differs between three scouting approaches due to the percent of field area covered in each approach (Figure 9a). Since it was assumed that the same amount of charges will be required to cover same size of a field across all scouting approaches, local-field and whole-field RL experienced the same charges at the same coverage setting, requiring 6 charges at 20% coverage rate and 12 charges at 40% coverage rate. This compares to 29 charges to map one hectare of land using the traditional exhaustive scouting method which was found to require considerably more labor and charges to complete the scouting mission.

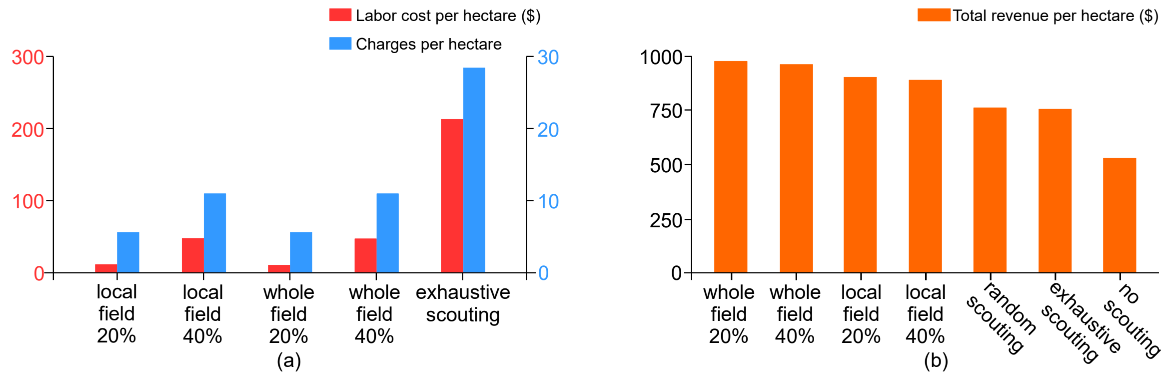

Figure 9. (a) Energy implementations and labor costs of autonomous scouting vs. exhaustive scouting, (b) the impacts that autonomous scouting have on revenue compared to state of the practice methods.

There were significant differences in labor costs for different scouting approaches (Figure 9a). Considering the economic data from the 2018 growing season [25], at hourly rates of $10 and $20 for unskilled and skilled workers respectively, autonomous scouting methods were roughly estimated to cost $29 and $44 for 20% and 40% coverage of an 80-acre crop field, which compares to the $212 mapping cost for exhaustive mapping using two labors.

According to recently published agricultural cost data [25], the revenue per acre for corn is $763.8 USD ($3.8/bushel * 200 bushels/acre). Revenue per management zone is calculated to be $0.8 for the size of 4.3 square meters per zone. The cost of fertilizer per acre is $130, which makes it $0.1 per management zone. Thus, one false negative management zone would cost a farmer $0.8 due to crop loss, while one false positive management zone would cost $0.1 in treatment cost. Based on this, when the field is not scouted, it is estimated to provide 36% less revenue than whole-field RL, and 27% less revenue than local-field RL (Figure 9b). Exhaustive scouting outperformed no scouting method by 20% but loses out to local-field and whole-field RL by 5% and 17% respectively. While exhaustive scouting will provide 100% accuracy, allowing farmers to properly treat their entire field, the labor costs of exhaustively mapping large fields is outweighed by lower coverage autonomous mapping with extrapolation. We also explored the effects of a random sampling approach at 40% coverage using local-field RL. This automated approach without RL pathfinding underperforms compared to exhaustive by 1.2%, demonstrating that not all automated and naïve autonomous approaches are superior.

Between autonomous approaches at 40% coverage, we found that whole-field RL garnered 13% more revenue than local-field RL. Despite similar labor costs, the accuracy improvements over local-field RL, particularly among the negative recall, provides a considerable increase in revenue for whole-field RL over local-field RL. By limiting false negatives, whole-field RL was found to reduce higher runoff risk by 12% compared to local-field RL. However, all the revenue data are generated from our simulation environment, which means that they are under the best case scenarios without considering other important factors such as climate, weather, market and insects.

4. Limitations and Future Work

Autonomous scouting methods inherently avoid surveying 100% of a field to save time, energy, and money. This, however, runs the risk of missing critical field health problems. This problem can be minimized if a field is regularly monitored for potential crop health problems during growing seasons. A problematic section of a field that might not have been picked up by autonomous scouting at one time is likely to be picked up if the field is regularly mapped.

The cost-benefit model used in this study considers fertilizer application as the only treatment for unhealthy zones. Zones may have stresses such as pest and water other than nutrients only, and thus need to be treated accordingly, which could in turn influence the treatment costs. Also, costs and benefits were estimated based on corn only. Some of these estimates can differ by crop and treatment types. Thus, future studies should be focused on evaluating some of these factors.

In the study, we used ExG as an indicator of crop health for simplicity. There are however other vegetation indices (e.g., NDVI, green index) and biophysical variables (e.g., soil organic carbon, pH, elevation) that are also reported to be good indicators of crop health [26]. Future studies could exploit a combination of these variables as indicators of crop health while developing models for autonomous scouting approaches.

Future work should also focus on training, reinforcement learning, and testing of the models based on data sets with a variety of crops collected from separate fields in separate conditions. The RL approaches used in this study, q-learning algorithm, can also be compared with other related sampling algorithms such as rapidly-exploring random trees. Similarly, future work should address how prediction window size and its effect on architectural latency affects overall cost and performance considering its effect on accuracy.

5. Conclusion

In this study we design and discuss a new fully autonomous aerial scouting approach, whole-field RL, as compared to local-field RL approach and the current naive UAS approach of exhaustive scouting. The performance of these two RL approaches along with other popular scouting methods were assessed in terms of accuracy, precision, recall, and execution time of crop health maps, and cost-saving potential across different field coverage ranging between 10% to 60% of the total field area. Compared to local-field RL, whole-field RL can boost accuracy of crop health maps by 9%. This approach produces accurate crop health maps after flying over only 40% of the field. Whole-field RL reduced labor cost by 4.8 times, increased agricultural profits by 36% and reduced runoff potential by 87%. We found that coverage rate offers diminishing improvements in accuracy after 40%. The considerable improvement in performance of whole-field scouting over local-field scouting can largely be attributed to its added CNN models to use surrounding ground truth data to predict health condition of management zones in flight. In-flight predictions allow the final crop health map to be iteratively updated and refined in flight, producing a more accurate final product. Our evaluation shows that fully autonomous aerial scouting can guide crop field management techniques that use less money, less agricultural product and achieve greater monetary return than the state of the practice.

Author Contributions: conceptualization,Z.Z., J.B. and C.S.; methodology, Z.Z. and J.B.; software, J.B.; validation, J.B. and Z.Z.; formal analysis, Z.Z., J.B. and C.S.; investigation, Z.Z.; resources, Z.Z. and S.K.; data curation, Z.Z. and J.B.; writing–original draft preparation, J.B.; writing–review and editing, Z.Z., J.B., C.S. and S.K.; visualization, Z.Z.; supervision, C.S.; project administration, C.S. and S.K.; funding acquisition, C.S. and S.K.

Funding: This research had no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Godfray, H.C.J.; Beddington, J.R.; Crute, I.R.; Haddad, L.; Lawrence, D.; Muir, J.F.; Pretty, J.; Robinson, S.; Thomas, S.M.; Toulmin, C. Food security: the challenge of feeding 9 billion people. *science* **2010**, *327*, 812–818.
2. Lal, R. Feeding 11 billion on 0.5 billion hectare of area under cereal crops. *Food and Energy Security* **2016**, *5*, 239–251.
3. Foley, J.A.; Ramankutty, N.; Brauman, K.A.; Cassidy, E.S.; Gerber, J.S.; Johnston, M.; Mueller, N.D.; O'Connell, C.; Ray, D.K.; West, P.C.; others. Solutions for a cultivated planet. *Nature* **2011**, *478*, 337–342.
4. Steensland, A.; Zeigler, D. Global Agricultural Productivity Report. *Global Harvest Initiative* **2017**.

5. Zhang, X.; Cai, X. Climate change impacts on global agricultural land availability. *Environmental Research Letters* **2011**, *6*, 014014.
6. Zhang, N.; Wang, M.; Wang, N. Precision agriculture—a worldwide overview. *Computers and electronics in agriculture* **2002**, *36*, 113–132.
7. Gebbers, R.; Adamchuk, V.I. Precision agriculture and food security. *Science* **2010**, *327*.
8. Pierce, F.J.; Nowak, P. Aspects of precision agriculture. In *Advances in agronomy*; Elsevier, 1999; Vol. 67, pp. 1–85.
9. Bell, J.; Butler, C.; Thompson, J. Soil-terrain modeling for site-specific agricultural management. Site-specific management for agricultural systems. Wiley Online Library, 1995, pp. 209–227.
10. Robert, P.C. Precision agriculture: a challenge for crop nutrition management. Progress in Plant Nutrition: Plenary Lectures of the XIV International Plant Nutrition Colloquium. Springer, 2002, pp. 143–149.
11. Boubin, J.; Chumley, J.; Stewart, C.; Khanal, S. Autonomic computing challenges in fully autonomous precision agriculture. 2019 IEEE International Conference on Autonomic Computing (ICAC). IEEE, 2019, pp. 11–17.
12. Candiago, S.; Remondino, F.; De Giglio, M.; Dubbini, M.; Gattelli, M. Evaluating multispectral images and vegetation indices for precision farming applications from UAV images. *Remote sensing* **2015**, *7*, 4026–4047.
13. Boroujerdian, B.; Genc, H.; Krishnan, S.; Cui, W.; Faust, A.; Reddi, V. MAVBench: Micro Aerial Vehicle Benchmarking. MICRO, 2018.
14. Vasisht, D.; Kapetanovic, Z.; Won, J.; Chandra, .R.; Kapoor, A.; Sinha, S.; Sudarshan, M.; Strãtman, S. FarmBeats: An IoT Platform for Data-Driven Agriculture. NSDI, 2017.
15. Larrinaga, A.R.; Brotons, L. Greenness indices from a low-cost UAV imagery as tools for monitoring post-fire forest recovery. *Drones* **2019**, *3*, 6.
16. Khanal, S.; Fulton, J.; Douridas, N.; Klopfenstein, A.; Shearer, S. Integrating aerial images for in-season nitrogen management in a corn field. *computers and electronics in agriculture* **2018**, *148*.
17. Boubin, J.; Babu, N.; Stewart, C.; Chumley, J.; Zhang, S. Managing edge resources for fully autonomous aerial systems. ACM Symposium on Edge Computing, 2019.
18. Lin, S.C.; Zhang, Y.; Hsu, C.H.; Skach, M.; Haque, M.E.; Tang, L.; Mars, J. The Architectural Implications of Autonomous Driving: Constraints and Acceleration. ASPLOS, 2018.
19. Song, M.; Zhong, K.; Zhang, J.; Hu, Y.; Liu, D.; Zhang, W.; Wang, J.; Li, T. In-Situ AI: Towards Autonomous and Incremental Deep Learning for IoT Systems. High Performance Computer Architecture, 2018.
20. Singh, S.P.; Jaakkola, T.; Jordan, M.I. Learning without state-estimation in partially observable Markovian decision processes. In *Machine Learning Proceedings 1994*; Elsevier, 1994.
21. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* **2014**.
22. Hughes, G.; Madden, L. Some methods for eliciting expert knowledge of plant disease epidemics and their application in cluster sampling for disease incidence. *Crop Protection* **2002**, *21*.
23. Kotsiantis, S.B.; Zaharakis, I.; Pintelas, P. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering* **2007**, *160*, 3–24.
24. Boubin, J.; Stewart, C.; Zhang, S.; Babu, N.T.; Zhang, Z. SoftwarePilot. http://github.com/boubinjg/softwarepilot, 2019.
25. Schnitkey, G. 2018 Crop Budgets: More of the Same. Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 2017.
26. Mulla, D.J. Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps. *Biosystems engineering* **2013**, *114*, 358–371.