

Wildfire Fuel Mapping with Convolutional Neural Networks for Remote Automated Exposure Assessment

by

Liam Bennett

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

in

Civil Engineering (Cross-Disciplinary)

Department of Civil and Environmental Engineering
University of Alberta

© Liam Bennett, 2023

Abstract

While beneficial to the natural environment in many cases, wildfires become hazardous when they intersect with the built environment. As such, there is an ongoing effort to understand the fire environment, the fuels it contains, and the way that wildfire interacts with the built environment. In particular, communities and structures in the wildland urban interface (WUI) are often assessed to determine their exposure to wildfire ignition processes. Exposure assessment workflows at both scales require the identification and classification of hazard fuels. Thus, a method is presented for the remote automated detection of hazardous wildfire fuels and the application of fuel detections to community and structure scale exposure assessments. The outputs of the automated processes are intended to supplement rather than replace existing practices and form a preliminary basis of information that is both simple and rapid to collect, process, and interpret.

Two workflows are devised to detect and classify large overstory trees from RGB imagery in the boreal, rocky mountain, and foothill natural regions of Alberta, Canada. The first workflow considers remotely piloted aircraft systems (RPAS), are used to collect RGB imagery. A convolutional neural network (CNN) is trained to detect trees in the overstory and to classify them as coniferous, deciduous, or snags. F1-scores reach 74.5% for tree detection and achieves classification F1-scores of 97.3%, 94.4%, and 90.9% for coniferous, deciduous, and snag classes respectively.

The second workflow uses RGB satellite imagery to detect individual trees using a second trained CNN. An R^2 of 0.76 is achieved comparing automated tree detection density to manual annotation density. A k-means clustering algorithm is used to determine winter 'leaf-off' imagery and classify trees as 'green-in-winter' or 'brown-in-winter', an indication of coniferous or deciduous trees. Classifications from satellite imagery reach an F1-score of 0.82.

Finally, tree detections and classifications from the RPAS model are visualized around a structure in the context of a FireSmart home assessment, and it is discussed how the workflow could be used to provide informative maps of tree vegetation around a structure. At a community scale, fuel maps derived from the satellite tree detection and classification

workflow are used in an existing community exposure assessment workflow to explore how fuel maps generated in this manner may be applied.

Preface

Some of the research conducted in this work has been previously published or is currently being considered for publication.

The satellite algorithm assessment procedure, manual annotation of images, survey flights, and machine learning code was done in collaboration with members of the lab team, including past and present undergraduate research assistants. Parts of chapters 1.3, 1.4, and 2.0 have been previously published in Bennett et al. (2022). Parts of chapters 3.0 and 4.0 are currently being formatted for journal submission. Co-authors on these papers include Dr. Jeff Boisvert, Dr. Jen Beverly, Brandon Wilson, Scott Selland, Linna Qian, Matthew Wood, Hannah Zhao, Zony Wu, Radomir Wasowski, and Camilla Da Silva.

Acknowledgements

I would like to give thanks to the mentors and colleagues I have met and worked with throughout the course of my degree. My supervisor, Dr. Jeff Boisvert, his academic knowledge and experience as a paid-on-call firefighter, and the team of undergraduate interns who worked on the project in its various stages.

Thanks are also extended to Dr. Jen Beverly, whose knowledge of the wildfire environment proved extremely valuable throughout the project, and to Dr. Li Cheng, who provided input and expertise regarding the machine learning / computer vision aspect of the work.

I would also like to thank the partners and experts in wildfire response and management, including Steve Otway (Yellowhead County Fire Department). Expertise, insight, support, and collaboration from our numerous partners in wildfire response continues to drive this project forward.

Many of the RPAS surveys were performed in cooperation with various homeowners and property owners in rural Alberta, without which the work would not be possible. Gratitude is extended to all.

Finally, I thank my friends and family, whose interest in my work and support in so many ways made this work happen. Thank you all, so much.

This research was funded in part by Alberta Agriculture and Forestry through the Canadian Partnership for Wildland Fire Science, grant agreement number 18GRWMB06.

Table of Contents

1.0	Introduction.....	1
1.1	Problem Statement	1
1.2	Problem Importance.....	1
1.3	Vegetation Mapping	2
1.3.1	Wildfire Fuels.....	2
1.3.2	Remote Sensing and Vegetation.....	3
1.3.3	Computer Vision.....	5
1.4	Exposure Assessment	7
1.5	Thesis Objectives and Outline.....	9
2.0	Tree Detection and Classification from RPAS RGB Imagery	10
2.1	Data Acquisition	10
2.2	RPAS Model.....	15
2.3	RPAS Model Results.....	18
2.4	RPAS Model Discussion	24
3.0	Tree Detection and Classification from Satellite RGB Imagery	27
3.1	Data Acquisition	27
3.2	Satellite Model.....	30
3.3	Satellite Model Results.....	33
3.4	Individual Tree Classification.....	37
3.5	Satellite Model Discussion	46
4.0	Remote Exposure Assessments	50
4.1	Structure Exposure Assessments	50
4.1.1	Methodology	50
4.1.2	Discussion	51

4.2	Community Exposure Assessments.....	54
4.2.1	Methodology	54
4.2.2	Discussion	59
5.0	Conclusion.....	64
5.1	Summary.....	64
5.2	Contributions	64
5.3	Limitations.....	64
5.4	Future Work	66
6.0	References	68
7.0	Appendices.....	75
	Appendix A: RPAS Detection Results	76
	Appendix B: Satellite Training Tiles	77

List of Tables

Table 1: Summary of RPAS systems utilized in this study	10
Table 2: Summary of RPAS cameras utilized in this study.....	11
Table 3: RPAS flight parameters for M30T.....	12
Table 4: Training survey summary	15
Table 5: RPAS dataset summary	17
Table 6: Specifications of satellite surveys used in this study	30
Table 7: Manual annotation dataset summary in satellite imagery	32
Table 8: Satellite detection assessment results	36
Table 9: Satellite tree classification accuracy results	44
Table 10: Exposure classes from Beverly et al. (2010)	58

List of Figures

Figure 1: Survey outputs for a home in rural Alberta, including (a) orthomosaic, (b) 3D mesh, (c) 3D pointcloud, and (d) DEM	13
Figure 2: RPAS survey locations, including training surveys and surveyed communities	14
Figure 3: Graphical representation of IoU	19
Figure 4: Training metrics for YOLOv5, folds 1 through 5	21
Figure 5: Sample tree detections and classifications on new RPAS survey	22
Figure 6: Sample RPAS property survey and classified tree detections.....	23
Figure 7: Automated RPAS image collection and detection workflow	24
Figure 8: Satellite survey locations in Alberta, Canada	28
Figure 9: Satellite imagery near Seba Beach, Alberta. Summer (left) and winter (right)	29
Figure 10: Example of annotation tiles, in urban, semi-urban, and forested environments	31
Figure 11: Satellite model pretraining procedure.....	33
Figure 12: Example of how a prediction box placed mere pixels away can result in a low IoU.....	34
Figure 13: An area without trees (a) could still be assigned a tree density value if a large radius is used alone. The 3 m circle (b) is used to set this cell to 0 if no trees are found.	35
Figure 14: Assessment of manual and predicted boxes in satellite imagery	37
Figure 15: Green-in-winter and brown-in-winter trees, and associated average RGB values, in winter and summer imagery	40
Figure 16: R, G, B distributions for trees detected in Lobstick imagery	41
Figure 17: Results of k-means clustering algorithm across the R, G, and B dimensions	42

Figure 18: RPAS manual annotation boxes and matching satellite model detection boxes overlaid on RPAS survey imagery at an IoU threshold of 0.5.....	43
Figure 19: Results of k-means classification at Seba Beach, AB, showing (a) summer imagery, (b) winter imagery, (c) detections classified as brown-in-winter, and (d) detections classified as green-in-winter.....	45
Figure 20: Fuel map comparison, including (a) summer imagery, (b) winter imagery, (c) Alberta 50 m fuel map (A. Beaudoin et al., 2014), and (d) satellite model detections ...	46
Figure 21: Tree detected in (a) summer imagery at high angle of incidence; and (b) tree detection box overlaid on winter imagery at different angle of incidence includes background pixels	48
Figure 22: RPAS model automated tree detections, colored by distance to inhabited structure based on FireSmart Alberta (2015) home assessment guidelines	52
Figure 23: Satellite image and fuel map near Lobstick, AB	55
Figure 24: Satellite tree detection boxes	56
Figure 25: Satellite classification results at Lobstick, AB, showing (a) summer imagery, (b) winter imagery, (c) detections classified as ‘deciduous’, and (d) detections classified as ‘coniferous’	56
Figure 26: Long-range spotting distance for a cell in the community	57
Figure 27: Remote community long-range ember exposure assessment, generated using satellite model detections	59
Figure 28: Remote community exposure assessment workflow.....	60
Figure 29: Tree count (left) and green-in-winter proportion (right) maps aggregated from individual tree detections.....	61

List of Abbreviations

Abbreviation	Term
CNN	Convolutional Neural Network
DEM	Digital Elevation Model
DSM	Digital Surface Model
DTM	Digital Terrain Model
FBP	Fire Behavior Prediction
GPS	Global Positioning System
GSD	Ground Sample Distance
IoU	Intersection over Union
LiDAR	Light Detection and Ranging
ML	Machine Learnings
RGB	Red-Green-Blue
ROS	Rate of Spread
RPAS	Remotely Piloted Aircraft System
RS	Remote Sensing
SAR	Synthetic Aperture Radar
SGD	Stochastic Gradient Descent
VAR	Value at Risk
WUI	Wildland-Urban Interface

List of Symbols

Symbol	Meaning
A_o	Area of Overlap
A_u	Area of Union
f	Focal Length
FN	False Negative
FP	False Positive
H	Flight Height
I_w	Image Width
S_w	Sensor Width
TP	True Positive

1.0 Introduction

1.1 Problem Statement

While wildfires are a natural part of the landscape, problems can arise when they intersect with human activity and the built environment. A number of strategies can be employed by communities to prepare for and mitigate the negative effects of wildfires and increase community resilience. This can include wildfire ignition exposure assessments at both the community and structure scale. This study explores leveraging technology in remote sensing (RS) and machine learning (ML) to supplement existing exposure assessment workflows with information that is automated, high-resolution, recent, and interpretable.

The work takes place within Alberta, Canada, including both the boreal, rocky mountain, and foothill natural regions. These regions are relevant to the work performed as all contain coniferous tree stands (Downing & Pettapiece, 2006) and experience wildfire; in particular, a recent study found that from 1959-2021 wildfires in the boreal region accounted for 87.7% of burned area in Alberta (Ahmed & Hassan, 2023). Models in this thesis are trained to detect and classify trees into broad categories of coniferous / deciduous rather than individual tree species using data within these regions. Application outside of these regions would require further study.

1.2 Problem Importance

Wildfires are natural events that are essential to maintain the health of boreal forests. However, wildfires can cause damage to infrastructure, property, and threaten the health and safety of community members when they intersect with human activity. Wildfires can cost the communities they intersect with in terms of insurance, community planning, wildfire response, wildfire preparedness, and in disaster recovery.

At the time of this thesis in 2023, news outlets speculate that this year Alberta could break its record area burned by wildfires, with Alberta Wildfire estimating over 1.4 million burned hectares so far (Mulcahy, 2023). In Alberta, \$3.5 billion CAD in insured losses were caused by the 2016 Fort McMurray wildfire, which also required the evacuation of 88,000 persons from the community and surrounding area (Mamuji & Rozdilsky, 2019). A study conducted post-fire revealed that, consistent with previous studies after other wildfires,

the majority of home ignitions were likely due to ember transmission from the wildfire to structures (Westhaver, 2017). As such, the investigation of new methods of fuel mapping may be of interest.

1.3 Vegetation Mapping

1.3.1 Wildfire Fuels

Wildfire behavior can be said to be governed by three factors: weather, topography, and fuels. This is referred to as the 'wildfire behaviour triangle'; these three factors interact and drive wildfire behavior (Countryman, 1972). Of these three factors, only fuel can be realistically modified by humans.

No fire can burn without fuel. The fuel a wildfire consumes can influence many of its properties. However, quantifying fuel is a difficult task. There are many properties that influence a fuel's propensity to burn, and more still that affect its behavior once it is burning. Fuel properties can vary on a small spatial scale, but information about those properties is required across large spatial extents - Brandt et al. (2013) estimate the boreal zone covers 552 million hectares in Canada, and Johnston & Flannigan (2018) estimate 32.3 million hectares of wildland-urban interface (WUI). Here, focus will be placed on 'crown fuels' – fires that burn in tree crowns are high intensity fires and can generate significant embers when accompanied by high winds (Van Wagner, 1983). In Canada, crown fires are generally exclusive to coniferous forests (Van Wagner, 1977). Thus, in order to quantify wildfire exposure as well as aspects of wildfire spread, fuel maps are required.

In Alberta, the Alberta Vegetation Inventory (AVI) exists to collect information about vegetation across the province (Forest Stewardship and Trade Branch, 2022). AVI data is generated through manual interpretation of aerial imagery by certified technicians, as well as a suite of ancillary data. Here, the minimum collection resolution for aerial imagery is specified to be 0.3 m / px. Outputs include manually delineated polygons containing a variety of attributes, for example tree species composition, tree crown closure, and tree height. Over the years, AVI has produced major inventories under evolving specifications, including AVI version 1.0 (1988), version 2.0 (1990), version 2.1 (1991), version 2.1.1

(2007), and the current version 2.1.5. Vegetation data generated by the AVI are then reorganized and used in a decision tree procedure to generate wildfire fuel maps across the province (Frederick, 2012).

Fuel maps are used to characterize the ‘fuel’ aspect of the wildfire behaviour triangle across the landscape. In Canada, the Fire Behaviour Prediction (FBP) system defines 5 groups of 16 fuel types (Forestry Canada Fire Danger Group, 1992). These groups consist of coniferous, deciduous, mixedwood, slash, and open fuels. In general, fuels are categorized into these groups based on species, form, size, arrangement, and continuity (Forestry Canada Fire Danger Group, 1992). These fuel maps can serve as inputs to FBP models such as Prometheus. Prometheus is an FBP system which uses topography, fuel types, and weather to simulate fire spread using a wave propagation model (Tymstra et al., 2010). Models such as Prometheus allow for predictions of fire growth, intensity, and rate of spread, allowing for the coordination of suppression activities. Fuel maps are also utilized when assessing community exposure to wildfire, which can in turn be used for planning fuel reduction treatments, creating firelines / firebreaks along vulnerable parts of the WUI, and prioritizing individual structures for risk reduction (Beverly et al., 2018).

In summary, fuel is the only aspect of the wildfire triangle that can be readily modified by humans, and thus it is of interest to map fuels across large spatial extents.

1.3.2 Remote Sensing and Vegetation

As there is a clear need for fuel maps across large spatial extents, the derivation of fuel attributes from RS data has been explored. There is considerable research into using RS for measuring vegetation attributes that would be impractical or impossible to densely sample across the domains they are needed. Satellites, aerial vehicles, and remotely piloted aircraft systems (RPAS) now offer data products that can be used for vegetation monitoring, ranging from high density pointclouds generated using light detection and ranging (LiDAR), to hyperspectral and multispectral imagery. Several methods have been explored for the extraction of vegetation and forestry characteristics from this data.

As RPAS become increasingly more accessible, interest in using these platforms for RS grows. To date, numerous works have been published on extracting forestry

characteristics from data collected via RPAS with a range of sensors. RGB imagery has been considered in Bennett et al. (2022), who use high-resolution RGB imagery collected via RPAS to train a convolutional neural network (CNN) to detect and classify trees as coniferous or deciduous, with F1-scores for tree detection averaging 72%. Mohan et al. (2017) performed similar work for the detection of trees using RPAS imagery; though in this work RGB imagery was used to generate a structure-from-motion (SfM) model of canopy height and a local maxima algorithm was used to detect trees. This was thus applied only to open canopy forests, where tree canopies do not overlap. Surovy & Kuželka (2019) provide a lengthy review on the use of RS for forest attributes, including satellite, aerial, and RPAS data. Authors review how both imagery and laser scanning (LiDAR) are used on RPAS platforms for plantation monitoring, individual tree segmentation, forest and tree height assessments, forest inventory, species classification, and forest health. Egli & Hopke (2020) collect imagery via RPAS in Germany and apply a CNN for species classification. The resulting model proved to be 92% accurate for classifying trees photographed at 1.6 cm / px resolution; authors also note the affordability and ease of collecting imagery in this manner.

Other sensors beyond RGB cameras are also common RPAS payloads. Aerial LiDAR has been used in forest inventory and vegetation mapping. Cameron et al. (2021) use aerial LiDAR over forest stands in Alberta to extract tree attributes relevant to wildfire behavior. Crown bulk density, crown fuel load, stem density, canopy height, and crown base height were predicted using the resulting data. It was found that all forest attributes, except for crown base height, had R^2 values greater than 0.81 compared to ground measurements when mapped at 40 m² resolution. RPAS-borne LiDAR was used by Rodriguez-Puerta et al. (2021) to detect individual trees and measure their height on tree plantations. RPAS-borne multispectral imagery was used by Gallardo-Salazar & Pompa-Garcıa (2020) to gather multiple tree characteristics, including height, area, crown diameter, and estimates of basal diameter, and diameter at breast height. Hyperspectral cameras have also been mounted on RPAS and used for this purpose: Nevalainen et al. (2017) use such a platform to detect and classify individual trees, and score detection accuracies between 40% and 95%, and an overall classification F1-score of 0.93. Together, the applicability of RPAS

and aerial RS platforms for forest inventory is apparent, with accurate results being achieved with low-cost RGB imagery, and with higher-cost multispectral, hyperspectral, and LiDAR sensors.

Numerous studies also exist investigating the extraction of forest attributes from satellite collected RS data. Surovy & Kuželka (2019) review the use of satellite data for stratification, plantation monitoring, forest height, inventory attributes, species classification, and forest health. Shtanchaev et al. (2021) use multispectral satellite imagery to classify tree crowns as birch, spruce, pine, and fir, and find that previously delineated tree crowns can be classified with F1-scores of up to 0.69. Classification using multi-spectral imagery collected via satellite is also considered by Immitzer et al. (2012), who used 8 spectral bands to classify trees into 10 species classes using a random forest algorithm with 82% classification accuracy. Synthetic Aperture Radar (SAR) has also been used, in combination with optical imagery, to produce maps of forest height in Finland, reaching an RMSE of 1.68 m (Ge et al., 2022). In Canada, Beaudoin et al. (2022) utilize a k-NN method to map forest attributes using a suite of input data, including multispectral satellite imagery, SAR, and environmental variables such as topography and climate indices. This work builds on previous work (Beaudoin et al. 2014), which utilized a similar method to map 127 forest variables including above ground biomass, species compositions, and landcover. This paper is of particular interest as it is the work cited by NRCAN as the basis for the national-scale FBP maps (Natural Resources Canada, n.d.), which are derived in part utilizing the forest attributes mapped by Beaudoin et al. (2014).

1.3.3 Computer Vision

With remote sensing platforms offering rapid collection of data across large spatial domains, methods must be considered for the extraction of information from the data. Satellite imagery can be obtained from a number of vendors, and RPAS are now easily accessible and often come with onboard high-resolution RGB cameras. Raw data collected is a set of many spatial high-resolution RGB images, and converting these images to useful information is not trivial.

The identification of individual trees from imagery is an ‘object detection’ problem; a class of problem concerned with localizing and classifying objects within an image (Zhao et al., 2018). This task is not straightforward, as objects can appear in many different contexts, in different lighting conditions, and there may be extensive variation within object classes (i.e., tree species, dog breeds, traffic signs, etc.). The goal is to detect objects in diverse contexts from RGB imagery; this is a task that machine learning is well-suited to. Neural networks are able to detect a variety of objects in a variety of settings by extracting feature representations. As discussed by Zhao et al. (2018), CNNs are one such type of network and have the ability to develop complex and deep feature representations during training; that is, feature representations do not have to be manually designed. Zhao et al. (2018) highlight numerous advantages of using CNN architecture for object detection over other methods, including hierarchical feature representation, increased expressive capability, ease of task combinations, and large learning capacities. There are numerous CNN architectures to select from. In this research, RetinaNet and YOLO are utilized and compared. RetinaNet leverages a pyramidal feature map structure to enhance scale invariance and a novel implementation of focal loss to account for background sample imbalance (Lin et al., 2016, 2017). YOLO uses a unified detection framework to both localize and classify objects, and operates rapidly and efficiently (Redmon et al., 2015).

The use of CNNs for tree detection has been explored in research. CNNs are commonly used for general image recognition, including tree detection with some extension into tree attribute modelling. Braga et al. (2020) implemented a mask R-CNN to detect and delineate tree crowns in high-resolution multispectral satellite imagery, reaching recall, precision and F_1 scores of 0.81, 0.91, and 0.86, respectively, in a tropical forest. The CNNs AlexNet and GoogLeNet were used to detect dead pine trees in southeastern China (Tao et al., 2020). When trained on 768 samples and tested in three regions, the algorithms obtained identical F_1 scores of 0.57. This score is depressed by errors of omission, with a sub 0.5 recall in both cases. Mubin et al. (2019) considered 0.3 m px^{-1} satellite imagery for palm tree detection, separating young and mature palms in the workflow. The CNN achieves overall accuracies of 95.11% and 92.96% for detection of the relatively well-delineated young and mature palm crowns. Ferreira et al. (2020)

detected palm species with an accuracy of $87.8 \pm 4.4\%$. This workflow incorporated ResNet-19 (He et al., 2016) in a DeepLabv3+ semantic segmentation architecture (L.-C. Chen et al., 2018). DeepForest is a tree detection convolutional model based on a RetinaNet structure with a ResNet50 backbone (He et al., 2016; Lin et al., 2017; Weinstein et al., 2019). It is a research model developed and trained on a database of annotated tree data and tested on various locations in the United States, reaching recall and precision of 0.69 and 0.61, respectively, with a stem detection of 0.82. Weinstein et al. (2019) overcame the low amount of available training data by relying on semi-supervised training that leveraged LiDAR data to improve the training set. Bennett et al. (2022) utilize the DeepForest API to detect and classify trees in Alberta boreal forest, and reach an average F1-score of 0.72, and classify deciduous and coniferous trees with F1-scores of 0.97 and 0.87 respectively. Briechle, Krzystek, and Vosselman (2021) use hyperspectral imagery, LiDAR pointclouds, and tree location polygons to train two CNNs (one per input modality) to classify trees as coniferous or deciduous and reach F_1 scores of over 96%. The use of hyperspectral and RGB imagery as inputs into CNNs for the purposes of tree species classification were compared by Fricker et al (2019); authors found that classification using hyperspectral offered a 23% higher accuracy compared to RGB data, though trees were classified into seven species and one dead class as opposed to coniferous / deciduous classification. As can be seen, the use of CNNs to detect trees from RS data is well-explored in research.

1.4 Exposure Assessment

Due to the cost of wildfire, it is in the interest of communities and governments to manage and mitigate its effects. This motivates the development of a community level exposure assessment.

At a community scale, it is of interest to assess exposure to wildfire ignition vectors: radiant heat, short-range firebrands, and long-range firebrands. Assessments at this scale can allow mitigation to be prioritized, communities to be compared to each other, and to prioritize home-scale assessments (Beverly et al., 2010). Originally developed by Beverly et al. (2010), FireSmart Alberta defines a community exposure assessment procedure (Beverly et al., 2018). The assessment procedure requires hazard fuels to be

mapped around the community. The definition of hazard fuels in Beverly et al. (2010) includes: “FBP fuel types C-1 (spruce–lichen woodland), C-2 (boreal spruce), C-3 (mature jack or lodgepole pine), C-4 (immature jack or lodgepole pine), C-7 (ponderosa pine–Douglas-fir), O-1 (grass), and M-2 (boreal mixedwood)” depending on the ignition vector considered. It is also noted that depending on the time of year, certain fuels types such as aspen could possibly be considered hazardous if they are in their spring dormancy or have low moisture content (Beverly et al., 2018). Once fuels are mapped, the community is divided into a grid. A moving circular window is used to calculate the area proportion of hazard fuel within the window at each grid cell. Three radii are used for the window, corresponding to different ignition types: 30 m is used to assess for exposure to radiant heat from wildfires, 100 m is used for short-range ember spotting, and 100 m to 500 m is used to assess exposure to longer range spotting (Beverly et al., 2010). The proportions of hazard fuel for each fuel type are then binned into nil, low, moderate, high, or extreme exposures. Finally, the community receives maps of exposure to each of the three ignition types. These exposure maps are useful for several reasons. Ember exposure has been shown to be the primary reason for structure loss in previous WUI fire events in Alberta (Westhaver, 2017). Thus, understanding where the community is most vulnerable to this ignition vector is a key aspect of preplanning. Further, it has been shown that wildfires primarily occur in areas with high exposure to hazardous fuels (Beverly et al., 2021). In Beverly et al. (2021), fuel maps were derived at a 100 m resolution using a combination of AVI (updated every 10-15 years), Alberta Ground Cover Characterization (generated with satellite data from 2000) (Sanchez-Azofeifa et al., 2004), and annually updated disturbance inventories. In Beverly et al. (2010), land cover was identified resolution via orthophoto interpretation at 1 m, field verification, and with support from local officials.

At the individual structure scale, FireSmart Alberta outlines a home assessment procedure to numerically score a structures exposure to wildfire (FireSmart Alberta, 2015). Attributes of both the structure itself and the surrounding area are assessed including topography, building materials, building condition, surface vegetation, and surrounding tree vegetation. Regarding tree vegetation, scoring is predicated on the establishment of vegetation management zones surrounding the structure. Within 10 m

from the structure, the surrounding tree type is classified as deciduous, mixed wood, or coniferous. At 10 m to 30 m from the structure, tree type is again assessed, and if coniferous, it is noted whether tree crowns are continuous or separated. Results of the assessment can be used to recommend modifications to the yard, surrounding vegetation, or structure itself that can reduce the hazard wildfire poses to the property.

1.5 Thesis Objectives and Outline

Based on the review of current practices in community exposure assessment, an opportunity exists for the leveraging of RS and ML in existing workflows to automate photo interpretation that supplements decision making. The objectives of this study are:

1. To develop a method of fuel mapping (automated interpretation of imagery) from satellite imagery and RPAS at an individual tree scale
2. To explore the possible use of these fuel maps in existing exposure assessment workflows

Thesis Statement: Automating tree detection and classification from RPAS and satellite imagery can produce useful outputs that can be integrated into exposure assessment workflows to supplement existing practices.

The work thus completed in this thesis is broken down into five chapters. Chapter 2 outlines the construction of a CNN model for the purpose of detecting and classifying trees from RPAS collected imagery, including RPAS and flight specifications, dataset summation, model architecture, assessment procedure, and overall performance. Chapter 3 is structured similarly and focuses on detecting and classifying trees from satellite imagery. Chapter 4 explores potential uses of these models in exposure assessments at both the individual structure and community scale, and notes both benefits and limitations. Chapter 5 includes a conclusion and outlines recommendations for future work.

2.0 Tree Detection and Classification from RPAS RGB Imagery

To aid in exposure assessment activities, it is useful to have high-resolution maps of fuels; part of the FireSmart Alberta Homeowners Assessment requires homeowners to assess local vegetation, including tree type and spacing up to 30 m from the home (FireSmart Alberta, 2015). To explore the benefits of high-resolution fuel mapping, a platform is sought that can rapidly collect very high-resolution imagery around areas of interest. This leads to the use of remotely piloted aircraft systems (RPAS), or ‘drones’, to collect imagery around areas of interest. These platforms allow for extremely high (> 2 cm) resolution maps to be created around the wildland-urban interface (WUI).

These systems are used to collect imagery around communities and values-at-risk (VARs), to process that imagery into high resolution orthomosaics, and to train a convolutional neural network (CNN) model to detect and classify trees from the collected imagery.

2.1 Data Acquisition

Over recent years, drone platforms have become increasingly accessible to both enterprise users as well as the general public. The market currently has many options, ranging from small, lightweight hobbyist drones to larger platforms tailored for use by governments, researchers, and emergency response agencies. In this research, three drones are used. All drones are manufactured by DJI. A table of drone and sensor parameters is included in Table 1.

Table 1: Summary of RPAS systems utilized in this study

RPAS Platform	Weight (g)	Wind Resistance (kph)	Onboard Sensors	Cost (\$USD)
DJI Mavic Mini	249	28	Wide Camera	\$400
DJI Mavic 2 Enterprise	905	38	Wide, Zoom	\$4,000
DJI M30T	4000	55	Wide, Zoom, Thermal Camera, Laser Rangefinder	\$14,000

The drones utilized in this project represent a range of what may be available to both researchers and agencies. Most drone platforms allow for surveys to be preplanned and flight paths automated to allow for efficient surveys and consistency between surveys. The Mavic Mini is a small, compact, relatively cheap drone that requires minimal piloting experience. At 249 g, the drone may be flown by any user under current regulations; drones 250 g and up require the user to have a drone pilot certificate to operate (Transport Canada, 2023). However, this drone must be manually piloted, and a drone of this size is also susceptible to wind. The DJI Mavic 2 Enterprise represents an intermediate drone, with additional onboard sensors and stronger wind resistance. Survey automation software is compatible with this drone, allowing flight parameters to be set and surveys to be automatically flown. The DJI M30T represents the high-end range of possible drones. This platform is specifically marketed for use by first responders, offers high wind resistance, a suite of onboard sensors, and increased weatherproofing. Due to its design for first responders, additional features including laser rangefinders for in-flight measurements and highly accurate GPS / surveying sensors are onboard to ensure captured products are accurate.

Despite the differences in drones used on this platform, onboard cameras are very similar. Table 2 summarizes the on-board wide-angle cameras used for imaging on each drone.

Table 2: Summary of RPAS cameras utilized in this study

RPAS Platform	Camera Sensor	Image Size
DJI Mavic Mini	1/2.3" CMOS: 12 MP	4000x3000
DJI Mavic 2 Enterprise	1/2.3" CMOS: 12 MP	4000x3000
DJI M30T	1/2" CMOS: 12 MP	4000x3000

A number of parameters must be set when flight planning to ensure consistency between surveys. First, the area to be flown is defined using a polygon on a map. Flight height must then be entered; the drone must fly high enough so as to be clear of all obstacles,

but low enough that the ground sample distance (*GSD*) will yield an acceptable resolution. Here, *GSD* is a relationship between camera sensor width, focal length, flight height, and output image width, and is calculated as follows.

$$GSD = \frac{H * S_w}{f * I_w} \quad \text{Eq. 1}$$

Where *GSD* is ground sample distance (cm / pixel), *H* is flight height (m), *S_w* is sensor width (mm), *f* is focal length (mm), and *I_w* is image width (pixels). Flight speed is also set; faster speeds can mean larger surveys in less time and battery life, though speeds that are too high will result in image blurring. Frontal and side overlap of the image footprints must also be set. This is the percentage to which adjacent images overlap; generally, higher overlap percentages allow for more detailed 3D models to be constructed by the stitching software. Parameters are chosen to reach a similar *GSD* to other studies that have been performed in tree detection and classification using CNNs: Egli & Höpke (2020) utilize 1.6 cm / px *GSD* for tree classification, while Schiefer et al. (2020) utilize 1.1 cm / px *GSD* imagery for the mapping of tree species. Table 3 summarizes the parameters commonly used on this project.

Table 3: RPAS flight parameters for M30T

Time	Midday, consistent cloud cover preferred
Flight Height (m)	50
Flight Speed (m/s)	2
Side Overlap (%)	80
Front Overlap (%)	90
Camera Resolution	3000 x 4000
Camera Focal Length (mm)	4
Camera Sensor Width (cm)	6.3
Final GSD (cm / px)	2.0

The metadata in the collected images is used to stitch the images into one large orthomosaic. Agisoft Metashape Professional is the software used to generate the orthomosaic from the collection of survey images (Agisoft, 2023). First, the images are added and aligned. Next, a digital elevation model (DEM), pointcloud, and 3D mesh model are generated of the survey area. Finally, an orthomosaic is generated and exported. Each survey now has an accompanying orthomosaic, 3D pointcloud, 3D mesh model, digital surface model (DSM), and digital terrain model (DTM). Figure 1 demonstrates these outputs for a home in rural Alberta.

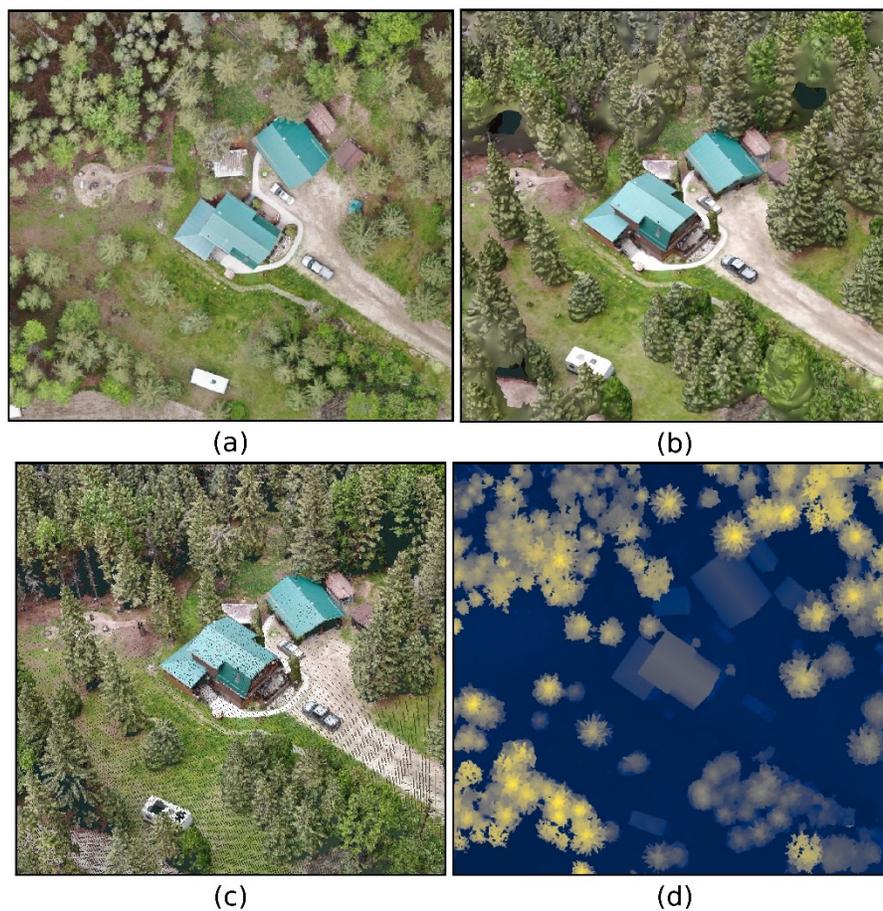


Figure 1: Survey outputs for a home in rural Alberta, including (a) orthomosaic, (b) 3D mesh, (c) 3D pointcloud, and (d) DEM

To train a CNN model to detect and classify trees, a training dataset of trees is required. To build this set, 11 surveys are flown at various locations around Alberta. These sites

contain several different tree species under various lighting conditions to ensure training dataset variability. Surveys are carried out at the parameters previously mentioned, ensuring GSD is maintained at 2 cm / px. The locations of these surveys are shown below in Figure 2, with additional information provided in Table 4.

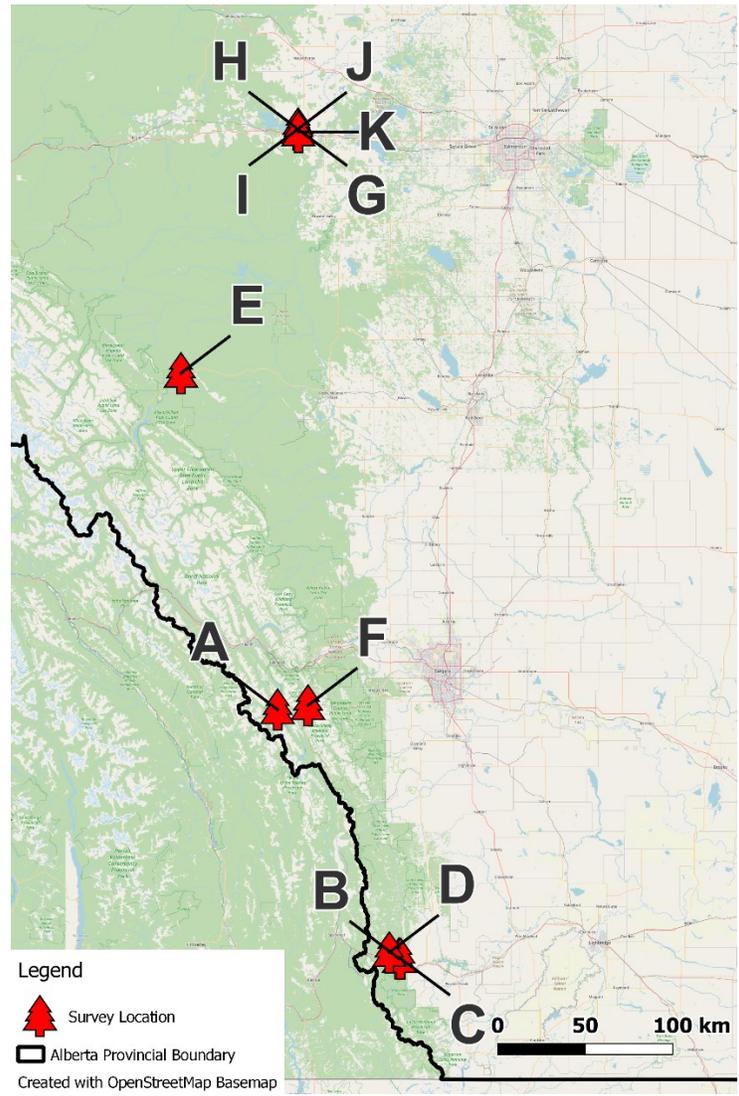


Figure 2: RPAS survey locations, including training surveys and surveyed communities

Table 4: Training survey summary

Survey ID Marker	Date Surveyed	Area (ha)	Conifer Count	Dec Count	Snag Count	Totals
A	10-Jun-20	2.69	682	0	39	721
B	12-Jun-20	1.89	154	155	28	337
C	12-Jun-20	2.52	374	20	5	399
D	12-Jun-20	3.36	581	97	45	723
E	19-Jun-20	2.28	693	0	89	782
F	22-Jun-20	2.68	1121	3	70	1194
G	07-Jun-22	1.15	1524	272	48	1844
H	07-Jun-22	1.54	122	611	17	750
I	07-Jun-22	4.2	79	667	57	803
J	07-Jun-22	4.61	1569	763	163	2495
K	07-Jun-22	0.94	29	145	1	175
Totals			6928	2733	562	10,223

2.2 RPAS Model

As mentioned previously, the need to detect trees from RGB imagery is a problem that lends itself well to the use of CNNs. Trees appearing in orthophotos can have different shapes and colors, while still sharing features that enable humans and machine algorithms alike to identify and classify the tree.

The field of computer vision is one that advances rapidly, and the optimal method for object detection from imagery is not yet settled. Researchers continually release models that use different techniques to identify objects in imagery, and additionally it is not certain that the model that currently best identifies cars will also best identify other objects, such as trees. As such, a suite of commonly used CNN models is assembled for the task of identifying tree locations and classifying trees in the stitched orthophotos. The models selected are RetinaNet and YOLOv5 (nano, small, and large configurations) (Jocher et

al., 2022; Lin et al., 2017). Both models are CNNs that score well in existing object detection benchmark competitions.

To detect trees in a set of imagery, the model must first be trained using a large number of examples. The more examples the model is allowed to train with, the more accurate the model is expected to be, as it has a deeper set of objects to extract features from. Generating these examples is still a human task – all surveys flown in the aforementioned set are manually annotated. A rectangular box is placed by a human over each tree in each survey, and the class of the tree is denoted as coniferous, deciduous, or snag (a standing dead tree). This allowed a dataset of 10,223 trees to be created, with 6928 conifers, 2733 deciduous, and 562 snags.

To fairly assess model performance, five datasets are created from the collected orthophotos and corresponding annotations. Five training and testing subsets are created, ensuring two things: that no model is both trained and tested using data from the same survey to preserve fairness, and that each model has approximately the same number of total tree annotations to train on. The total dataset is thus split using these two rules, so that models could be trained using four of the subsets and tested using the remaining one. This is commonly known as a k-fold cross validation and is repeated 5 times so that each subset is omitted from the training data and used as a test set. A summary of the dataset, as well as the five subsets, is included in Table 5.

Table 5: RPAS dataset summary

Split Group ID	Training Surveys	Testing Surveys	Training Conifer	Training Deciduous	Training Snag	Testing Conifer	Testing Deciduous	Testing Snag	Training Total	Testing Total
1	A, B, C, D, E, F, G, H, I, K	J	5359	1970	399	1569	763	163	7,728	2,495
2	A, B, C, D, E, F, H, I, J, K	G	5404	2461	514	1524	272	48	8,379	1,844
3	D, E, F, G, I, J, K	A, B, C, H	5596	1947	473	1332	786	89	8,016	2,207
4	A, B, C, D, E, G, H, J, K	F, I	5728	2063	435	1200	670	127	8,226	1,997
5	A, B, C, F, G, H, I, J	D, E, K	5625	2491	427	1303	242	135	8,543	1,680

With the datasets created, the models are then trained. Python packages have been created by various groups to allow common models to be modified, trained, and tested. The RetinaNet model API used is provided by MMDetection (K. Chen et al., 2019), and the YOLOv5 model API by Ultralytics (Jocher et al., 2022).

Training and testing are performed in an Anaconda Environment on a Windows 10 desktop computer. The computer is equipped with 256 GB of RAM and an NVIDIA GeForce RTX 4090 GPU. First, for each of the five above datasets, training surveys and their corresponding annotations are tiled to allow training to fit into RAM, and to allow varied image augmentations. The surveys are sliced into 800 px tiles that overlap with each other by 30%, to ensure trees on tile borders appear whole in at least one image. The tiles are then split into a training and validation set at a 90/10 ratio. Upon training, image augmentations are applied, including random image flips, image scaling, and alterations to the image's hue, saturation, and value characteristics. Training consists of

200 epochs with an initial learning rate of 0.01 using the stochastic gradient descent (SGD) optimizer. Training time depends on the size of the dataset as not all are exactly equal but takes around 5 hours to complete. Upon completion, the training metrics are analyzed, and the model is tested on the withheld surveys to assess its performance. Inference time is negligible when compared to training time, as it is on the scale of seconds to minutes per survey.

2.3 RPAS Model Results

Model performance is assessed via a k-fold testing strategy. Using the manual annotations as truth and the model outputs as predictions, performance metrics are calculated. Metrics are divided into two categories: identification performance and classification performance.

Precision, recall, stem recall, and F1-score are used to determine the model's ability to identify and properly box trees regardless of class. Here, 'properly boxed' means that the box predicted by the model sufficiently overlaps the true box. This is determined by calculating the intersection-over-union (IoU) of the predicted box and the true box. IoU is calculated as follows.

$$IoU = \frac{A_o}{A_U} \quad \text{Eq. 2}$$

Here, A_o is the area of overlap, and A_U is the area of union. Figure 3 provides a graphical representation. In the figure, two boxes are represented; one being ground truth, the other being the model detection. The shaded area represents the area calculated: in the numerator, the area of intersection only, and in the denominator the area of union. Dividing the two results in IoU .

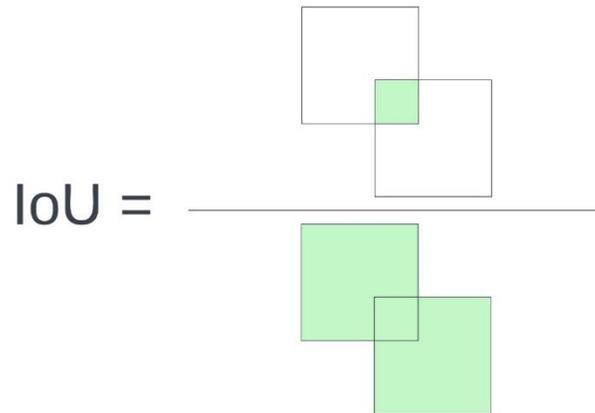


Figure 3: Graphical representation of IoU

A cutoff threshold is then used – a prediction box is said to have properly boxed the true tree if the IoU is greater than 0.5. With a ‘correct box’ defined, precision and recall are then calculated. Precision is a measure of how often an estimated tree box output by the model correctly boxed a true tree and is calculated as follows.

$$Precision = \frac{TP}{TP + FP} \quad \text{Eq. 3}$$

Here, a true positive (TP) refers to a box drawn by the model that correctly overlaps a true tree, and false positive (FP) refers to a box drawn by the model that does not properly overlap any true tree. Recall measures the model’s ability to properly identify all true examples in the testing set and is calculated below.

$$Recall = \frac{TP}{TP + FN} \quad \text{Eq. 4}$$

Here, false negative (FN) refers to the number of true trees that are not properly boxed (i.e., missed) by the model. F1-score is the harmonic mean of precision and recall, combining them into one number so that model performances can be more directly compared. F1-score is calculated as follows.

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall} \quad \text{Eq. 5}$$

Stem recall is the final metric used for tree detection in this study. Due to the nature of the study, there are circumstances where tree size may be irrelevant, such as when calculating tree stem density. As such, true trees must not always be ‘properly’ boxed (i.e., IoU > 0.5) to count as a true positive. It is of interest to know how often the model detects a tree at all, regardless of if the box is properly sized. Stem recall is then defined as the percentage of times the stem of the tree, defined as the centre point of the true box, is within a prediction box regardless of size.

Classification of the identified trees into the three defined classes must also be assessed. Similar metrics are used. Here, classification metrics are calculated using only detection boxes that properly boxed (IoU > 0.5) a true box. This is to ensure the classification assessment is fair, as the model can only be asked to classify trees that are properly boxed. Further, no ground truth can be assigned to a box that does not properly contain a tree. Again, precision and recall are calculated and combined into the F1-score for classification of each class. The formulas used are included below.

$$Precision_{class\ 1} = \frac{TP_{class\ 1}}{TP_{class\ 1} + FP_{class\ 2} + FP_{class\ 3}} \quad \text{Eq. 6}$$

$$Recall_{class\ 1} = \frac{TP_{class\ 1}}{TP_{class\ 1} + FN_{class\ 2} + FN_{class\ 3}} \quad \text{Eq. 7}$$

A FP for a class occurs when, for instance, a tree of class 2 or 3 is classified as class 1. A FN occurs when a tree of class 1 is classified as class 2 or 3. During training, the precision, recall, and F1-score are calculated on the validation set to ensure training is progressing properly and the model is learning. Figure 4 demonstrates the learning curve of the model on the validation set during training.

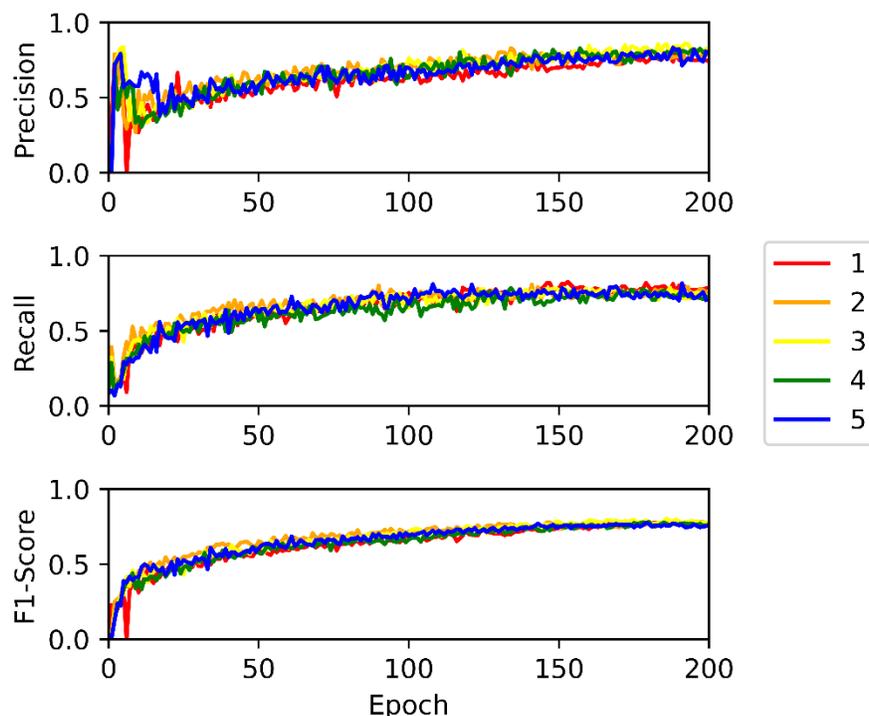


Figure 4: Training metrics for YOLOv5, folds 1 through 5

Once all five models are trained, they are each tested on their corresponding test set. As stated, the surveys included in the test set of each data split were not used for training, and thus represent a fair assessment of model performance. A table of the overall results for both identification and classification is included in Appendix A: RPAS Detection Results. In the instance that multiple surveys are withheld for testing in a fold, such as in sets 3, 4, and 5, a weighted average of the metrics is calculated corresponding to the number of trees contained in that survey so that a final overall metric can be determined for that fold. It is seen that the large configuration of YOLOv5 has the best combined F1-score, as well as the best classification accuracies of coniferous and deciduous trees. With an F1-score of 74.5% and a stem recall of 87.1%, the model identifies trees fairly consistently. Classification F1-scores reveal that the model is also highly accurate at classifying identified trees as coniferous, deciduous, or snag: the large configuration of the model achieved classification F1-scores of 97.3%, 94.4%, and 90.9% respectively.

Results can also be assessed visually. Boxes output from the model are drawn on the imagery and compared to the manual annotations. Figure 5, generated using the YOLOv5L model, demonstrates the model in action, with the prediction boxes also color coded by predicted class.

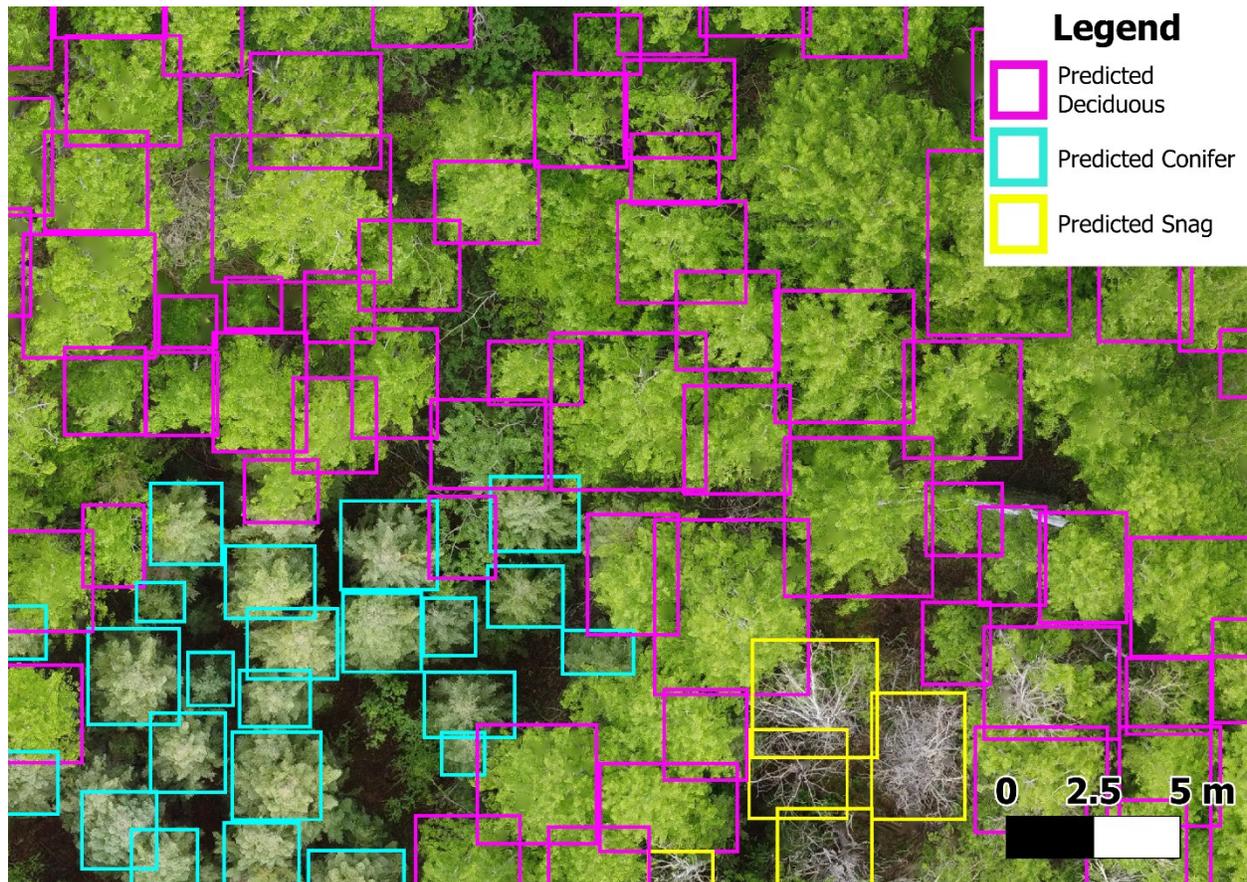


Figure 5: Sample tree detections and classifications on new RPAS survey

As can be seen visually, the model is able to detect trees of all classes, and to classify the trees into their respective types. The survey above was previously unseen by the model; it was not part of the training set. Individual tree crowns are clearly well-delineated, especially in the case of conifers. The predictions made by the model in this sample are not error free; dense and overlapping deciduous stands remain challenging for individual crown detection. However, as discussed previously, deciduous trees are generally not considered to be hazardous fuels. Coniferous trees are correctly detected and classified with an F1-score of 79.9% - compared to the global average for detection alone at 74.5%,

it is clear the model excels at coniferous detection and classification. Detections around a structure can also be shown, as in Figure 6.

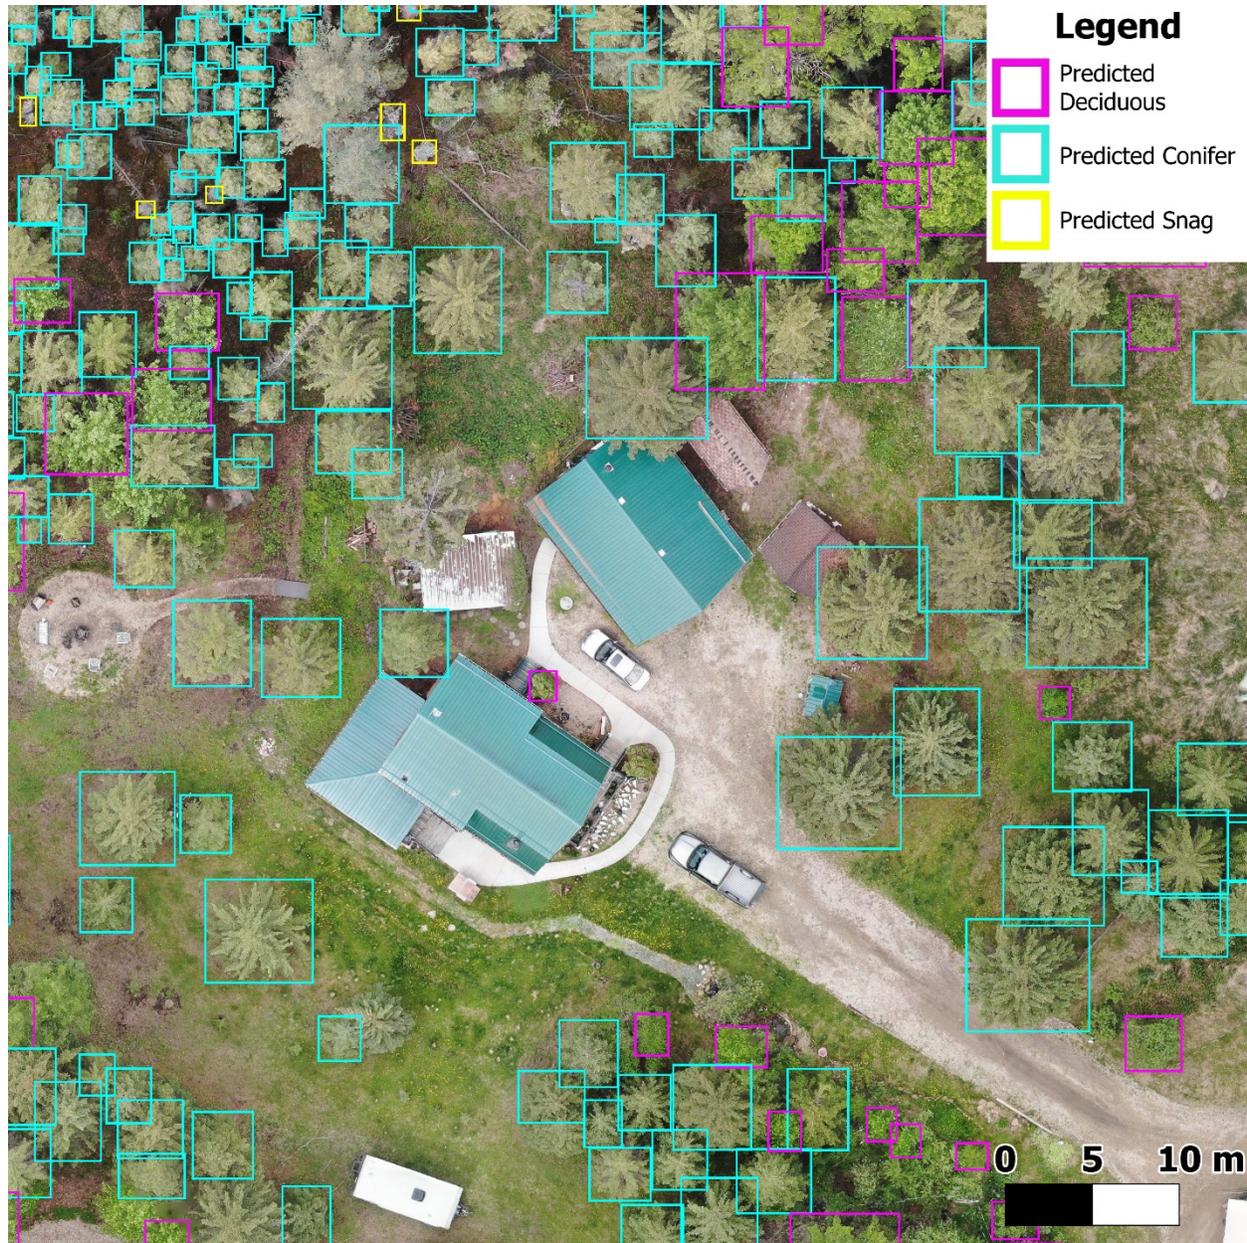


Figure 6: Sample RPAS property survey and classified tree detections

Figure 6 demonstrates tree detections around an inhabited structure. As can be seen the majority of vegetation, both coniferous and deciduous, is well-boxed and correctly classified around the structure.

2.4 RPAS Model Discussion

In summary, a workflow is presented in which aerial imagery is flown using an RPAS and collected nadir RGB imagery, processed into a suite of data products, and a CNN is developed and applied to detect and classify trees. The workflow developed is fully automated and shown in Figure 7.

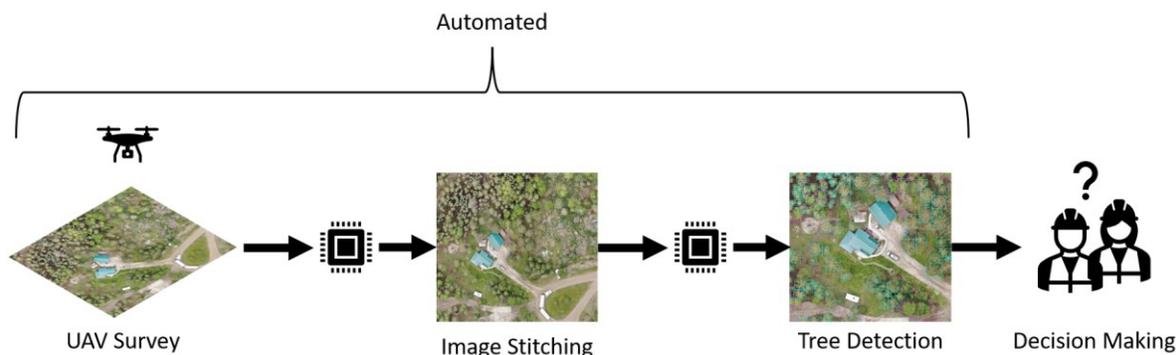


Figure 7: Automated RPAS image collection and detection workflow

Once survey bounds are decided, they can be uploaded to the RPAS and the survey is automatically collected. Data is transferred onto a computer for processing, which is automated and requires only a few keypresses. The resulting outputs can thus be collected and generated rapidly and with little work to the individual performing the survey. Outputs such as these could be generated on timescales as little as an hour or two from flight takeoff to output generation and visualization. This demonstrates the workflow's ability to rapidly gain information about the vegetation surrounding a VAR.

It is important that YOLOv5 detection results be understood and used in context of their limitations. Primarily, the fact that a perfect RPAS model that correctly detects and classifies every single tree is an extremely difficult goal to achieve and may be functionally impractical to develop. Most significantly in Figure 6, a conifer just to the north of the structure was missed by the model. Reasons as to why can only be speculative, but it is worth reiterating that the F1-score of 74.5% and the stem recall of 87.1% means errors still occur. There are a multitude of variables that can change between surveys and affect how trees appear in the final survey imagery, including variations in tree species, tree growth, weather conditions, lighting, and RPAS camera or flight parameters. Model

performance is expected to only improve as additional surveys are carried out for an increased number of training examples, and as model architecture continues to be developed in the field of computer vision. RPAS surveys that collect RGB imagery are also limited to detecting trees in the overstory. As visible light cannot pass through thick canopies, the survey collected by the drone will capture only the top layer of the trees. Smaller trees and undergrowth, which can often represent a significant portion of the wildfire fuel in an area, cannot be detected using this method.

RPAS data collection can also be limited due to the upfront cost of the drone system. While a number of options are presented for use that range from \$400 to \$14,000 USD, drone purchasing, accessories (including additional batteries), software, and training can represent a significant up-front investment. While much of the process in both data collection and stitching is automated, it still requires the training and certification of personnel piloting the drone. Surveys may also be considered time consuming depending on the scenario – while the majority of the work is in preplanning and the flight itself is automated, someone must still be present to monitor the drone while it is flying.

Survey flights can also be impeded by a number of factors. Weather conditions, including heavy rains and high winds, can ground the drone. Different drones have different weather and wind resistance ratings, but moderate to high winds can limit takeoff and can limit the area that can be surveyed due to increased battery consumption. Flights can also be impeded by members of the public concerned about privacy and by local bylaws that regulate drone flight and imagery collection that vary by jurisdiction.

The time-lag between collection and data processing is minimal, though not zero. The time to plan a survey, fly, process, and interpret the data into useful products are all affected by various factors, though each step itself is on the time scale of minutes to hours. While instantaneous turn-around is not the intention of this study, the workflow represents a means to which fuel information can be collected same-day.

The proposed workflow can detect and classify individual overstory trees and achieves an F1-score for detection of 74.5% and classification accuracies above 90%. Many wildfire processes do not operate on the tree-to-tree scale; rather, they are larger scale

processes that rely on patches of fuels, including both overstory tree fuels and ground / surface fuels. As such, future work should be performed in the aggregation of individual tree properties to the stand level. Further, the extraction of other fuel attributes such as tree height, crown base height, and crown bulk density from other data products output from the survey (such as the pointcloud seen in Figure 1) should be explored. Such work could allow for generated fuel maps to be used in additional workflows such as Fire Behaviour Prediction.

The outputs of the proposed workflow offer high resolution overstory fuel detections that can be collected and processed rapidly and in an automated fashion. Within the bounds of the stated limitations, these outputs could be used as an informational supplement to structure scale exposure assessments.

3.0 Tree Detection and Classification from Satellite RGB Imagery

While remotely piloted aircraft system (RPAS) imagery can offer extremely high-resolution information over multiple hectares, the boreal zone alone in Canada covers 552 million hectares (Brandt et al., 2013). As such, it is infeasible to collect 2 cm resolution RPAS imagery across its entire extent. Even ignoring outlying forests and focusing on the wildland-urban interface (WUI), it is estimated that Canada has 32.3 million hectares of WUI, 10.5 million hectares of wildland-industrial interface, and 109.8 million hectares of wildland-infrastructure (including railways and roads) across the country (Johnston & Flannigan, 2018). It is then of interest to utilize a new source of data across these massive areas. Satellite imagery is increasingly accessible to interested buyers at increasing resolutions. Satellites offer reasonably high-definition imagery with high spatial coverage and within recent timespans. Satellites such as WorldView-2 offer high-resolution (0.5 m panchromatic and 1.8 m multispectral) data with revisit times of up to 1 day (DigitalGlobe, 2009). The utilization of an algorithm to detect and classify fuels in satellite imagery across massive spatial extents is thus of interest.

Satellite imagery is purchased covering 5 communities within Alberta, as well as their surrounding area. A convolutional neural network (CNN), similar to what was developed in the RPAS workflow, is utilized to detect trees within the satellite imagery after being trained using manual annotations. Model performance is assessed by comparing estimated tree densities to manual annotations. A classification scheme to classify trees into a proxy for coniferous / deciduous fuel types is then developed and assessed. The result is an estimate of wildland fuel coverage and type at high resolutions and at large spatial extents. The possible use of these results in a community exposure assessment is later demonstrated.

3.1 Data Acquisition

Satellite imagery is ordered for 5 communities in Alberta. Seba Beach, Lobstick, and Tomahawk are in the boreal forest region of Alberta, while Bragg Creek and Redwood

Meadows fall within the rocky mountain natural region. All communities contain significant amounts of WUI. A map of these communities is provided in Figure 8.

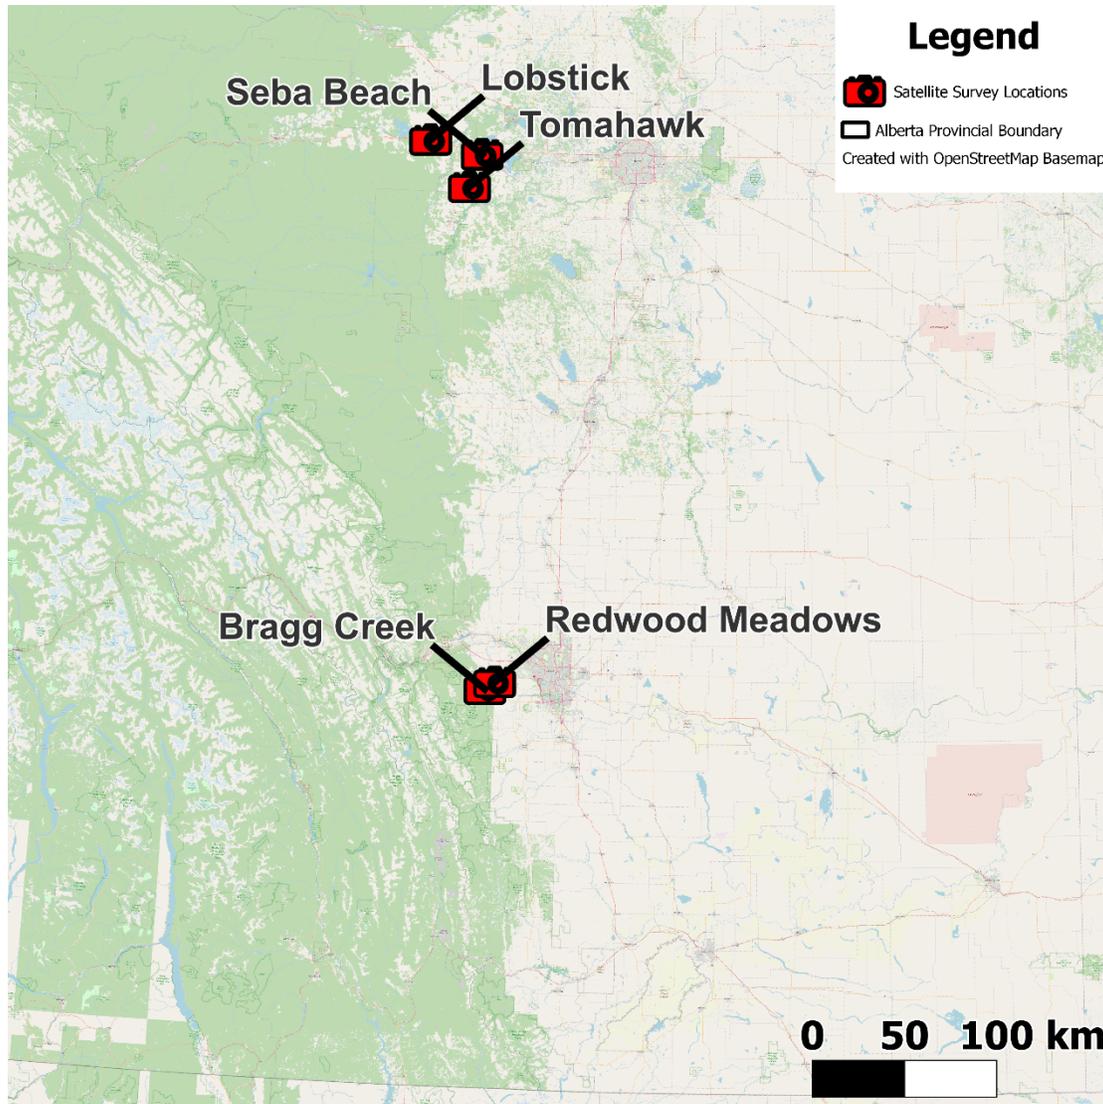


Figure 8: Satellite survey locations in Alberta, Canada

The goal is to acquire high-resolution RGB imagery covering the community and its surrounding area. A number of specifications are important when collecting satellite imagery to ensure it is useful for fuel identification. Resolution is specified to be 0.5 m / px: the highest resolution available at the time of purchase across the areas of interest. Imagery selected must also be free of cloud cover, as trees are then obscured and undetectable. The angle off-nadir is also held to less than 30°; this ensures that trees are

viewed at approximately the same angle of incidence. Views that are considerably off-nadir can impact model performance as trees begin to look significantly different as more of the side profile is captured. Finally, the most recent imagery that can be acquired that meets all specifications is ordered. An image is ordered in the summer, when all trees are green and visible, and a fall/winter image is also ordered, after the point when deciduous trees have gone dormant but without significant snow coverage, for later use in fuel classification. A side-by-side of satellite imagery near Seba Beach, Alberta, is shown in Figure 9 to demonstrate the resolution of the imagery as well as the difference in summer / winter imagery. Table 6 summarizes all acquired surveys, which are then used to develop a tree detection model.



Figure 9: Satellite imagery near Seba Beach, Alberta. Summer (left) and winter (right)

Table 6: Specifications of satellite surveys used in this study

Location	Satellite	Date Surveyed (Summer)	Date Surveyed (Winter)	Area (ha)	Angle Off-Nadir (Summer)	Angle Off-Nadir (Winter)
Bragg Creek, AB	Pleiades-1	29-Jun-21	17-Oct-18	4700	21.4	10.6
Seba Beach, AB	Pleiades-1	12-Sep-19	25-Oct-15	3900	15.9	25.0
Lobstick, AB	Pleiades-1	15-Jul-19	25-Oct-15	3600	9.0	11.4
Redwood Meadows, AB	Pleiades-1	17-Jul-18	30-Dec-15	4000	18.9	12.6
Tomahawk, AB	Pleiades-1	12-Aug-17	12-Oct-15	3700	9.8	15.8

3.2 Satellite Model

The large configuration of the YOLOv5 model is preferred for RPAS imagery (Chapter 2). Therefore, this model is also considered for satellite imagery. The same distribution provided by Ultralytics is utilized for model training and inference (Jocher et al., 2022).

A set of annotations is needed to train the model to predict tree locations. Training CNN models is most successful when the models are supplied with many examples of the object of interest under various lighting conditions, orientations, shapes, and sizes. This allows the model to learn the features that constitute the object, allowing for a robust detection model to be built. However, manually annotating satellite imagery is time consuming, difficult, and prone to errors as trees are difficult to distinguish at the resolution available and are closely packed. Further, even small sections of forest can contain hundreds or thousands of visible trees. When generating an annotated dataset, all examples of the object must be annotated; if trees are left in the dataset but are not annotated, they will be interpreted by the model to be negative examples during the training process, which can result in instability and an error-prone model.

To begin training, large satellite images are thus split into more manageable tiles – 150 m x 150 m (2.25 ha) crops taken from the greater survey image. Thirteen tiles are then completely annotated by a human: each example of a tree in the image is manually boxed and assigned the class of ‘tree’. These surveys include Bragg Creek, Redwood Meadows,

and Seba Beach. Survey tiles are chosen so that trees are included as well as typical surrounding environments; dense woods, trees bordering grass fields, and trees in the urban and semi-urban environment are all represented in the dataset. Figure 10 demonstrates a sample of tiles selected for human annotation, and Table 7 summarizes the manual annotation dataset for satellite imagery. All 13 tiles are included in Appendix B: Satellite Training Tiles.



Figure 10: Example of annotation tiles, in urban, semi-urban, and forested environments

Table 7: Manual annotation dataset summary in satellite imagery

Survey	Image ID	Boxed Trees
Bragg Creek	BC1	576
	BC2	389
	BC3	486
	BC4	928
	BC5	118
Redwood Meadows	RW1	412
	RW2	284
	RW3	258
	RW4	295
	RW5	201
Seba Beach	SB1	609
	SB2	894
	SB3	465

Human annotated tiles are then used to train a blank YOLOv5 model to detect trees. To assist in the generation of the many annotations required for training mentioned previously, an iterative process is devised, shown in Figure 11. First, a model is trained using only manual annotations. This first newly trained model is then used to generate predictions over the entirety of the greater satellite image. Model predictions are then utilized as a ‘pretraining’ set for a new, blank YOLOv5 model. Pretraining allows weights to be initialized and to allow the model to build a level of familiarity with the object prior to the regular, full course of training on high-quality human annotated data. The model is then refined using human annotated data, and the process is iterated: the model is used to generate predictions over the entire satellite image again, and these predictions are in-turn used as pretraining for another new model. Iteration continues until negligible performance increases are observed. This process took 11 iterations to complete. The final model is then saved and used for prediction.

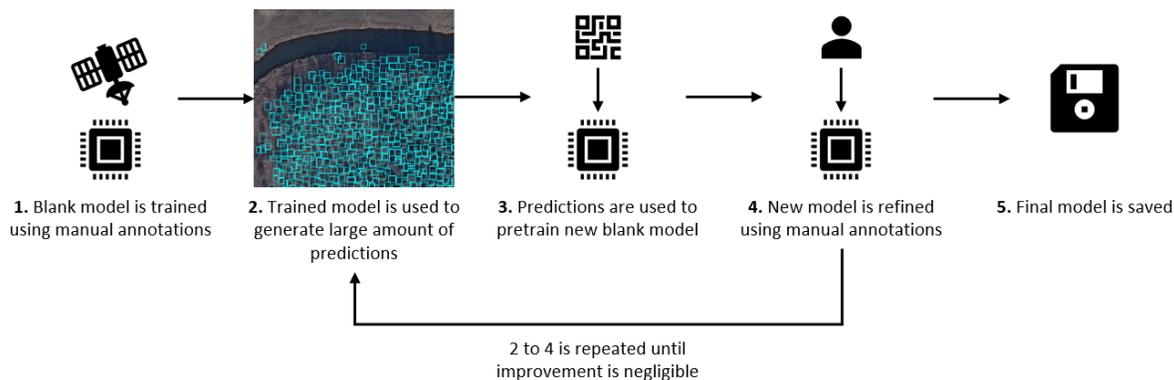


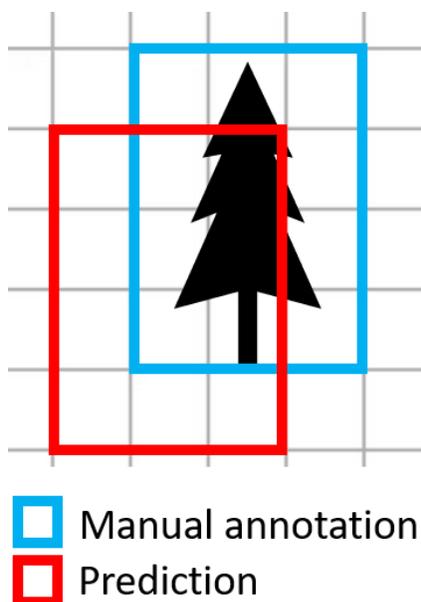
Figure 11: Satellite model pretraining procedure

Satellite model training took place on the same machine utilized for RPAS model training. In pretraining, satellite surveys and their annotations are tiled in a similar manner, resulting in a dataset of 12688 tiles and 5,901,510 corresponding tree annotations. To initialize weights, the model is pretrained for 100 epochs at a learning rate of 0.01 with the stochastic gradient descent (SGD) optimizer. This is followed by 100 epochs of training using only the human-annotated dataset at a learning rate of 0.01. Due to the low resolution of the imagery as well as the small size of the objects attempting to be detected, limited image augmentation is considered, only random flips were used. The pretraining of the satellite model took 4 hours to complete, while the training process took approximately 30 minutes due to the low number of human-annotated tiles.

3.3 Satellite Model Results

Model performance is assessed via k-fold cross-validation. To ensure fairness, in each fold an entire satellite survey is withheld from the pretraining and training set and used for testing. For example, when testing on manually annotated tiles in the Bragg Creek set, the Bragg Creek survey and annotations are removed from the pretraining set and the training set to ensure an unbiased assessment. However, the assessment metric utilized for satellite tree detection differs from the precision, recall, and F1-score used in the RPAS assessment. Trees in the 0.5 m / px resolution satellite imagery are typically only a few pixels in size. This means that the bounding boxes that are drawn are very small. Further,

drawn bounding boxes must align with the pixel grid as bounding boxes cannot be described in terms of fractions of a pixel. Thus, the minimum step a bounding box can be placed is 0.5 m. As the formula for intersection-over-union (IoU) is calculated as a fraction of area of overlap over area of union, shifting a corner of a bounding box by even a singular pixel can cause the area of union to grow significantly, or the area of overlap to shrink significantly, even though the tree is still approximately bounded. Figure 12 demonstrates how a tree detection shifted merely by a single pixel left and downwards can result in an IoU of 0.33 – well below the 0.5 threshold normally used for scoring a ‘correct’ box.



$$IoU = \frac{\text{Area of Overlap}}{\text{Area of Union}} = \frac{6}{18} = 0.33$$

Figure 12: Example of how a prediction box placed mere pixels away can result in a low IoU

The result is that a tree must be boxed with near-perfect accuracy to count as a positive case. This assessment style is unfair to the model, as there may be the discussed human error or disagreement between human and machine understanding of where the exact bounds of the tree are interpreted to be. Thus, a metric of stem density is introduced and utilized for assessment.

Here, the stem of a tree is defined as the geometric centre of the bounding box that defines it. Once predictions are generated for the tile in the testing set, measures of density of tree stems are then taken from both the predicted results and the manually annotated results. Recall that training tiles are square tiles with side lengths of 150 m. To calculate density, a circle with a 20 m radius is centered on each grid location and the number of stems in each circle is counted. The grid location is then assigned a density value ($n_{\text{trees}} / \text{area}$) in units of trees per square metre. However, this means that a grid location in a treeless field 20 m away from the nearest trees will still record the grid location as having some trees per square metre. To correct for this ‘blurring’ effect, a second circle is used with a radius of 3 m. If no tree stems are within this circle, the grid location is set to 0. Figure 13 demonstrates both the issue of an area far from trees being assigned a density value, as well as the 3 m circle used to address the issue. Finally, the density grids of both manual and predicted trees can be compared.

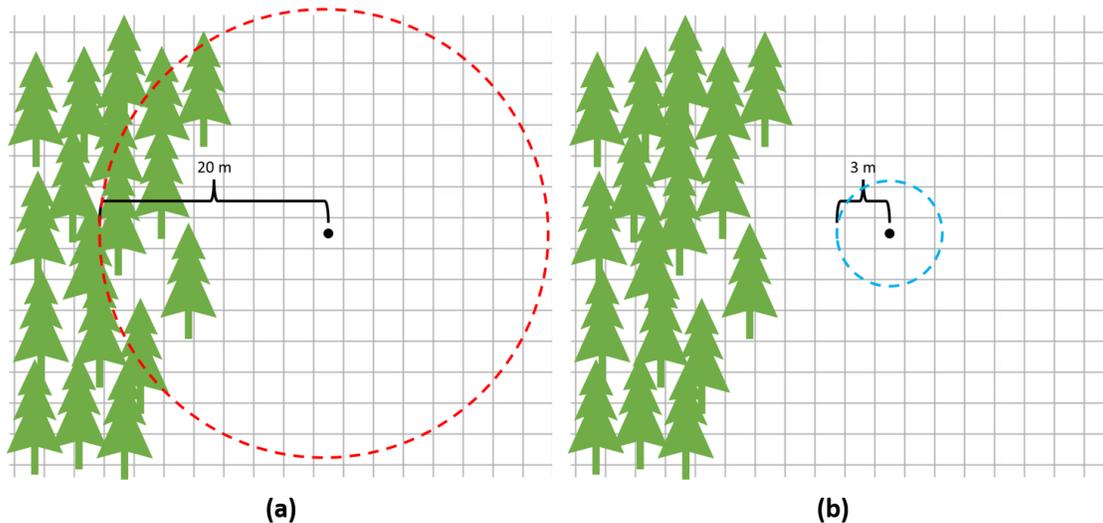


Figure 13: An area without trees (a) could still be assigned a tree density value if a large radius is used alone. The 3 m circle (b) is used to set this cell to 0 if no trees are found.

The manual and machine-predicted densities of the grids overlaid on each tile are compared, and R^2 calculated. Table 8 summarizes the results of the k-fold.

Table 8: Satellite detection assessment results

Survey	Image ID	Image ID	Actual Tree Boxes	Predicted Tree Boxes	R ²
Bragg Creek	1	BC1	576	558	0.73
	2	BC2	389	471	0.60
	4	BC3	486	539	0.80
	7	BC4	928	797	0.58
	9	BC5	118	105	0.75
Redwood Meadows	2	RW1	412	331	0.66
	3	RW2	284	311	0.82
	7	RW3	258	222	0.83
	8	RW4	295	232	0.75
	10	RW5	201	304	0.94
Seba Beach	2	SB1	609	697	0.71
	3	SB2	894	800	0.85
	4	SB3	465	514	0.83
Average					0.76

Densities correlate with an R² of 0.76. Sample plots in Figure 14 show the manual and predicted tree boxes, along with the generated density comparison chart, and a cross-plot of manual and predicted tree densities.

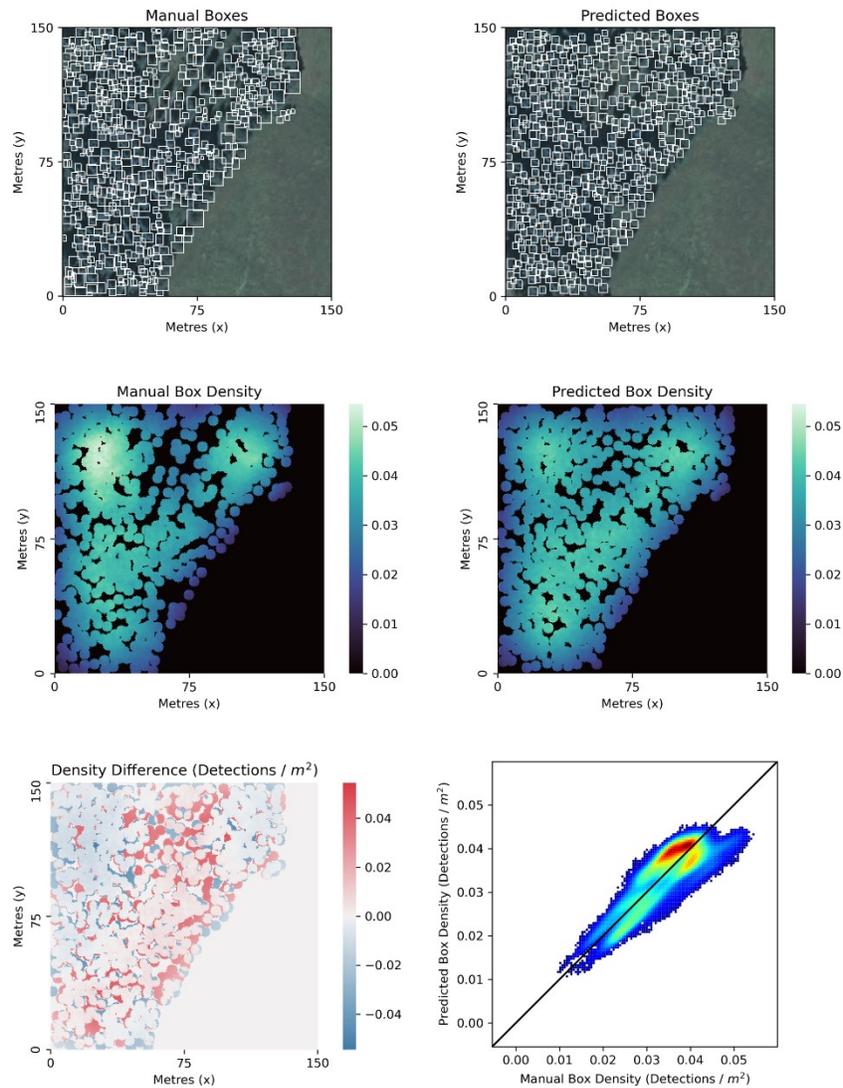


Figure 14: Assessment of manual and predicted boxes in satellite imagery

3.4 Individual Tree Classification

A number of methods for classifying coniferous and deciduous trees from remote sensing (RS) data have been explored in research. Persson et al. (2018) used multiple sensing bands, including shortwave infrared and near infrared to classify tree species in Sweden with a random forest model. Immitzer et al. (2012) also use a random forest to classify tree species, here considering up to 8-band WorldView-2 imagery and achieving an overall accuracy of 82% when performing object-based classification. Bolyn et al. (2022) use a deep learning algorithm to derive tree species proportions from satellite imagery

captured via the Sentinel-2 satellite, which acquires imagery at 10 to 60 m resolution, depending on the selected band. Researchers used this data to estimate tree species proportions in tree stands at a resolution of 2.5 m and achieved an overall R^2 score of 0.50 when assessing species proportions, though when mapping majority proportion an overall accuracy of 0.73 is achieved. Authors note that combining remote sensing and deep learning is advantageous in scenarios such as forestry. Gazzea et al. (2022) use imagery from WorldView-2 at 0.5 m / px to classify pixels as pine, spruce, or deciduous, and achieve F1-scores of 0.883, 0.648, and 0.761 respectively. Overall, classifying tree species from satellite-collected imagery is a problem that has been explored in research. A method for classification for tree classification in this study is outlined.

The classification of fuels into coniferous and deciduous categories is a useful fuel property. As mentioned, coniferous fuels drive wildfires with increased severity and intensity, and are considered to be hazard fuels (Van Wagner, 1977). Alternatively, deciduous fuels can often limit the spread of crown fires (Quintilio et al., 1991). Tree classification, at least into coniferous / deciduous categories, is a requirement for exposure assessment. To this end, winter and summer imagery is purchased for the same area. While all trees are detected utilizing summer imagery, winter imagery can be utilized for tree categorization. Winter imagery is intentionally purchased after deciduous trees have gone dormant, but prior to significant snow accumulation on trees. Tree boxes generated by the model utilizing the summer imagery are then drawn on the winter imagery; as the imagery is georeferenced, it is assumed that boxes generated using summer imagery will still box trees in the winter imagery. The fact that deciduous trees go dormant, lose leaves, and 'turn brown' is leveraged for classification. Trees are categorized into 'brown-in-winter' and 'green-in-winter' categories, a proxy for coniferous and deciduous vegetation.

An unsupervised classification method, k-means clustering, is utilized for tree classification in satellite imagery. Unsupervised classification refers to the ability of the model to classify without the use of training data. Comparing to the supervised YOLOv5 CNN model utilized for tree detection, which is fed a set of imagery and corresponding annotations to train model weights to identify trees, unsupervised classification methods

make classifications directly from the set of data they are provided. This is advantageous, as it does not need a training set of annotations to be manually generated, which is often a highly time-consuming process, and difficult to perform for two reasons. Firstly, it is difficult to distinguish between a coniferous and deciduous tree in summer imagery due to the 0.5 m resolution of the satellite imagery. While in high-resolution RPAS imagery the different characteristics of each tree type are clear, many of the differentiating features of coniferous and deciduous trees are not discernable at 0.5 m / px resolution. Secondly, it is difficult to distinguish 'bare' or 'dormant' deciduous trees in winter imagery at the available resolution. Dormant deciduous trees at 0.5 m / px resolution are extremely difficult to reliably identify, especially as surrounding vegetation also enters dormancy, loses leaves, and 'turns brown'. Recalling Figure 9, the figure demonstrates an area in both winter and summer imagery; the issues with visually classifying trees as coniferous or deciduous in the summer imagery and with identifying individual deciduous trees in the winter imagery can be clearly seen.

These two reasons mean it is an impractical task for a human to build a training dataset, as it would require the annotator to either classify trees using summer imagery, or to identify trees in winter imagery; both tasks that are likely to result in high error rates. However, once tree boxes are drawn on summer imagery, transferring the boxes to winter imagery and calculating their color values is a simple task. Finally, it is useful to detect all trees in summer imagery and later classify as coniferous or deciduous, rather than simply to detect all green trees in winter imagery, as deciduous trees can still impact fire spread and danger.

Once tree detection boxes are transferred to the winter imagery, average color values of each boxed tree are calculated. The red (R), green (G), and blue (B) values of pixels within each bounding box are averaged, so that each tree has an average R, G, and B value, shown in Figure 15 for a green-in-winter and brown-in-winter tree in winter and summer imagery.

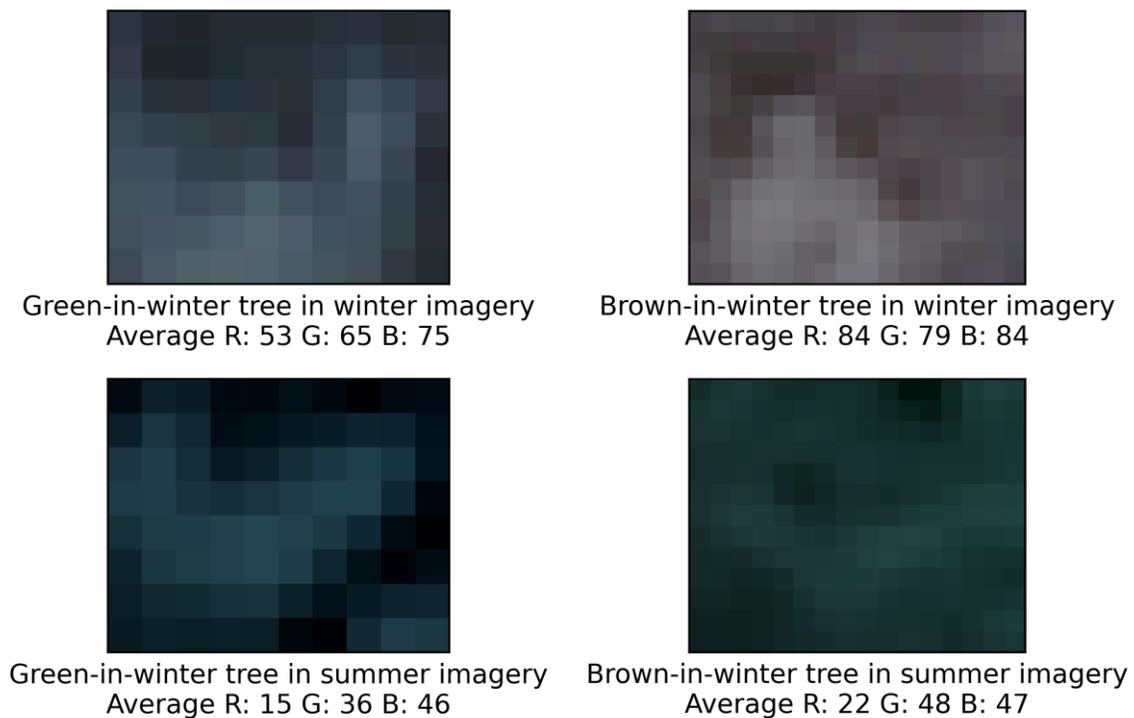


Figure 15: Green-in-winter and brown-in-winter trees, and associated average RGB values, in winter and summer imagery

The k-means clustering algorithm is then applied. Clustering algorithms separate data into groups based on the data itself (that is, there is no comparison or training set). In particular, k-means clustering iteratively calculates the 'centre' of each cluster, assigns data nearest to the centre to that cluster, calculates the new centre, and repeats (Muller & Guido, 2016). This is done until all data is assigned to a cluster. Inputs consist of the data itself, in this case the average RGB values of each tree, and the number of clusters anticipated for this dataset, in this case two (green-in-winter and brown-in-winter). Figure 16 visualizes the input data for the Lobstick, AB survey.

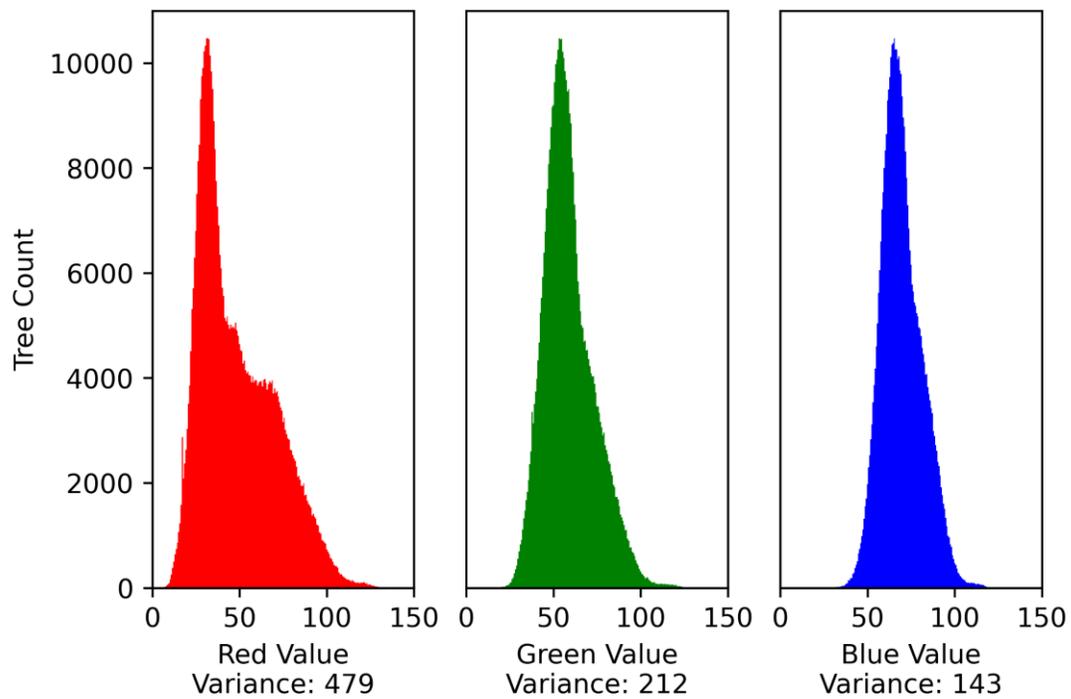


Figure 16: R, G, B distributions for trees detected in Lobstick imagery

As can be seen, much of the variation appears in the red value of the trees. This is logical as in RGB color space, increasing a red value while holding green and blue values similar and constant results in a color approaching brown. Finally, the results of the k-means clustering algorithm applied to the Lobstick imagery across each of the R, G, and B dimensions are plotted in Figure 17.

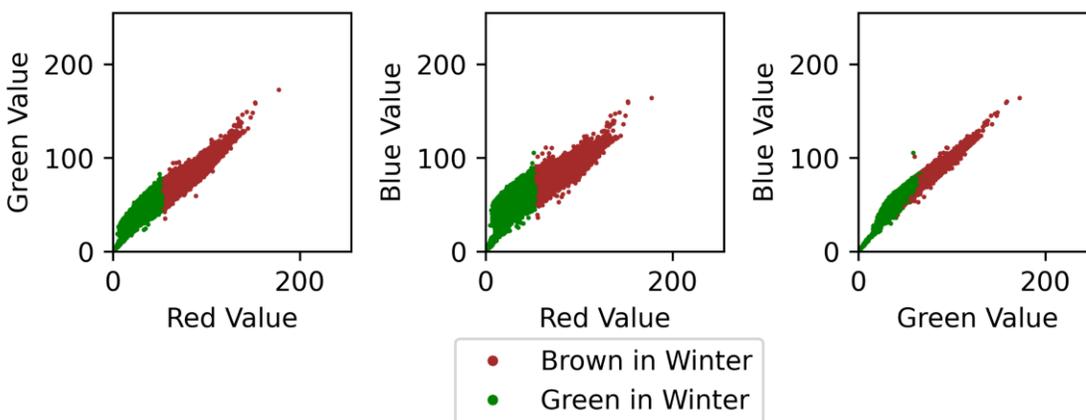


Figure 17: Results of k-means clustering algorithm across the R, G, and B dimensions

As can be seen in Figure 17, cluster boundaries are difficult to distinguish. An assessment procedure is defined to quantify the accuracy of the k-means algorithm at classifying individual trees.

To assess the results of the classification scheme, boxes are compared to manual annotations generated using RPAS survey data. As both RPAS surveys and satellite imagery are georeferenced, boxes from one can be overlaid on the other. This is taken advantage of for the purposes of satellite tree classification assessment. As previously shown, coniferous and deciduous trees can be clearly distinguished by a human annotator in RPAS imagery. Thus, manual RPAS annotations can be taken as the true class of each tree, and the classifications generated by the k-means algorithm are compared. Manual annotations performed in RPAS imagery and detections generated by the satellite model are overlaid. A search is performed for satellite detection boxes and manual RPAS annotation boxes that box the same tree (defined at various IoU thresholds). Figure 18 depicts a manual annotation box made using RPAS imagery as well as its corresponding satellite model detection box.

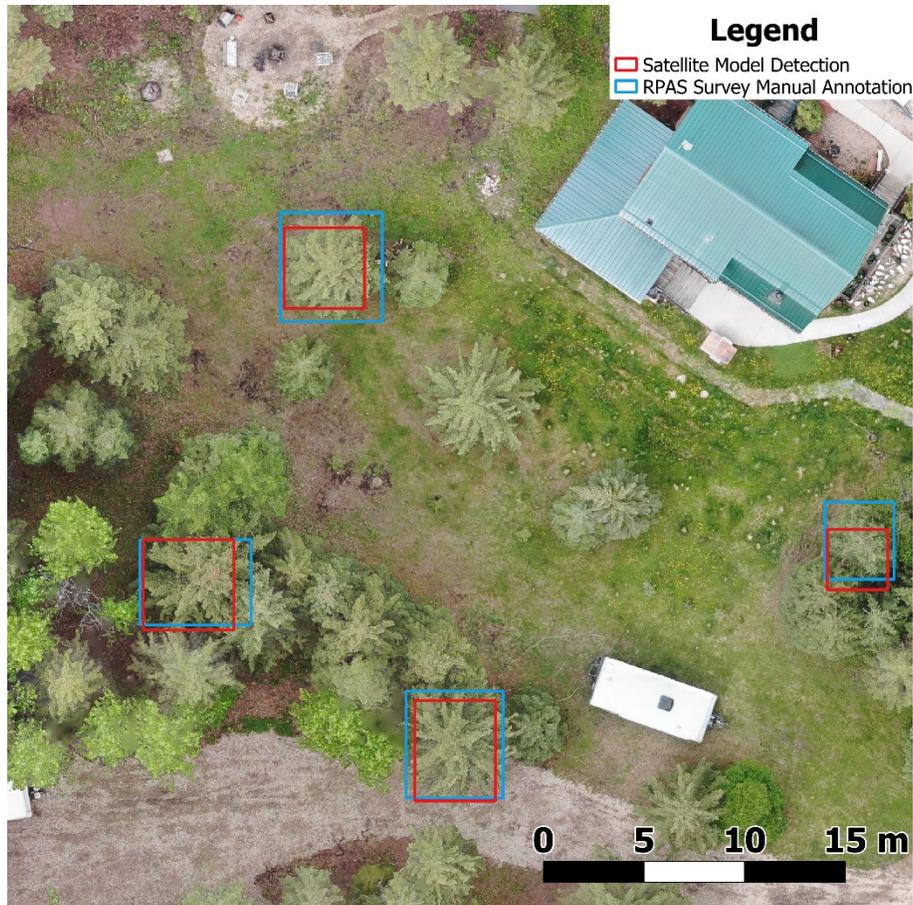


Figure 18: RPAS manual annotation boxes and matching satellite model detection boxes overlaid on RPAS survey imagery at an IoU threshold of 0.5

As mentioned, manual RPAS boxes and satellite model prediction boxes are said to box the same tree if they overlap with a certain IoU. Due to the issues mentioned above with measuring IoU in satellite detection boxes, this is performed at multiple IoU thresholds. In Figure 18, the IoU threshold used to define a match is set at 0.5; as this is the most strict threshold, only trees that are precisely boxed by both the RPAS and satellite models will be considered. Taking the manual RPAS annotations as the ground truth, the accuracy of the satellite classification scheme is assessed. Results are included below in Table 9.

Table 9: Satellite tree classification accuracy results

IoU	Total RPAS Annotations	Total Matches Found	Conifers Correct	Conifers Incorrect	Deciduous Correct	Deciduous Incorrect	Conifer F1	Deciduous F1
0	15023	15022	9559	474	3298	1221	0.92	0.80
0.1	15023	4368	1852	193	1570	581	0.83	0.80
0.2	15023	1924	670	96	784	298	0.77	0.80
0.3	15023	814	278	42	352	111	0.78	0.82
0.4	15023	358	110	17	163	57	0.75	0.82
0.5	15023	140	46	6	71	14	0.82	0.88
Average:							0.81	0.82

As can be seen, the results of the classification scheme are promising. Overall, the classification procedure yields an average F1-score of 0.81 and 0.82 respectively. Results of the clustering are also assessed visually. The image at Seba Beach shown earlier is examined in Figure 19 with prediction boxes separated the two classes.

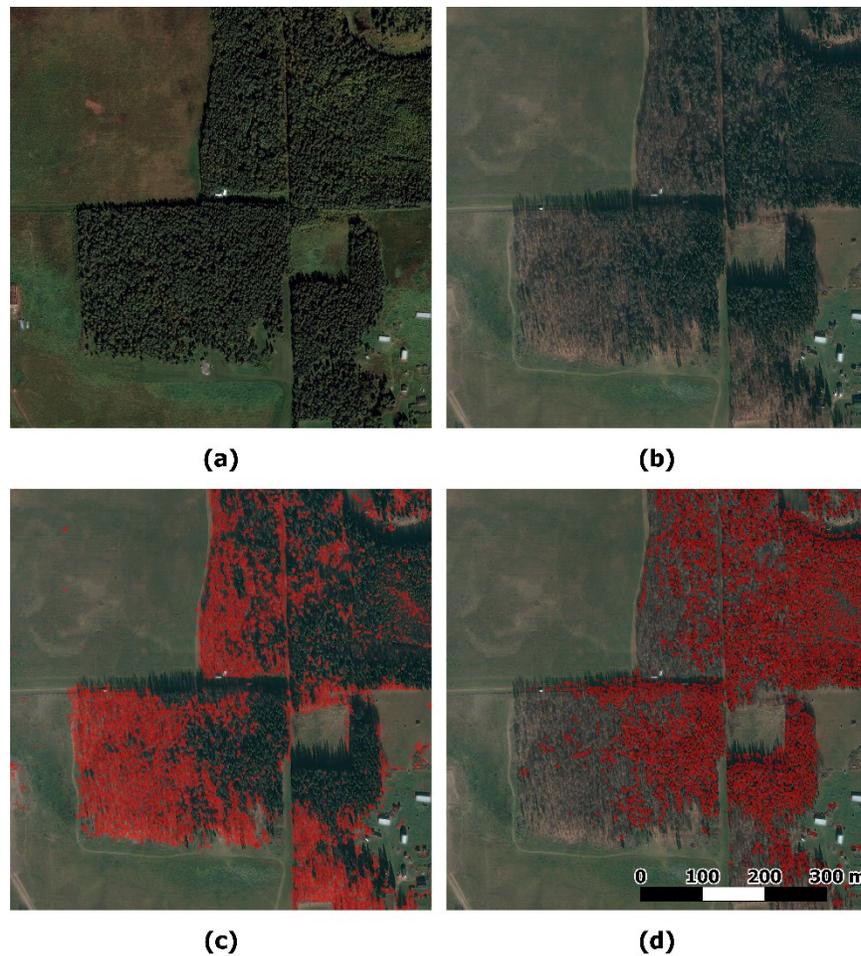


Figure 19: Results of k-means classification at Seba Beach, AB, showing (a) summer imagery, (b) winter imagery, (c) detections classified as brown-in-winter, and (d) detections classified as green-in-winter

As can be seen, the method separates detections into green / brown classes. These detections can then be used to determine where hazardous fuels may exist.

Comparisons can be drawn between the model outputs and a fire behavior prediction (FBP) system fuel type map generated by (A. Beaudoin et al., 2014), shown in Figure 20. Both maps in Figure 20 identify coniferous fuels in the west, north, and east portions of the image, while trees in the south, northwest, and southeast portions can be classified as deciduous. The detections output by the satellite tree detection model are shown to be at a much greater resolution than the FBP fuel map. There is some disagreement in the northeast (highlighted blue); closer examination of the winter satellite image suggests the

fuel map may be in error as the area is classified as fuel type O-1 (grass) in an area where trees are seen to exist. This error may be due to an error in the generation of the fuel map, or the time difference between the creation of the fuel map and the collection of satellite imagery used in this work.

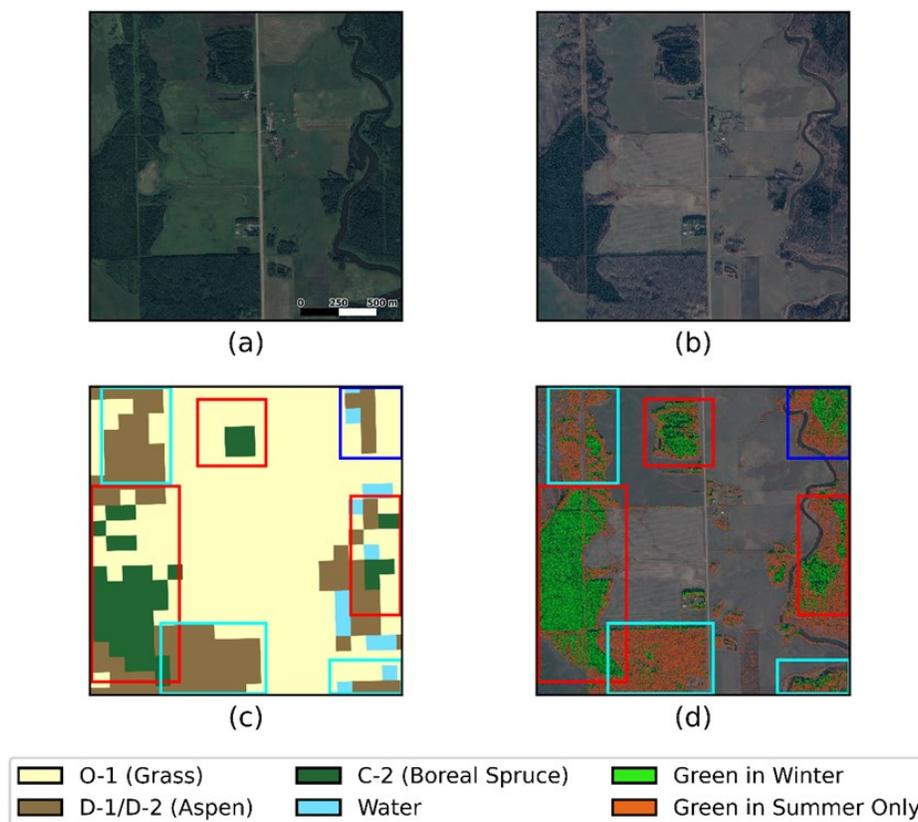


Figure 20: Fuel map comparison, including (a) summer imagery, (b) winter imagery, (c) Alberta 50 m fuel map (A. Beaudoin et al., 2014), and (d) satellite model detections

3.5 Satellite Model Discussion

A number of considerations must be made when viewing or utilizing the results of tree detection from satellite imagery. Fuel detections must be understood in the context in which they are generated. Satellite imagery resolution is limited, and low when compared to aerial or RPAS systems. A resolution of 0.5 m / px means that even large trees that have crown widths of multiple metres will only be represented by a few pixels. This makes detections susceptible to changes in the environment, such as lighting. This also means

that trees at or smaller than this resolution are impossible to detect. Thus, detections represent only large overstory trees.

The accuracy of the model must also be considered. Even under the conditions of the study, the model struggles to reliably detect every tree individually. However, as shown in the density evaluation, detections considered at a large area extent correlate well with manual annotations, reaching an R^2 of 0.76. Tree detections from satellite imagery using this process should then be used in applications where the scale of the task being performed covers large areas, and not in circumstances where the detection of every individual tree is required or can significantly change the output of the task or assessment.

The classification of trees output from the process must also be considered in context. While a logical proxy for coniferous or deciduous classes, 'brown-in-winter' and 'green-in-winter' classifications may not necessarily represent the reality of the fuel environment under examination. Trees may turn brown for various reasons, not simply because they are seasonal deciduous species. For example, coniferous trees afflicted with disease or pests (such as the mountain pine beetle) may shift color towards red or brown as the disease progresses, while still being a hazardous wildfire fuel. Misclassifications can also occur when the area under examination contains deciduous conifers such as tamarack which shed their needles in the fall. Finally, classifications are made based on the dominant colour in the overstory; this work does not detect understory fuels.

There is also an associated error in tree classification from satellite imagery. Though an F1-score above 0.8 is achieved, not every tree is correctly classified when compared to manual annotations done on high-resolution RPAS imagery. Error can come from several sources. In the assessment performed, the RPAS imagery was captured in 2022, the summer imagery in 2019, and the winter imagery in 2015. As all are used and compared in the classification and assessment procedure, the time lag between surveys could be a source of error. As trees grow, age, and are felled, their appearance can change drastically, especially over a time period of 7 years (2015 to 2022). This could result in trees being detected in summer imagery that do not appear in winter imagery, meaning the classification algorithm has incorrect data to base its classification on. This could be

resolved with updated satellite imagery, should it be available. Another source of error in classification could be the difference in viewing angle between summer and winter imagery. While the summer and winter imagery are of the same area, it is not guaranteed they are captured from the same viewpoint; the imaging satellite may capture the same area on a different route. This could cause issues with detection boxes made in summer imagery aligning with trees in winter imagery: the detection box will still align with the trees geospatial location; however, the tree may be viewed from a different angle of incidence, resulting in background pixels being included in the detection box and used in the color-based classification scheme. Figure 21 demonstrates this effect.

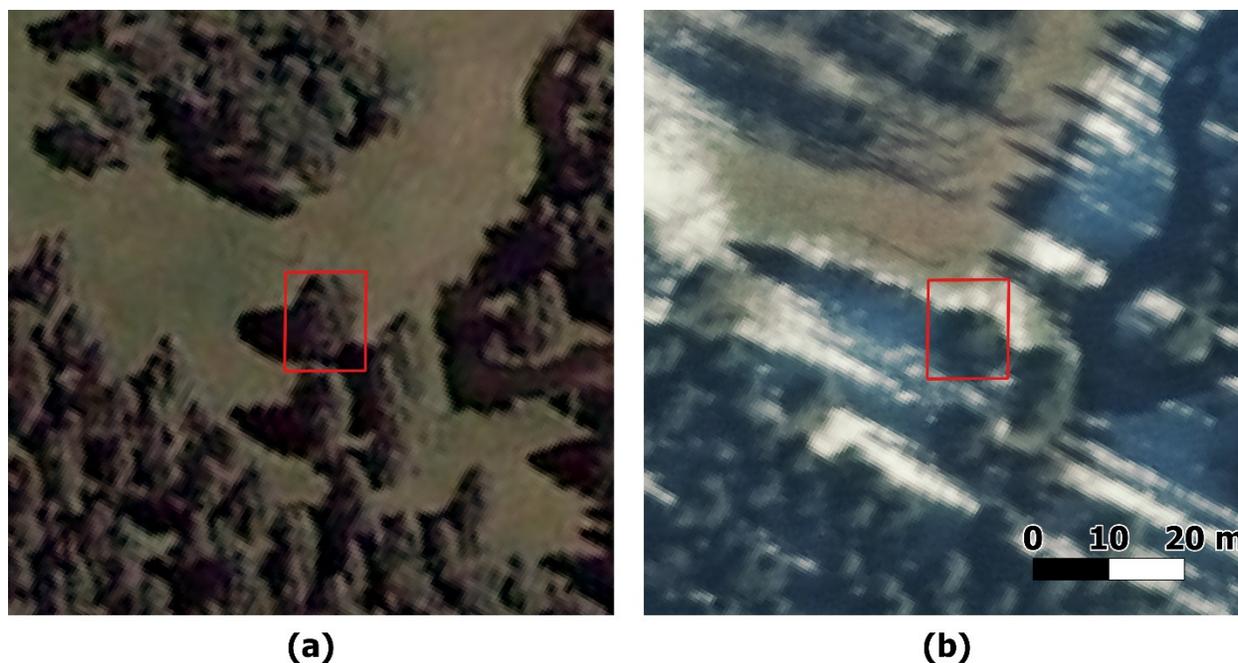


Figure 21: Tree detected in (a) summer imagery at high angle of incidence; and (b) tree detection box overlaid on winter imagery at different angle of incidence includes background pixels

Figure 21 depicts a tree in the Bragg Creek survey. This survey has the most pronounced effect as the difference between viewing angles is the greatest. As can be seen, more background pixels are included in the box when overlaid on the winter imagery. While classification still reaches F1-scores of above 0.8, classifications should be interpreted and used only when the nature of the area and forest under examination is well-

understood. Finally, while RPAS imagery offers very high-resolution images of trees for human classifiers to interpret, it is suggested that field work be performed in the future to validate the classification procedure outlined above.

Field work should also be performed to validate the satellite tree detections. The R^2 score of 0.76 was reached when comparing detection density to manual annotation density. As the satellite tree detection workflow can only detect large overstory trees, surveyed areas should be measured to assess the relationship between satellite detections and field-measured forest densities.

Future work should include the aggregation of individual tree detections into stand-level fuel maps. As discussed in Chapter 2, the proposed workflow offers individual tree detection, while certain applications require stand-level attributes. As such, future work should explore the aggregation of individual tree detections into fuel maps at any scale. Utilizing additional spectral bands for tree classification is also of interest as a possible way to classify trees by species rather than 'green-in-winter' or 'brown-in-winter'.

While subject to a number of limitations, tree detections from satellite imagery are still a useful and informative basis of information when these limitations are understood. The method proposed offers a way to generate general inferences about hazardous wildfire fuels rapidly and across massive extents. This could be used to generate rough maps of hazard fuels across any area where satellite imagery is available; particularly around the 32.3 million hectares of WUI (Johnston & Flannigan, 2018). With the limitations in mind, one possible application of this method of fuel mapping is the generation of remote community level wildfire ignition exposure assessments for areas of WUI.

4.0 Remote Exposure Assessments

4.1 Structure Exposure Assessments

4.1.1 Methodology

A structure-level exposure assessment determines a value-at-risk's (VAR) exposure to wildfire. VARs can take many forms: dwelling structures, infrastructure, or other assets that are deemed to have value. Determining a VARs exposure to wildfire is of interest as it allows hazards to be preemptively mitigated and can aid in emergency planning before actual wildfire encroachment. FireSmart Alberta, a wildfire preparedness initiative in Alberta, defines an exposure assessment procedure for individual structures to improve wildfire outcomes. There are many factors that go into an exposure assessment, including local vegetation, topography, and building materials (FireSmart Alberta, 2015). Structures assessed using the FireSmart assessment are typically along the wildland-urban interface (WUI) – the border where wildlands and urban areas meet. Information is gathered by visiting the structure and assessing both the structure itself and its surrounding environment. The structure then receives a numerical score reflective of its exposure to wildfire.

The FireSmart Alberta home assessment procedure defines several zones around the structure and considers vegetation within these zones. Exposure is elevated if there are any coniferous trees within 10 m of the structure. In the 10 m to 30 m zone, the surrounding forest is assessed as deciduous, mixed wood (coniferous and deciduous), or coniferous. Further, if the surrounding forest is coniferous, it is determined whether the forest is separated, with wide gaps between tree crowns, or continuous, where trees are tightly spaced (FireSmart Alberta, 2015). Scores from each zone are then combined and the structure receives a final assessment rating as low, moderate, high, or extreme. In the case of high exposures, FireSmart can provide homeowners with advice on how to lower their exposure rating by modifying vegetation or making changes to their home, resulting in a house less threatened by encroaching wildfires.

The applicability of the remotely piloted aerial system (RPAS) tree detection and classification workflow to a structure scale exposure assessment is explored. This

workflow is more appropriate for application at this scale as it is more accurate in both detection and classification and can detect smaller trees due to the increased resolution when compared to satellite imagery. Application is straightforward: flights are conducted in a portion of the WUI, and surveys are stitched using the same method defined in Chapter 2. Once stitched, the model developed in Chapter 2 is applied to detect and classify trees within the survey. The location of the inhabited structure is manually digitized. As the RPAS survey is geospatially projected, detections can then be automatically binned based on distance to structure. Distance is measured from the edge of the structure to the edge of the tree bounding box. Finally, structures can be examined visually from the imagery with detections overlaid, and the number of coniferous trees in each of the FireSmart zones can be counted.

4.1.2 Discussion

The workflow suggested above is applied. FireSmart recommends no vegetation within 1.5 m of the VAR, only deciduous within 10 m, and only deciduous (with a limited amount of widely spaced conifers also being acceptable) within 10 m to 30 m. Figure 22 is generated for a home in the WUI and coniferous tree detections in each zone are counted.

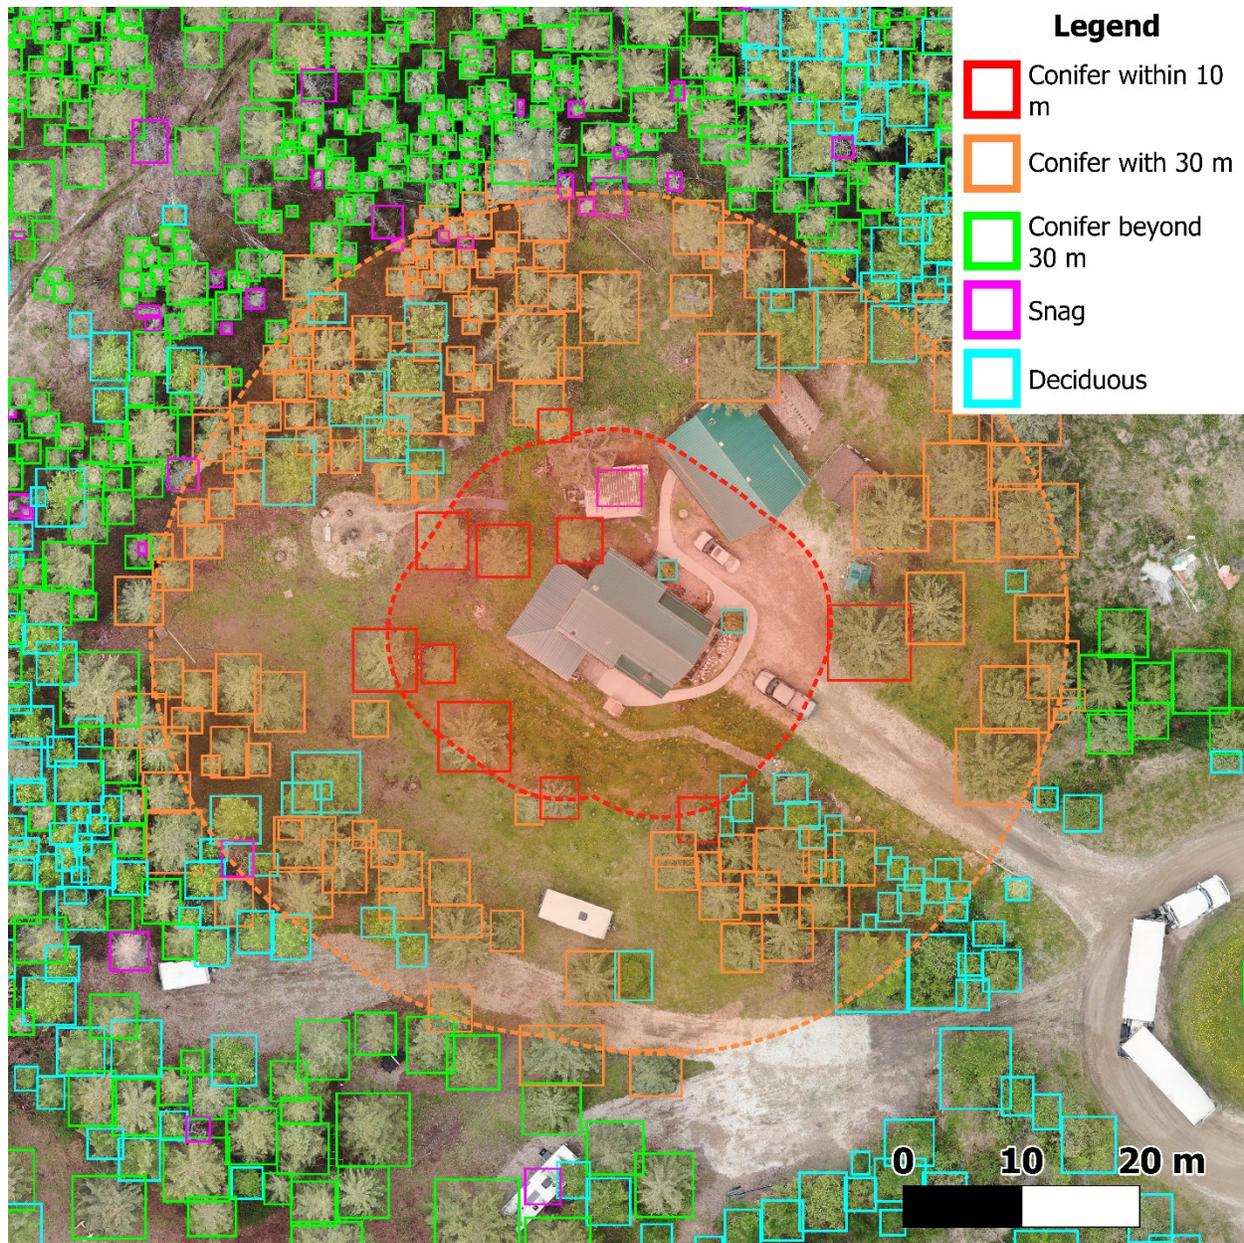


Figure 22: RPAS model automated tree detections, colored by distance to inhabited structure based on FireSmart Alberta (2015) home assessment guidelines

In Figure 22, the model detects 1 conifer within 1.5 m, 9 between 1.5 m and 10 m, and 133 conifers between 10 m and 30 m. This could be quickly performed for every structure in collected surveys. As demonstrated, even preliminary analysis of model outputs can result in information that may be useful for VAR hazard assessments.

The proposed workflow could be used to quickly assess overstory tree vegetation around a structure in the context of a structure scale exposure assessment. The developed model can be applied to imagery collected at any distance from the structure. Further, the coniferous / deciduous classification allows for hazard fuels to be identified. If an entire community is surveyed, this workflow presents a remote and automated option for counting hazardous trees around structures in each of the FireSmart zones. Spacing between tree crowns is another consideration in the FireSmart Alberta home assessment procedure that could be assessed using this method. The use of orthographic and geospatially projected imagery allows the viewer to assess the spacing between trees as well as tree size.

A number of limitations are present in the proposed workflow. Notably, this procedure remains ignorant to aspects in the structure scale exposure assessment other than tree vegetation. Structure scale exposure assessments include factors such as building materials, yard tidiness, deck and roof condition, and topography (FireSmart Alberta, 2015). These are not assessed by the proposed workflow, and thus it is suggested that the proposed workflow supplement rather than replace existing methods. The model achieves a detection F1-score of 74.5% and classification accuracies above 90%, meaning that errors still occur. In the case of structure level exposure assessments, the omission of even one hazardous tree in close proximity to the VAR can affect its rating. This is especially relevant as this is a safety-critical-setting, and it is reiterated that this study explores the automation of one aspect of exposure assessment and does not make the case for replacing existing practices. Errors in detection and classification will also occur if the model is applied outside of the domain of its training data (in this study, surveys were collected in boreal, rocky mountain, and foothill regions in Alberta, Canada). As convolutional neural network (CNN) architectures continue to advance, and as more data is added to the training dataset as additional surveys are collected, accuracies in both identification and classification are expected to increase. Another important part of the exposure assessment workflow is the opportunity to connect with the homeowner and open a dialogue about best practices in preparation for a wildfire. Even if some aspects

of the assessment process can be automated, it remains important to keep humans in-the-loop in both the assessment and decision-making processes.

4.2 Community Exposure Assessments

4.2.1 Methodology

As previously discussed, exposure assessments are used by communities to determine areas in the WUI that may have elevated exposure to wildfire. Once these areas are identified, preplanning activities and even mitigation can take place that is informed using fuel maps and exposure maps. In this section, the application of the satellite tree detection and classification model (Chapter 3) to exposure assessments for communities in the WUI is explored.

Beverly et al. (2010) outline the first method wildfire exposure assessment method for the built environment using spatial analysis. As mentioned, community bounds are defined, and a 5 m grid is overlaid on the community. Hazard fuels are identified and used to generate wildfire exposure maps for three wildfire ignition types (radiant heat, short-range embers, and long-range embers) by measuring the proportion of hazard fuel in a defined radius around a grid cell. Fire behavior prediction (FBP) fuel types C-1 (spruce–lichen woodland), C-2 (boreal spruce), C-3 (mature jack or lodgepole pine), C-4 (immature jack or lodgepole pine), C-7 (ponderosa pine–Douglas-fir), O-1 (grass), and M-2 (boreal mixedwood) are defined as hazard fuels (Beverly et al., 2010). While the satellite model cannot identify FBP fuel types to this degree of specificity, it can classify detected trees into a green-in-winter and brown-in-winter proxy for coniferous / deciduous. It is also worth noting that, aside from grass, the FBP fuel types mentioned previously are coniferous or mixedwood fuels. Beverly et al. (2010) also note that deciduous forests are generally not considered as fuels of concern. Finally, it is noted that Beverly et al. (2010) only consider grass to be a fuel of concern for radiant heat exposure and short-range ember transfer and not for long-range ember ignitions. The exposure map for long-range ember ignitions is considered because the satellite model is designed to identify and classify overstory trees. Ember ignitions are of particular interest; previous post-wildfire investigations have shown that the majority of destroyed structures are ignited via ember exposure, such as

in the 2016 Fort McMurray wildfire (Westhaver, 2017). Moreover, it has been shown that high wildfire exposure correlates well with wildfire burn maps (Beverly et al. 2021).

Applying the method described by Beverly et al. (2010) using the data output by the satellite model is straightforward. First, a WUI community in Alberta is selected, and satellite imagery purchased. To this end, Lobstick, AB will be used to demonstrate remote exposure assessment generation. Lobstick is a community in western Alberta that is surrounded by C-2 and D-1 fuels, shown in Figure 23. Further, the portion of the community that is selected consists of rural acreages that have significant intermingling of structures and the forest. Figure 23 shows both the satellite image of the area, as well as an FBP fuel map (André Beaudoin et al., 2022).

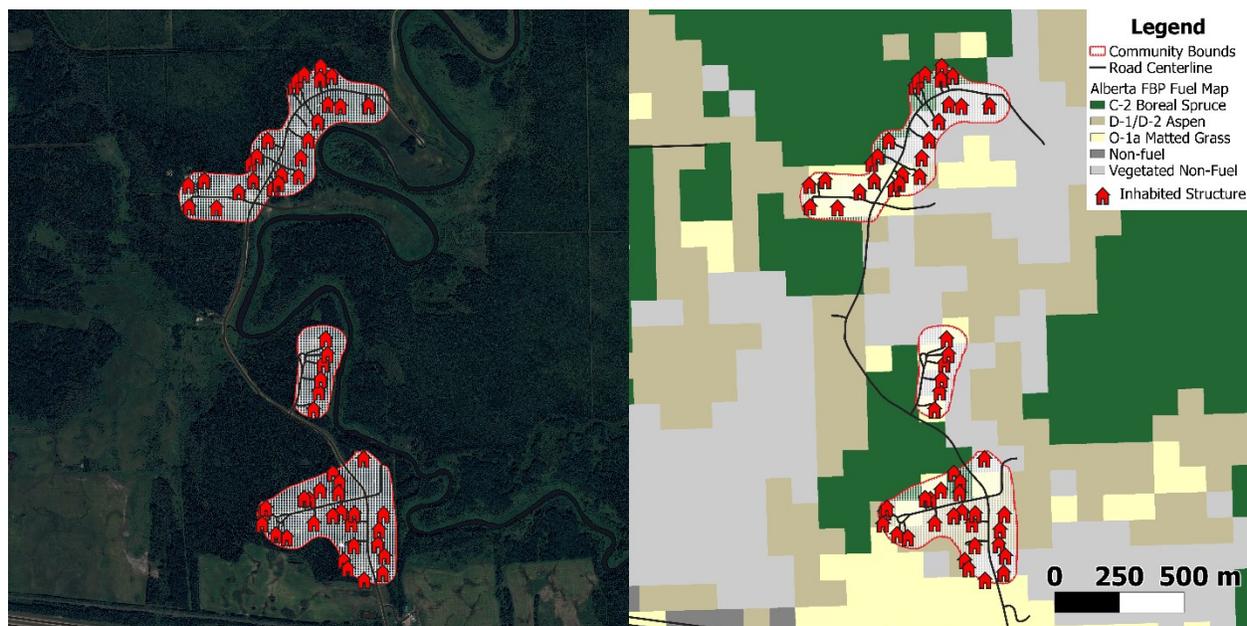


Figure 23: Satellite image and fuel map near Lobstick, AB

The model (Chapter 3) can be applied to the satellite image obtained. The tree detection boxes can be drawn on the satellite image (Figure 24). The satellite image and detections have been rotated to display the trees more clearly to the reader due to the off-nadir viewing angle. Next, tree classification is performed using winter imagery. Figure 25 demonstrates the winter image, along with the tree detections separated by class.



Figure 24: Satellite tree detection boxes

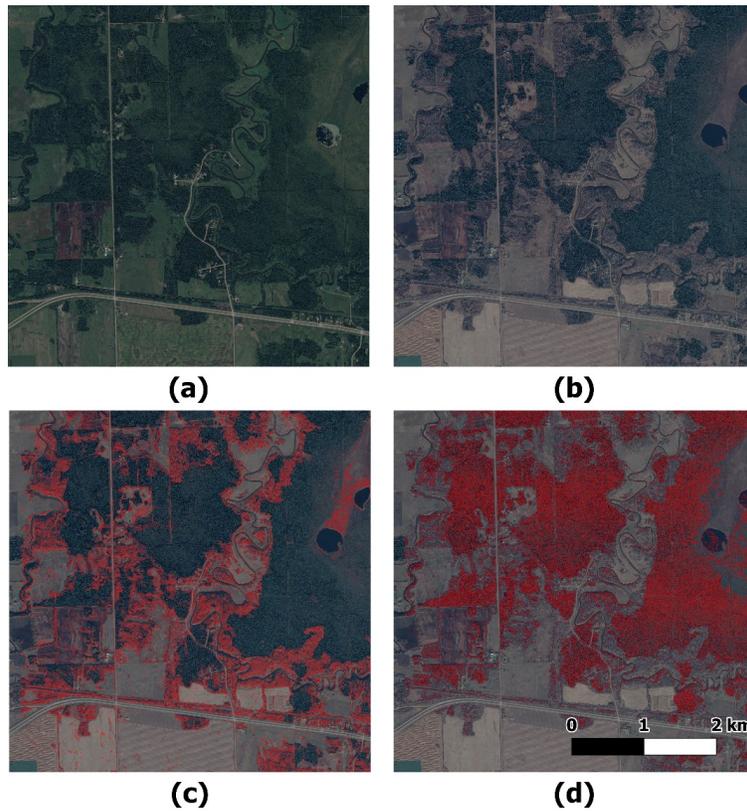


Figure 25: Satellite classification results at Lobstick, AB, showing (a) summer imagery, (b) winter imagery, (c) detections classified as 'deciduous', and (d) detections classified as 'coniferous'

Once classified, the tree detections are utilized as a fuel map for the exposure assessment procedure. All 'green-in-winter' tree detections are taken as hazardous fuels. The wildfire exposure assessment can be performed using the workflow in Beverly et al (2010). Figure 26 demonstrates how the assessment is calculated for long-range ember transfer. A circle with a radius of 500 m is placed at each grid cell in the community and the fuel groundcover proportion is calculated.

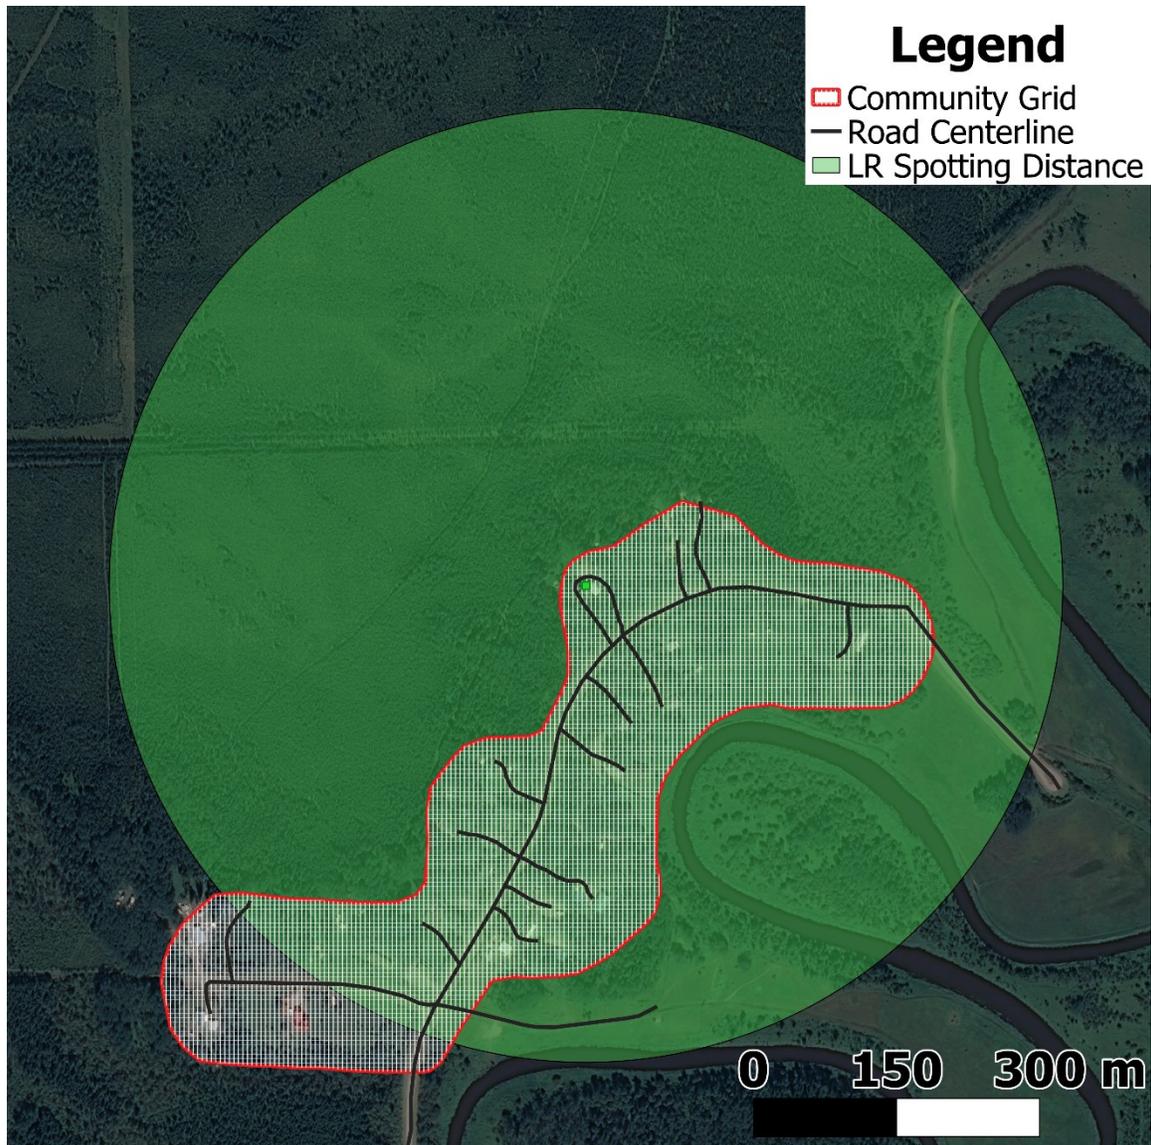


Figure 26: Long-range spotting distance for a cell in the community

This can then be repeated for each grid cell within the community. Once nearby fuel groundcover proportions are calculated for each grid cell, thresholds included in Beverly et al. (2010) are used for converting fuel proportions to a hazard level. Table 10 lists the thresholds used. The remote exposure assessment is performed using these thresholds (Figure 27).

Table 10: Exposure classes from Beverly et al. (2010)

Land-Cover Occupied by Detected Coniferous Fuels from 100 m to 500 m	
Nil	0
Low	> 0 - 0.15
Moderate	0.15 - 0.30
High	0.30 - 0.45
Extreme	> 0.45

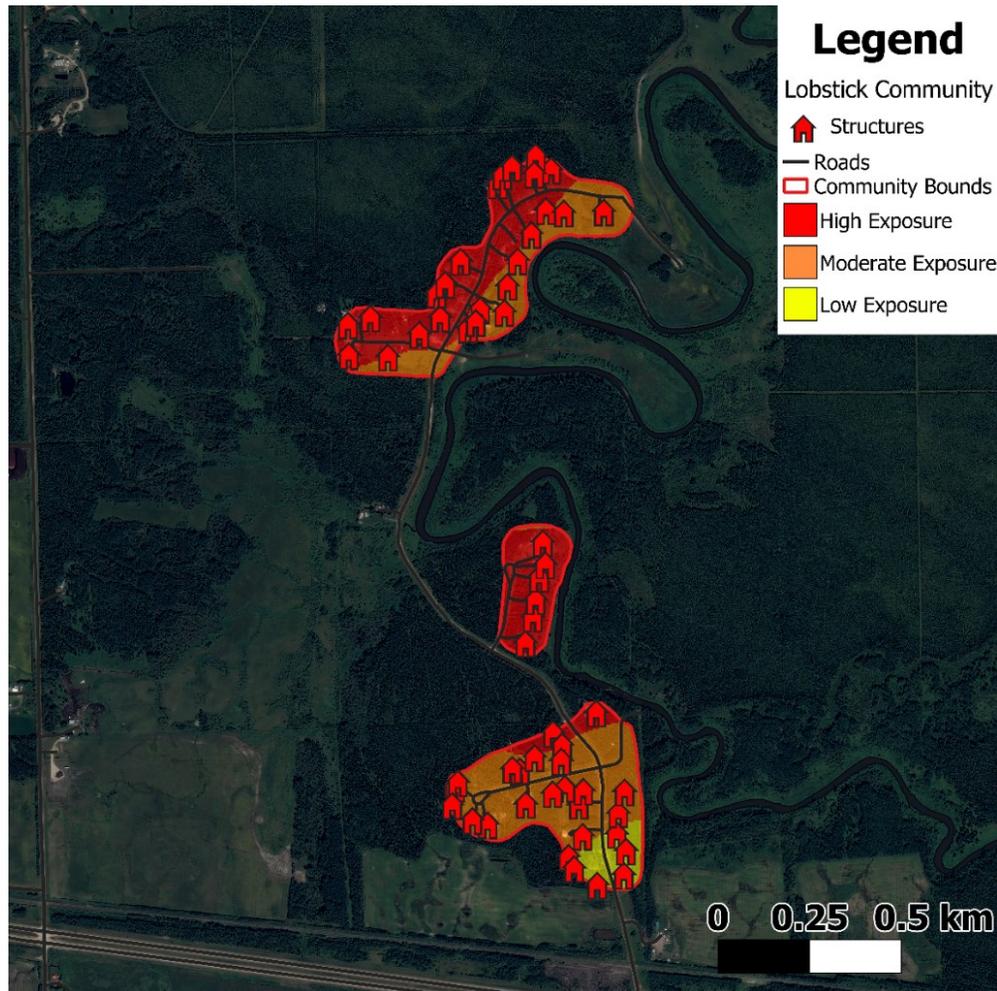


Figure 27: Remote community long-range ember exposure assessment, generated using satellite model detections

4.2.2 Discussion

The application of the satellite tree detection and classification workflow to community scale exposure assessments is explored. This workflow is fully automated (Figure 28), meaning information can be generated rapidly for use as input to existing decision-making processes.

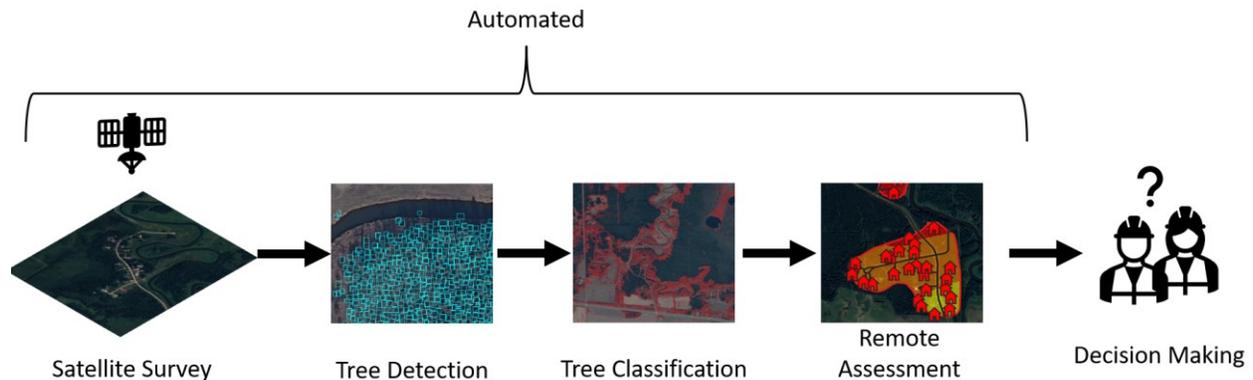


Figure 28: Remote community exposure assessment workflow

This workflow explores a first-step towards both the generation of extremely high-resolution fuel maps from satellite imagery, as well as the derivation of high-resolution exposure maps from said fuel maps. Compared to manual interpretation the proposed workflow (Figure 28) is fully automated while also being remote, generating results quickly from RGB satellite imagery. In its current form, this process could be applied to any satellite imagery that is both within the domain of the study (i.e., in the boreal, rocky mountain, and foothills regions of Alberta, Canada) and within the specifications set in this study (0.5 m / px, > 30° off-nadir, clear imagery). This means that individual trees could be detected and classified as coniferous / deciduous at any time period in which acceptable satellite imagery is available, allowing the change in the landscape over time to be assessed if historical imagery is obtained for an area. The workflow can input the latest in-specification satellite imagery, guaranteeing up-to-date results.

The proposed workflow has several limitations. Stand attributes are omitted in the detection and classification of individual trees. The hazard fuels used in the above assessment are individual coniferous tree detections; these detections are not aggregated into stand-level patches. As such, mixedwood stands are not classified as they were in Beverly et al. (2010). However, a benefit of individual tree detection is the possibility for aggregation into fuel maps of any resolution. In Figure 29, a 30 m grid is overlaid on the area surrounding the community, and tree count is calculated within each cell. The proportion of green-in-winter trees within each cell is also calculated. Future

work should investigate the use of calculated metrics to aggregate detections into stand-level patches.

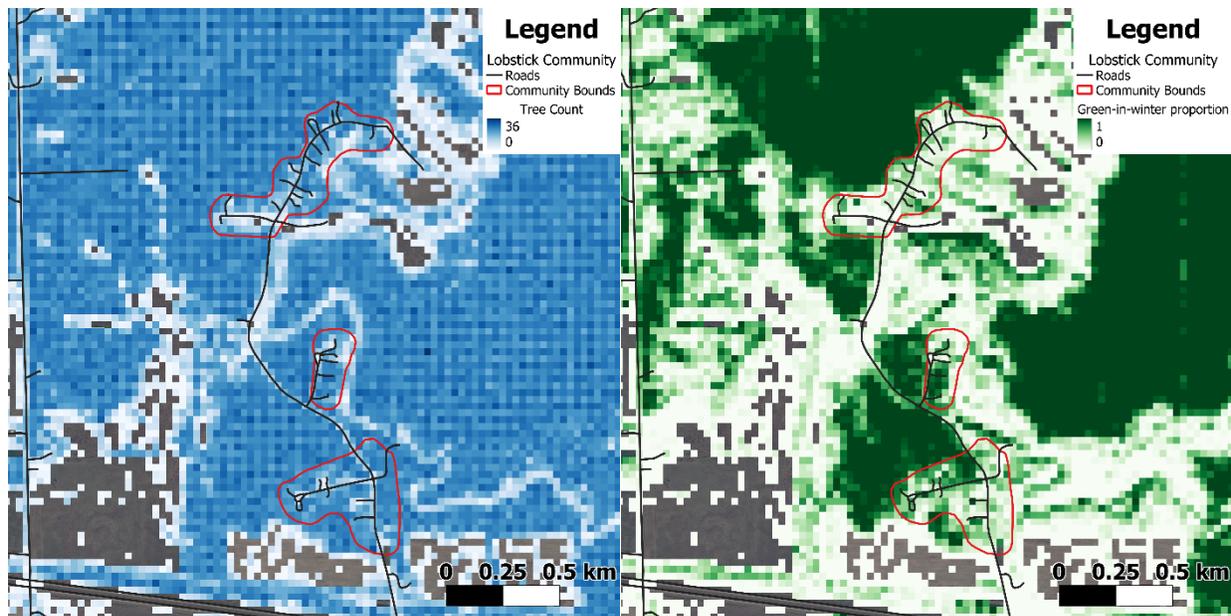


Figure 29: Tree count (left) and green-in-winter proportion (right) maps aggregated from individual tree detections

The mapping of grass as a hazardous wildfire fuel was omitted from this study. Methods for mapping grass have been explored in other research and could be incorporated into the proposed workflow. The limitations of the satellite tree detection and classification workflow are reiterated – this workflow only detects large overstory trees and classifies as coniferous or deciduous based on a proxy ('green-in-winter' or 'brown-in-winter') rather than coniferous or deciduous categories directly. This could influence exposure ratings when coniferous trees appear brown in the winter because of pests or disease, or difficult to classify species such as tamarack. These may biasedly lower the perceived community exposure should this proposed method be relied upon alone. The accuracy of the satellite tree detection and classification model can also limit the applicability of this fuel mapping method. The model developed in Chapter 3 correlated tree detections to human interpretation of the same imagery at an R^2 of 0.76, with classification scores above 80%. As wildfire planning is a safety critical task, all sources of error and limitations must be carefully considered, understood, and communicated when new methods are explored.

This section explores a wildfire exposure assessment, rather than a full threat or risk assessment. The R11 Forest Management Plan, a document that outlines forest management activities for the R11 Forest Management Unit in Alberta Canada, provides an example of a wildfire threat assessment (Alberta Sustainable Resource Development, 2007). The threat assessment contains four components – fire behaviour potential, fire occurrence risk, values at risk, and suppression capability. While the proposed method of mapping wildfire fuels identifies and classifies trees, its outputs are not sufficient to inform fire behaviour potential mapping, which considers additional fuel attributes, topography, and climate. Finally, work has recently been performed investigating directional vulnerability to wildfire (Beverly & Forbes, 2023). The work in this study consisted of omnidirectional exposure assessments and did not consider how wind trajectories can influence community exposure.

The availability and recency of satellite imagery should be considered. Satellite imagery is not always available at the resolution and specifications required. Despite having return periods of up to 1 day, this varies significantly depending on the location of the area of interest. Further, conditions may be unfavorable for satellite imagery capture on the day the satellite passes the area of interest. In Canada, the boreal forest may be snow covered or dormant for much of the year, leaving only a few months where ‘summer’ imagery can be collected. The collection of this imagery may be further impeded by clouds and other atmospheric conditions. Depending on the orbit of the satellite capturing the area, the angle off-nadir may also be too significant for the algorithm to be reliably applied. These limitations can result in a time-lag between the present day and the date at which suitable satellite imagery can be collected. Imagery used in this chapter was purchased between 2017 and 2021.

The use of the satellite tree detection and classification workflow (Chapter 3) is explored. This study represents a first-step towards automated near real-time high-resolution fuel and exposure maps. This workflow is not intended to replace existing procedures and has limitations. There are multiple areas for future work. Additional validation through comparison with ground measurements and manual classification would improve the assessment of classification accuracy. The generated remote community exposure

assessments should be compared to existing validated community assessments to determine differences. This may be challenging to perform, as a fair comparison requires that a community assessment be performed at the same time acceptable satellite imagery is available. Further, the detection of other hazard fuels, such as grass, should be explored to extend the analysis of the fuel environment – the current assessment is limited to long-range ember exposure (a process that only includes tree fuels). The aggregation of individual tree detections into stand-level attributes should be further explored.

5.0 Conclusion

5.1 Summary

A workflow is presented wherein tree vegetation in the overstory can be detected from remotely sensed (RS) data products including remotely piloted aircraft system (RPAS) and satellite RGB imagery. Detections are performed utilizing a convolutional neural network (CNN); classifications are also generated for detected trees. The CNN trained on RPAS imagery is able to detect trees and to classify as coniferous, deciduous, or snag with detection F1-scores reaching 74.5% and classification accuracy above 90% across all classes. Satellite imagery is used to train a separate CNN to detect trees, and classifications are then made utilizing a k-means clustering algorithm. Satellite tree detection densities reproduce manual annotation densities with an R^2 of 0.76, and classification accuracy is assessed to be above 81% for all classes. The potential use of these detections is demonstrated for exposure assessment. Satellite detections are utilized in a community scale wildfire ignition exposure assessment, showing how satellite tree detections can offer high-resolution and recent data for remote automated assessments. RPAS detections are shown around a value-at-risk (VAR), and vegetation in each of FireSmart Alberta's priority zones is highlighted.

5.2 Contributions

A framework is presented for the automated processing of RS imagery into overstory fuel maps. A method of collecting fuel information using RPAS is shown to offer a high-resolution method for fuel detection at the individual tree scale. A similar method using satellite imagery is also proposed; a model is developed and used to convert RGB satellite imagery into overstory fuel maps. Finally, a possible use of the generated fuel maps in exposure assessments at both the community and individual structure scale is demonstrated.

5.3 Limitations

Several limitations to the work performed remain evident. The accuracy of the CNN tree detection models from RPAS and satellite imagery could be improved. As the field of computer vision progresses, models continually become more accurate and stable. As

such, new models and CNN architectures should be implemented as they become available to ensure tree detections and classifications are as accurate as possible. Accuracy could also be improved through further data collection containing diverse tree species in a variety of conditions to enrich the training dataset with additional examples.

It is also noted that wildfire fuels consist of more than overstory tree vegetation. All work in this study is performed utilizing overstory tree detections and classifications; surface fuels such as grass, litter, and other tree fuels in the understory are not considered. Further, validation is performed by comparing detections to manual annotations of the same imagery. It may be of interest to perform site measurements of stand species and densities to validate tree classification and tree detection accuracy.

This study took place in the boreal, foothills, and rocky mountain regions of Alberta, Canada. As such, application to new regions would require the collection of new datasets for training the tree detection models to detect and classify tree species that may not have been seen within this region. Application of the workflow in regions beyond the domain of this study was not explored.

Limitations also exist in the community exposure assessment workflow. The lack of grass detection limits the community exposure assessment to long-range embers. Grass detection should be added, though this task is not trivial as grass is only considered a hazardous fuel at certain lengths and moisture contents. Fuel maps generated via the satellite tree detection and classification workflow define hazard fuels as any location where a coniferous tree was detected and did not aggregate detections to the scale of tree stands, meaning hazard fuels consisted of overstory coniferous fuels only.

The resolution of satellite imagery limits tree classification, which in this work is performed using color information. This study used a 'green-in-winter' and 'brown-in-winter' proxy for coniferous and deciduous classification. While classifications were shown to be fairly accurate when compared to high-resolution RPAS annotations, if classifications are applied in areas where coniferous trees appear brown in color (for example, due to disease or pest), they will either be interpreted as deciduous or not detected at all. This proxy for coniferous and deciduous classification must be understood and communicated

when applying the algorithm to new areas. The model should only be applied in cases where the nature of the fuel landscape is understood so that any errors or issues can be identified. Additionally, the availability, recency, and quality of satellite imagery can be a limiting factor, as specific cloud cover and viewing requirements should be satisfied when purchasing imagery.

The use of RPAS detections in structure scale exposure assessments is subject to several limitations. Utilizing an RPAS platform can be a convenient way to collect aerial data but can be subject to flight limitations. Depending on the RPAS platform, battery life can limit the size of surveyed areas if no power is available for recharging at the survey site. Weather, site conditions, or legal restrictions can limit where and when RPAS surveys can be performed. Finally, the proposed exposure assessment workflow considers only local overstory tree vegetation, and thus should not be used as a substitute for complete FireSmart exposure assessments.

5.4 Future Work

Future work is proposed. Pace should be kept with advancements in CNN architecture to continue to increase detection and classification accuracy. Further, field validation for RPAS and satellite tree detections should be performed in order to improve accuracy assessment. The RPAS model is shown to be accurate in classifying coniferous, deciduous, or snag trees. The expansion of these classes into actual tree species, a useful characteristic in fuel mapping, should be explored with the collection of additional surveys to increase the size of the training dataset. Similarly, detecting grass from satellite imagery is also of interest; this would allow for a short-range ember and radiant heat exposure assessment to be performed. These exposure maps should be compared to existing community exposure maps to validate the procedure and contextualize any errors.

Future work could explore the generation of more detailed fuel maps from other data products. In the case of RPAS, a pointcloud is also generated when stitching the survey using photogrammetry. This could prove useful data to describe the 3D aspect of fuels. In the satellite workflow, other research has shown that additional spectral bands can be

used to classify tree species – incorporating this data into the proposed workflow could allow for the classification of trees as coniferous or deciduous directly, rather than the ‘green-in-winter’ or ‘brown-in-winter’ proxy. Finally, the aggregation of individual tree detections into stand-level fuel patches should be further investigated – detecting attributes on the individual tree scale could possibly allow for fuel maps of any resolution to be generated if properties are correctly aggregated.

Finally, the work presented is still subject to the resolution of the data input. While the RPAS workflow utilizes extremely high-resolution imagery (~2 cm / px), the satellite imagery used was 0.5 m / px. In the future, this resolution is expected to continue to improve; thus, work should be performed to keep pace with improvements in RS technology to ensure the most accurate and detailed information can be delivered.

The automation of aerial and satellite photo interpretation is explored, and fuel maps are generated. Use of these fuel maps in existing exposure assessment workflows is explored in the context of the workflow’s limitations. Overall, this study represents promising possibilities for the use of machine learning and remotely sensed data in the field of wildfire fuel mapping.

6.0 References

- Agisoft. (2023). *Agisoft Metashape Professional* (1.7.3). Agisoft.
- Ahmed, M. R., & Hassan, Q. K. (2023). Occurrence, Area Burned, and Seasonality Trends of Forest Fires in the Natural Subregions of Alberta over 1959–2021. *Fire*, 6(3), 96. <https://doi.org/10.3390/fire6030096>
- Alberta Sustainable Resource Development. (2007). *R11 Forest Management Plan*.
- Beaudoin, A., Bernier, P. Y., Guindon, L., Villemaire, P., Guo, X. J., Stinson, G., Bergeron, T., Magnussen, S., & Hall, R. J. (2014). Mapping attributes of Canada's forests at moderate resolution through k NN and MODIS imagery. *Canadian Journal of Forest Research*, 44(5), 521–532. <https://doi.org/10.1139/cjfr-2013-0401>
- Beaudoin, André, Hall, R. J., Castilla, G., Filiatrault, M., Villemaire, P., Skakun, R., & Guindon, L. (2022). Improved k-NN Mapping of Forest Attributes in Northern Canada Using Spaceborne L-Band SAR, Multispectral and LiDAR Data. *Remote Sensing*, 14(5), 1181. <https://doi.org/10.3390/rs14051181>
- Bennett, L., Wilson, B., Selland, S., Qian, L., Wood, M., Zhao, H., & Boisvert, J. (2022). Image to attribute model for trees (ITAM-T): individual tree detection and classification in Alberta boreal forest for wildland fire fuel characterization. *International Journal of Remote Sensing*, 43(5), 1848–1880. <https://doi.org/10.1080/01431161.2022.2048914>
- Beverly, J. L., Bothwell, P., Conner, J. C. R., & Herd, E. P. K. (2010). Assessing the exposure of the built environment to potential ignition sources generated from vegetative fuel. *International Journal of Wildland Fire*, 19(3), 299. <https://doi.org/10.1071/WF09071>
- Beverly, J. L., Braid, G., Chapman, L., Kelm, S., Pozniak, W., Stewart, L., & Johnston, K. (2018). *FireSmart Wildfire Exposure Assessment*. Government of Alberta.
- Beverly, J. L., & Forbes, A. M. (2023). Assessing directional vulnerability to wildfire. *Natural Hazards*, 117(1), 831–849. <https://doi.org/10.1007/s11069-023-05885-3>

- Beverly, J. L., McLoughlin, N., & Chapman, E. (2021). A simple metric of landscape fire exposure. *Landscape Ecology*, *36*(3), 785–801. <https://doi.org/10.1007/s10980-020-01173-8>
- Bolyn, C., Lejeune, P., Michez, A., & Latte, N. (2022). Mapping tree species proportions from satellite imagery using spectral–spatial deep learning. *Remote Sensing of Environment*, *280*, 113205. <https://doi.org/10.1016/j.rse.2022.113205>
- Brandt, J. P., Flannigan, M. D., Maynard, D. G., Thompson, I. D., & Volney, W. J. A. (2013). An introduction to Canada's boreal zone: ecosystem processes, health, sustainability, and environmental issues. *Environmental Reviews*, *21*(4), 207–226. <https://doi.org/10.1139/er-2013-0040>
- Briechle, S., Krzystek, P., & Vosselman, G. (2021). Silvi-Net – A dual-CNN approach for combined classification of tree species and standing dead trees from remote sensing data. *International Journal of Applied Earth Observation and Geoinformation*, *98*, 102292. <https://doi.org/10.1016/j.jag.2020.102292>
- Cameron, H., Schroeder, D., & Beverly, J. L. (2021). Predicting black spruce fuel characteristics with Airborne Laser Scanning (ALS). *International Journal of Wildland Fire*, *31*(2), 124–135. <https://doi.org/10.1071/WF21004>
- Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., ... Lin, D. (2019). *MMDetection: Open MMLab Detection Toolbox and Benchmark* (3.0.0). <https://doi.org/1906.07155v1>
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In V. Ferrari, M. Hebert, C. Sminchisescu, & Y. Weiss (Eds.), *Computer Vision -- ECCV 2018* (pp. 833–851). Springer International Publishing.
- Countryman, C. M. (1972). *The fire environment concept*. USDA Forest Service, Pacific Southwest Forest and Range Experiment Station.

- DigitalGlobe. (2009). *WorldView-2*.
<https://www.spaceimagingme.com/downloads/sensors/datasheets/WorldView2-DS-WV2-Web.pdf>
- Downing, D. J., & Pettapiece, W. w. (2006). *Natural Regions and Subregions of Alberta*.
- Egli, S., & Höpke, M. (2020). CNN-Based Tree Species Classification Using High Resolution RGB Image Data from Automated UAV Observations. *Remote Sensing*, 12(23), 3892. <https://doi.org/10.3390/rs12233892>
- Ferreira, M. P., Almeida, D. R. A. de, Papa, D. de A., Minervino, J. B. S., Veras, H. F. P., Formighieri, A., Santos, C. A. N., Ferreira, M. A. D., Figueiredo, E. O., & Ferreira, E. J. L. (2020). Individual tree detection and species classification of Amazonian palms using UAV images and deep learning. *Forest Ecology and Management*, 475, 118397. <https://doi.org/10.1016/j.foreco.2020.118397>
- FireSmart Alberta. (2015). *FireSmart Homeowner's Assessment*. Alberta Government.
<https://open.alberta.ca/dataset/d8cc1bbd-7a69-4b04-89b1-9ced448ab43b/resource/dd2e85d9-662c-4a41-aeec-479a8319e2fd/download/2015-06-firesmart-homeownersassessment.pdf>
- Forest Stewardship and Trade Branch. (2022). *Alberta Vegetation Inventory Standards, Version 2.1.5*. Alberta Agriculture, Forestry and Rural Economic Development.
- Forestry Canada Fire Danger Group. (1992). *Development and Structure of the Canadian Forest Fire Behaviour Prediction System*. Forestry Canada.
- Frederick, K. (2012). *Revised Process to Convert Alberta Vegetation Inventory (AVI) Data to Canadian Forest Fire Behaviour Prediction System (FBP) Fuel Types*.
- Fricke, G. A., Ventura, J. D., Wolf, J. A., North, M. P., Davis, F. W., & Franklin, J. (2019). A Convolutional Neural Network Classifier Identifies Tree Species in Mixed-Conifer Forest from Hyperspectral Imagery. *Remote Sensing*, 11(19), 2326. <https://doi.org/10.3390/rs11192326>
- G. Braga, J. R., Peripato, V., Dalagnol, R., P. Ferreira, M., Tarabalka, Y., O. C. Aragão,

- L. E., F. de Campos Velho, H., Shiguemori, E. H., & Wagner, F. H. (2020). Tree Crown Delineation Algorithm Based on a Convolutional Neural Network. *Remote Sensing*, 12(8), 1288. <https://doi.org/10.3390/rs12081288>
- Gallardo-Salazar, J. L., & Pompa-García, M. (2020). Detecting Individual Tree Attributes and Multispectral Indices Using Unmanned Aerial Vehicles: Applications in a Pine Clonal Orchard. *Remote Sensing*, 12(24), 4144. <https://doi.org/10.3390/rs12244144>
- Gazzea, M., Kristensen, L. M., Pirotti, F., Ozguven, E. E., & Arghandeh, R. (2022). Tree Species Classification Using High-Resolution Satellite Imagery and Weakly Supervised Learning. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–11. <https://doi.org/10.1109/TGRS.2022.3210275>
- Ge, S., Gu, H., Su, W., Lönnqvist, A., & Antropov, O. (2022). A Novel Semisupervised Contrastive Regression Framework for Forest Inventory Mapping with Multisensor Satellite Data. <http://arxiv.org/abs/2212.00246>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- Immitzer, M., Atzberger, C., & Koukal, T. (2012). Tree Species Classification with Random Forest Using Very High Spatial Resolution 8-Band WorldView-2 Satellite Data. *Remote Sensing*, 4(9), 2661–2693. <https://doi.org/10.3390/rs4092661>
- Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., NanoCode012, Kwon, Y., Michael, K., TaoXie, Fang, J., imyhxy, Lorna, Yifu), 曾逸夫(Zeng, Wong, C., V, A., Montes, D., Wang, Z., Fati, C., Nadar, J., Laughing, ... Jain, M. (2022). *ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation (7.0)*. Zenodo. <https://doi.org/10.5281/zenodo.7347926>
- Johnston, L. M., & Flannigan, M. D. (2018). Mapping Canadian wildland fire interface areas. *International Journal of Wildland Fire*, 27(1), 1. <https://doi.org/10.1071/WF16221>

- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2016). *Feature Pyramid Networks for Object Detection*.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). *Focal Loss for Dense Object Detection*.
- Mamuji, A. A., & Rozdilsky, J. L. (2019). Wildfire as an increasingly common natural disaster facing Canada: understanding the 2016 Fort McMurray wildfire. *Natural Hazards*, 98(1), 163–180. <https://doi.org/10.1007/s11069-018-3488-4>
- Mohan, M., Silva, C., Klauberg, C., Jat, P., Catts, G., Cardil, A., Hudak, A., & Dia, M. (2017). Individual Tree Detection from Unmanned Aerial Vehicle (UAV) Derived Canopy Height Model in an Open Canopy Mixed Conifer Forest. *Forests*, 8(9), 340. <https://doi.org/10.3390/f8090340>
- Mubin, N. A., Nadarajoo, E., Shafri, H. Z. M., & Hamedianfar, A. (2019). Young and mature oil palm tree detection and counting using convolutional neural network deep learning method. *International Journal of Remote Sensing*, 40(19), 7500–7515. <https://doi.org/10.1080/01431161.2019.1569282>
- Mulcahy, K. (2023, June 12). Alberta could break decades-old record for amount of area burned during 2023 wildfire season. *CTV News*. <https://edmonton.ctvnews.ca/alberta-could-break-decades-old-record-for-amount-of-area-burned-during-2023-wildfire-season-1.6438194>
- Muller, A., & Guido, S. (2016). *Introduction to Machine Learning with Python* (D. Schanafelt (ed.); 1st ed.). O'Reilly.
- Natural Resources Canada. (n.d.). *FBP Fuel Types Map*. Retrieved April 8, 2022, from <https://cwfis.cfs.nrcan.gc.ca/background/maps/fbpft>
- Nevalainen, O., Honkavaara, E., Tuominen, S., Viljanen, N., Hakala, T., Yu, X., Hyypä, J., Saari, H., Pölonen, I., Imai, N., & Tommaselli, A. (2017). Individual Tree Detection and Classification with UAV-Based Photogrammetric Point Clouds and Hyperspectral Imaging. *Remote Sensing*, 9(3), 185.

<https://doi.org/10.3390/rs9030185>

- Persson, M., Lindberg, E., & Reese, H. (2018). Tree Species Classification with Multi-Temporal Sentinel-2 Data. *Remote Sensing*, 10(11), 1794. <https://doi.org/10.3390/rs10111794>
- Quintilio, D., Alexander, M. E., & Ponto, R. L. (1991). *Spring fires in a semimature trembling aspen stand in central Alberta*. Forestry Canada, Northern Forestry Centre.
- Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2015). *You Only Look Once: Unified, Real-Time Object Detection*. <https://doi.org/1506.02640v5>
- Rodríguez-Puerta, F., Gómez-García, E., Martín-García, S., Pérez-Rodríguez, F., & Prada, E. (2021). UAV-Based LiDAR Scanning for Individual Tree Detection and Height Measurement in Young Forest Permanent Trials. *Remote Sensing*, 14(1), 170. <https://doi.org/10.3390/rs14010170>
- Sanchez-Azofeifa, G. A., Chong, M., Sinkwich, J., & Mamet, S. (2004). *Alberta Ground Cover Characterization (AGCC) Training and Procedures Manual*. Earth Observation Systems Laboratory, Department of Earth and Atmospheric Sciences, University of Alberta.
- Schiefer, F., Kattenborn, T., Frick, A., Frey, J., Schall, P., Koch, B., & Schmidlein, S. (2020). Mapping forest tree species in high resolution UAV-based RGB-imagery by means of convolutional neural networks. *ISPRS Journal of Photogrammetry and Remote Sensing*, 170, 205–215. <https://doi.org/10.1016/j.isprsjprs.2020.10.015>
- Shtanchaev, A., Bille, A., Sutyryna, O., & Elelimy, S. (2021). *Automated Remote Sensing Forest Inventory Using Satellite Imagery*.
- Surovy, P., & Kuželka, K. (2019). Acquisition of Forest Attributes for Decision Support at the Forest Enterprise Level Using Remote-Sensing Techniques—A Review. *Forests*, 10(3), 273. <https://doi.org/10.3390/f10030273>
- Tao, H., Li, C., Zhao, D., Deng, S., Hu, H., Xu, X., & Jing, W. (2020). Deep learning-based dead pine tree detection from unmanned aerial vehicle images. *International Journal*

of Remote Sensing, 41(21), 8238–8255.
<https://doi.org/10.1080/01431161.2020.1766145>

Transport Canada. (2023). *Canadian Aviation Regulations: Part IX — Remotely Piloted Aircraft Systems*.

Tymstra, C., Bryce, R. W., Wotton, B. M., Taylor, S. . W., & Armitage, O. B. (2010). Development and structure of Prometheus: the Canadian wildland fire growth simulation model. *Natural Resources Canada, Canadian Forest Service, Northern Forestry Centre, Information Report NOR-X-417*.(Edmonton, AB).

Van Wagner, C. E. (1977). Conditions for the start and spread of crown fire. *Canadian Journal of Forest Research*, 7(1), 23–34. <https://doi.org/10.1139/x77-004>

Van Wagner, C. E. (1983). Fire behaviour in northern conifer forests and shrublands. In R. Wein & D. MacLean (Eds.), *The Role of Fire in Northern Circumpolar Ecosystems* (pp. 65–80). John Wiley and Sons.

Weinstein, B. G., Marconi, S., Bohlman, S., Zare, A., & White, E. (2019). Individual Tree-Crown Detection in RGB Imagery Using Semi-Supervised Deep Learning Neural Networks. *Remote Sensing*, 11(11), 1309. <https://doi.org/10.3390/rs11111309>

Westhaver, A. (2017). *Why some homes survived: Learning from the Fort McMurray wildland/urban interface fire disaster*. Institute for Catastrophic Loss Reduction.

Zhao, Z.-Q., Zheng, P., Xu, S., & Wu, X. (2018). *Object Detection with Deep Learning: A Review*. <http://arxiv.org/abs/1807.05511>

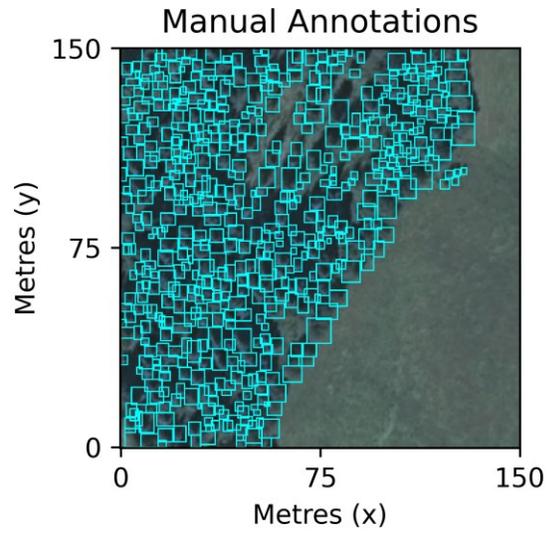
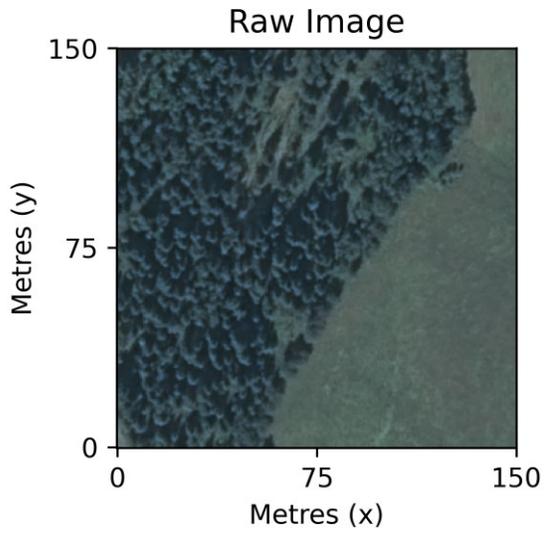
7.0 Appendices

Appendix A: RPAS Detection Results

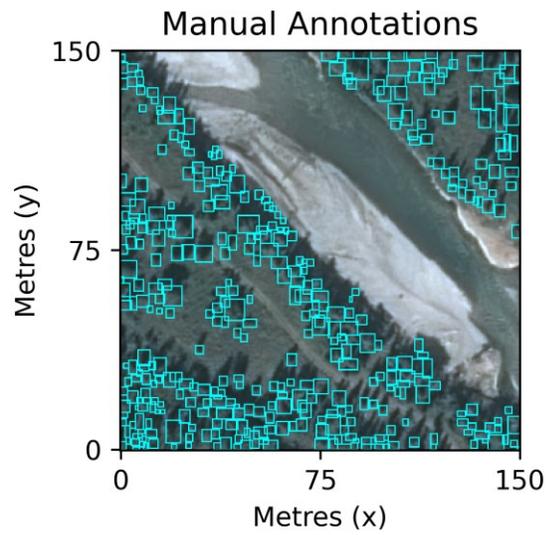
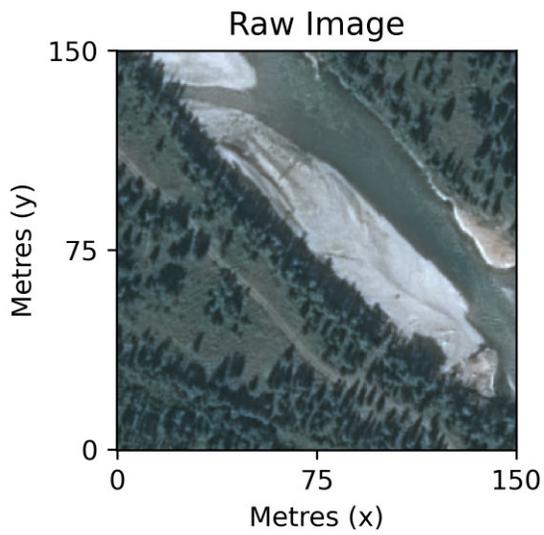
Model	k-fold Code	Training Cons	Training Decs	Training Snags	Testing Cons	Testing Decs	Testing Snags	Precision	Recall	Stem Recall	F1-Score (Detection)	F1-Score (Con)	F1-Score (Dec)	F1-Score (Snag)
YOLOv5 Small	1	5359	1970	399	1569	763	163	59.2%	76.8%	92.6%	66.9%	95.6%	89.6%	84.7%
	2	5404	2461	514	1524	272	48	78.6%	73.1%	91.8%	75.7%	98.5%	90.9%	90.3%
	3	5596	1947	473	1332	786	89	71.6%	76.5%	87.0%	73.8%	96.9%	93.1%	91.3%
	4	5728	2063	435	1200	670	127	76.5%	77.3%	87.3%	76.7%	96.2%	97.7%	96.7%
	5	5625	2491	427	1303	242	135	59.9%	78.6%	91.4%	67.9%	96.8%	84.8%	95.2%
RESULTS								68.9%	76.5%	90.0%	72.1%	96.8%	92.3%	91.5%
YOLOv5 Nano	1	5434	1977	399	1569	763	163	58.3%	78.5%	93.0%	66.9%	95.7%	90.6%	86.0%
	2	5479	2468	514	1524	272	48	75.2%	73.8%	93.7%	74.5%	98.5%	89.9%	80.0%
	3	5671	1954	473	1332	786	89	66.3%	77.8%	90.3%	71.5%	97.6%	94.9%	94.2%
	4	5803	2070	435	1200	670	127	71.4%	80.8%	89.9%	75.8%	96.4%	97.4%	94.8%
	5	5625	2491	427	1378	249	135	55.8%	71.1%	94.6%	66.4%	94.8%	80.1%	97.6%
RESULTS								65.2%	76.7%	92.2%	70.9%	96.6%	92.5%	91.6%
RetinaNet	1	5434	1977	399	1569	763	163	58.1%	57.1%	77.6%	57.6%	97.9%	95.9%	82.5%
	2	5479	2468	514	1524	272	48	65.0%	45.4%	79.3%	53.5%	91.7%	40.2%	0.0%
	3	5671	1954	473	1332	786	89	55.7%	63.3%	85.3%	58.6%	89.7%	4.3%	0.0%
	4	5803	2070	435	1200	670	127	70.5%	78.7%	91.7%	74.3%	92.0%	0.5%	0.0%
	5	5625	2491	427	1378	249	135	52.3%	78.0%	95.2%	62.4%	96.3%	90.1%	3.2%
RESULTS								60.3%	64.0%	85.2%	61.1%	93.7%	40.3%	24.7%
YOLOv5 Large	1	5434	1977	399	1569	763	163	72.0%	71.7%	85.0%	71.9%	96.5%	91.3%	87.9%
	2	5479	2468	514	1524	272	48	78.8%	69.1%	90.1%	73.7%	99.2%	95.8%	77.8%
	3	5671	1954	473	1332	786	89	73.4%	74.2%	88.1%	73.7%	97.7%	95.3%	93.1%
	4	5803	2070	435	1200	670	127	80.2%	75.7%	85.0%	77.7%	95.2%	98.1%	95.7%
	5	5625	2491	427	1378	249	135	74.9%	78.5%	88.1%	76.6%	97.8%	89.8%	93.1%
RESULTS								75.6%	73.7%	87.1%	74.5%	97.3%	94.4%	90.9%

Appendix B: Satellite Training Tiles

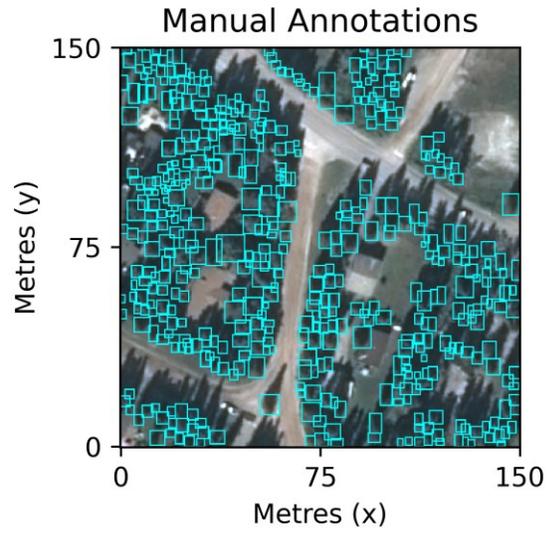
BC1



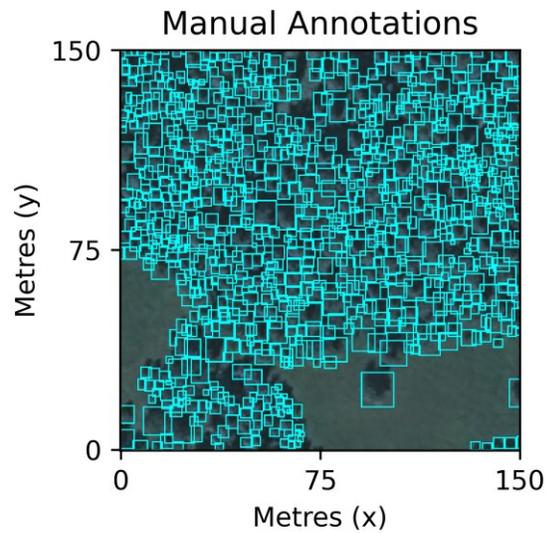
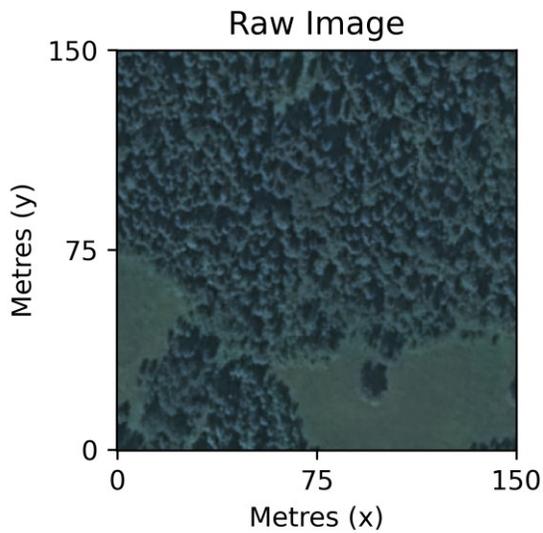
BC2



BC3



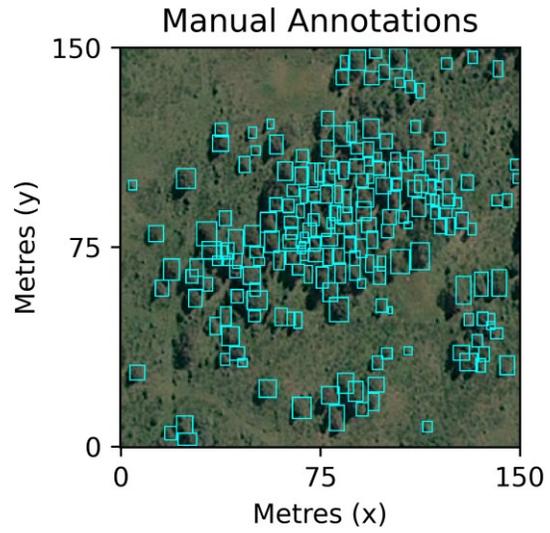
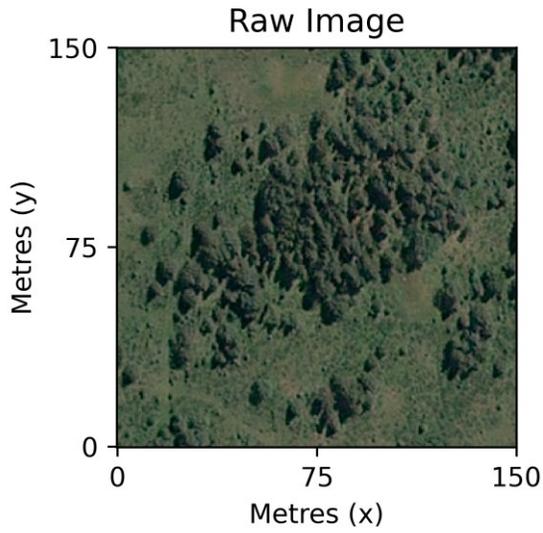
BC4



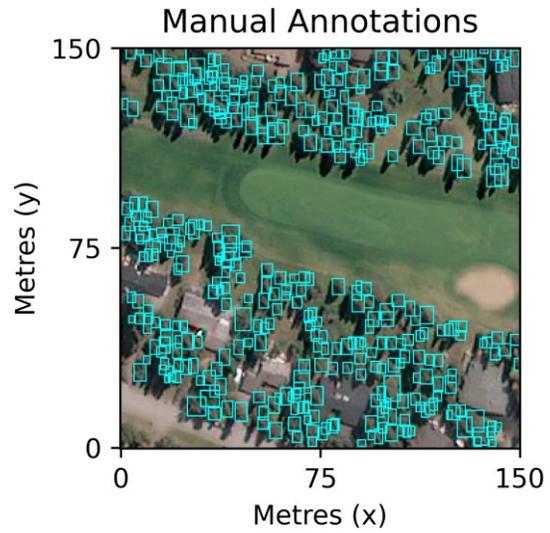
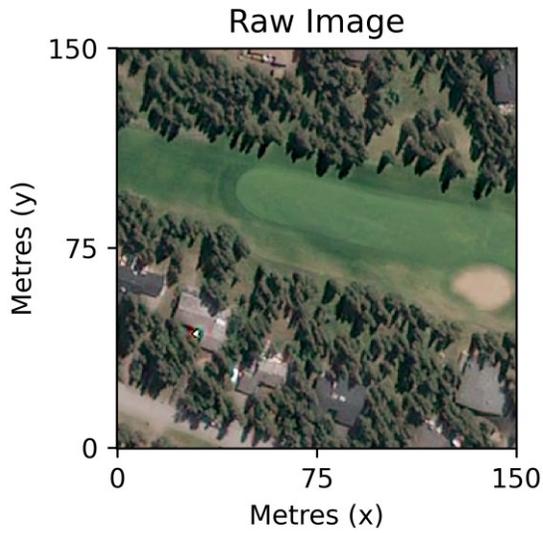
BC5



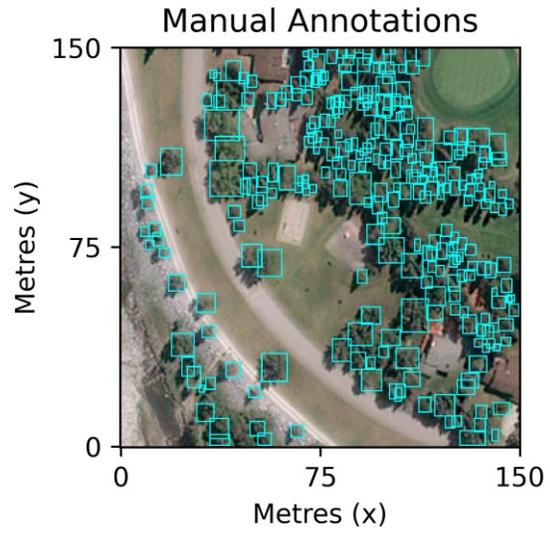
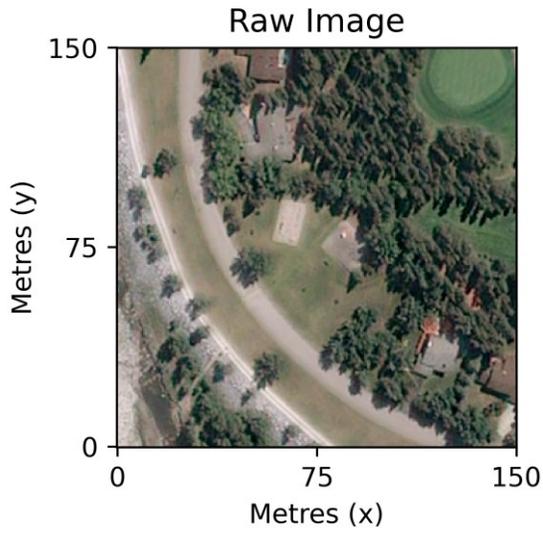
RW1



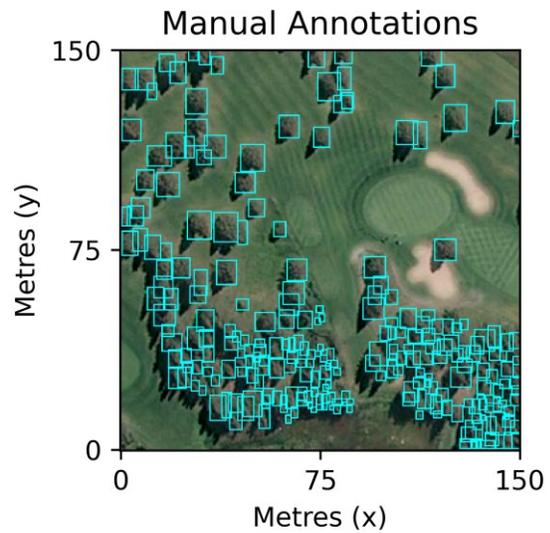
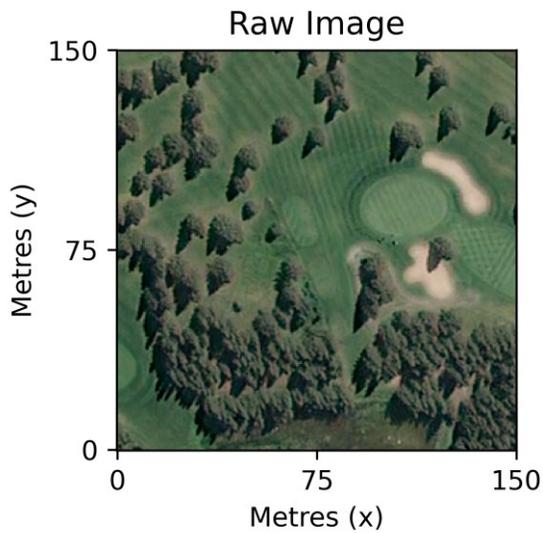
RW2



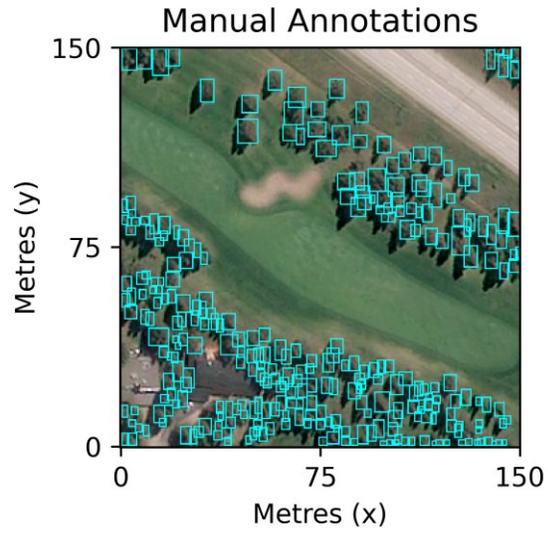
RW3



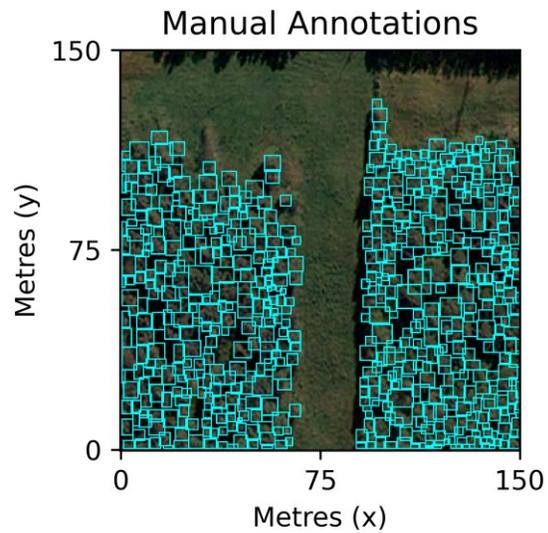
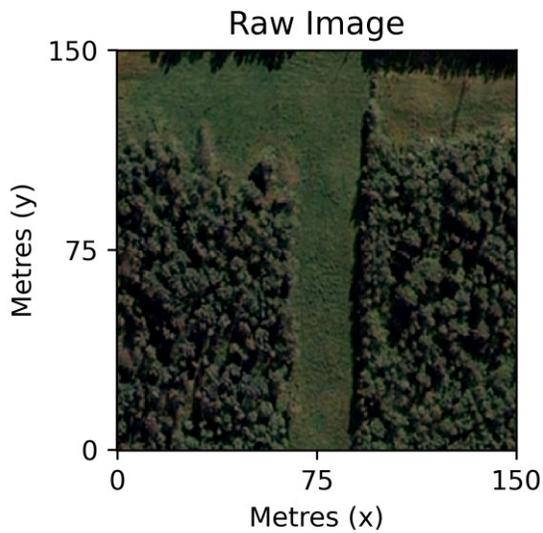
RW4



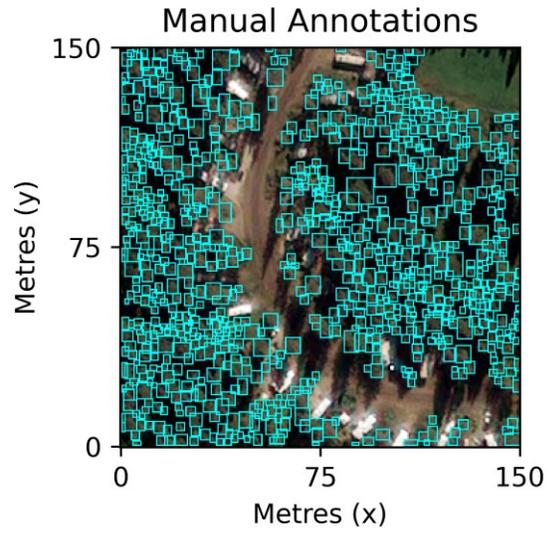
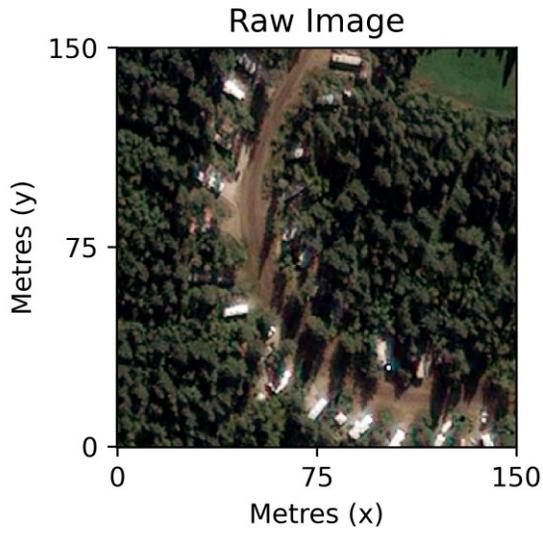
RW5



SB1



SB2



SB3

