Weeds Out!

DEVELOPING AN AI-BASED WEED MONITORING PROGRAM FOR HUME CITY COUNCIL

March 1, 2024

Students: Jake Bowen, Mia Gilmore, Kimberly Huang and Vivek Voleti
Advisors: Stephen McCauley and Sara Saberi

Developing an AI-based weed monitoring program for Hume City Council

An Interactive Qualifying Project submitted to the Faculty of Worcester Polytechnic Institute in partial fulfilment of the requirements for the degree of Bachelor of Science.

Written by: Jake Bowen, Mia Gilmore, Kimberly Huang, and Vivek Voleti

Advised by: Stephen McCauley and Sara Saberi

Community Sponsor: Hume City Council

Submitted On: March 1, 2024

Acknowledgements

Special thank you to all who made this project possible.

We would like to thank our sponsor, Hume City Council, and their representatives for this project, Liz Turner, and Julian Thompson, as well as Ravi Krishnamurthi and other employees of the council whom we worked with. This project was made successful through all of their help and resources.

We also would like to thank our advisors, Professors Stephen McCauley, Sara Saberi, and Jason Davis. This project would not have been possible without their help, support, and excellent feedback throughout the last two terms.

Our project was completed on the lands of the Wurundjeri people and we wish to acknowledge them as Traditional Owners. We would also like to pay our respects to their Elders, past and present, and Aboriginal Elders of other communities.

Contributions

Jake Bowen

Jake was the team member responsible for communicating with external resources outside of Hume City Council. Jake additionally led the research for learning about other case studies and software similar to our project and was the group's point of contact for connecting with the other programs. Jake was a prominent editor for other sections and contributed his well thought out ideas to the group.

Mia Gilmore

Mia was the leading team member responsible for following through with our participant consent during research, focus groups and interviews. Mia also significantly aided the team in gathering field imagery and external factors to be used for presentation and report purposes. By observing the fields on drone days, Mia was able to record key climate conditions which were used to construct the Deliverables. Mia co-conducted the focus group and was an excellent note taker throughout the entire project. Mia was a strong supporter of all sections and was essentially a very flexible team member with her ability to be knowledgeable with a variety of tasks.

Kimberly Huang

Kimberly was the main organizer and primary editor for all assignments turned in. With her organizational skills, Kimberly was the main coordinator of the group in setting up meetings and group communication to Hume City Council and advisors. Kimberly co-conducted the focus group which led to her being in charge of the background research of stakeholders and social aspects of the project. Kimberly naturally fell into the role of being the delegator of tasks and kept the group on task.

Vivek Voleti

Vivek was the essential team member responsible for learning and using various software for the project. Vivek dedicated much time outside of group meetings to research and dedicated his computer for training and testing training models on ArcGIS. Vivek's strength in robotics and computer science led to him being the optimal team member to learn and employ different machine learning models. Vivek's devotion to this project created a positive team atmosphere and helped the team accomplish its technical goals.

Abstract

Invasive noxious weeds pose a significant economic and environmental threat to Victoria, Australia. This is especially the case in Hume City, where environmental and agricultural productivity is important. Our project aided Hume City Council in developing an artificial intelligence (AI) and drone-based protocol to automatically detect and map weeds to help monitor the spread of weeds over time. The data gathered from the analyses completed as a part of this study will allow Hume to enforce and adapt their weed control strategies, helping them combat noxious weeds throughout the municipality. Through social context assessment, analysis of existing solutions, and comprehensive experimentation with weed mapping conditions, we developed a protocol and a base set of AI tools to form the backbone of the Council's weed monitoring program.

Executive Summary

Invasive species of weeds are a serious threat to landscapes in Victoria, Australia. These noxious weeds are disruptive to the environment and animals that inhabit it. With their ability to spread with ease, they increase the risk of bushfires, can harbor pests and disease, and interrupt the natural processes of native flora and injure wildlife that grazes on these weeds. The municipality of Hume, which begins just 15 km outside of downtown Melbourne, has been particularly vulnerable to noxious weeds due to the large area of green space encapsulated within the city. The presence and spread of these noxious weeds, specifically Serrated Tussock and Artichoke Thistle, is one of the biggest challenges that Hume's rural community faces.

To address this problem, Hume City Council (HCC) wants to encourage and enforce the treatment of these noxious weeds. However, they require a reliable and efficient method of identifying and mapping weed hotspots to begin with. Our project aimed to test methods for an artificial intelligence (AI) based weed mapping protocol that Hume City Council can employ to map invasive weeds, which will allow the council to measure the effectiveness of their weed control measures over time and help enforce treatment of weeds on private lands. Working with HCC's sustainability team, we pursued four key objectives that helped us to accomplish our goal.

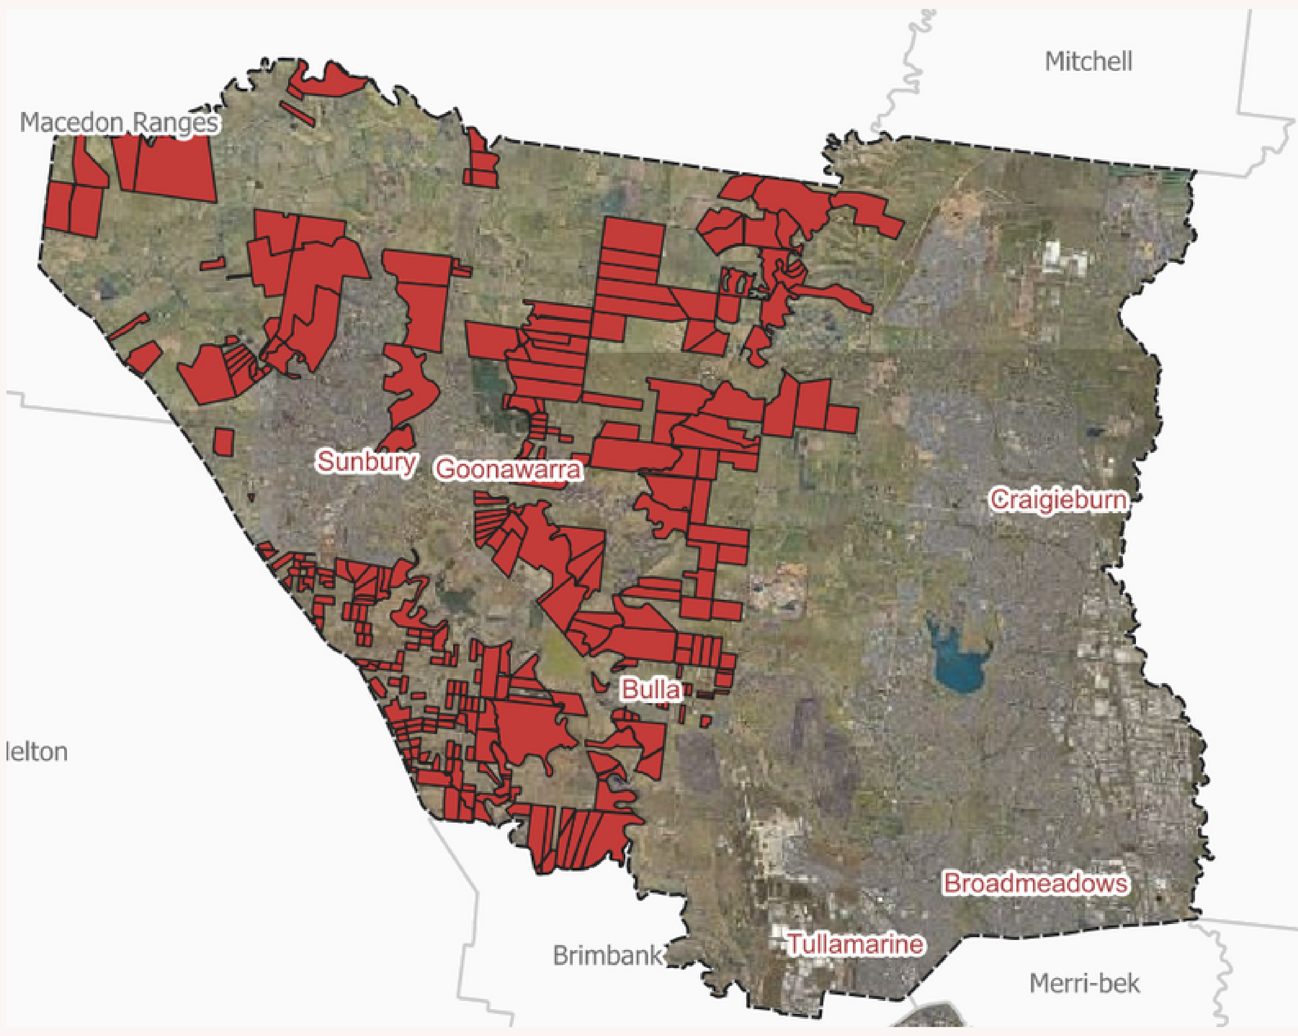

Figure ES-1. Hume City Council's map of artichoke thistle (red) in the rural areas of Hume.

Assessing the Social Context of Weeds

Our first objective was to assess the social context involved in the weed compliance program, through the use of causal research, focus groups, and archival research. During an informal discussion with a HCC Rural Officer we discovered that neglected weeds can spread to neighboring properties, sparking disputes among neighbors due to the substantial expenses associated with treating and managing the weeds. The focus group we conducted led to an understanding that HCC has not begun enforcement in the weed compliance program and is focused on the goal of being able to detect the noxious weeds on properties. Lastly, from exploring the different documents in the HCC shared Microsoft folder, an understanding of the drone laws was built with the laws of drones cannot exceed over 120 meters, has to fly 30 meters over humans, and 5.5 kilometers away from airports (“no flight zones”).

Figure ES-2. Weed seeds from an unmaintained property blown onto a maintained property

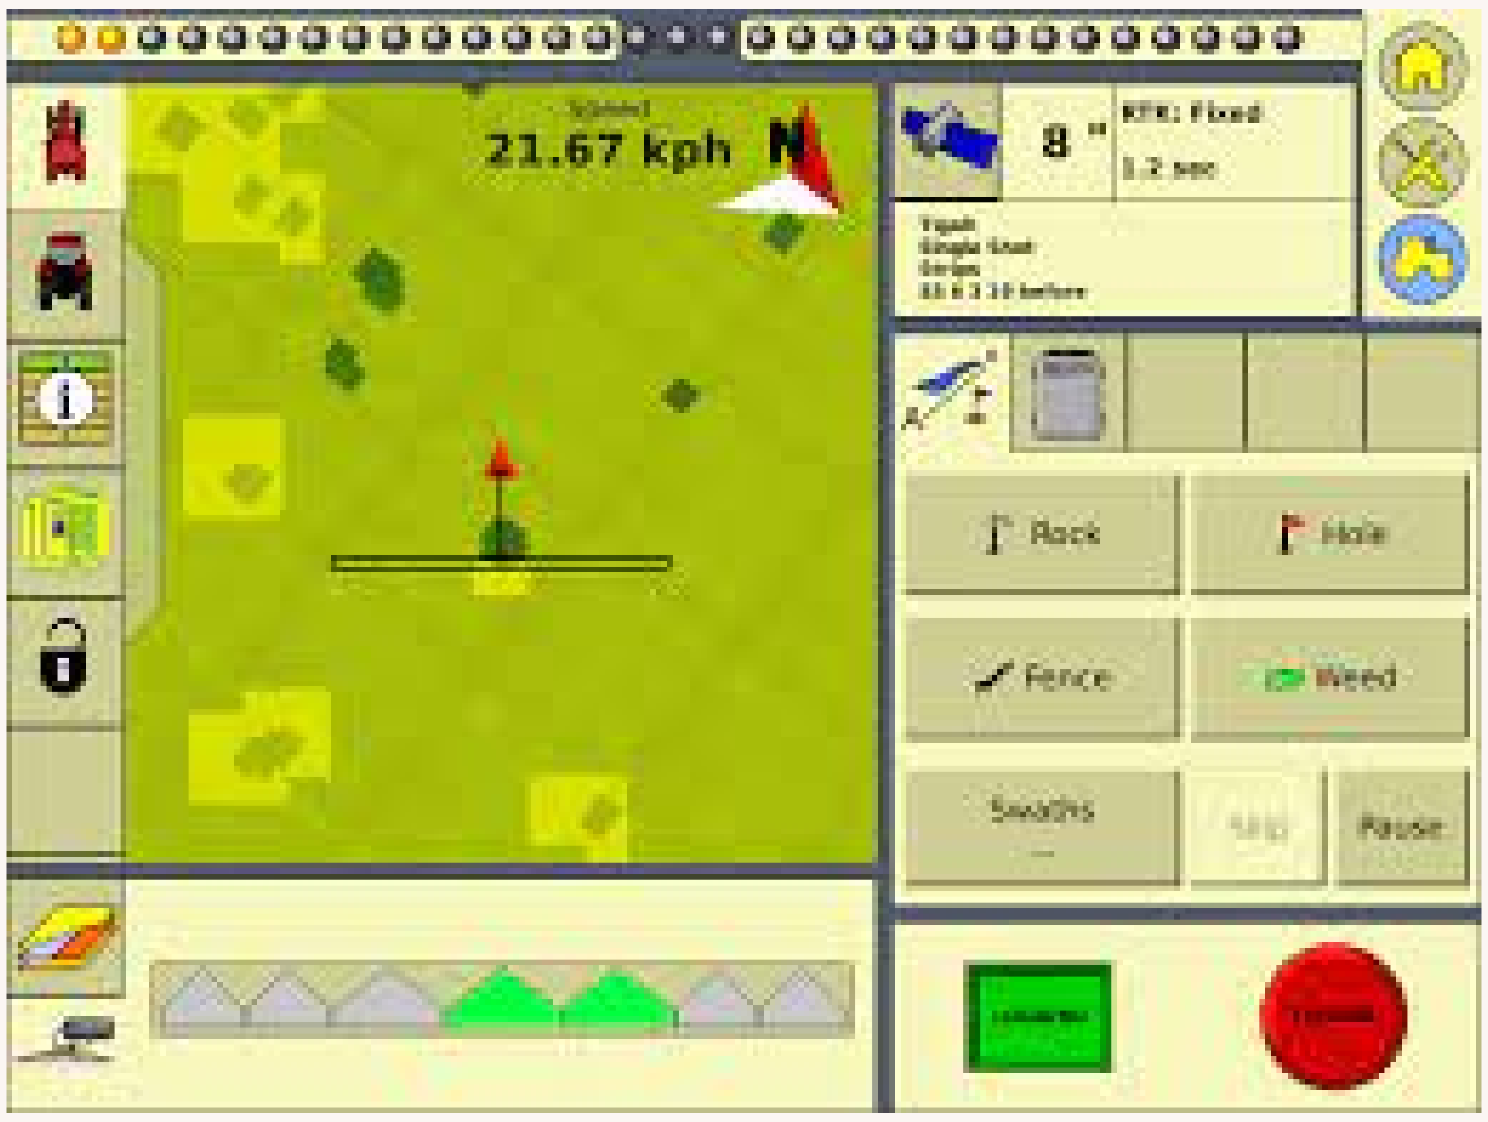

Figure ES-3. Image of software program used by SingleShot to map their drone

Exploring Effective AI-based Methods of Weed Detection

Our second objective was to explore effective AI-based methods conducted by neighboring organizations through archival research and personal communication lines. During the research process, we found that the technology is relatively new to Australia yet three distinct types of organizations are utilizing it. Firstly, there are private firms such as SingleShot, which have developed and patented AI-driven weed mapping technology. These firms offer services for land surveying and weed detection, capable of pinpointing weeds as small as a cork and plotting their GPS coordinates for treatment. The company does have the resources to treat the weeds the same day through sprayer attachments on the drone or the drone sends the coordinates to a tractor that will drive sound the land spraying the weeds. The second category comprises academic institutions, notably Central Queensland University and the University of Queensland, Australia.

These universities have gathered a group of professors with backgrounds in computer science related fields and have developed two-year programs for software development, data training, along with field studies. The universities are extremely protective over this information because they have not produced any published reports at this time. Lastly, local governments constitute the last category, although their noxious weed problems are comparatively less intense. Although they lack the specialized technology, our group conducted an interview with a member of Whittlesea Council where our group was pointed in the direction of third party software that can allow us to stitch our images more accurately and quickly.

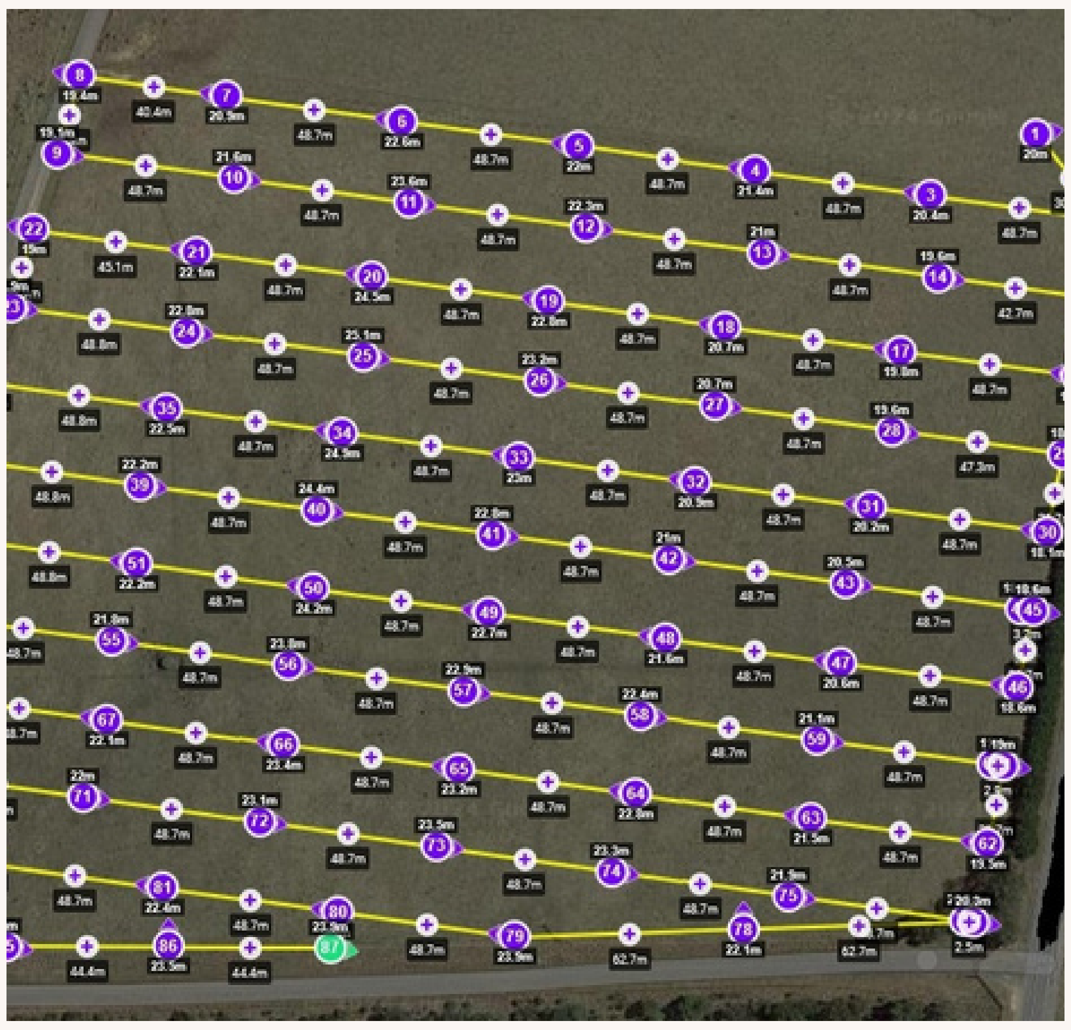

Figure ES-4. Flight path generated by DJI Air 2S software that includes image intervals

Collecting and Processing Field Imagery to Train Different AI Models

Our third objective involved collecting and processing field imagery to serve as training data for the different AI models offered by GIS. To accomplish this objective, we conducted geographic sampling through the use of drones that collected aerial imagery of various landscapes throughout Hume. During data collection, we documented multiple environmental factors influencing the imagery, such as time of day, light conditions (including sun and cloud cover), wind speed, terrain characteristics (such as hills or flat terrain), and the vegetation composition, including weed-to-native flora ratios and specific weed species present. Another important aspect of our data collection was leveraging the automatic flight planning system available on HCC's DJI Air 2S drone. This program enabled us to highlight geographic areas and customize image parameters, including camera height, gimbal angle, and image overlap percentage. With this information, the DJI drone's software was able to calculate a flight path that would take images at predetermined intervals. This drone program provided us with the necessary field imagery to train our AI models effectively.

Figure ES-5. Flight path generated by DJI Air 2S software to demonstrates the full flight path

Testing DIfferent AI Models to Determine the Highest Performing Model

Our fourth objective entailed training our own custom machine learning models using ArcGIS, and evaluating our trained models to measure their effectiveness at weed detection. This objective was achieved by generating an orthomosaic image—a composite image stitched together from drone-captured images—to establish a training dataset using the imagery obtained in objective three.

Figure ES-6. Property with artichoke thistle and serrated tussock

We iterated on several model architectures, but eventually trained two YOLOv3 object detection models (one for Serrated Tussock, and one for Artichoke Thistle). We tested these models against various different properties with various levels of weed infestation, and explored each model's strengths and weaknesses by analyzing the weeds the models predicted correctly, the weeds they did not detect, and the weeds that were detected but misclassified. One such property we analyzed is shown in figures ES-6 and ES-7. Ultimately, we determined that while our initial machine learning models produce errors (as all AI models do), they display a high accuracy given the small size of our training dataset. We determined that an AI-based weed monitoring approach can be viable at scale, and we provided further methods and recommendations to improve on our work.

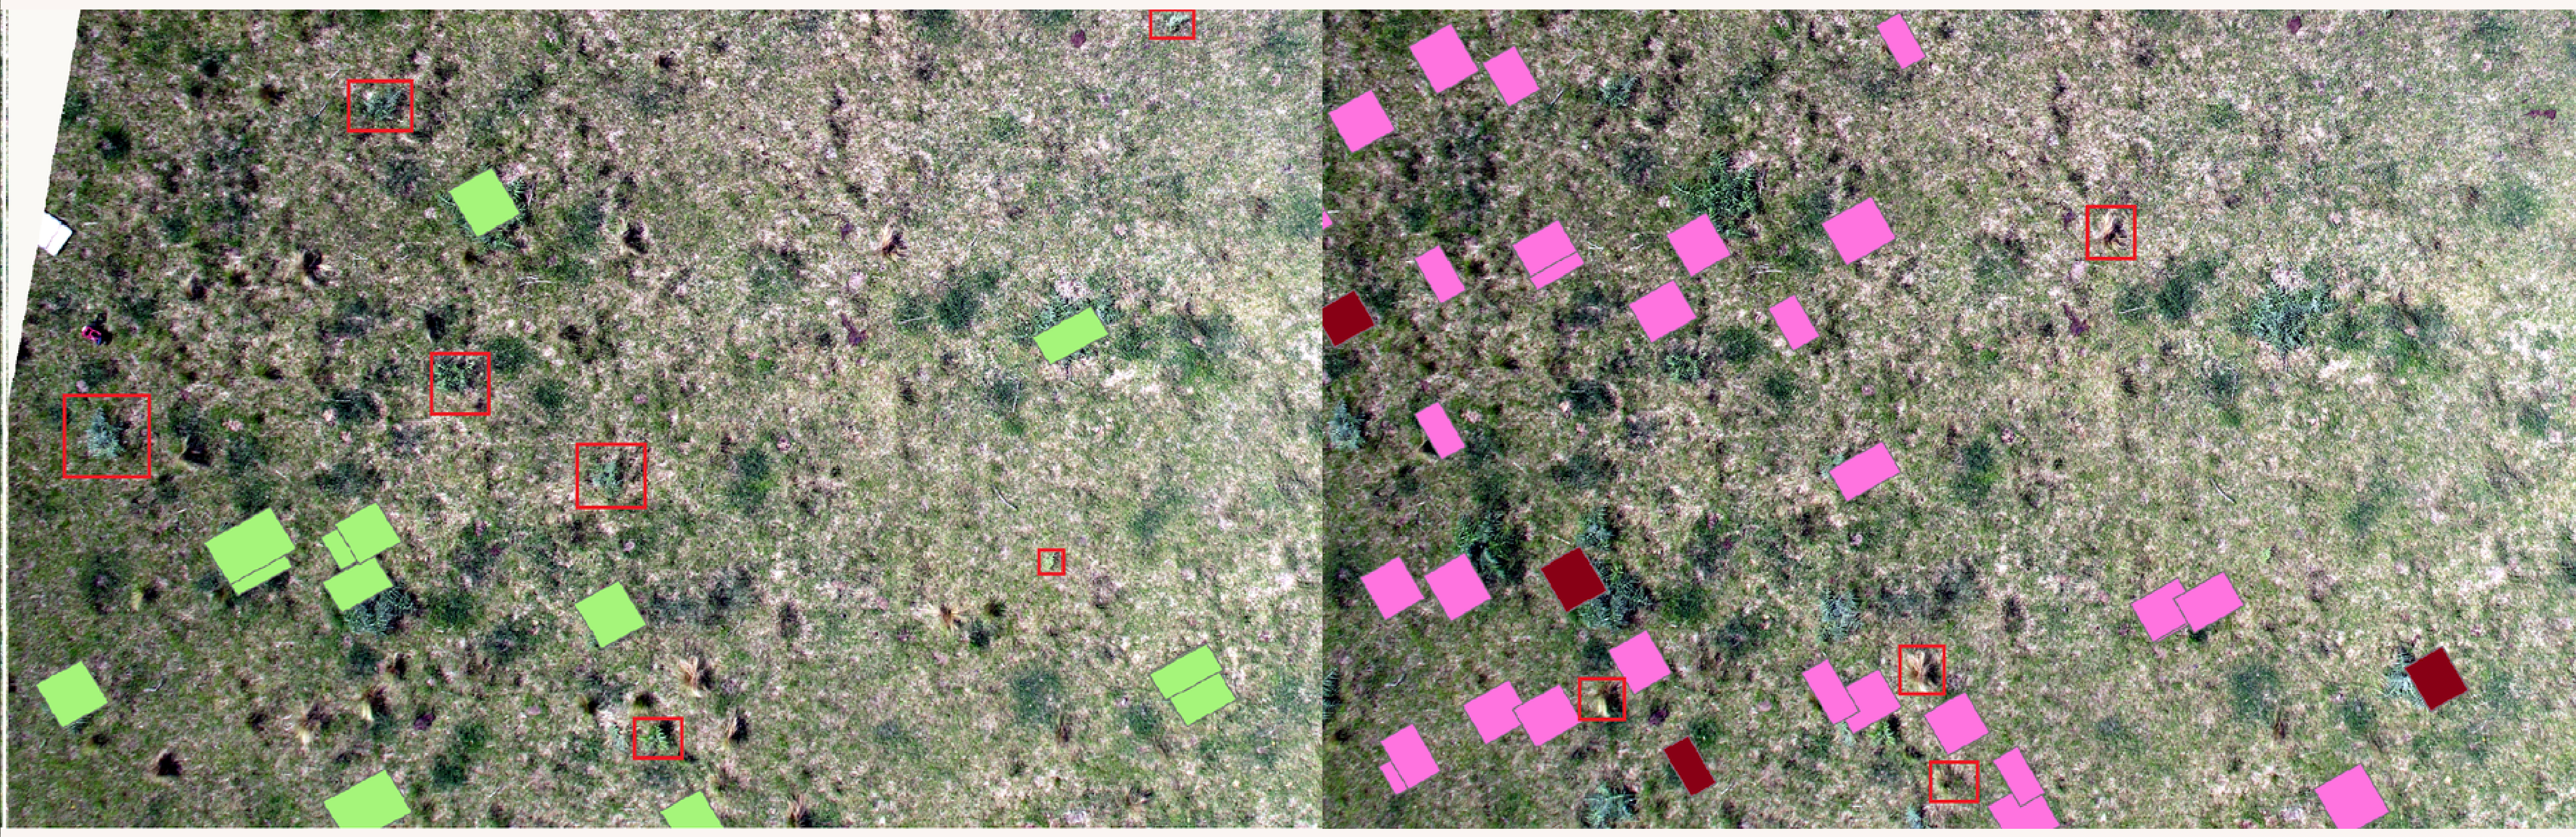

Figure ES-7a & ES-7b. The model's predicted artichoke thistle outputs (light green), and serrated tussock outputs (pink). Missed weeds are outlined in red boxes, and vegetation that was mistaken as Tussock is dark red.

Upon completion of our project, we had multiple recommendations for Hume City Council to continue building on the team's work. It is important to understand the social context regarding such a program, which is why we recommend conducting interviews with landowners to further comprehend their thoughts and apprehensions on using AI, and to strengthen lines of communication between Council officers and the general public. Additionally, questions regarding landowner data privacy may arise regarding who will have access to their data, and what will be done with it. Regarding the technical aspects of this project, there will likely always be errors that arise with the use of AI for weed detection. As a result, we recommend HCC to continually improve on our base models by training on a wider variety of training data under different weather conditions and different stages in the weeds' lifecycles. We further recommend human supervision to catch errors throughout deployment of models, and to provide methods of communication for landowners to HCC in case of errors. We finally recommend using supplementary hardware and software to increase model accuracy and cut down on processing time. For example, through use of multispectral imagery to improve model accuracy or exploring additional computing resources, the AI-based protocol we developed will be easier to scale for use across the entire municipality.

It is evident that the increasing presence of noxious weeds in Hume is harmful to the environment, and poses an economic burden on those who are tasked with managing them. Making use of AI technology with aerial imagery to map weeds can help HCC enforce and adapt their weed management strategies as needed to effectively combat noxious weeds throughout the municipality.

Figure ES-8. Photograph of land in Hume with a scattered presence of noxious weeds.

Table of Contents

List of Figures

.

List of Figures

List of Tables

Introduction

Invasive weed species have been taking over the landscape in Victoria, Australia. These weed species (which we will refer to as “noxious weeds”) are known to be disruptive to the environment, and are easily spreadable. They harm native flora by consuming space, light, water, soil nutrients and pollinators. Noxious weeds also endanger the public, as they interfere with fire protection, harbors pests and diseases, and spoils crops (Hume City Council, n.d.). As a result of..., Australian grain growers have lost ~$3.5 billion (AUD) within one year (Australian Broadcasting Corporation, 2017)

Noxious weeds have been rapidly spreading particularly in Hume city, a municipality in the northern Melbourne metropolitan region which is largely covered in agricultural lands, both private and publicly-owned. Recently, the Hume City Council (HCC) has instituted a new local law known as the Weed Compliance Program that requires landowners to manage noxious weeds on their private property to avoid being fined. While the new program provides a basis for managing weeds, HCC currently lacks the ability to monitor the spread of these weeds.

Previously, the council has mapped the locations of noxious weeds in a laborious fashion, which involved officers manually mapping weeds by foot on the ground or by manually viewing airplane imagery. HCC desired to automate this process, namely by using drones to capture images of land from above and artificial intelligence (AI) to detect and label weeds. AI use in agricultural settings (particularly in weed detection) is still an emerging field - while drone imagery in combination with AI-based classification is increasingly used in large commercial farming applications, it is still an emerging technology for municipal governments. Much of the commercial work is proprietary, meaning details on existing successful implementations may be rare. However, the potential time and accuracy gains achievable through an automated process made this a desirable approach to pursue.

Our project goal was to test methods for an AI-based weed mapping protocol which Hume City Council can employ to map different species of weeds, which will allow the council to measure the effectiveness of their weed control program over time. As a group, we have worked closely with the Council's sustainability team, building on their ongoing efforts to construct a weed monitoring method by using AI. We have achieved our goal by following four objectives: (1) assess the social context surrounding the noxious weeds in Hume; (2) explore the state of AI mapping technologies in the surrounding area (particularly around Hume) which the HCC can reasonably replicate; (3) collect and process training data for the artificial intelligence weed detection models; and (4) test several models on unseen imagery, and identify the most viable high performing models. This process has allowed us to test different weed detection and mapping systems for HCC. To improve the efficiency of the previous approach, we have determined standardized AI-based procedures to map and monitor weed locations. This procedure has allowed us to establish an approach which will allow HCC to map weeds at regular time intervals and compare the data to determine the effectiveness of their weed control methods.

Figure 1. Noxious Weeds on Agricultural Property

Background

This chapter provides an in-depth background on the noxious weeds spreading across Victoria, particularly in the city of Hume. These invasive weeds have the ability to negatively affect crop growth, livestock, and the health of humans in the area. We examine the ways the seeds of these invasive species are spread, how they affect the landscape and landowners, and some of the most promising methods of monitoring these weeds. This chapter will also cover tools that are used in automated weed monitoring, including AI models, software to visualize weed locations on a map, and machine learning.

Noxious Weed Overview

Australian landscapes have been geographically isolated from the evolution of plants on other continents since the split of Pangea about 200 million years ago. As a result, Australian landscapes are left vulnerable to invasion by species from elsewhere (Cullen et al., 2023). There are various ways for seeds to travel to and throughout Australia. For instance, both aquatic and land weeds can travel on the bottom of boats that travel to Australia from surrounding areas. As a result, these weeds may arrive on shore, where their seeds are dispersed by human movement. Many weeds are also able to spread their seeds through the wind and by human travel or animal travel. It has been noted by a HCC officer that Hume has a greater volume of weeds compared to neighboring municipalities. The HCC officer has also noted that a large portion of the rural areas of Hume are infested with either Serrated Tussock, Artichoke Thistle or both (J. Thompson, 2024).

Serrated Tussock (Nassella Trichotoma), Figure 2, is a weed of interest for Hume City Council, as it is widespread due to its ability to spread several kilometers in the wind, and even farther when attaching to humans (Department of Jobs, 2023). Serrated tussock is a perennial tussock-forming grass that grows in dense clumps. Due to its diffused and fibrous roots, serrated tussock has a long lifespan. Dense infestations of this weed can pose a serious fire hazard with a recorded burn intensity of up to 7 times greater than native grasslands. Serrated tussock is so problematic because its seed production is abundant; a hectare of dense tussock growth can produce more than 2 tons of seed annually. Larger weeds can produce 100,000 seeds a year, which can remain dominant in the soil for over 15 years (Department of Jobs, 2023).

Figure 2. Serrated Tussock

Figure 3. Artichoke Thistle

Artichoke Thistle (Cynara Cardunculus), as seen in Figure 3, is the second weed of interest for the HCC. This weed can extend its roots up to 8 feet underground, which allows it to keep water for itself long after it rains (Los Angeles Times, 1998). The plant's twisty branches also prevent sunlight from reaching other plant species that are native to the area. This weed also tends to grow long thorns on its stem which will pierce whatever may come across it. Its seeds easily travel in the wind, where they can move up to 20 meters from the original plant. Livestock, birds, and human clothing all also have the ability to transport the seeds to other landscapes, which further increases the surface area they cover (Department of Jobs, 2023).

Weed Monitoring and Mapping Tools

For weed detection, HCC's goal is to use artificial intelligence to automatically classify a plot of land as either: containing one of the targeted noxious weed species, or not containing any noxious weeds. One subclass of AI models, Machine Learning (ML), was found to be useful by other organizations utilizing AI in our research . Machine Learning models strive to learn some pattern or behavior from various examples given during a "training" phase (MIT, 2021). For example, if one were to create a model to identify any weeds present in a given image of some agricultural land, training a model would entail feeding it hundreds of images (potentially more) along with the corresponding ground truth labels, which dictate whether or not a particular image contains a targeted weed species. Over a model's training cycle, it consumes an image, makes a prediction, compares this prediction with the image's corresponding ground truth label, and adjusts its parameters ("learns") as necessary. After training, ML models are typically tested on data that was not used during training to evaluate their accuracy and potential faults (MIT, 2021).

Machine Learning models have proven to be effective in agricultural applications. In addition to aerially captured imagery, ML is already being used for several different agricultural purposes, including detection of diseases in food and monitoring of plant water levels, among other things (US Dept. of Agriculture). Specifically, researchers have been attempting to apply this technology to weed detection, aided by computer vision and machine learning techniques. Bullock et al. (n.d.) mounted a camera on a drone, and trained a Convolutional Neural Network (CNN) - a type of ML model - based on training data taken from the drone at different altitudes. They trained models with three different data processing methods, and the model with the highest accuracy correctly detected the presence of weeds in 97.1% of the training images (Bullock et al., 2019).

Automatic weed detection is still relatively new in Australia's technology scene, however some programs are developed and possibly outperform Bullock et al.'s model when put in the field. When using an image-capturing drone, ML models (both existing and self-trained) can be used to process thousands of images, identifying which of them have weeds. Based on the location of the image taken, this information can be used to build a complete geographic profile of weed locations. It is important to note that there are other approaches for AI-based image classification, but ML approaches seem feasible due to our time and resource limitations.

HCC had already used Esri's ArcGIS software to visualize geographic weed hotspots. ArcGIS gives users powerful computational capabilities, which include both image processing and map visualization. The software's map visualization features have allowed HCC officers to manually construct a map of Artichoke Thistle weeds around Hume city using data collected from their manual inspection approaches, which can be seen in Figure 4. In this figure, the red zones depict lands that contain noxious weeds. ArcGIS also provides AI models, which can be used in conjunction with its mapping capabilities to develop an efficient and automated approach to detect weeds on land properties.

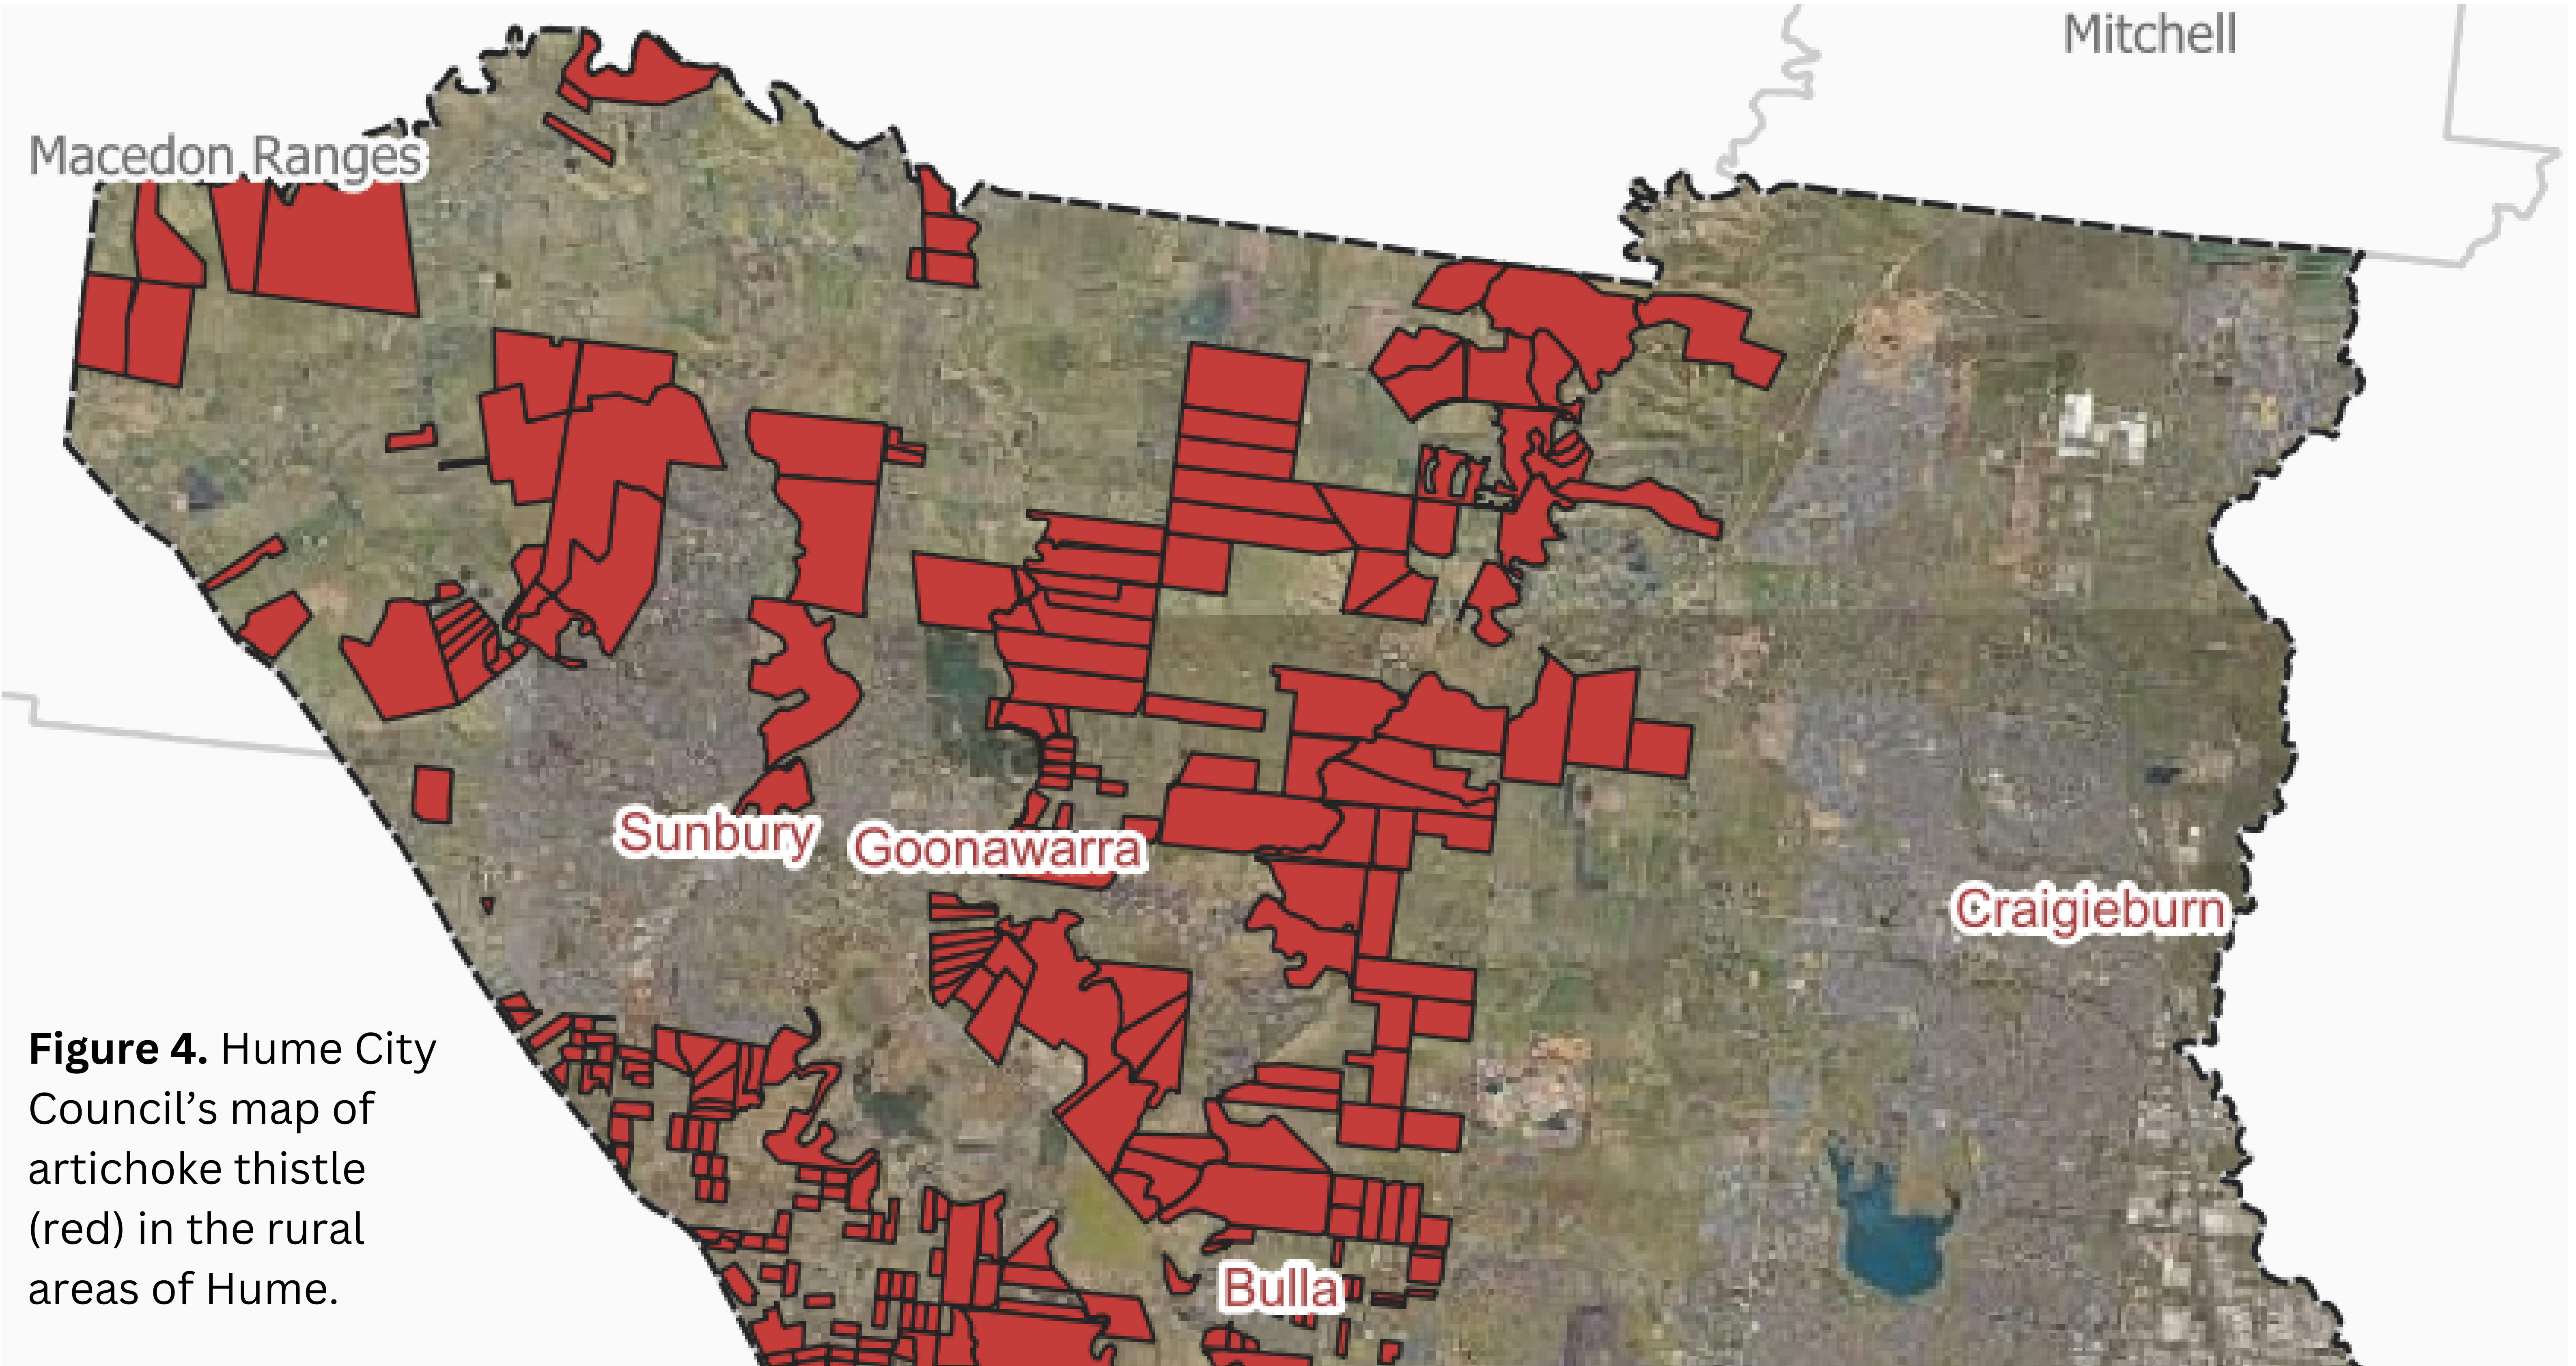

Figure 4. Hume City Council's map of artichoke thistle (red) in the rural areas of Hume.

Relevent Examples of AI-based Weed Detection

Other government organizations across the world have implemented AI based weed detection and mapping softwares into their own weed monitoring programs. As weed monitoring programs have been recently being developed, these initiatives are included within emerging research where science is always advancing and efforts are starting to be applied in land management programs.

Regarding automatic weed detection, the Australian government had funded the Centre For Invasive Species Solutions for the development of the WeedScan app. “WeedScan is powered by a CSIRO AI identification model trained on more than 120,000 weed images and tested across Australia by scientists, farmers, community groups, agronomists, rangers, weeds and natural resource management officers'' (CISS, 2023). Central Queensland University has also implemented a two year plan where research teams will use light-weight drones to capture ultra-high-resolution images of weeds that will be processed to create GPS locations (CQU, 2023).

The Oregon Department of Agriculture draws data from even more sources, resulting in a comprehensive map of noxious weed locations. The government uses an ArcGIS-based tool called WeedMapper (Figure 5), which plots the locations of noxious weeds across the state. It displays data drawn from individual reports of weed locations, and from the online tracking tool iMapInvasives (ODA, 2020). Though this data is mainly collected through manual crowdsourcing approaches, this application presents a way to use spatial data on multiple different noxious weed species, which has been useful while constructing HCC’s monitoring program.

Figure 5. The WeedMapper tool, showing locations of noxious weeds (Oregon Dept. of Agriculture)

Methodology

The goal of this project was to test different weed detection and mapping strategies, which will allow the council to measure the effectiveness of their weed control enforcement strategies over time. To allow for continued development and implementation of the Weed Compliance Program we have compiled our findings to a document which HCC will use to advance their weed monitoring systems. To construct our weed mapping approach, we determined appropriate AI-based weed detection tool(s) that align with HCC's current resources to be used in conjunction with GIS' mapping capabilities to build a mapping protocol that may be repeated over regular time intervals in the future. We completed four main objectives to achieve this goal:

1. Assess the social context surrounding the noxious weeds monitoring program
2. Explore AI-based methods of automatic weed detection used by other organizations
3. Construct a data sampling protocol, and collect training data for the AI weed detection models
4. Train and test AI models on collected data, and identify strengths, weaknesses, and next steps

Objective One: Understanding the social context and the needs and interests of relevant stakeholders regarding the noxious weed monitoring program

Our group's first objective was to understand important aspects of the social context for introducing an AI and drone-based approach to weed monitoring. This objective involved learning about HCC's weed monitoring program, drone regulations, privacy laws, and the view of landowners about noxious weeds.

HCC's weed monitoring program

Landowner's Perspective

Drone Regulations

To start off our social context research, we had to understand HCC's current weed monitoring program and what they hoped for us to accomplish. With the use of a focus group, we met with HCC officers to learn about the Weed Compliance program. Focus group questions, found in Appendix A, helped us gain an idea on what weed monitoring methods have been tried and what has or has not worked so far. Throughout the whole duration of this project, our team was accompanied with HCC sustainability team members including the weed compliance manager who contributed to numerous helpful conversions regarding our research. By discussing HCC's past monitoring methods, we were also able to grasp a better understanding of HCC's resources for monitoring noxious weeds.

Sustainability Team

Rural Officers

Weed Program Officers

To continue our research with the social context of this project, we had to understand the effect of the noxious weeds on a social and environmental level. By casually chatting with a HCC officer, we were able to learn about the extent to which noxious weeds have affected Australia's lands. The knowledge learned from casual conversations all term, allowed us to understand the motivations behind the Weed Compliance Program and the impact of these noxious weeds on Australian landowners. The combination of the focus group and the chatting allowed us to recognize the effect of the noxious weeds on a social and environmental level and how these insights could inform the implementation of an AI and drone-based approach.

To finish our social context research we had to understand the legal abilities of intaking drone footage. We conducted research on privacy laws and drone regulations to ensure that our work would not be a breach of privacy. The HCC officers have been able to clearly define any problems that would have occurred regarding privacy breaches due to data collection and accessibility. Reliable resources have been provided to us which describe the drone regulations of our data collection and privacy laws that keep landowners safe from privacy aspects. This research allowed us to gain more consideration of ethical and legal dimensions in implementing an AI and drone-based weed monitoring system. Understanding privacy laws and drone regulations was pivotal in ensuring that our approach aligns with legal frameworks, and safeguarding the rights and privacy of individuals. The insights gathered from the HCC officers provided clarity on potential issues related to data collection and accessibility. This comprehensive understanding of the legal aspects surrounding our project did not only ensure ethical practices but also established a framework that creates a sense of confidence in our AI and drone-based weed monitoring framework.



Objective Two: Explore AI-based methods of automatic weed detection used by other Organizations

Our second objective was to explore effective methods of weed detection or related technology used by other organizations which the HCC could explore and replicate in their own monitoring program. Other councils around Hume were given special consideration, since their agriculture landscape is likely to be similar to that of Hume, and connections between the HCC officers and surrounding governments had been helpful in obtaining tools and information. We particularly focused on Agencies that use AI-based methods in their monitoring programs.

To meet this objective, we performed detailed case study research on surrounding governments' such as Whittlesea Council AI-based methods, along with the current two year program in place at Central Queensland University. We specifically explored drone camera specifications (e.g. what drone model they use, which spectral bands are used, etc.), image sampling protocols for training and testing, the software and algorithms used for training models and weed classification, and how effective their models are. Our team had tried reaching out to these programs and had gotten in touch with some of the representatives briefly, but was unable to set up a zoom meeting. We also obtained administrative details regarding any legal barriers, staff and expertise needed to manage the program, and the overall purpose of their specific program.

Our group analyzed the different methods used by these programs and assessed which methods were the most promising and practical to test, given the regulations and resources determined from objective one. If it was realistically available, we tested off-the-shelf applications that have been utilized by other organizations or solo practitioners. To better understand the methodologies along with their steps of implementation, we interviewed a member of Whittlesea Council in an informal manner resulting in better conversation flow and more time to explore third party softwares which allowed us to gain specialized information to help implement them in Hume.

Objective Three: Construct a data sampling protocol, and collect training data for the AI weed detection models

Our third objective was to collect and process field imagery to be used as input data for the ML models we trained. To complete this objective, we performed geographic sampling, which involved capturing field imagery across various locations around Hume along with any needed metadata (e.g. location) using drone images from the DJI Air 2S drone provided by Hume City Council. Along with our sampling, we documented different environmental conditions of our field days.

Before we did this, it was important to first develop a sampling protocol to eliminate data bias and abide by the law. Our sampling locations were decided with respect to variety in weed density (e.g. no weeds present vs. low presence of weeds vs. high presence of weeds), the mechanism used to capture this data (e.g. drone camera, satellite imagery, etc.), the terrain the data collection occurred on, and the time of day and weather conditions under which to capture images (NIST/SEMATECH, 2012). Once enough imagery was collected, we cleaned, labeled, and separated it as necessary to construct a training, validation, and test dataset for the models we evaluated.

While working in agricultural fields, it is essential to prioritize personal protective equipment (PPE) as the presence of noxious weeds can pose various health risks. The lands we have explored contain potentially hazardous nature factors such as invasive plant species that pose health risks through contact or inhalation (Department of Health). Our specific PPE included wearing sturdy work boots, long socks and thick pants. An image of the group's PPE can be viewed in Figure 6. The combination of this set of PPE kept us safe from any sharp plants as well as any harmful wildlife, such as a venomous snake. By prioritizing the use of PPE, individuals can safeguard their health and well-being while navigating through the Australian environments.

Figure 6. Group Members Kimberly Huang(left) and Mia Gilmore(right) wearing proper PPE

Figure 7. Agricultural lands

A typical day out in the field would consist of HCC representatives flying the drone with the guidance of WPI team member Vivek Voleti. WPI team member, Jake Bowen, was in charge of keeping track of the parameters of the testing day in an excel sheet. WPI team members Kimberly Huang and Mia Gilmore were in charge of documenting environmental conditions and surveying the land by foot. The documented environmental conditions led to a set list of parameters such as light conditions, terrain type, weed sparsity, and more

Date of Field Day

Time of day

Light Level

Wind Speed

Image collecting height

Sample Type

Objective Four: Train and test AI models on collected data, and identify strengths, weaknesses, and next steps

Our last objective was to test all trained models on imagery not used to train the model, and evaluate the models to find which performs the best. To accomplish this, we established a testing dataset, by separating a portion of the collected imagery from objective three to use specifically for testing. The selected models were then tested in a controlled field experiment. The controlables we dictated while collecting data include selecting the terrain type, the density of the weeds in the area, the time period of data collection, and the height at which the drone was flown, while the wind speed, and level of sunlight were conditions that were considered as well.

Our models were employed against various plots of land not included in the training set, and we analyzed each model's strengths, weaknesses, and ways to improve. This process allowed us to compare the results of our trained ML models and determine whether or not an AI-based approach was an effective approach for weed detection. If so, Hume City Council would be able to improve on our base models and eventually use them to extract locations of detected weeds. We planned for this standardized procedure to form the basis of the AI-based weed monitoring program

Findings

We have obtained data through various sources, including drone imagery and input from Hume City Council employees. Our findings have consistently reinforced the notion that noxious weeds have a detrimental impact on the ecosystem and stakeholders.

Social context in relation to the noxious weeds in Hume

In our investigation into the social context surrounding noxious weeds in Hume, we explored two key areas. Firstly, we explored HCC's previous approaches concerning weed detection. Secondly, we navigated the complexities of local laws and drone regulations to better understand the impacts they had on our research.

Hume City Council's resources and goals

Part of our methodology to achieve our goal was to conduct a focus group to understand HCC's status with monitoring the noxious weeds. Prior to the on-site research term, we were aware of the existence of the Weed Compliance program but was unsure if it had reached its enforcement stage. WPI team members conducted this focus group on a virtual call with several HCC officers from different departments ranging from conservation reserves to rural roadside departments. From this focus group, we determined that HCC essentially does not have an actionable weed monitoring process - the previous monitoring process included field officers going out by land and recording by hand the presence of weeds. The officers claimed that the process has been extremely "time consuming and laborious, especially with some properties being 100 hectares." The officers made it known that at this point they are simply looking for an efficient method for detecting weeds on a property and that the next step will be to establish a method to determine the percentage of noxious weeds to native fauna on a property.

Effects of weeds on landowners

Objective one started to understand the reason why we were conducting this project in the first place and to see how landowners are affected by these weeds. Throughout the term we had casual conversations with one of our main sponsors in which we were able to learn more about the presence of noxious weeds through his many connections through landowners and other councils. We learned that lands that are not maintained are burdens to their neighboring lands. With the winds and movement from animals, weeds scatter across land fences and end up on the maintained lands. Figure 8 shows weeds that had blown from one property to another. When weeds are not kept up with, they become a financial burden for those landowners who treat their weeds. The noxious weed problem is also affecting landowners mentally leaving them to feel like they should give up on their situation. Essentially, the Weed Compliance Program is not only helping the native Australian flora but is also helping to lower weed costs for landowners.

Figure 8a&b. Weeds blown to a maintained farm

Laws and regulations for intaking drone footage

As part of the first objective, one of the first tasks that the group had to conduct was to understand local laws and regulations for intaking drone footage. This was an important research aspect of our project as everything we were doing in the fields had to follow privacy laws. On the first day of working, the team was provided access to Hume City Council's team drive which included legal specifications for field officers as well as drone flight regulations.

The Local Government Act 1989 gives council officers the authority to enter land or buildings to enforce relevant laws. This act essentially outlines the procedures and responsibilities of authorized officers from the council. This act explains officers' enforcement of Acts, regulations and local laws. Essentially, the Act specifically details the requirements for identity cards, powers of authorized officers, penalties for non-compliance and provisions allowing police officers to enforce specific local law upon notification by the Council. Relating to the Weed Compliance Program, this Act establishes the framework for the appointment and responsibilities of authorized officers by local councils. This act grants authorized officers the authority to enter land or buildings within the municipal district to enforce relevant laws. Also regarding the Weed Compliance Program, this act highlights the fact that local councils may use drones as a tool for monitoring and identifying weed infestations on private properties. Essentially, authorized officers may employ drones for aerial surveillance to detect weeds and ensure compliance with local regulations regarding weed control. This act does not go into depth about the drone regulations that would need to be followed to ensure that the use of the drone aligns with legal requirements and respects the rights and privacy of property owners. To gather information about drone regulations, we spoke to the HCC officer who currently has a drone license and is authorized to fly drones over private property.

The Civil Aviation Safety Authority (CASA) establishes guidelines for piloting drones in Australia with the intention to keep the users and surroundings safe on the ground and in the air. The main law regarding our drone footage that we had to consider was that the drone could not exceed 120 meters from the ground and could not be flown 30 meters within people. This led to our group flying the drone at heights from 20 to 40 meters in the air, and at least 30 meters within the flight team. With the fact that Hume is located around an airport, we had to keep in mind the “no flight” zones. The law regarding airports is that the drone must stay at least 5.5 kilometers away from these airports. To re-enforce this law, we checked out property locations with websites that show the no flight zones, OpenSky and ok2fly. Our HCC drone officer also operated the drone with a CASA approved app which showed the areas where we were not allowed to fly the drone due to aviation legislation. Another main law that our group followed was that only personnel with a CASA drone permit were authorized to operate the drone. With the cooperation and availability of HCC’s resources, we had a licensed drone operator working at HCC fly the drone for our footage.

The State of Automated Weed Monitoring in Australia

AI based weed mapping technology has arrived on the Australian environment conservation scene. The technology of AI based weed mapping systems is still relatively new to the area. However, we explored various released technologies and programs from private organizations and universities that developed (or are developing) AI weed mapping systems

From the time we had begun preparatory work for this project to the time we arrived in Australia, there have been multiple private firms that have released technology to both map the weeds and even spray the weeds using AI. One of the most accurate o f these systems was released by SingleShot. They are a contract-based service that goes to various properties with high end industrial drones equipped with ~$15,000 thermal imaging cameras. They then fly over the property in a grid-like fashion (similar to figure 10). During the flight, the AI can identify and map the locations of weeds. Subsequently, the drone, equipped with a weed sprayer, autonomously flies to the locations where weeds were detected and sprays them ("spot spraying"), or it sends the data to a tractor. Since SingleShot and similar companies are private, their methods are proprietary, making it difficult for municipalities to learn the details of their innovations. Furthermore, their services can be extremely costly, especially for the amount of land Hume City Council wants to survey.

The academic sector is currently witnessing numerous breakthroughs in the realm of drone AI weed mapping, with several institutions making notable strides. Two of these institutions that have similar programs are the University of Queensland (UQ) and Central Queensland University (CQU). Both universities have teams of professors and students, mostly with a background in computer science related fields. Since their projects are still in development, we were not able to obtain many specifics regarding their methodology. However, from released information, we found that their drones are equipped with thermal imaging cameras, which may make weed detection easier for AI models.

Figure 9. Example SingleShot Drone

During the group’s time in Australia, a company named 2pi Software gave a demo to HCC employees about a weed mapping software named WeedRemeed. This software presented an AI-based solution for weed detection at scale that requires no training data. By uploading a set of images and providing a color profile of the targeted weed, WeedRemeed runs clustering algorithms across the provided images and selects groups of pixels corresponding to the given color profile. Since clustering is an unsupervised machine learning algorithm, no training data is needed (Xu et al., 2015). The demo presented scalability as another advantage of WeedRemeed; since processing would be performed on cloud servers, the processing time and resources for HCC would meaningfully decrease. However, their clustering approach would only match groups of pixels that matched a certain color, meaning it would likely miss weeds with colors similar to the native vegetation. We ultimately determined that their software would be ideal in situations where a weed significantly pops out of its environment, but was likely to fall short if weeds were to blend into their environment (e.g. Serrated Tussock).

City councils serve as the foremost avenue for disseminating information in the locality, constituting the primary means of information sharing within the area. Often, if one city council is facing a problem that another council has either faced or is familiar with, there will be collaboration between both councils to arrive at a solution. Our contact at the HCC sustainability team put us into contact with a member of Whittlesea Council - a neighboring council to Hume. During the meeting, we discussed the capabilities of the drone model Hume has, advantages and disadvantages regarding the capabilities of ArcGIS and QGIS, and additional softwares such as One3d and Pix4d which may perform image preprocessing more efficiently than in ArcGIS. We discuss these softwares in more detail in section 5.2.

Data sampling protocol and captured imagery

While collecting data via drone imagery, we made note of various environmental conditions that may have affected our data collection process. The conditions we observed included the time of day, brightness, wind speeds, terrain type, and an approximate ratio of local to noxious flora in the sample. It was important to try to vary these parameters to provide different data samples for the AI model to learn from, and to identify the model’s weaknesses.

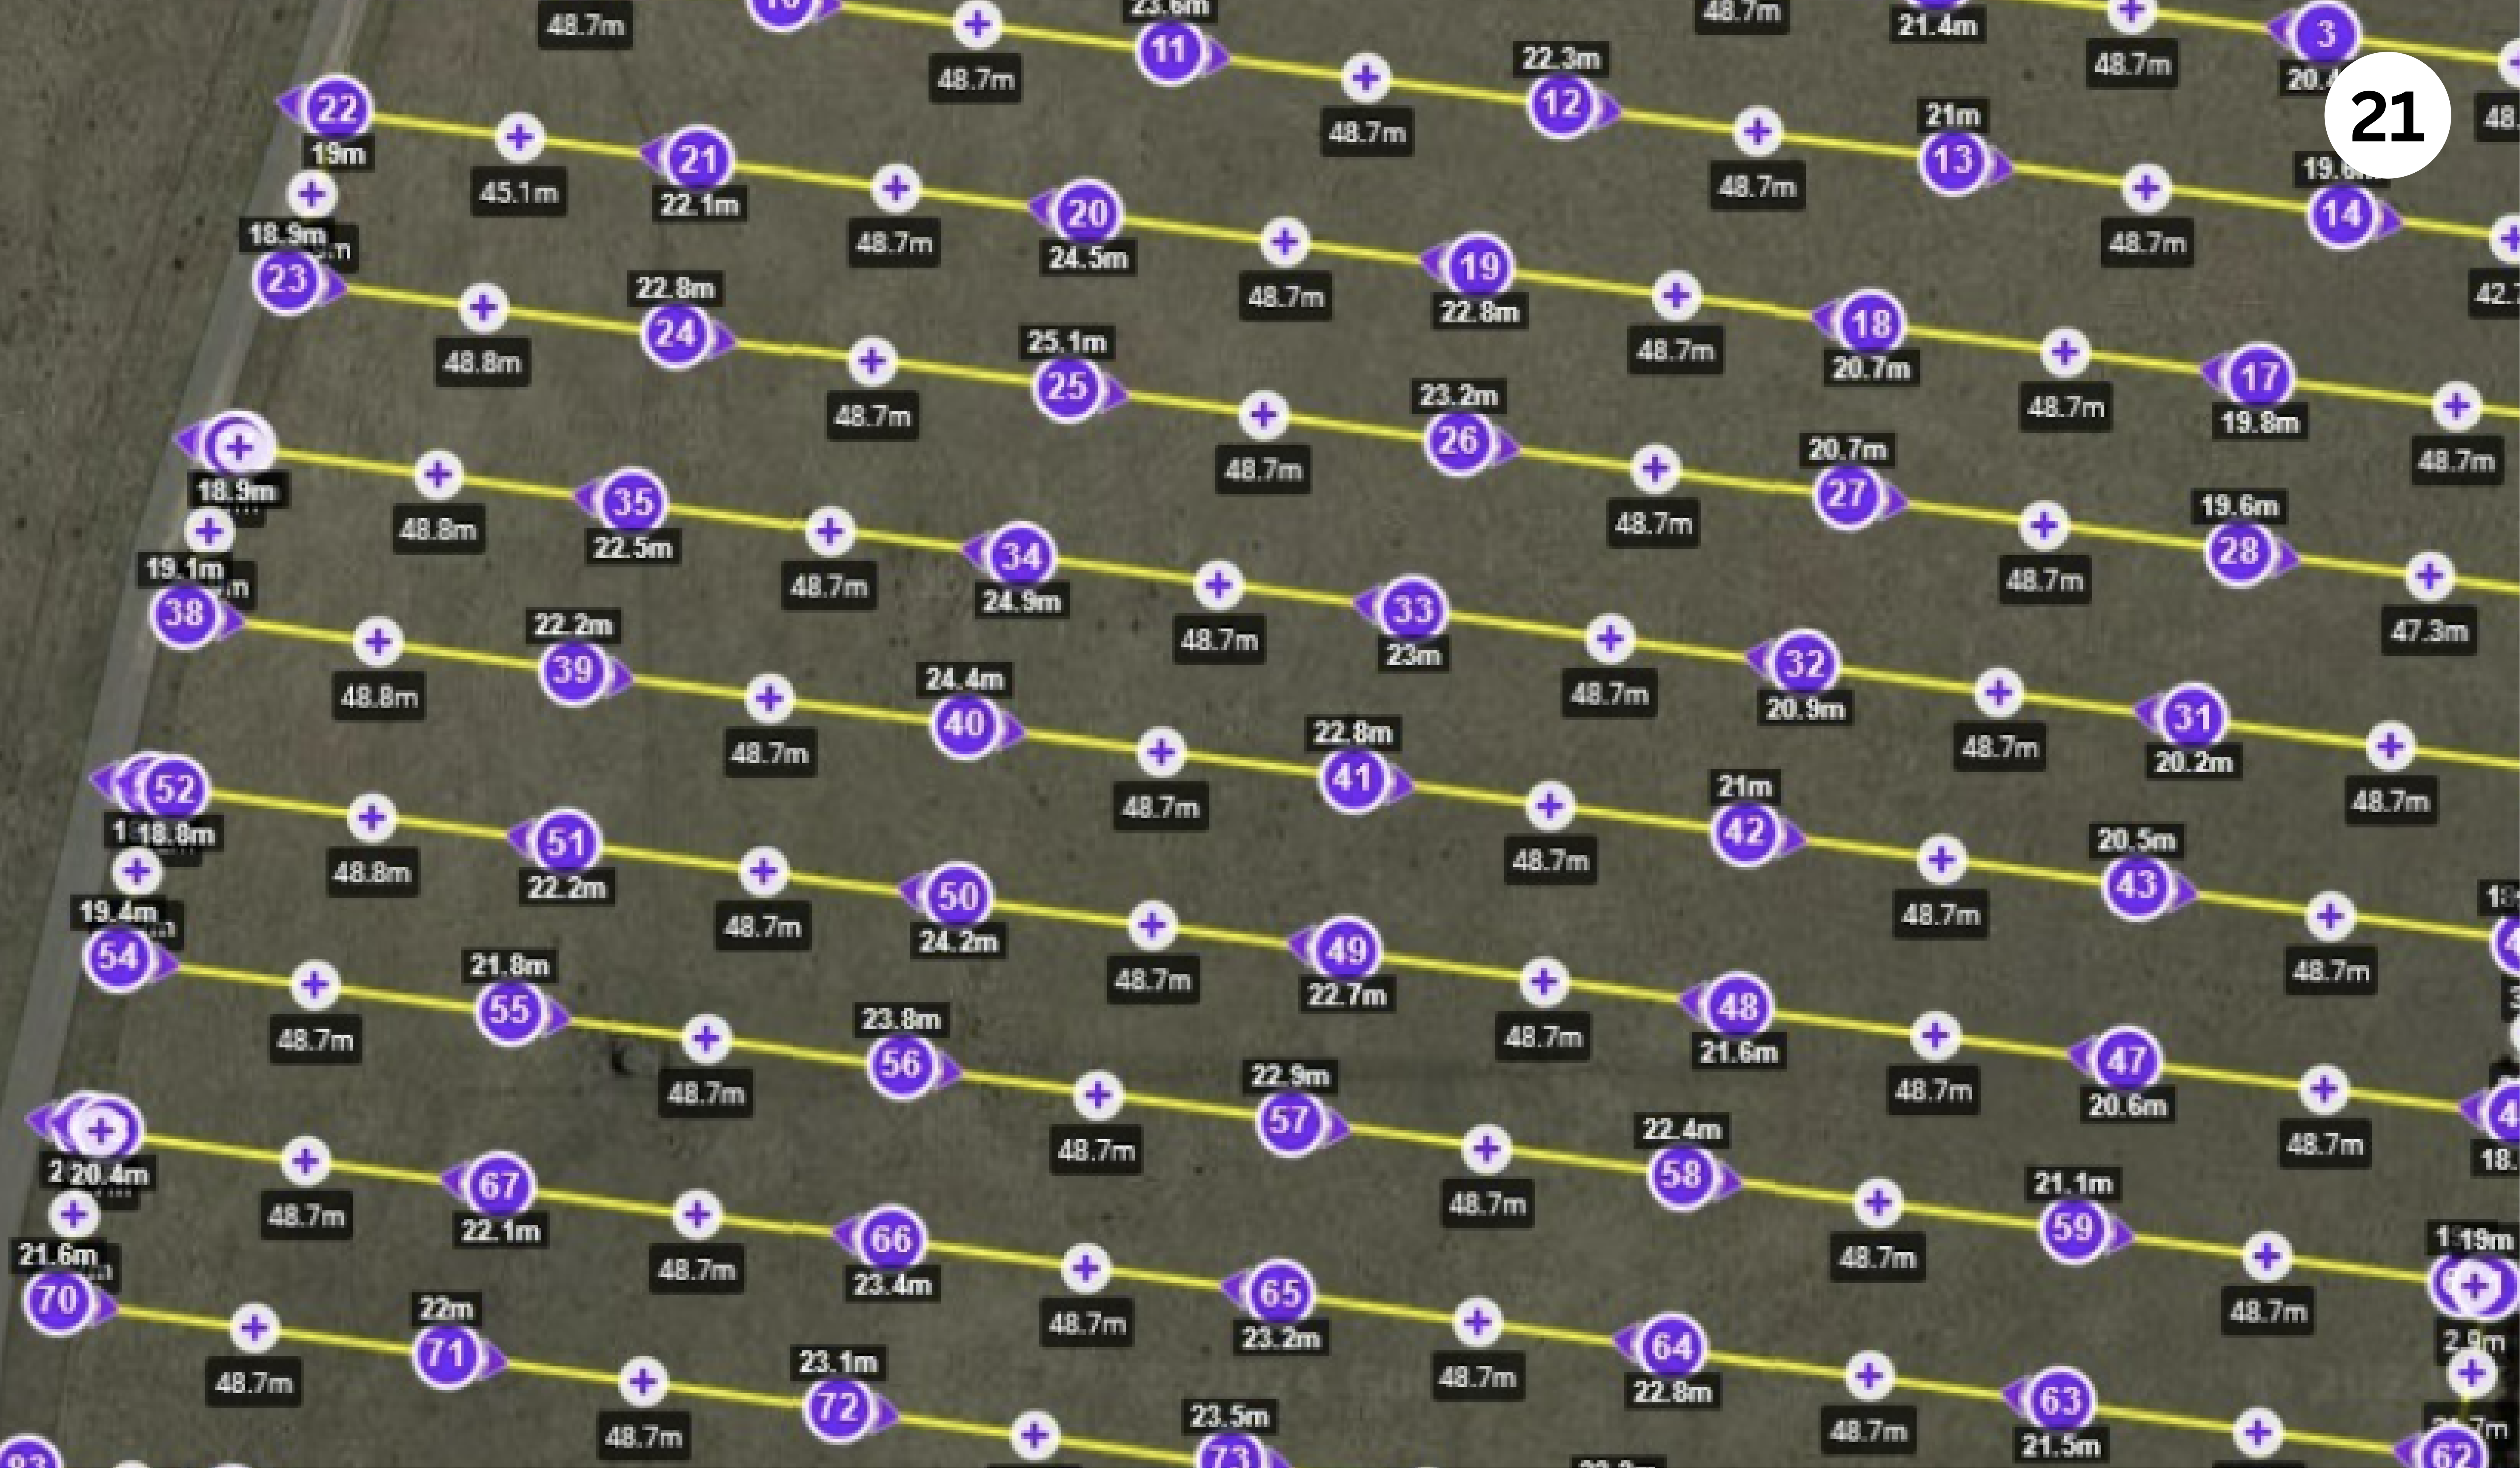

Figure 10. An example flight path (partially cut off to eliminate location identifying information)

We found that the most effective method of imagery collection on HCC's available drone (DJI Air 2S) was to use the built-in flight planning capability. This tool allowed for us to highlight a geographic area, enter height, camera gimbal angle, and image overlap percentage parameters, and the DJI software calculated the exact drone's flight path to collect images with the requested conditions (Figure 10). This setup allowed for the drone to take images at regularly programmed intervals, which increased the effectiveness of orthomosaic generation in ArcGIS (more in section 4.4) as compared to when drone flight and image collection was done manually by the drone operator.

We went out into the field for data collection a total of three times after all practical considerations were taken into account (e.g. temperatures in the range of 90°F to 100°F increase risk of fires in the field, which prevented data collection on hot days). Across all of our data collection days, we varied parameters (aside from the natural weather changes) to get a variety of data (Table 1). We collected imagery of land that contained pure vegetation with a low quantity of weeds present, land that had a mix of both Artichoke Thistle and Serrated Tussock, and we varied the terrain type (flatland vs hills) as well (Table 2).

Figure 11.
No Weeds Present

Figure 12.
Thistle and Tussock

Figure 13. Slope surface with Thistle

Figure 14.
Serrated Tussock

Figure 15. Native and Noxious Tussock

Date	Time of Day	Light Level	Wind Speed	Image Collection Height	Sample Type	Terrain Type
22.1.24	10AM-12PM	Bright sun	24 kph	20m	50/50 artichoke thistle to serrated tussock	Mostly flat
6.2.24	11AM-1PM	Slightly cloudy	32 kph	20m & 40m	60/40 native grass to artichoke thistle	Hilly
20.2.24	1:15PM-2:15PM	Bright sun	21-24 kph	20m	70/30 serrated tussock to native grass	Hilly

Table 1. Table of environmental conditions observed each field day.

Image	Description	Height (m)	Camera Angle (°)
Figure 11	Flat pasture with almost no weeds present	20	90
Figure 12	Flat pasture with a mix of Thistle and Tussock	20	90
Figure 13	Sloped surface containing Thistle (far in the background; may be difficult to see)	20, 40	90
Figure 14	Flat pasture with mainly Tussock	20	90, 85, 80
Figure 15	Sloped surface with a mix of Tussock and a similar-looking native grass: Poa Labillardieri	20	90, 85, 80

Table 2. Table of descriptions and flight parameters for each area sampled

Data processing, model training, and evaluation

Orthomosaic Generation

The first important stage in our workflow was to create an "orthomosaic" of all the images taken - one large image composed of all the drone imagery "stitched" together overlaid onto a map. We used the ArcGIS OrthoMapping software in an attempt to keep the entire image processing workflow within the same software for ease of use for HCC. While functional, however, we found that the ArcGIS orthomosaic generation was slow and processor intensive. It took approximately 35-40 minutes to stitch together a series of 230 images on a laptop with an Intel Core i7 11th generation processor. The 230 images only covered a small portion of the property, meaning it may take longer to create the orthomosaic if HCC was to survey a larger plot of land in the future.

Figure 16. An orthomosaic created using ArcGIS (yellow = image centers; orange = drone flight path)

Dataset Creation

Regardless of the time taken, creating an orthomosaic allowed us to visualize the drone imagery as one contiguous unit overlayed on a map. Using ArcGIS' dataset creation wizard, we then could create several labeled datasets of both Artichoke Thistle and Serrated Tussock. It is important to note that these datasets were several iterations of (approximately) the same data, since different ML model architectures offered by ArcGIS required different metadata formats. The final dataset created was composed of 200 training "chips" (a subimage of the map selected by the user to be a training example for a model) of each Artichoke Thistle and Serrated Tussock taken from property b) (Table 2 and figure 11-15).

Separating the data into distinct training and testing sets was crucial for accurately evaluating model accuracy. If we were to use the same set of images to test the model as we did to train it, model accuracy would falsely appear higher than reality since it was trained on this data. By reserving some images exclusively for testing, we ensured that we could measure model accuracy on imagery it has never seen before, thereby testing the model's ability to generalize learned patterns as opposed to its ability to "memorize" its training set.

Trained Models and Results

Using the labeled data, we trained several models in ArcGIS. We initially experimented with using the RetinaNet object detection model (Lin et al., 2017). Training took about 9 hours on an Intel Core i7 11th gen processor (no GPU acceleration), with a dataset of 140 examples of each weed type. Running inferencing on the original training imagery took about 1.5 hours, Although achieving decent accuracy, some weed instances were still missed, suggesting the need for more training data and fine-tuning for improvement.

We eventually transitioned to using the YOLOv3 model (Redmon et al., 2016), since it generally has faster training and inference speed while maintaining good accuracy. Using this architecture, we trained two separate models (one for Artichoke Thistle, and one for Serrated Tussock) to streamline model fine-tuning in the future. After training one YOLOv3 model on 200 examples of Artichoke Thistle and the other on 200 examples of Serrated Tussock, training time was reduced to 1-1.5 hours per model. YOLOv3 offered faster training and inference with generally good accuracy. However, it can benefit from additional data and fine tuning to address missed weeds and misclassified non-weed vegetation as weeds.

While the team did not have sufficient time to rigorously ground truth the images and measure accuracy metrics for each model, figures 17-24 qualitatively detail the models' strengths and weaknesses across four properties, and give insight about each models' potential in a larger scale weed detection setting. These properties are labeled as "Property 1" through "Property 4", anonymized for landowner privacy. Figures 17-19 show each model's results on Property 1, which has a mix of Artichoke Thistle and Serrated Tussock. Figure 20 shows Property 2, which contains Artichoke Thistle on a sloped surface, which our initial models did not identify. Figures 21-22 show the model's predictions on Property 3, which has no observable weeds present. Finally, figures 23-24 show a Serrated Tussock infested Property 4, along with the model's outputs, It is important to note that any model errors labeled in the following images may not be comprehensive - they are only the errors that can be confidently identified as such.

Despite the errors present, our initial models generally perform well, especially considering the nature of their training data. Due to time constraints, all data in the training datasets were captured from one orthomosaic image. This likely resulted in both models being somewhat overfit to a particular appearance of their respective weed. For example, many of the Artichoke Thistle weeds in the dataset have a similar appearance given that they were on the same property experiencing the same weather and sunlight conditions. This is likely one of the main reasons why the models may misclassify weeds, miss them entirely, or hallucinate nonexistent weeds. We hypothesize that the likeliest solution would be to collect a wider variety of training data, with varying lighting conditions and weeds at different stages in their lifecycle.

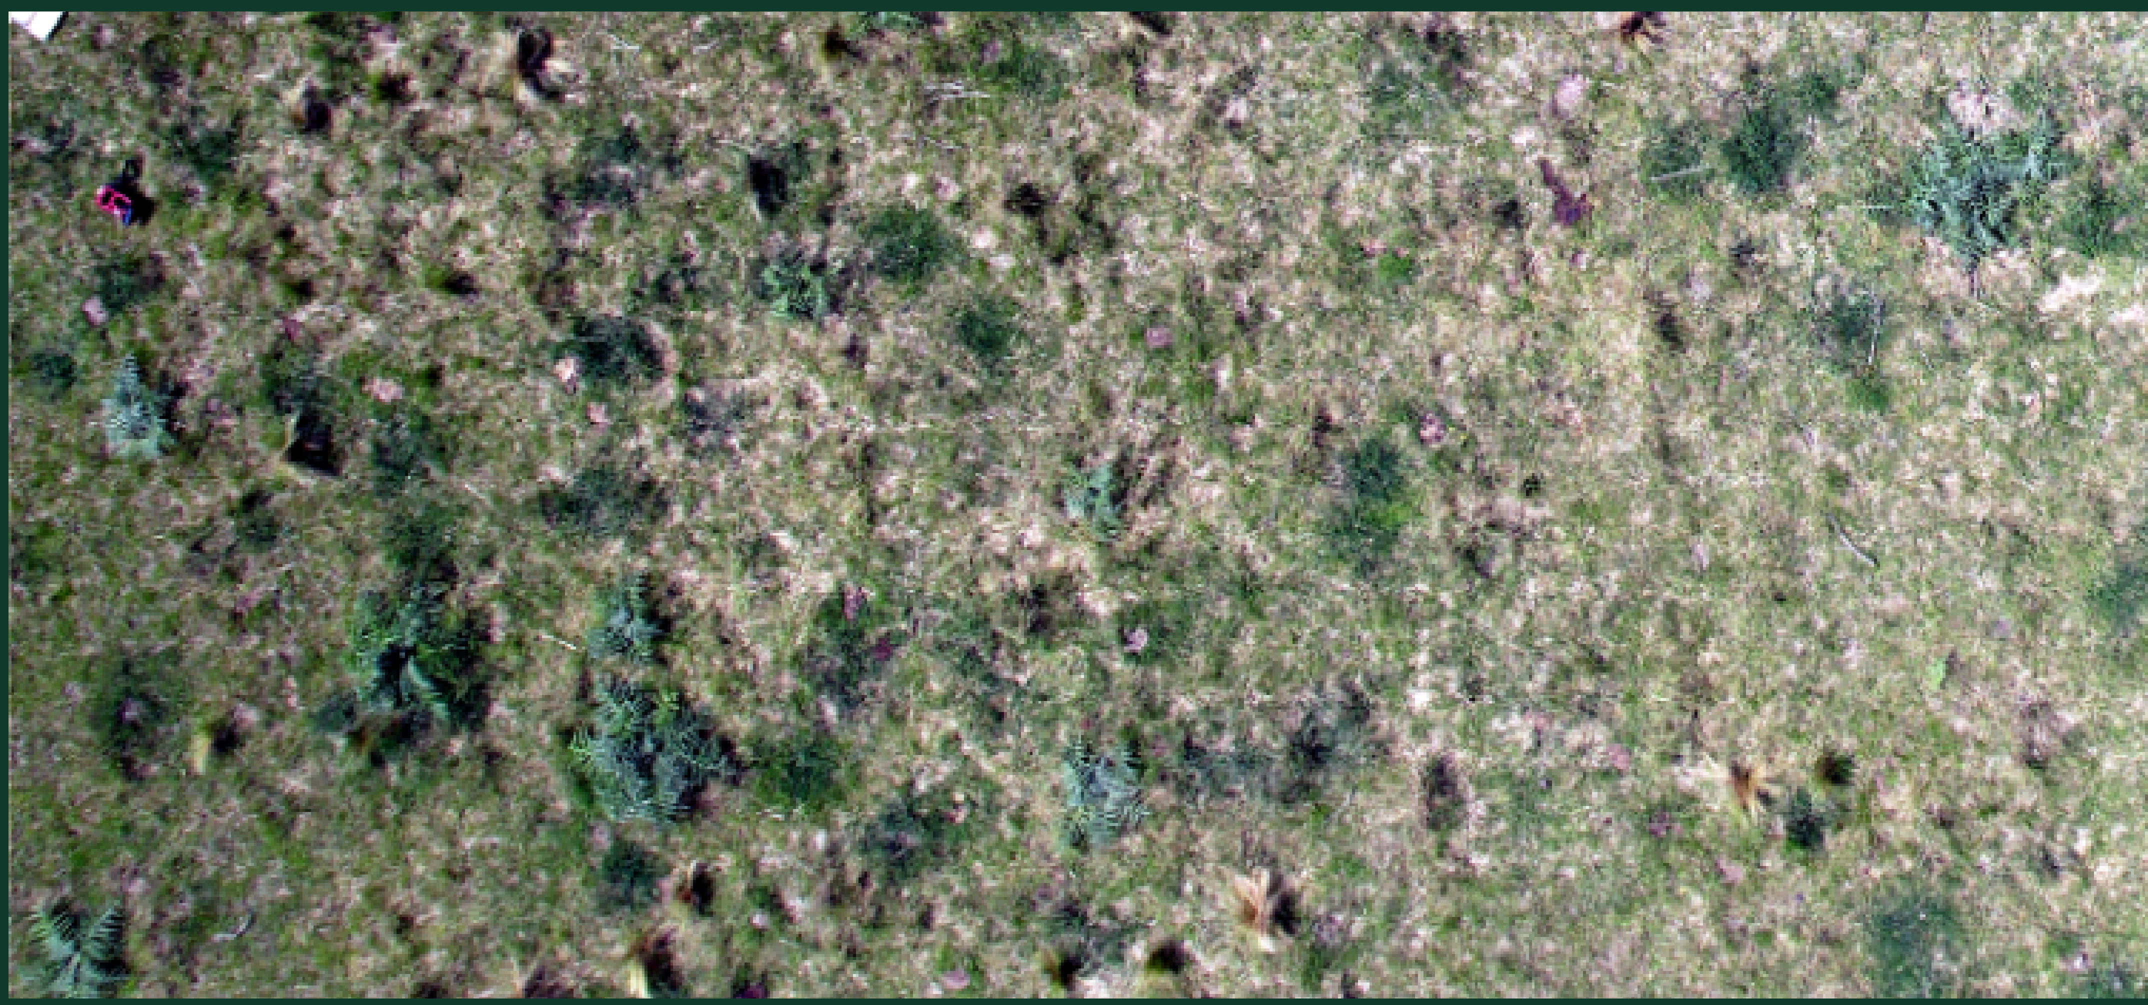

Figure 17. Property 1, which has a mixed amount of Artichoke Thistle and Serrated Tussock

Figure 18. Property 1 with the model's predicted Artichoke Thistle outputs (light green) and missed weeds (red outlines)

Figure 19. Property 1 with the model's predicted Serrated Tussock outputs, (pink) with missed weeds (outlined in red boxes) and vegetation that was mistaken as Tussock (dark red)

Figure 20. Sloped surface on Property 2 with Artichoke Thistle (light green spots). The model did not identify any Thistle with these lighting conditions.

Figure 21. Property 3 (no weeds present) with the model's incorrect Artichoke Thistle predictions outlined

Figure 22. Property 3 (no weeds present) with the model's incorrect Serrated Tussock predictions outlined

Figure 23. Property 4 infested with significant amounts of overlapping Serrated Tussock. Any black spots surrounded by beige vegetation are likely patches of Tussock.

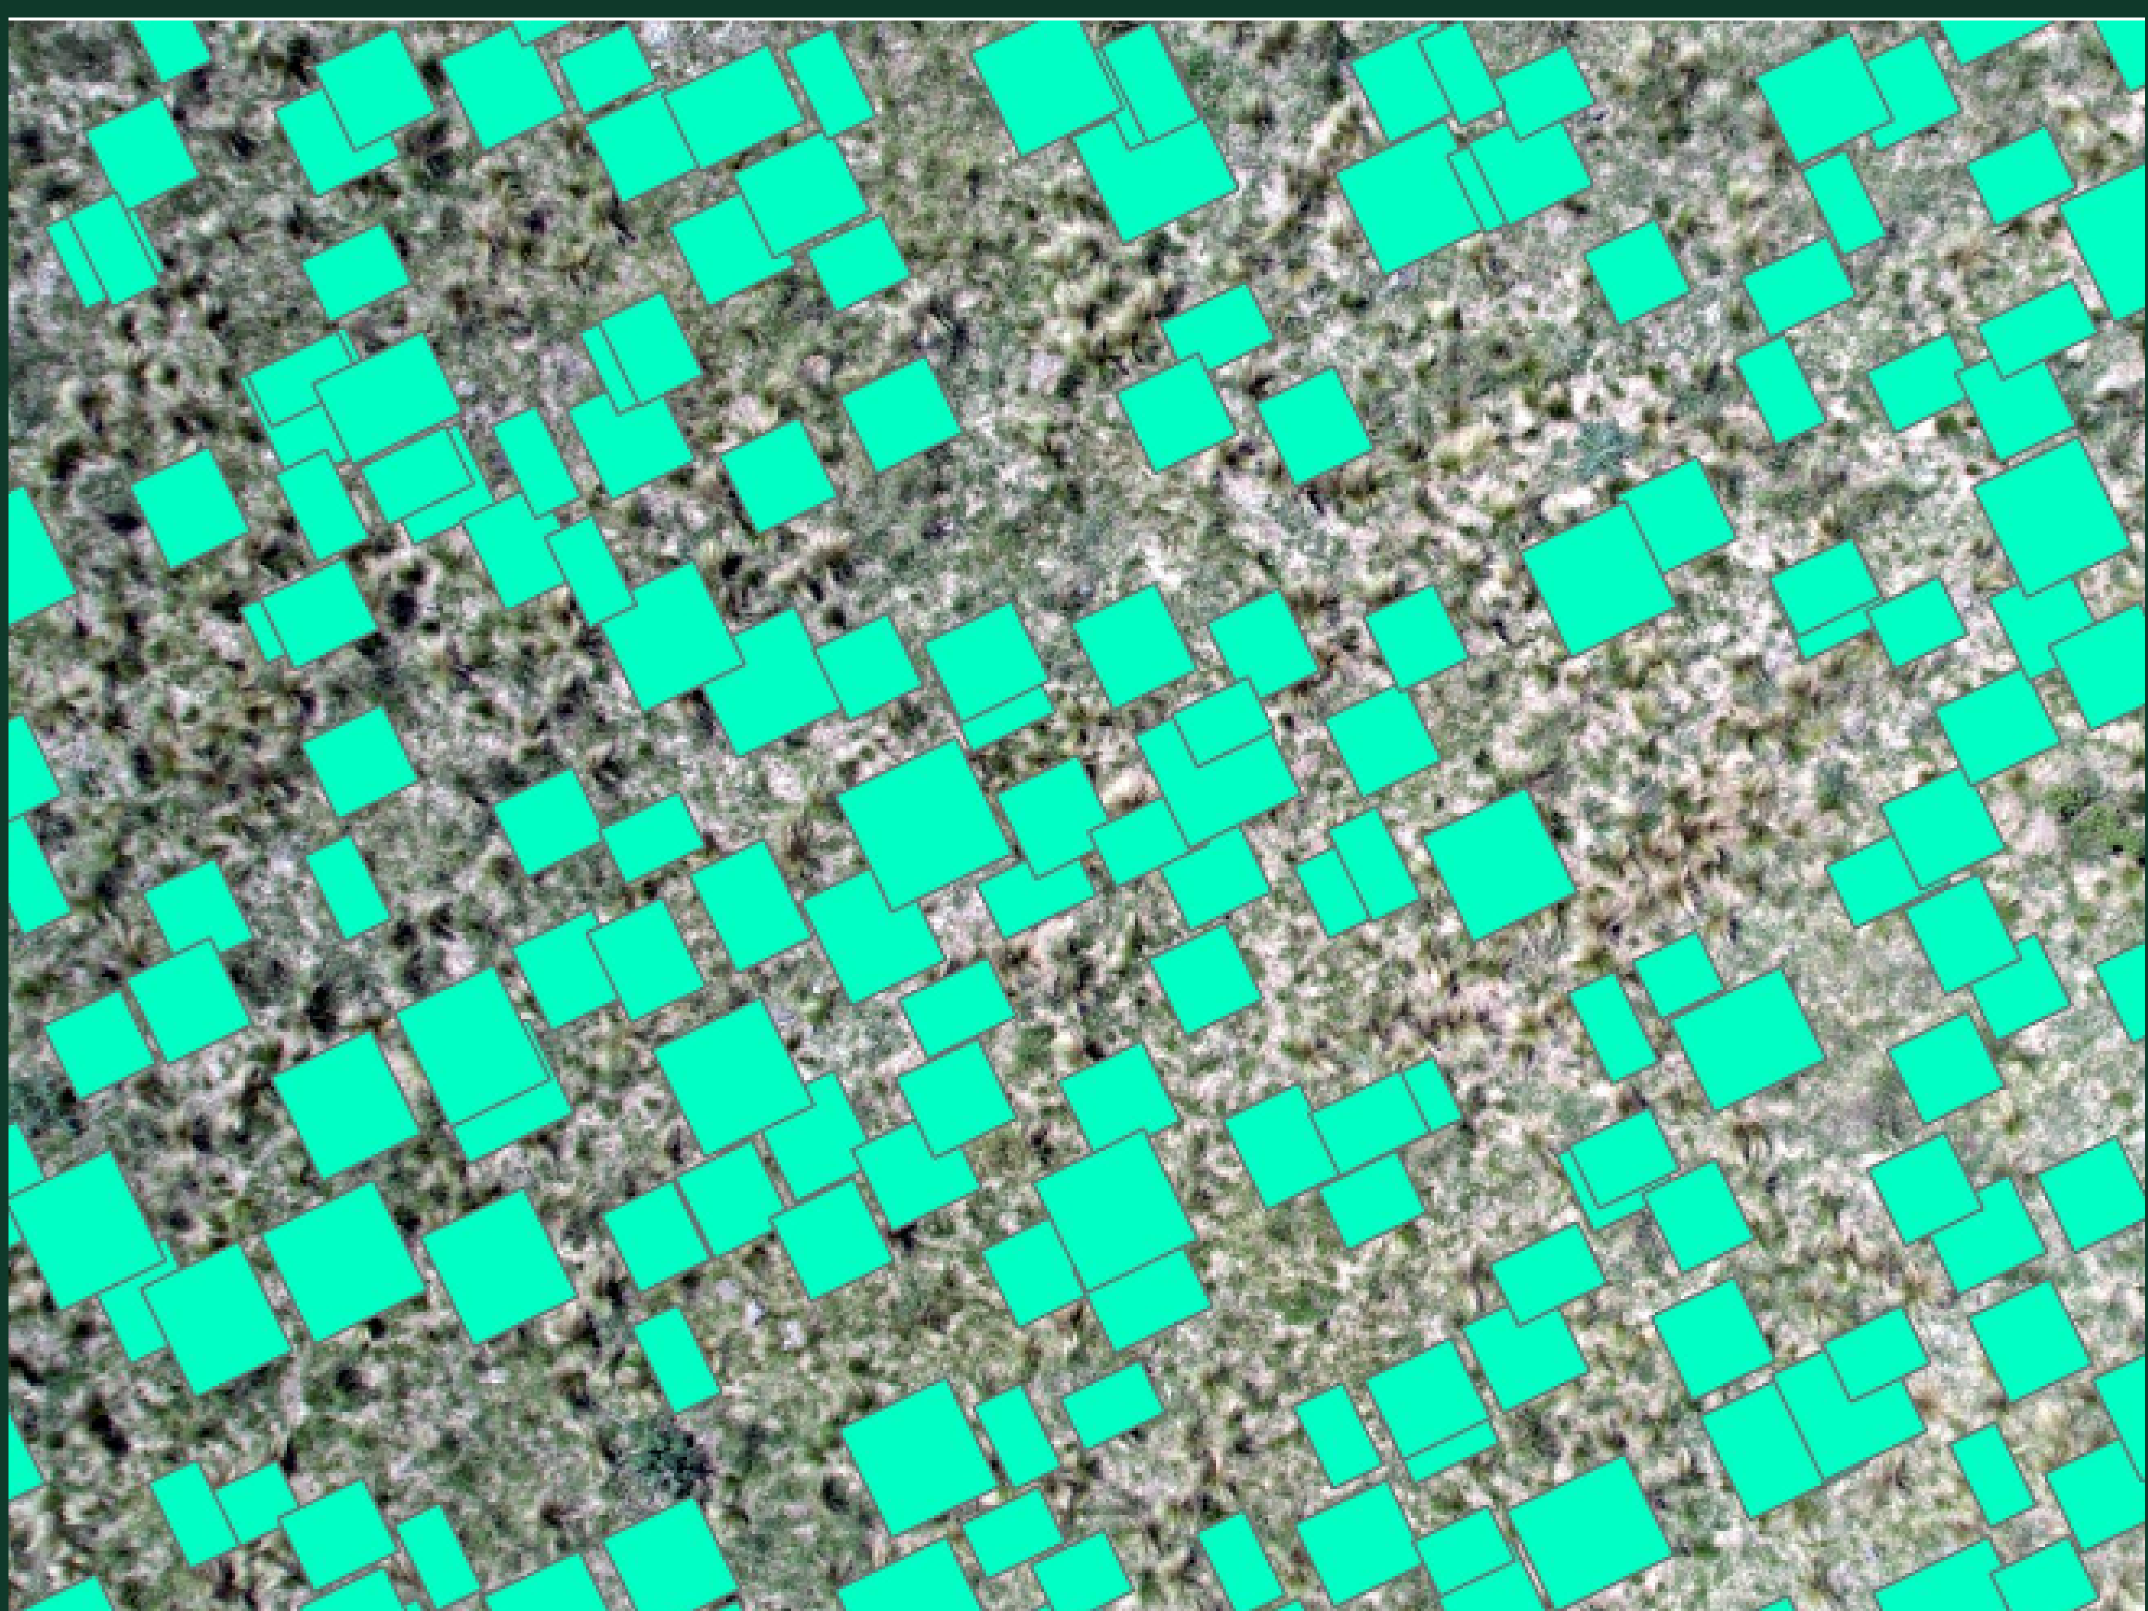

Figure 24. Property 4 with the predicted Serrated Tussock locations.

Nevertheless, from a qualitative standpoint, our initial models clearly capture the general trend of weed infestation across the various properties. Property 3, which contained no weeds, yielded only a few model inferences for both Artichoke Thistle and Serrated Tussock. Property 1 contained a moderate presence of both weeds, and the amount of identified weeds from each model also indicates this when compared to Property 3. On Property 4, which was characterized by a significant Tussock infestation, the Serrated Tussock model’s predicted outputs covered a large portion of the image. The model did miss some overlapping patches of Tussock, but the model very clearly conveyed that this patch of vegetation contained large amounts of Tussock. In general, even the initial models show a rough correlation between the quantity of weeds in a set of images and the model's ability to infer weed quantity. With further fine-tuning of our models with additional data, we hypothesize that AI-based tools can be effective for weed detection at scale.

Aside from model performance itself, we found that the biggest other limiting factor is processing time and compute space. All model training and inferencing was done on an Intel Core i7 11th generation processor. The training time with the final YOLOv3 models was not excessive with a small to moderately sized dataset of 200 images per dataset (1-1.5 hours per model), but inferencing took significantly longer. For example, inferencing on the entirety of Property 2 (approximately 70,000 square meters) took ~8 hours. Scaling this to large properties may not be feasible without parallelization of processing across separate computers, or potentially even dedicated server compute space with Graphics Processing Units (GPUs), which machine learning models generally benefit from.

Recommendations

This section concludes our findings by outlining various recommendations and next steps for Hume City Council to implement this protocol. We provide both non-technical and technical recommendations to consider.

Non-Technical

Based upon our experience with completing this project, we recommend that landowners are interviewed during this process. A landowner interview would be beneficial to gain an understanding of the social context regarding the situation of noxious weeds. These invasive plants impose an economic burden on landowners while also wreaking havoc on local flora and fauna. We hypothesize that landowners may be apprehensive about capturing aerial imagery using drones, with uncertainties regarding the privacy or public accessibility of these images. We recommend interviewing landowners (Appendix B) to examine how well understood the Weed Compliance Program is, and landowners thoughts on data privacy.

Additionally, there may be growing concern regarding landowner compliance following the implementation of AI. Previously, Hume City Council's engagement with landowners regarding weeds, which could be followed by enforcement if there is no compliance, may encourage landowners to resist compliance. We therefore recommend that the council takes an educational approach when communicating with landowners. HCC officers can encourage landowners to control their noxious weeds by educating them on the impacts of weeds on other crops on their land, animals that may graze on the weeds, and the economic burden that their unmaintained weeds put on themselves and their neighbors.

Machine learning is difficult, and even the most advanced models have errors. For this reason, it is important to follow up on sparse weeds detected on properties. While AI decreases the need for human supervision, we recommend that human supervision is used when the weeds detected on properties are scant, due to the possibility of error. If AI is to be deployed to enforce the Weed Compliance Program, we further recommend that AI use is communicated with landowners, and to provide a method to contact HCC in case of error. This will help illustrate the models' weaknesses, and possibly provide training data to improve performance.

Technical

Due to time constraints, the team was not able to fully explore all hardware and software possibilities outside of ArcGIS. In this section, we detail some supplemental hardware and software to replace the hardware and software used in this project. HCC may substitute some of the team's workflow as they see fit, but the high level process of: (1) collect imagery; (2) create orthomosaic; (3) label training data; (4) train model; and (5) evaluate model can stay roughly the same.

Regarding capturing imagery, the team only had access to drones with an RGB camera, which did not capture complete spectral information. Considering that our initial models showed promise with only three spectral bands (red, green, and blue), we hypothesize that multispectral multispectral cameras may significantly increase model accuracy. While expensive, multispectral cameras capture spectral bands invisible to the human eye (e.g. infrared and ultraviolet), which will add more features for the models to learn from.

Alternatively, HCC may also find multispectral satellite imagery useful. Rasmussen et al. (2021) show that Sentinel-2 satellite imagery was not a reliable detector of Cirsium Arvense (Creeping Thistle, similar in appearance to Artichoke Thistle), and they thus determined that image resolution is a crucial factor for automatic weed detection. At its finest resolution, Sentinel-2 only achieves a coarse 10x10 meter spatial resolution, but additional satellites exist to capture much finer imagery. Though these are paid services, we recommend that HCC explores commercial satellite imagery to significantly cut down on time spent collecting data. While many options exist, the most important factors are to have a very high resolution (preferably 0.5 meters or less) and spectral bands beyond what the human eye can capture (a good starting point for satellites that meet these criteria is Maxar's WorldView series).

With ArcGIS orthomosaic capabilities being somewhat volatile and very slow, we recommend exploring supplemental software to create orthomosaic images. Three potential options include Pix4D, One3D, and 3DF Zephyr. These softwares will not change the procedure of training a machine learning model, but would serve as an expedited way to generate an orthomosaic image. If training and running inferences with machine learning models become too resource heavy, QGIS is another potential software to explore. Being an open-source alternative to ArcGIS, QGIS may be a more cost efficient and lightweight tool for image analysis. However, HCC may incur extra overhead if switching to QGIS, as much of the functionality does not come out-of-the-box as it does with ArcGIS, and the council may need to search for custom plugins from the QGIS community, or write them from scratch. Table 3 summarizes the various softwares discussed and their approximate prices

Software	Purpose	Approximate Price
Pix4D	Orthomosaic reconstruction	~$292.67 AUD per month
One3D	Orthomosaic reconstruction	First 2500 images free; 0.01€ per image
3DF Zephyr	Orthomosaic reconstruction	€199.00 for Lite (minimum capability for HCC's applications)
WebODM	Orthomosaic reconstruction	~$229.48 AUD one time
QGIS	Image processing, deep learning, and geographic analysis	Free

Table 3. Summary of supplemental software to explore

Variation in data collection is important to help machine learning models generalize. Collecting data in different types of weather (sunny, cloudy, rainy, etc.) will provide different views of each weed that will allow the machine learning model to more accurately identify weeds across various settings. Additionally, with weeds having different appearances during different stages of their life cycle, it is important to capture data across weed lifespan so that machine learning models can more easily generalize that particular weed's appearance.

In addition to varying the data itself, it may also be useful to further explore manipulating drone parameters - specifically drone height and camera gimbal angles. At different drone heights, weeds may differ in size and clarity, which may result in the model not identifying them. Exploring the models' limits at various heights may be informative regarding model limits. While the team did not have sufficient time to explore height, we did begin exploring altering the camera gimbal angle. Serrated Tussock is specifically hard to detect from a top down angle, so we hypothesize that the model can more easily identify the weed with an angled camera. Figure 25-27 shows Tussock from a straight down camera view, a 10° tilt off-nadir, and (approximately) a 45° tilt off-nadir. With the higher off-nadir angles, a side profile of the Tussock becomes more visible, which may give the model more spatial information about the weed. Note: the 45° image was not explicitly captured by angling the drone camera. This sub image was taken from the edge of a larger image captured by the drone with the camera facing straight down. Given the field of view of the DJI Air 2S camera, we estimate that the drone captured these patches of Tussock at approximately 45°.

Figure 25. Serrated Tussock viewed from a straight down angle

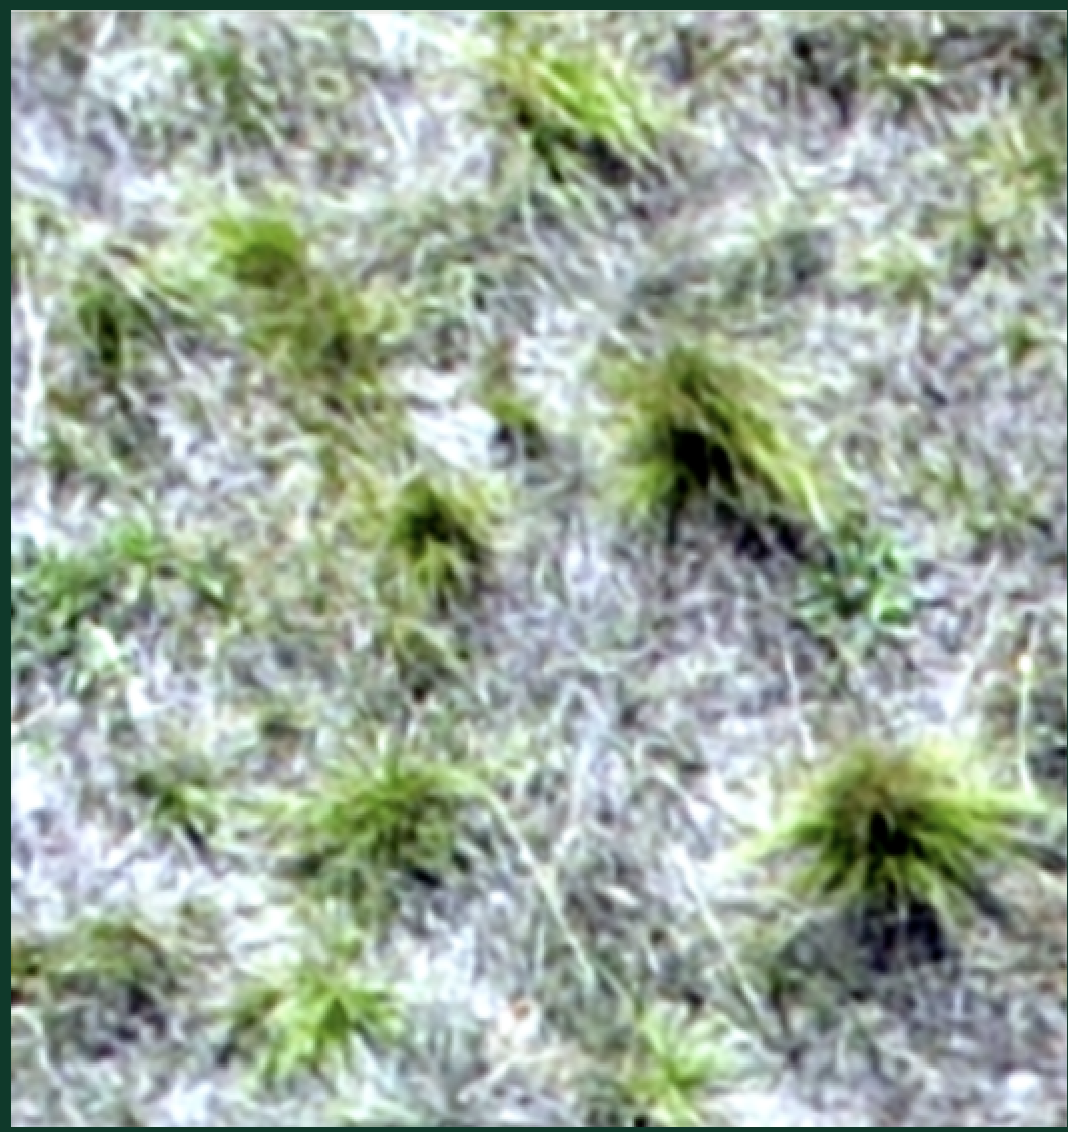

Figure 26. a 10° tilt off-nadir

Figure 27. a ~45° tilt off-nadir

The final set of recommendations aim to cut down on time spent processing to help scale our AI-based protocol to larger areas of Hume. When processing images, we recommend splitting processes across different computers where possible. For example, if multiple computers are each running analyses on separate properties, this has the potential to significantly cut down on the time spent training and running models. HCC may also explore cloud computing options; while it is difficult to obtain a price estimate, running machine learning processes on an external server with specialized hardware - namely GPUs - will also cut down on processing time. Purchasing any of the proposed orthomosaic creation softwares (One3d, Pix4d, WebODM, etc.) effectively does this for the orthomosaic generation, which is a part of the price for these softwares.

Conclusion

The presence of noxious weeds in Hume municipality is progressively increasing and becoming more problematic. After conducting focus groups with employees of Hume City Council, and multiple field days to gather data regarding the prevalence of noxious weeds, we have concluded that these invasive species are detrimental to the health of both the local flora and fauna, as well as the economy. Motivated by this knowledge, we started development of AI technology to identify and flag noxious weeds using drone footage we collected. Given our initial results and recommendations to continue our work, our protocol can enable Hume City Council to detect and map noxious weeds on a large scale. Ultimately, we hope this work allows the council to enforce their weed control strategies and adapt them as needed aiding in their efforts to combat noxious weeds, and we hope that our AI-based protocol may be of use to municipalities across Victoria for both agricultural and (potentially) non-agricultural purposes.

Figure 28. Group Photo, taken by drone

References

Adams, P. (2017). Resistant weeds: Battling Australia's multi-billion dollar weed problem: An international program based in Germany is researching new herbicides to combat resistant strains of weeds, and an engineering solution conceived by a WA farmer is attracting interest from around the world. Ultimo: Australian Broadcasting Corporation. Retrieved from http://ezproxy.wpi.edu/login?url=https://www.proquest.com/other-sources/resistant-weeds-battling-australias-multi-billion/docview/2121355114/se-2

Agriculture Victoria. (n.d.). Victorian Weed Spotters. https://www.agriculture.gov.au/sites/default/files/documents/weed-spotters-victoria.pdf

Benjamin, C. (2023, March 24). Wayne Zeisemer, Bongeen Qld. WeedSmart. https://www.weedsmart.org.au/content/wayne-zeisemer-bongeen-qld/

BONNIE, H. H. (1998, Jul 28). Thorns in the side; despite botanists' efforts to control the foreign invader, artichoke thistle is pushing out native species and spreading: [orange county edition]. Los Angeles Times, pp. 1. Retrieved from http://ezproxy.wpi.edu/login?url=https://www.proquest.com/newspapers/thorns-side-despite-botanists-efforts-control/docview/421485796/se-2

Bullock, D., Mangeni, A., Wiesner-Hanks, T., DeChant, C., Stewart, E., Kaczmar, C., Kolkman, J., Nelson, R., Gore, M., & Lipson., H. (2019). Automated Weed Detection in Aerial Imagery with Context. arXiv. https://arxiv.org/ftp/arxiv/papers/1910/1910.00652.pdf

Centre for Invasive Species Solutions. (2023). Bringing AI smarts to tackling Australia’s priority weeds. WeedScan. https://weedscan.org.au/#:~:text=WeedScan%20is%20powered%20by%20a,and%20natural%20resource%20management%20officers.

Combatting chilean needle grass. (2021, Dec 16). *Glen Innes Examiner,* pp. 0. Retrieved from http://ezproxy.wpi.edu/login?url=https://www.proquest.com/newspapers/combatting-chilean-needle-grass/docview/2610097356/se-2

Commonwealth of Australia, (2017). Australian Weeds Strategy 2017-2027. Retrieved from: https://www.agriculture.gov.au/sites/default/files/sitecollectiondocuments/pests-diseases-weeds/consultation/aws-final.pdf

CQUniversity, (2023). CQUs IOT-augmented drones are taking war on weeds to the skies: Research with impact. https://www.cqu.edu.au/research/impact/cqus-iot-augmented-drones-are-taking-war-on-weeds-to-the-skies.

Cullen, J. M., W. A. Palmer, and A. W. Sheppard. “Biological Control of Weeds in Australia: The Last 120 Years.” *Austral entomology* 62.2 (2023): 133–148. Web.

Department of Health. Victoria, A. (n.d.). *Pesticide use and personal protective equipment*. Retrieved January 30, 2024, from https://www.health.vic.gov.au/environmental-health/pesticide-use-and-personal-protective-equipment

Department of Jobs, Precincts and Regions. (2023, October 26). *State prohibited weeds*. Agriculture Victoria. https://agriculture.vic.gov.au/biosecurity/weeds/state-prohibited-weeds

He, K., Zhang, X., Ren, S. & Sun, J. (2016). Deep Residual Learning for Image Recognition. Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (p./pp. 770--778), June, : IEEE. https://arxiv.org/pdf/1512.03385.pdf

Hume City Council. Implementation of the local law for control of noxious weeds on private property. (2023). Hume City Council. https://docs.google.com/document/d/1RXkUcocKYeMcnkEgJ_FLkcNv-wIMcS_u/.

Hume City Council. Privacy policy. (2021). Hume City Council https://www.hume.vic.gov.au/Your-Council/Governance/Council-Plans-Reports-and-Policies/Council-Policies/Privacy-Policy.

Hume City Council . (n.d.). *Weeds, vegetation and pests*. Greater Hume Council. https://www.greaterhume.nsw.gov.au/Living-in-Greater-Hume/Weeds-Vegetation-and-Pests

Li, Y., Al-Sarayreh, M., Irie, K., Hackell, D., Bourdot, G., Reis, M., & Ghamkhar, K. (2021). Identification of Weeds Based on Hyperspectral Imaging and Machine Learning. National Library of Medicine. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7868399/

Lin, T.-Y., Goyal, P., Girshick, R. B., He, K., & Dollár, P. (2017). Focal Loss for Dense Object Detection. CoRR, abs/1708.02002. Retrieved from http://arxiv.org/abs/1708.02002

Llewellyn, R., Ronning, D., Clarke, M., Mayfield, A., Walker, S. & Ouzman, J.(2016) Impact of weeds on Australian grain production: the cost of weeds to Australian grain growers and the adoption of weed man-agement and tillage practices. Report for GRDC, CSIRO, Australia.

Machine learning, explained. (2021). Massachusetts Institute of Technology. https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained

Maher, J., Stringham, O. C., Moncayo, S., Wood, L., Lassaline, C. R., Virtue, J., et al. (2023). Weed wide web: Characterising illegal online trade of invasive plants in australia. *NeoBiota, 87*, 45-72. https://doi.org/10.3897/neobiota.87.104472

NIST/SEMATECH. (2012). E-handbook of statistical methods. National Institute of Science and Technology. https://www.itl.nist.gov/div898/handbook/ppc/section3/ppc33.htm

Oregon Department of Agriculture. (2020). Invasive Noxious Weed Control Program Annual Report. https://www.oregon.gov/oda/shared/Documents/Publications/Weeds/NoxiousWeedProgramAnnualReport.pdf

Parikh, R., Mathai, A., Parikh, S., Chandra Sekhar, G., & Thomas, R. (2008). Understanding and using sensitivity, specificity and predictive values. Indian journal of ophthalmology, 56(1), 45–50. https://doi.org/10.4103/0301-4738.37595.

Raja, R., Nguyen, T., Slaughter, D., & Fennimore, S. (2020). Real-time weed-crop classification and localisation technique for robotic weed control in lettuce. Biosystems Engineering. https://www.sciencedirect.com/science/article/pii/S1537511020300428?via%3Dihub

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 779–788). https://doi.org/10.1109/CVPR.2016.91

Rural notebook | the prickly matter of chilean needle grass.(2022, Apr 29). *Western Advocate,* Retrieved from http://ezproxy.wpi.edu/login?url=https://www.proquest.com/newspapers/rural-notebook-prickly-matter-chilean-needle/docview/2656295443/se-2

Appendix A: Focus Group Questions for HCC Members Prior to Implementation of our Methodology

1. What is the current state of HCC’s Weed Compliance Program?
2. What current methods are being used for data collection?
3. Would you consider those methods to be either efficient or effective?
4. Is there data available from this program?
5. What results would you like to see out of a new weed management program?
6. Is there anything in the current program that you believe is important we implement into this new program?
7. What methods have you tried that didn't work?

Appendix B: Questions for Landowners

1. Are you aware of Hume City Council's noxious weed monitoring program and if so what do you know about it?
2. Are you aware of Hume City Council's right to monitor your property if there is suspicion of a noxious weed problem?
3. What are your current or future plans for your land? Ex: farming, livestock, renovations, real estate, etc.
4. Do you have any concerns involving privacy with the HCC weed management program?
5. What other concerns do you have about images of your land being collected?
6. Do you know if other landowners would have apprehension of land data being collected? If so, why might they be apprehensive?
7. Do you believe a fine for noncompliance is justifiable?