Dense neural network outperforms other machine learning models for scaling-up lichen cover maps in Eastern Canada

Galen Richardson

A thesis submitted to the University of Ottawa in partial fulfillment of the requirements for the degree of Master of Science in Geography

Department of Geography, Environment and Geomatics
Faculty of Arts
University of Ottawa

Supervisor:
Dr. Anders Knudby

Thesis Committee:
Dr. Wenjun Chen
Dr. Michael Sawada

Abstract

Lichen mapping is vital for caribou management plans and sustainable land conservation. Previous studies have used Random Forest, dense neural network, and convolutional neural network (CNN) models for mapping lichen coverage with remote sensing data. However, to date, it is not clear how these models rank in the performance of this task. In this study, these machine learning models were evaluated on their ability to predict lichen percent coverage in Sentinel-2 imagery covering Québec and Labrador, NL. The models were trained on 10-m resolution lichen coverage (%) maps created from 20 drone surveys collected in July 2019 and 2022. The maps were divided into quadrant blocks and then split into train, validation, and test datasets. The quadrant-blocking approach exposed the models to a variety of different landscapes and reduced spatial autocorrelation between the training sites. All three models performed similarly when evaluated on the test set. However, the dense neural network achieved a higher accuracy than the other two, with a reported Mean Absolute Error (MAE) of 5.2% and an R^2 of 0.76. By comparison, the Random Forest model returned an MAE of 5.5% (R^2: 0.74) and the CNN had an MAE of 5.3% (R^2: 0.74). The models were also evaluated on their ability to predict lichen coverage (%) for larger quadrant blocks consisting of, on average, 400 Sentinel-2 pixels. The Random Forest and dense neural network had an R^2 of 0.93, while the CNN had an R^2 of 0.90. The MAE in this assessment for the dense neural network, Random Forest, and CNN were 2.1%, 2.3%, and 3.1% respectively. A regional lichen map was created using the dense neural network and a Sentinel-2 image mosaic. Model predictions have larger errors for land covers that the model was not exposed to in training, such as mines and deep lakes. While the dense neural network requires more computational effort to train than a Random Forest model, the 5.9% performance gain in the test pixel comparison and 9.1% performance gain in the quadrant block comparison renders it the most suitable for lichen mapping. This study represents progress toward determining the appropriate methodology for generating accurate lichen maps from satellite imagery for caribou conservation and sustainable land management.

Acknowledgments

I would like to thank my partner, Elisha Richardson, for her continued support and generous love throughout my life and especially during this thesis. I would also like to extend my gratitude towards my family and friends; Luciana Frigerio, Eric Richardson, Olivia Richardson, Karen Richardson, Edie Zorychta, Ryan Farr, and Mickey. The SWEOL, EO4CE, and GGSA are associations for which I am grateful for participating in. Finally, the support and willingness to take a chance on me from Anders Knudby, Julie Lovitt, Wenjun Chen, and Sylvain Leblanc have allowed me to pursue such ambitious concepts in this thesis.

Preface

This thesis was written in article format and consists of an introduction (Chapter 1), a manuscript submitted to a peer-reviewed journal (Chapter 2), a chapter addressing misclassified pixels in the regional lichen map (Chapter 3), and a conclusion (Chapter 4). The manuscript in Chapter 2 has been submitted PLOS ONE, with the proposed authorship as follows:

> Richardson G, Knudby A, Chen W, Sawada M, Lovitt J, He L, and Naenic LY, (submitted) Dense neural network outperforms other machine learning models for scaling-up lichen cover maps in Eastern Canada. PLOS ONE.

Author contributions: Conceptualization, Galen Richardson, Anders Knudby, Wenjun Chen, and Michael Sawada; methodology, Galen Richardson and Anders Knudby; software, Galen Richardson, Anders Knudby, and Liming He; validation, Galen Richardson, and Anders Knudby; formal analysis, Galen Richardson, and Anders Knudby; investigation, Galen Richardson, and Anders Knudby; resources, Galen Richardson, Anders Knudby, Wenjun Chen, Julie Lovitt, Liming He, and Leila Yousefizadeh Naeni; data curation, Galen Richardson, and Leila Yousefizadeh Naeni; writing—original draft preparation, Galen Richardson, and Anders Knudby; writing—review and editing, Wenjun Chen, Michael Sawada, Julie Lovitt, Liming He, and Leila Yousefizadeh Naeni; visualization, Galen Richardson; supervision, Anders Knudby, Wenjun Chen, and Michael Sawada; project administration, Anders Knudby and Wenjun Chen; funding acquisition, Galen Richardson, Anders Knudby, and Wenjun Chen

Table of Contents

List of Figures

Chapter 1: Introduction

1.1 Background

Lichen is a symbiotic life form consisting of a fungus and an autotrophic photobiont such as green algae or cyanobacteria (Asplund & Wardle, 2017; Petzold & Goward, 1988). Lichens are slow-growing and occur in most terrestrial ecosystems such as boreal forests, grasslands, and tundra (Asplund & Wardle, 2017; Kuusinen et al., 2020). They are reliant on atmospheric moisture and nutrition for growth and are typically found growing on soil, tree trunks, and rocks (Asplund & Wardle, 2017; Kuusinen et al., 2020). Aside from being a source of food, shelter, and camouflage for organisms, lichen is also an essential component of global nitrogen and carbon cycles (Allen et al., 2019; Asplund & Wardle, 2017). There has been a decline in global lichen diversity and abundance for more than a century (Allen et al., 2019; Joly et al., 2009). Global threats such as air pollution disproportionally harm lichens compared to other types of organisms (Allen et al., 2019; Kuusinen et al., 2020). A warming climate in northern latitudes has facilitated a widespread expansion of vascular plants, which are experiencing longer growing and snow-free seasons (Myers-Smith et al., 2020; Nelson et al., 2013; Thompson et al., 2015). This harms lichen populations since vascular plants outcompete lichens for ground space and access to sunlight (Kuusinen et al., 2020; MacAnder et al., 2020). While deciduous shrubs might provide additional nutritional sources for grazing animals in summer, the shedding of leaves limits winter grazing (Joly et al., 2009).

In Canada, lichen is the primary food source for caribou, comprising 75% of the caribou diet in the winter months (Théau et al., 2005; Thompson et al., 2015). Since lichen is almost exclusively consumed by caribou, evidence of lichen over-grazing is accepted as a positive indicator of caribou activity (Nordberg & Allard, 2002; Théau et al., 2005). However, the consumption of lichen by caribou is one of the many factors that alter lichen abundance in Canada. The expansion of managed forests, increased frequency and severity of forest fires, and cumulative impacts of human development have reduced the abundance of lichen in caribou ranges (Jozdani et al., 2021; Kuusinen et al., 2020; MacAnder et al., 2020; Myers-Smith et al., 2011; Thompson et al., 2015).

Coincidentally, over the past 30 years Canadian caribou populations have declined by an average of 52% to an estimated 1.3 million caribou in 2015 (Gunn, 2016). Gunn (2016) reported that in 37 of the 52 Canadian boreal caribou subpopulations where data was available, 81% were in decline since 2002. Many negative impacts on caribou are linked to a warming climate, such as heightened insect harassment, increased frequency and intensity of weather events, and changes in foraging plant phenology (Vors & Boyce, 2009). The earlier emergence and increased abundance of insects indirectly reduce the summer forging by reducing the amount of undisturbed time grazing (Vors & Boyce, 2009). Freezing rain events may hinder winter grazing by locking pastures under impenetrable ice (Vors & Boyce, 2009). In addition to these climate-related threats, the dynamics of the caribou ranges are changing due to anthropogenic landscape changes, particularly an increase in forest harvests and petroleum infrastructure construction (Gunn, 2016; Théau et al., 2005; Thompson et al., 2015; Vors & Boyce, 2009). Forest management reduces the habitat available to woodland caribou, since the lichen mats are commonly destroyed and take up to 50 years to regenerate (Nordberg & Allard, 2002; Théau et al., 2005; Thompson et al., 2015). Caribou tend to avoid managed habitats, often using them as regions for travel instead of grazing (Thompson et al., 2015). Climate change and habitat loss are primary drivers of the loss in caribou populations, and it is critical to understand where these changes are occurring.

Mapping landcover change is essential for creating Canadian caribou recovery plans, and a crucial aspect of monitoring caribou habitat change is mapping their primary food source, lichen. These maps are essential to creating informed sustainable land management policies and caribou recovery plans. Lichen has been mapped across North America, Northern Europe, and Antarctica using remotely sensed imagery (Casanovas et al., 2015; He et al., 2021; Nordberg & Allard, 2002; Théau et al., 2005). Bright lichens containing usnic acid, such as *Cladonia Stellaris*, are distinguishable in remote sensing imagery due to their high blue reflectance and unique spectral response (Kuusinen et al., 2020; Nelson et al., 2013). Recent studies on lichen mapping have focused on using machine learning models that leverage the spectral signature of lichen to inform predictions on specific caribou ranges (He et al., 2021; MacAnder et al., 2020; Théau et al., 2005; Zimmermann et al., 2018). This thesis will contribute to this body of work by determining which type of machine learning models work best for this task.

1.2 Overview of machine learning models

Machine learning is the use of computer systems and algorithms to analyze data and provide predictions. The objective of this section is to introduce machine learning and elaborate upon machine learning models that have been previously used for lichen mapping. Additional details on non-machine learning approaches as well as discussion of machine learning applications to lichen mapping will be elaborated upon in Chapter 2.1.

Machine learning models tend to excel at regression and classification tasks with complex data and many input variables. Implementing machine learning in remote sensing classification applications has tended to produce a higher accuracy than traditional parametric classifiers (Maxwell et al., 2018). Random Forest, dense neural networks, and convolutional neural networks (CNNs) are machine learning models that have been used in lichen mapping, and their predictive performance was evaluated in this thesis.

One of the most common machine learning models implemented in remote sensing classification and regression models is Random Forest. Developed by Breiman (2001), Random Forest is an ensemble classifier consisting of a collection of decision trees that are each created from random subsets of training data. The model predictions in Random Forest models are based on the average/majority of individual tree predictions (Belgiu and Drăguţ, 2016; Fawagreh et al., 2014). The user can define the number of trees included in the model, however each decision tree within the Random Forest operates independently (Belgiu and Drăguţ, 2016). Random Forest models are considered easy to optimize compared to other machine-learning algorithms due to the smaller amount of user-defined parameters (Maxwell et al., 2018). Another benefit is that Random Forest models are interpretable, and can provide assessments of the relative importance of predictor variables, which can provide a greater understanding of the contributing factors to a prediction (Maxwell et al., 2018).

Dense neural networks are conceptualized as biologically-inspired computational models designed to simulate human neurons (Maxwell et al., 2018; Nazari & Yan, 2021). Dense neural networks consist of self-optimizing neurons that are fully connected to each other by weights. These neurons learn by first randomly guessing values for the weights, and then iteratively altering the weights to optimize the predictive performance of the network (Maxwell et al., 2018). Relative to Random Forest models, dense neural networks are substantially slower to train and require proper guidance to produce optimal results. Additionally, dense neural networks and

CNNs are ‘black box’ classifiers which are not easily understood, making it challenging to understand the factors leading to a prediction (Maxwell et al., 2018).

CNNs are similar to dense neural networks, but are primarily developed to analyze and make predictions of images (Nazari & Yan, 2021). There are different implementations of CNNs in image-related applications, such as image segmentation, feature detection, feature counting, and image classification. Each convolutional layer in a CNN consists of a convolutional operation where a series of filters sweeps over an entire image resulting in feature maps (Bhatnagar et al., 2020). These feature maps are then passed through an activation function to introduce nonlinearity and generate an output (Bhatnagar et al., 2020). The filters used to generate the feature maps are self optimized with backpropagation to improve model performance over training.

In the study of lichens, CNNs have been implemented for pixel classification and image segmentation (Jozdani et al., 2021; Lovitt et al., 2022; Richardson et al., 2021). There is an important distinction between fully convolutional networks and CNNs that also include dense layers; the former can be implemented in image segmentation and produces a prediction 2-dimensional array, while the latter is typically used to produce a single predicted value such as lichen coverage percentage. This study focused on using CNNs to predict a single value representing lichen percent coverage since the desired output for regional lichen maps is a lichen coverage (%); a difficult task for a fully connected convolutional network to predict due to image variation and noise when calculating the loss function. The CNNs assessed in this study include dense layers that generate a single continuous pixel prediction, similar to models proposed by Al Najar et al., (2022), Aptoula & Ariman, (2022), and Erlandsson et al., (2022).

1.3 Research questions

The objective of this thesis is to generate high accuracy regional lichen maps which can be used to inform decisions on sustainable land management. Previous studies have created lichen coverage maps using Landsat imagery due to the vast available archive of such imagery (He et al., 2021; MacAnder et al., 2020; Nelson et al., 2013; Nordberg & Allard, 2002; Théau et al., 2005; Zimmermann et al., 2018). However, with a pixel resolution of 30 meters, it is challenging to align Landsat imagery with high resolution datasets (He et al., 2021). Sentinel-2 imagery is freely available, has a higher resolution of 10 m in some bands for more accurate alignment to higher resolution datasets, and can be used to create current lichen maps with a higher spatial resolution than Landsat imagery.

Past studies have shown that Random Forest models, dense neural networks, and deep CNNs are all effective at lichen mapping. However, to date, it is unclear which type of model performs best at generating regional lichen maps. In this thesis, these models were assessed in a study region comprising the Manicouagan Caribou Range in Québec and the Red Wine Mountain Caribou Range in Labrador, where the lichen species *Cladonia stellaris* is dominant. This thesis will address the following questions:

1. Is it feasible to create regional lichen maps with Sentinel-2 imagery?
2. Which regression machine learning methodology is most effective at lichen mapping?

These objectives are addressed throughout the remaining chapters of this thesis. The discussion of the results from this thesis can inform future lichen and landcover mapping applications.

1.4 Thesis structure

In the following, I present the paper "Dense neural network outperforms other machine learning models for scaling-up lichen cover maps in Eastern Canada", submitted for publication in PLOS ONE (Chapter 2), followed by a chapter addressing misclassified pixels in the regional lichen map (Chapter 3), and a conclusion to the thesis (Chapter 4).

Chapter 2: Dense neural network outperforms other machine learning models for scaling-up lichen cover maps in Eastern Canada

2.1. Introduction

In 2016, *Rangifer tarandus*, commonly known as the caribou, was categorized as a vulnerable species due to an observed decline of 40% between the 1990s and 2015 (Gunn, 2016). Caribou are threatened by changes to their habitat, unregulated hunting, and herd fragmentation due to human-made infrastructure (Gunn, 2016; Théau et al., 2005). Terricolous (ground-dwelling) lichen is a primary food source for caribou, comprising 75% and 25% of their winter and summer diets respectively (Théau et al., 2005). In recent decades, ecological changes and climate change impacts on the caribou habitat, such as the expansion of shrublands and managed forests, have reduced the abundance of lichen in caribou ranges (MacAnder et al., 2020; Myers-Smith et al., 2011; Thompson et al., 2015). The increasing frequency, severity, and spatial extent of forest fires in the boreal region have also impacted the availability of lichen, as lichen requires decades to fully recover (MacAnder et al., 2020; Myers-Smith et al., 2011). Mapping and monitoring lichen availability is thus critical for informing sustainable land management policy and caribou recovery plans.

The base data for lichen mapping generally consists of field measurements, digital photographs, UAV surveys, and satellite imagery. These data are used to train models to detect lichen at different spatial resolutions. Early lichen mapping models focused primarily on linking the spectral characteristics of lichen to remotely sensed imagery (Casanovas et al., 2015; Nordberg & Allard, 2002; Olthof & Pouliot, 2009). For example, the normalized difference vegetation index (NDVI) was used to detect lichen in Landsat TM imagery located in tundra regions with no tree canopy coverage (Nordberg & Allard, 2002). However, many caribou ranges exist below the tree line, and models based on NDVI struggle to distinguish lichen from other vegetation in regions with canopy cover (Théau et al., 2005). To address this issue, Théau et al., (2005) improved the classification of lichen in Landsat TM imagery by implementing Spectral Mixture Analysis. This approach was able to correctly detect lichen in 75% of the 24 field-verified sites and tended to overestimate lichen coverage (%) when compared to lichen maps from aerial surveys (Théau et al., 2005).

A wide range of studies have broadly found that machine learning models for classification tend to produce higher accuracy compared to traditional parametric models (Maxwell et al., 2018). To generate sufficient training data for machine learning models, researchers have implemented multi-scaled approaches to model development. This involves training models on higher-resolution data and then applying them to lower-resolution data to make predictions (maps) over large areas. Multi-scaled approaches have been used to map tree mortality (Campbell et al., 2020), detect invasive plant species (Alvarez-Taboada et al., 2017), and for lichen mapping (Fraser et al., 2021; MacAnder et al., 2020; Zimmermann et al., 2018).

One of the most common machine learning models implemented in multi-scaled lichen mapping is Random Forest. Developed by Breiman (2001), Random Forest is a non-parametric ensemble classifier consisting of a collection of decision trees that are each created from random subsets of training data, with model predictions based on the average/majority of individual tree predictions (Belgiu and Drăguţ, 2016; Fawagreh et al., 2014). The user can define the number of trees included in the model, however each decision tree within the Random Forest operates independently (Belgiu and Drăguţ, 2016). Random Forest models have been successfully implemented in multi-scale lichen mapping applications (MacAnder et al., 2020; Zimmermann et al., 2018). MacAnder et al. (2020) developed a methodology to classify lichen in Landsat 7 imagery, using a Random Forest model and resampled UAV orthomosaics as training data. They achieved an R2 of 0.71 and a Mean Average Error (MAE) of 5.2% in interior Alaska and Yukon datasets (MacAnder et al., 2020). According to the authors of this study, the model tended to overestimate lichen coverage in areas with lower true cover, and underestimate lichen coverage in areas with high cover (MacAnder et al., 2020). Other research has shown that estimating lichen coverage through direct scaling of UAV orthomosaics to Landsat imagery is associated with high uncertainties due to poor co-registration between datasets of significantly different spatial resolutions (He et al., 2021). Therefore, extra care must be taken to minimize errors arising from co-registration issues.

Recent approaches to lichen mapping involves implementing neural networks, to create lichen maps. He et al. (2021) used binary lichen maps from WorldView (WV) imagery (50 cm) to train a regression dense neural network to predict lichen coverage (%) in Landsat 8 OLI imagery (30 m). This model had an R^2 value of 0.60 and an MAE of 6.6% when comparing the lichen cover from WV imagery to the lichen cover in Landsat 8. The training data in this study used binary lichen maps at the 50 cm resolution, and it was reliant on at-cost WV imagery (He et al., 2021). Another type of model that has been used for lichen mapping is CNNs. Unlike single-pixel classifiers, CNNs can create inferences based on image texture, intensity, and other spatial patterns, and use these along with spectral information to classify pixels in an image (Bhatnagar et al., 2020; Lovitt et al., 2022). CNNs have many different remote sensing applications, such as detecting trees and buildings (Maggiori et al., 2016; Zhao et al., 2018), segmenting ecotypes (Bhatnagar et al., 2020), and detecting landcover change (Shi et al., 2021). Jozdani et al. (2021) and Richardson et al. (2021) have explored the use of CNNs for lichen mapping with high resolution imagery. Jozdani et al. (2021) implemented a CNN for binary classification of lichen in WV imagery (50 cm), with an overall accuracy of 85.28% and an F1-score o 84.38%, outperforming a Random Forest classifier. Richardson et al. (2021) developed UAV LiCNN, a multi-scaled CNN trained on vegetation plot photos to classify pixels in UAV orthomosaics (2 cm). This model had mean user and producer accuracies of 85.84% and 92.93% respectively when classifying pixels containing a high lichen percentage but had lower accuracies at classifying lichen in shadows (Richardson et al., 2021).

A small body of published work has focused on implementing CNNs for regression applications. Regression CNN models have been implemented to map freshwater shortages (Panahi et al., 2020), bathymetry (Al Najar et al., 2022), chlorophyll-a (Chl-a) in lakes (Aptoula & Ariman, 2022), and lichen (Erlandsson et al., 2022). Aptoula et al. (2022) developed a six-layer CNN which used 3x3 tiles with 12 Sentinel-2 bands to estimate Chl-a in Lake Balik. The authors discovered that CNNs outperformed non-deep learning methods and that a two-dimensional CNN outperformed a three-dimensional CNN. Erlandsson et al. (2022) developed a Landsat-based lichen biomass model that used Scandinavian field records collected over twenty years to inform a deep CNN. This model had an R^2 of 0.57, and an MAE of 7.5% when predicting lichen volumes. Erlandsson et al. (2022) reported that the design of the datasets was likely unable to compensate for spatial autocorrelation and that the results are suspected to be overly optimistic. The low number of training samples (3,290) and complex structure of this model may also indicate an overdetermined CNN, further supporting the authors' conclusion.

An issue for many lichen mapping models is the availability of transferable true observations to open-source satellite data. Previous studies have created fractional lichen coverage maps using Landsat imagery due to the vast archive of imagery available (He et al., 2021; MacAnder et al., 2020; Nelson et al., 2013; Nordberg & Allard, 2002; Théau et al., 2005; Zimmermann et al., 2018). However, with a pixel resolution of 30 meters, it is challenging to align Landsat imagery with high resolution datasets. Sentinel-2 imagery is freely available, has a higher resolution of 10 m in some bands for more accurate alignment to UAV datasets, and could serve as an intermediate scaling-up step to Landsat datasets. Additionally, Sentinel-2 imagery can be used to create current lichen maps with a higher spatial resolution than Landsat imagery.

Previous research has shown that Random Forest models, dense neural networks, and deep CNNs are all effective at lichen mapping. However, to date, it is unclear which type of model performs best when compared on the same dataset. In this present study, these models were assessed in a study region with a dominant lichen species of *Cladonia stellaris*, consisting of the Manicouagan Caribou Range in Québec and the Red Wine Mountain Caribou Range in Labrador. This study assesses which regression machine learning methodology is most effective at lichen mapping. The results serve as guidance for future lichen mapping approaches and other landcover mapping applications.

2.2 Materials and methods

2.2.1 Lichen map datasets

The lichen maps used for model training and evaluation were created from 20 orthomosaics derived from UAV imagery collected between July 24 - July 31, 2019, and June 25 - July 2, 2022. The imagery was collected at sites located along the Québec Route-389 or the Trans-Labrador Highway and were accessible by foot or vehicle (Figure 1). The sites represented a variety of different land covers and topographies in the Québec and Labrador region. All sites were included in this study to maximize the different land cover conditions the machine learning models were exposed to.

The 2019 UAV imagery was collected using either a DJI Inspire 1 with a Sentra Double 4 K Normalized Difference Red Edge camera or a DJI Mavic 2 pro with its stock RGB camera. One frame per second was extracted from the 4 K video for the DJI Mavic 2 Pro, while one photograph per second was captured by the Sentra Double 4 K camera, following the methodologies outlined by Fernandes et al. (2018) and Leblanc (2020). One of the 2019 UAV image datasets were flown with a DJI Mavic 2 Pro and captured images with a 1 cm resolution over an area of 2.5 ha. The 11 other sites visited in 2019 had a spatial resolution of 2 cm resolution and areas of approximately 17 ha. The 2022 UAV imagery consisted of eight sites and was collected using a DJI Phantom 4 Pro with its stock RGB camera. One photograph was collected every two seconds during operation and all resulting 2022 orthomosaics had a 1.5 cm resolution over areas of approximately 10 ha.

Pix4D mapper version 4.4.12 and 4.75 was used to process the 2019 and 2022 UAV orthomosaics respectively. All UAV orthomosaics used a minimum of three high-accuracy Global Navigation Satellite System (GNSS) data ground control points to improve the locational accuracy of the UAV orthomosaics.

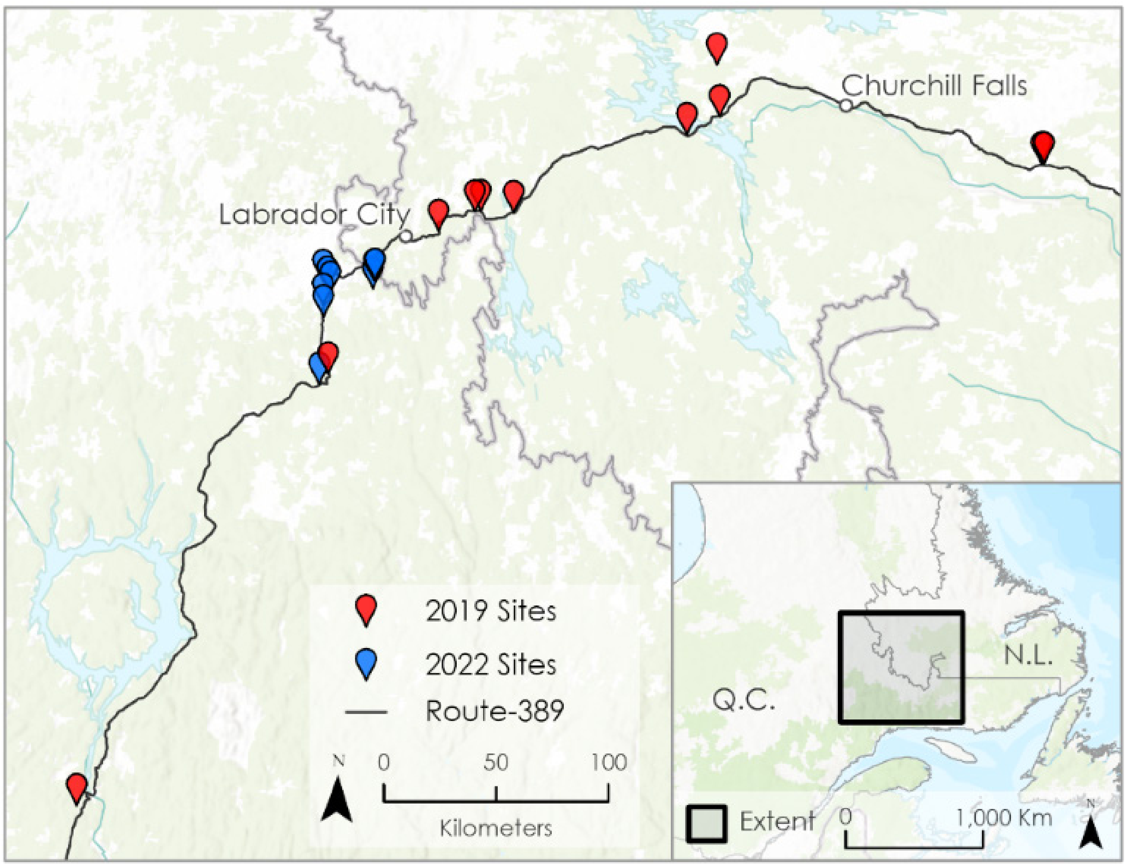

Figure 1. Map of the sites visited in this study.

The 1 cm and 1.5cm resolution orthomosaics were resampled to 2 cm pixels with bilinear resampling before classification. UAV-LiCNN, a methodology developed by Richardson et al. (2021), was used to classify lichen in orthomosaics into three classes, which represented 0% lichen coverage, 50% lichen coverage, and 100% lichen coverage. The classification outputs were manually cleaned using the pixel editor in ArcGIS Pro to correct substantial misclassifications of shadows, bright logs, and sand. Lichen coverage (%) maps at 10 m resolution were generated from the UAV lichen maps by using the block statistics tool in ArcGIS Pro (Figure 2). The mean lichen coverage (%) from the UAV lichen maps were calculated for each Sentinel-2 pixel, creating 10-m lichen maps where pixel values ranged from 0 to 100.

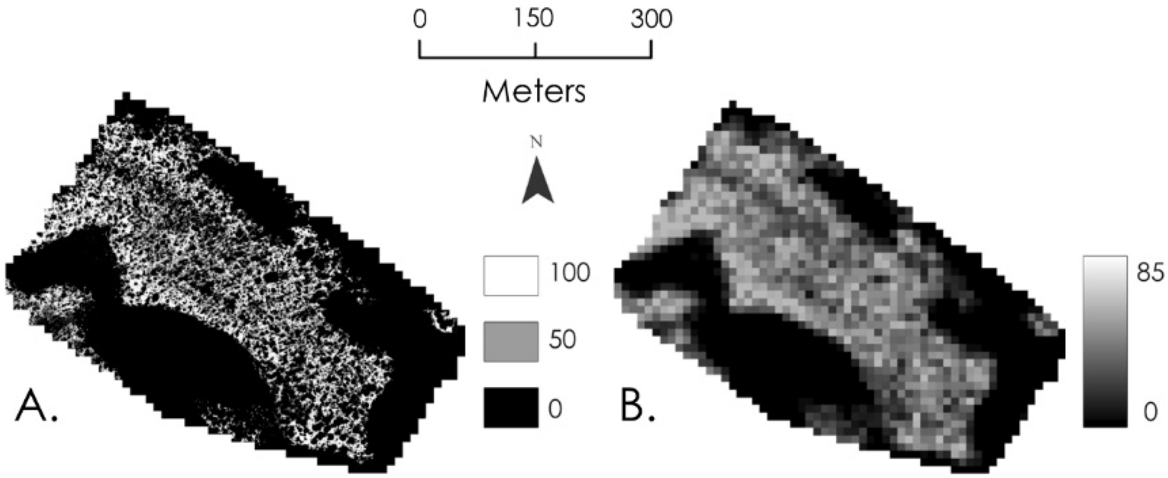

Figure 2. A. Example of a UAV lichen map, B. Example of a 10-m lichen map.

2.2.2 Sentinel-2 imagery

The Sentinel-2 surface reflectance imagery used in this study was from the Harmonized Sentinel-2 MSI: Multispectral Instrument, Level-2A collection accessed through Google Earth Engine (GEE). The harmonized surface reflectance collection was selected to reduce pre-processing and to ensure that the DN values in the 2022 imagery were in the same range as those from previous years. Sentinel-2 images that did not contain clouds or cloud shadows over any of the sites visited that year were selected for this study. The Sentinel-2 imagery used in this study for the 2019 and 2022 sites was captured on August 26th, 2019, and August 18th, 2022, respectively. The imagery was downloaded from GEE, with each band at a resolution of 10 m in a GeoTIFF format.

For this study, we excluded all natively 60-m bands due to their low spatial resolution. Band 8A (20 m) was also excluded due to the spectral closeness with Band 8 (10 m) and lower spatial resolution. Figure 3 outlines a workflow for processing the Sentinel-2 imagery. The remaining nine bands were composited and reprojected to their respective UTM projection with bilinear resampling. The 12.5 m geolocation accuracy of Sentinel-2 required the imagery to be georeferenced to the UAV orthomosaics (Stumpf et al., 2018). This was done in ArcGIS Pro with one tie point to translate Sentinel-2 imagery into alignment. Then the Sentinel-2 imagery was clipped and saved as three 16-bit portable network graphic (PNG) files for further processing in TensorFlow.

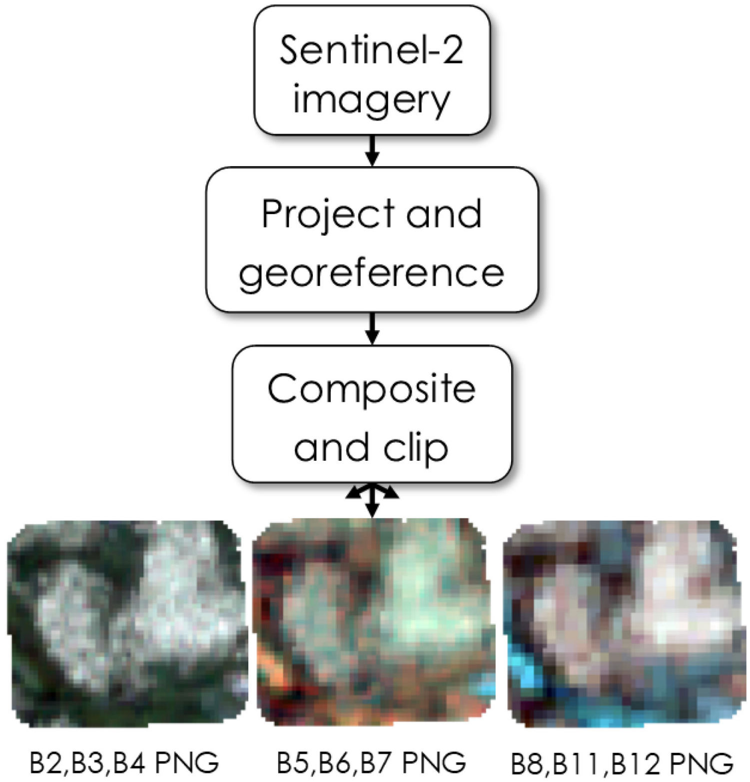

Figure 3. Sentinel-2 imagery processing workflow.

2.2.3 Dataset preparation

It was essential to create a training dataset that exposed the machine learning models to the diversity of landscapes found in the UAV orthomosaics. Therefore, portions of all sites, excluding the smaller 2019 2.5 ha site, were incorporated into the train, validation, and test datasets. An important consideration was the spatial autocorrelation found in the Sentinel-2 imagery and 10-m lichen map. Previous studies that have presented seemingly outstanding performances of CNNs commonly use random cross-validation, which does not account for spatial autocorrelation between training and testing datasets (Kattenborn et al., 2022). However, mapping studies that do not address spatial autocorrelation can have overoptimistic assessments of model predictive power (Kattenborn et al., 2022; Ploton et al., 2020; Wadoux et al., 2021). To avoid this, we implemented spatial partitioning (also known as “blocking”), where we split the sites into quadrant blocks, assigning blocks rather than individual pixels to the train, validation, and test dataset (Figure 4). Following guidance from Ploton et al., (2020), this strategy produced sufficiently large blocks to address spatial autocorrelation within each of the orthomosaics by largely eliminating pseudo-replication issues and the interdependence between train, validation and test data that would ensue with a random sampling of spatially autocorrelated data. However, at the same time, the model is still exposed to many different landscapes. All orthomosaics were subjected to blocking except for the small 2.5 ha site which was completely included in the test dataset.

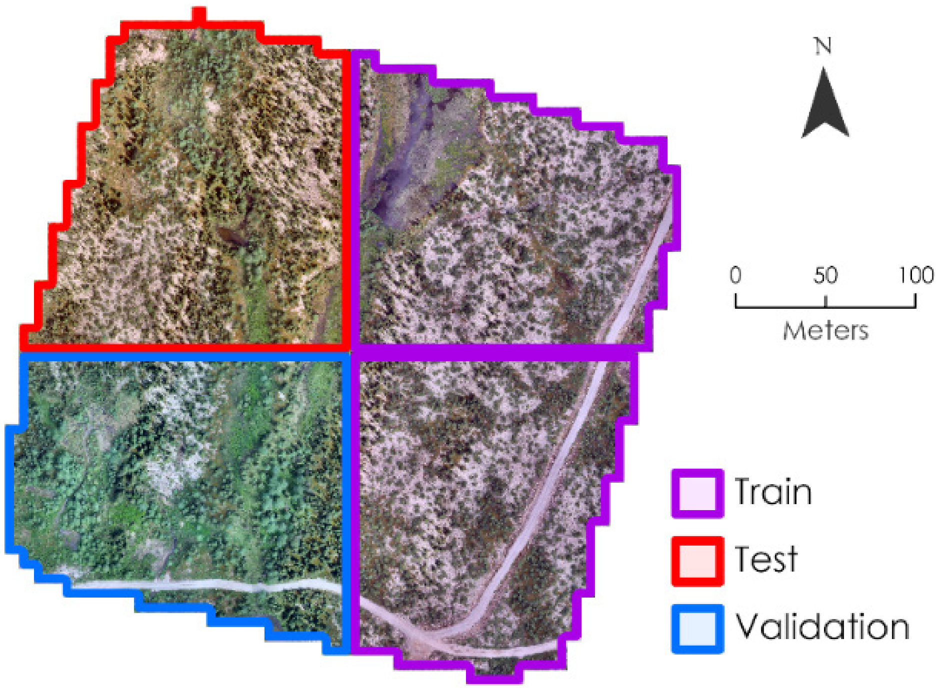

Figure 4. Example of quadrant blocking used to minimize spatial autocorrelation between the datasets.

The Sentinel-2 imagery and 10-m lichen maps were clipped to the quadrant block extents and split with 60%, 20%, and 20% of the blocks allocated to the train, validation, and test datasets respectively. We initially tested randomized allocation of blocks to the train, validation, and test datasets, but this resulted in uneven distribution of different landscapes and lichen percent coverages. Additionally, we wanted to ensure that the test and validation datasets contained data from as many different sites as possible. To address these issues, we manually allocated the blocks to ensure that each dataset had a balanced representation of different landscapes and lichen coverages. Lichen near sand, recent forest fire burns, lichen regrowth, high and low lichen coverage patches, roads, lichen in forests, lichen in wetlands, lichen surrounding downed trees, and rocks with lichen were some of the landscapes considered when deciding which block went into which dataset. As an additional step, before proceeding we verified that the datasets had similar lichen percent coverage histogram distributions (Appendix A).

2.2.4 Convolutional neural network dataset preparation

The inputs for the CNN model were square tiles generated from Sentinel-2 imagery blocks, and the output was a single value which represented the lichen percent coverage of the centre pixel of each tile. PNG files were used as a file format in this study for CNN training as TensorFlow does not yet offer native support for Tag Image File Format (TIFF) files with more than three bands. We tested neural network training with 9-band TIFF files in TensorFlow using the *tifffile* python package and found that it was significantly slower than loading three PNG with the same data. The quadrant blocks designated for each dataset were iterated through, and tiles with sides of 3 pixels (px), 5 px, and 7 px were generated from the three Sentinel-2 PNG files and the 10-m lichen maps (Figure 5). There was overlap between each tile in their respective dataset to maximize the number of unique training inputs.

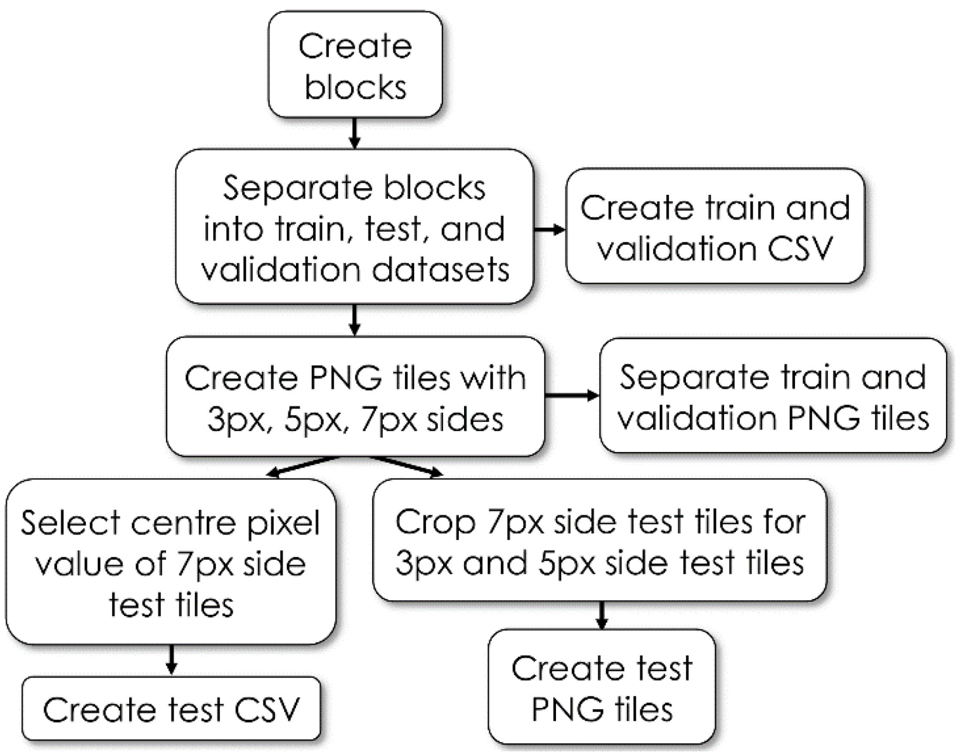

Figure 5. Workflow depicting dataset preparation for the CNN, dense neural network, and Random Forest models.

Datasets with larger tile sides had fewer possible tiles created within each block since all tiles containing NoData values were excluded from the datasets. The number of training and validation inputs differed with the tile size, with 11,621 training inputs for the 7 px tiles and 19,594 training inputs for the 3 px tiles. To ensure consistent test datasets for comparing models, only the central values of the 7 px tiles were used. We cropped the 7 px test dataset tiles to generate 3 px and 5 px tiles for their respective test datasets. This resulted in 3,963 test pixels used to compare the different models in this study.

2.2.5 Random Forest and dense neural network dataset preparation

For each site, the pixel values within the three Sentinel-2 PNG files and the 10-m lichen map quadrant blocks were read and extracted into a comma-separated value (CSV) file for the Random Forest and dense neural network (Figure 5). The training dataset and validation dataset consisted of 24,998 and 7,794 pixels respectively. To ensure consistency in model comparison, we limited the test dataset for the Random Forest and dense neural network to only contain the central pixel also found in the CNN test dataset, resulting in 3,963 test pixels.

2.2.6 Random Forest training

The Random Forest model was created using the Sci-kit Learn Random Forest regressor. The training and validation datasets were used to optimize model hyperparameters *n_estimators* (number of trees) and *max_features* (number of subsets to consider in node splitting). Model performance was assessed with MAE because the average performance of these models over large regions was the main measure of model efficacy.

We tested values of 30, 100, 300, and 1000 for *n_estimators*, and auto, sqrt, and log2 for the *max_features* parameters. The model that performed the best on the validation dataset had *n_estimators* = 1000 and *max_features* = log2, taking 5 minutes to train on an Intel i5-12600k CPU. The best performing model had an R^2 value of 0.73, and an MAE of 5.53% when compared to the validation dataset.

2.2.7 Neural network parameter and model selection

For the dense neural network and CNN models, variations of dense, drop out, batch normalization, and one-dimension convolutional layers were tested in different arrangements.

Various forms of CNN model architecture were tested including combinations of 2D convolutions, cropping, concatenations, and flattening layers. The Adam optimizer, randomized initial weights, and a combination of rectified linear unit (ReLU) and linear activation functions were used in our final models. Appendix B outlines our selection of activation functions, loss functions, and optimizers in model selection.

Hyperparameter tuning was completed as an iterative process using a Python script that cycled through various settings and logged results. We then evaluated the loss curves of the best performing models to determine if they showed signs of overfitting. Classic overfitting is seen when the model does not generalize well from the training data to the validation data (Salman & Liu, 2019; Ying, 2019). Neural networks are prone to classic overfitting due to the large number of parameters that can learn to map inputs to outputs without being able to generalize beyond the training dataset (Salman & Liu, 2019). To avoid selecting a model that has overfit, we implemented an early stopping function. We also disregarded models which developed significant gaps between the training and validation loss curves, because that is a sign of learned noise patterns in the training dataset (Salman & Liu, 2019). Models which had the lowest MAE on the validation dataset and did not show signs of overfitting were evaluated in this study.

2.2.8 Dense neural network training

The dense neural networks assessed in this study were created in TensorFlow 2.8 and consisted of dense layers and input layers. The dense neural network design started from the three-layer model He et al. (2021) proposed, which was improved upon by adding layers and altering parameters. The input layer consisted of a tensor containing the nine Sentinel-2 band values, and the output value was the estimated lichen percent coverage of the input pixel. The first four dense layers of this model used a Rectified Linear Unit (ReLU) activation function, and the last dense layer used a linear activation function. The model contained 2,337 trainable parameters and was compiled with a learning rate of 0.001 using the Adam optimizer. It had a batch size of 1000 with 225 steps per epoch and took 66 seconds to train on an Nvidia RTX 3070ti graphics card. The model had a patience of 15 epochs; if the validation loss did not decrease over this time the best model weights were saved, and the training would end. The model weights were saved at epoch 54 since the validation loss did not improve after that point

(Figure 6). The model had an R^2 of 0.76, and an MAE of 5.14% when compared to the validation dataset.

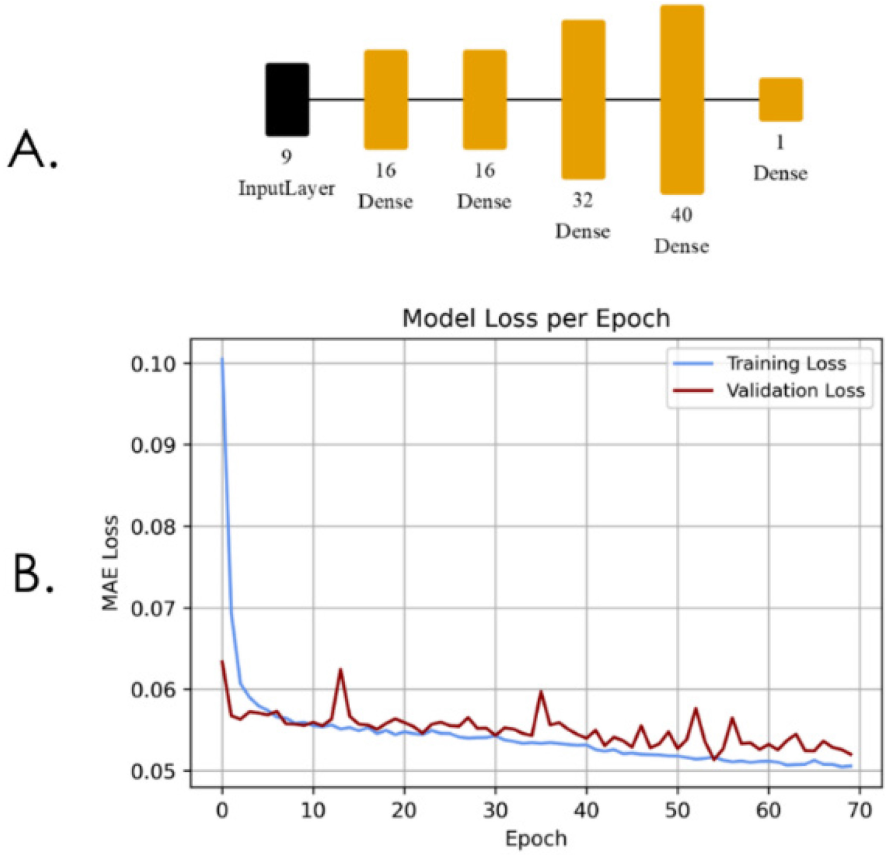

Figure 6. A. Visual representation of the dense neural network model using Net2Vis (Bäuerle et al., 2021), B. training and validation loss curves.

2.2.9 Convolutional neural network training

The objective for the CNN models was to interpret a tile of Sentinel-2 imagery and make an estimate of the lichen percent coverage of the centre pixel. The training tiles were passed through a data augmentation step, which randomly flipped and rotated the tensor to expose it to different orientations. To reduce the number of features, the CNNs in this study use valid padding strategies to discard the data at the edges of the tensor (Bhatnagar et al., 2020; Chen et al., 2020). Valid padding ensures that convolutional filters are only applied to values in the tensor and that the height and width of the output tensor is eroded (Alsallakh et al., 2020; Bhatnagar et al., 2020; Chen et al., 2020). Same padding maintains the same output tensor size as the input by adding zeros to the edges, and is used to maintain the size of the feature map

(Alsallakh et al., 2020; Bhatnagar et al., 2020; Chen et al., 2020). We created a CNN which used valid padding to erode the tensor until the feature map resembled a shape of 1x1 with the number of bands equal to the number of filters used. At this point, the feature map was flattened and passed through a series of dense layers that returns a single percent lichen coverage prediction.

The best performing CNN in this study was created in TensorFlow 2.8 and consisted of two-dimension convolutional, flatten, and dense layers. The two-dimension convolutional layers and the first three dense layers of this model used a ReLU activation function, and the last dense layer used a linear activation function. The model contained 9,001 trainable parameters and was compiled with a learning rate of 0.001 using the Adam optimizer. It had a batch size of 500 with 80 steps per epoch and took 13 minutes to train on an Nvidia RTX 3070ti graphics card. The model had a patience of 10 epochs and the weights were saved at epoch 34 since the validation loss did not continue to decrease after that point (Figure 7). The selected model had an R^2 of 0.75, and an MAE of 5.21% when compared to the validation dataset.

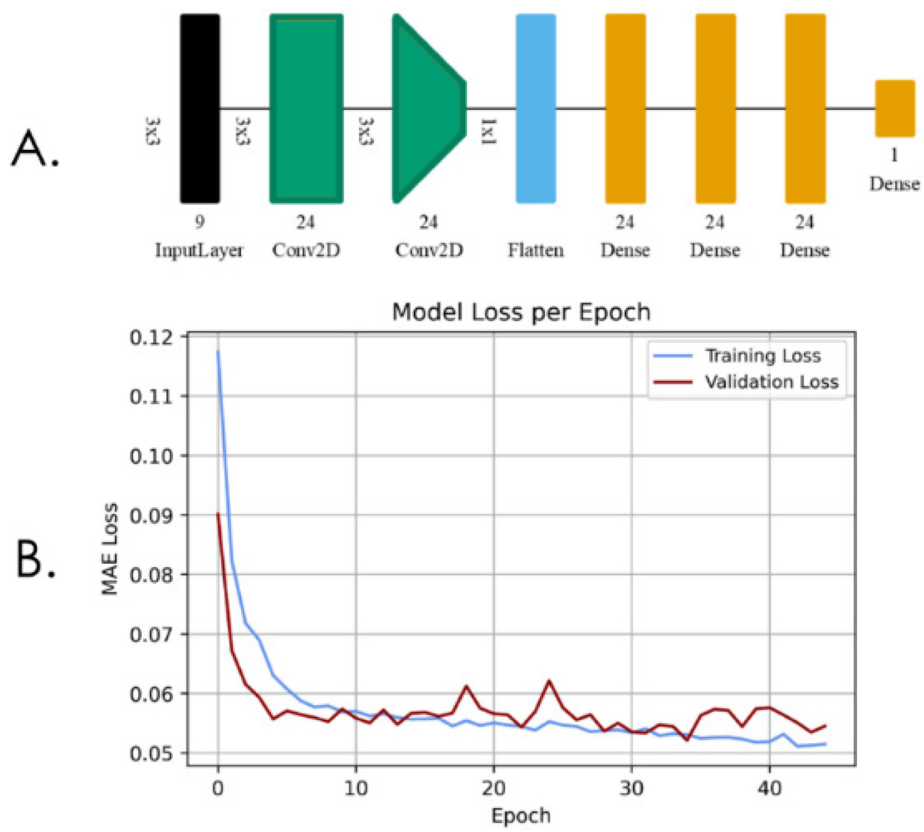

Figure 7 A. Visual representation of the CNN model using Net2Vis (Bäuerle et al., 2021), B. training and validation loss curves.

2.3 Results

2.3.1 Test dataset pixel comparison

The dense neural network outperformed the Random Forest and CNN models in the test dataset pixel comparison, with an R^2 value of 0.76 and an MAE of 5.2% (Figure 8). The Random Forest and CNN had MAEs of 5.5% and 5.3% respectively, and each had R^2 values of 0.74. All three models tended to under-predict pixels with high lichen coverage.

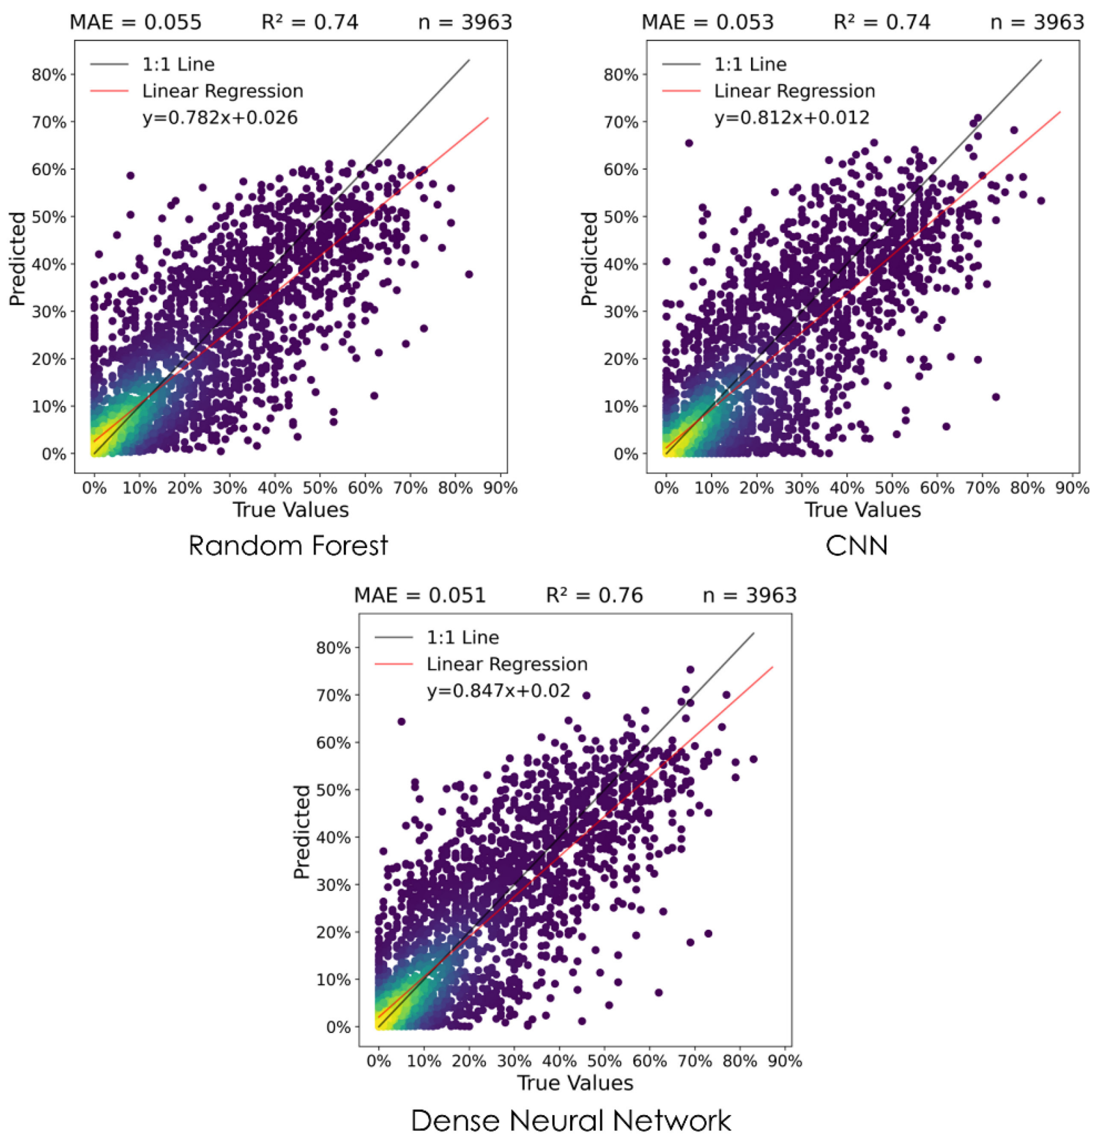

Figure 8. Pixel lichen percent coverage scatterplots, MAE, and R^2 values for the model performance on the test dataset in this study. Brighter yellow colours indicate a higher frequency of points.

The Random Forest model seemed to misclassify more non-lichen pixels (true value of zero) as having lichen than the CNN and dense neural network. Conversely, the CNN tended to predict no lichen (predicted value of zero) for more pixels that had lichen than either other model.

2.3.2 Test dataset blocks comparison

This comparison assessed the performance of the models over larger areas, rather than comparing the individual pixels, by predicting the mean lichen percentage of each block within the test dataset. The test dataset consisted of 17 blocks from different sites, with a mean block size of 400 Sentinel-2 pixels. The Random Forest and dense neural network had an R^2 of 0.93, showing a very high coefficient of determination (Figure 9).

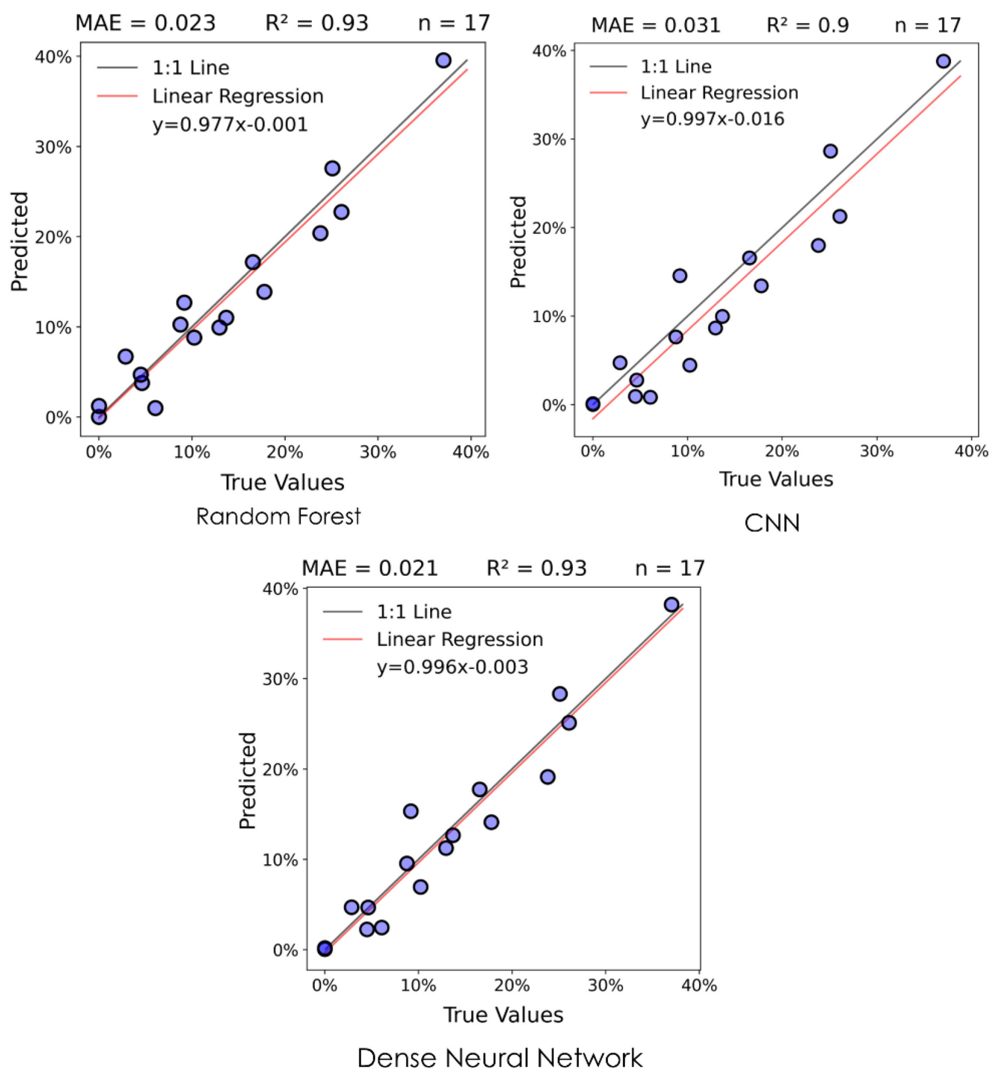

Figure 9. Graphic depicting mean block lichen percent coverage scatterplots, MAE, and R^2 values for the models in this study.

The MAE for the dense neural network and Random Forest were 2.1% and 2.3%, giving the dense neural network a slight edge. The CNN had a lower R^2 (0.90) and a higher MAE (3.1%) than the other two models and showed a stronger systematic tendency to underpredict mean lichen percentage across blocks.

2.3.3 Lichen map using the dense neural network

We created a Sentinel-2 harmonized surface reflectance mosaic that covered most of the sites used in this study. The mosaic was created from the most recent cloud-free pixel from an August image. There were a few isolated instances, primarily over bodies of water, where the pixels had no imagery available. These pixels were filled in by using median focal statistics. The dense neural network model took 14 hours using an Intel i5-12600k CPU to create a 180 km^2 (15878 columns, 11396 rows) regional lichen map (Figure 10).

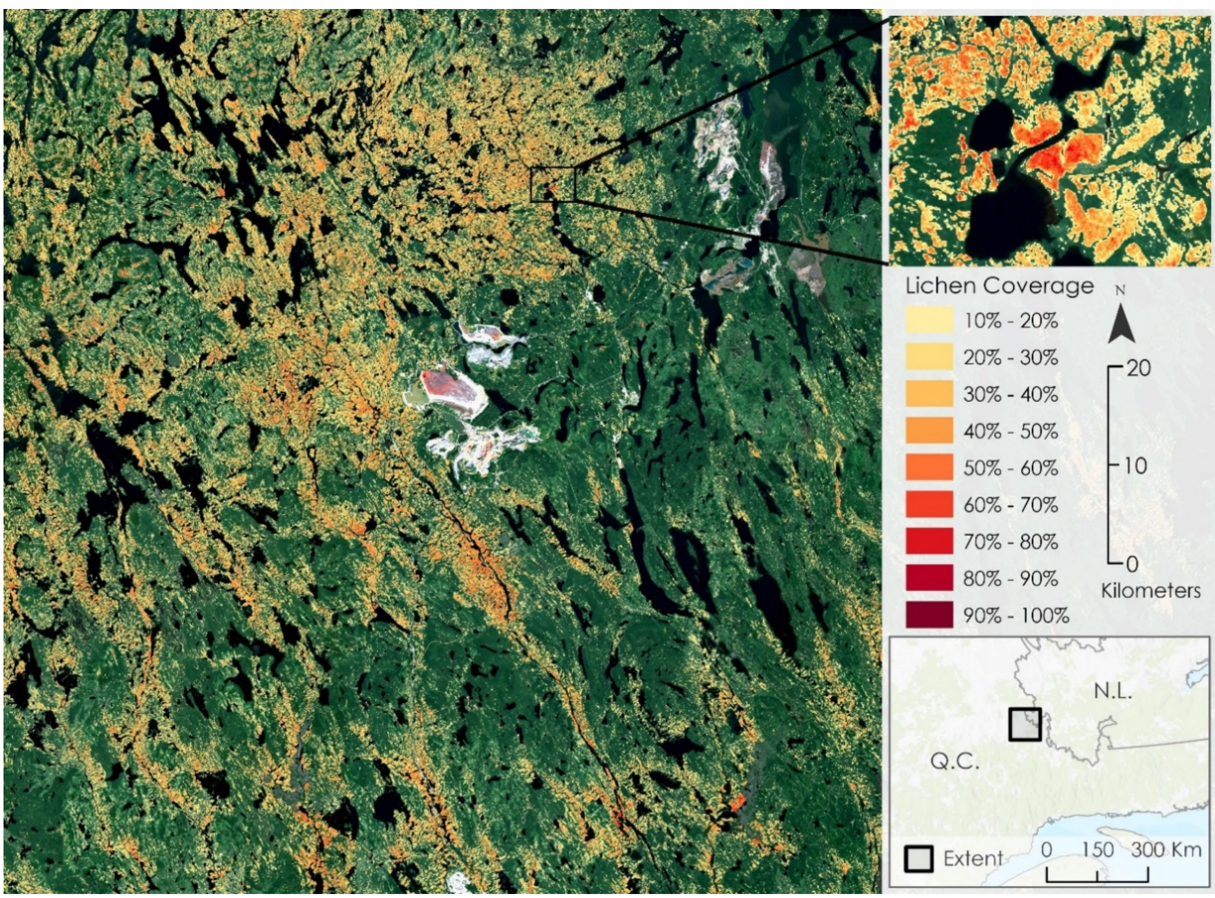

Figure 10. Lichen map of Fermont Quebec with inset of high lichen region. A lichen coverage of less than 10% is transparent in this map.

2.3.4 Comparison of lichen map and true lichen values

To better understand how the models performed at creating lichen maps, we compared the model generated lichen maps from the Sentinel-2 imagery mosaic with the true values for two of the 20 sites used in this study (C8 and S4). Site C8 consisted of very high lichen coverage interspersed between black spruce and Site S4 consisted of lichen interspersed between trees with a distinct parking lot surrounded by lichen. Figure 11 shows the Random Forest, dense neural network, CNN, and true lichen maps for these sites.

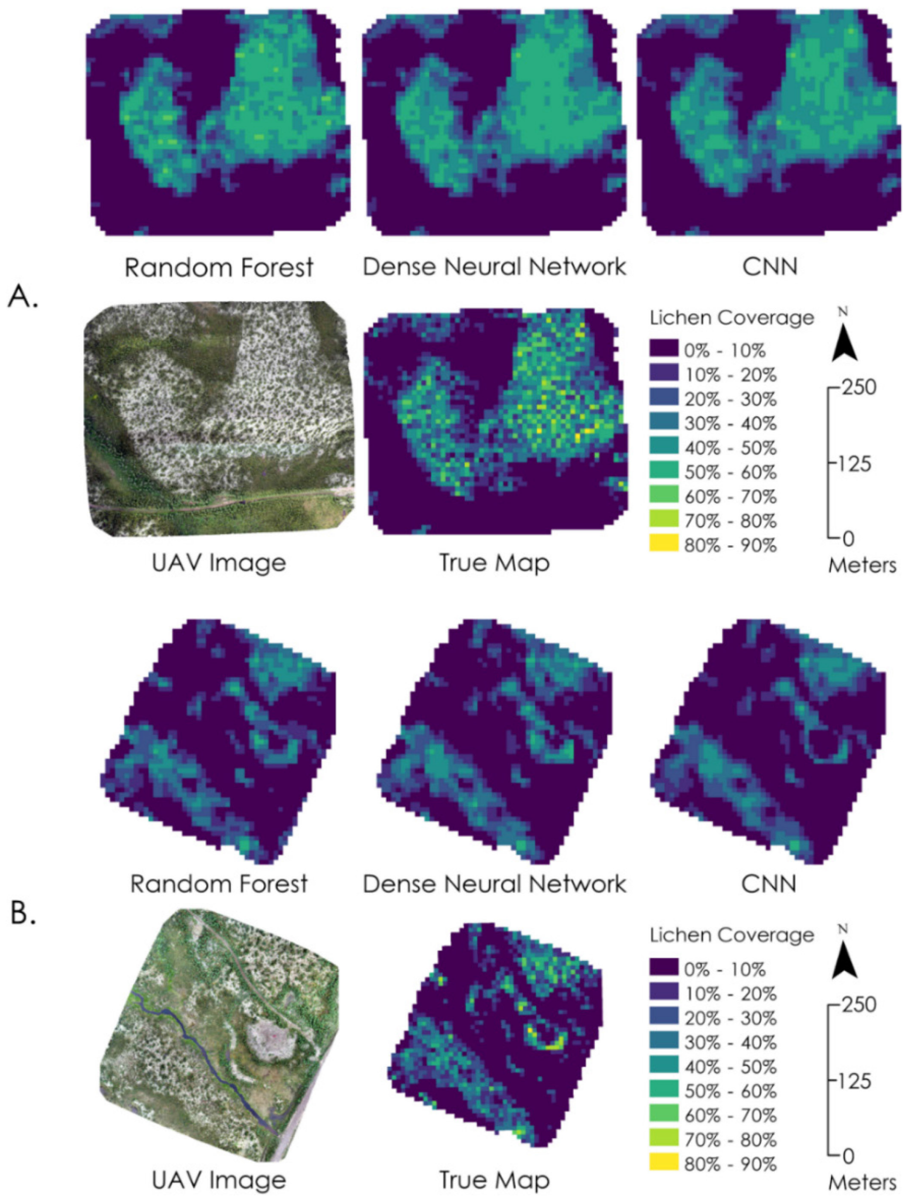

Figure 11. A. Comparison of lichen maps for Site C8, B. Comparison of lichen maps for site S4.

Across both sites, the predicted maps appear smoother and miss local variability found in the true lichen maps. This could be due to the spatial resolution of the Sentinel-2 sensor compared to the UAV imagery which was used to create the 10-m lichen maps. As observed in the scatterplot comparisons (Figures 8 & 9), the CNN had a greater tendency to underpredict lichen coverage percentage in these sites than the dense neural network and Random Forest model. This was apparent in the pattern of the lichen values in site C8, and the lower lichen values in the lichen pixels across site S4. The Random Forest model and the dense neural network produced similar lichen maps for site C8, with the Random Forest predicting more lichen pixels in the 60-70% bin than both the dense neural network and CNN.

We created difference maps between the true lichen maps and the predicted lichen maps to identify locations with large prediction errors. The difference map for both sites had mean errors of approximately 0 with standard deviations of approximately 10 percentage points (p.p.) difference between the predicted and true values. To highlight pixels that show substantial overestimation or underestimation, we thresholded each difference map at +/- 20 p.p. difference. Figure 12 shows the resulting difference maps where red and blue pixels represent substantial overestimates and underestimates respectively.

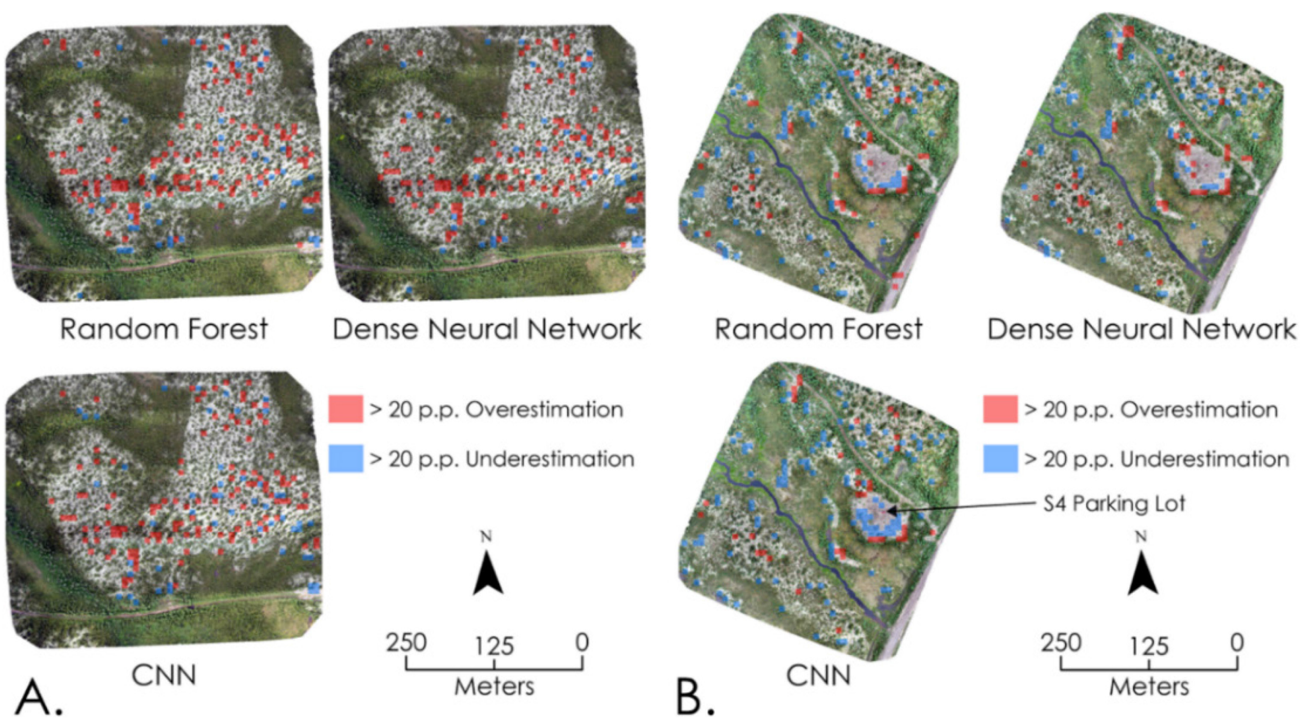

Figure 12. A. Map of overestimations and underestimations in site C8, B. Map overestimations and underestimations in site S4 with an arrow pointing to the parking lot.

Across all models, site C8 has a significant number of overestimations in the high lichen patches. This could be due to the difference in local variability found in the true lichen map and model predictions. The CNN systematically underestimates lichen in site S4, a pattern that is not replicated by the Random Forest and dense neural network. There are noticeable differences in model predictions of the site S4 parking lot surrounded with lichen, a feature consisting of high lichen coverage near exposed ground. Here, all models have numerous substantial overestimations and underestimations, with the CNN and Random Forest models having more underestimations than the dense neural network. Unlike the Random Forest model, neither the dense neural network nor the CNN have substantial misclassifications on the road in site S4 as lichen.

2.4 Discussion

In this study, we compared three different types of machine learning models on their ability to predict lichen percent coverage from Sentinel-2 imagery. Each type of model had different predictive performance, complexity, and computing power required for training. The dense neural network model outperformed the Random Forest and convolutional neural network models in terms of lower MAE in both test dataset comparisons, and a higher R^2 in the test pixel comparison.

The dense neural network was significantly faster to train and did not require as many pre-processing steps as the CNN. A substantial challenge to training CNNs using TensorFlow is the reading of nine-band imagery. While opening three 16-bit PNG files was faster than opening one 9-band TIFF, the limiting factor in CNN training speed was the opening of files, which could potentially be resolved by accessing all the data from memory or using TFRecord files.

Herein, the most common hyperparameters for each model were tuned using grid search, with the sole purpose of finding the combination of parameters that minimized loss and maximized performance. Doing so was a valid attempt to undertake model tuning for the Random Forest, dense neural network, and CNN using the most common hyperparameters that are known to affect performance. While no claims are made that all models were tuned perfectly, all models were at the least optimized. The optimization provides a good degree of fairness among the three models that were developed.

A natural question that arises from this work is why the CNN did not perform better given the consensus that CNNs generally outperform dense neural networks and Random Forest models for image analysis problems. Often very deep CNNs such as VGG19, ResNet, and Xception are used in image analysis but using a transfer learning approach, the most common being a pre-trained network with ImageNet weights (Kornblith et al., 2019). Given the nature of the training data in this research, there were no well-proven CNN architectures or benchmark datasets that were considered sufficient to represent the problem at hand. Instead, we trained all models from scratch, with random weights initially assigned and used minimally deep networks for training.

We found that increasing the window size of the CNN tiles did not reduce the MAE for lichen mapping. Larger window sizes had a stronger tendency to overfit than smaller window sizes. The MAE on the test pixel dataset for the 7 px and 5 px tile models was 5.5%, while the best 3 px tile model had an MAE of 5.3% lichen coverage. Larger tile sizes could increase the noise introduced into the model and potentially lower the accuracy of predictions. Different types of CNNs were tested on all tile sizes, with some excluding dense layers and other models passing the central pixel values along with the flattened CNN feature map into dense layers. Overall, the CNNs that performed the best had a series of CNN layers, followed by a series of dense layers.

The dense neural network was applied to a Sentinel-2 surface reflectance harmonized mosaic to create the lichen map shown in Figure 10. A small number of pixels located inside mines, along route 389, or in the middle of lakes were predicted to have lichen coverage or NoData values. The models were not exposed to pixels with a reflectance similar to these land covers in the training since this study used lichen maps from vegetation surveys. Regional lichen maps could be improved by integrating land covers not found in UAV vegetation surveys into the datasets and either retaining the model or leveraging transfer learning, a neural network method of refining an existing model with new variables (Astola et al., 2021).

There are significant limitations to the Sentinel-2 surface reflectance harmonized image collection found in GEE at the time of publishing, with no imagery located over our area of interest earlier than 2019. This complicates the process of creating lichen maps from earlier Sentinel-2 imagery or conducting regional time series analysis. However, the regional lichen maps generated in this study could be used as an intermediate step for training a Landsat-based lichen model.

2.5 Conclusion

All machine learning models assessed in this study were accurate at predicting lichen percent coverage in Sentinel-2 imagery covering a variety of different landscapes. This study shows that for lichen mapping in this region, the dense neural network outperformed the CNN and Random Forest models. While the dense neural network requires more computational effort to train than a Random Forest model, the 5.9% performance gain in the test pixel comparison and 9.1% performance gain in the quadrant block comparison renders it the most suitable for lichen mapping. The dense neural network trained on Sentinel-2 imagery had a pixel comparison MAE of 5.2% and R^2 of 0.76, and a block comparison MAE of 2.1% and R^2 of 0.93, making it appropriate for lichen mapping. Further research on retraining or transfer learning with the dense neural network is necessary for a regional lichen map to account for land covers not seen during the vegetation surveys used as training data in this study, such as mines and deep lakes. This study represents substantial progress toward determining the appropriate methodology for generating accurate lichen maps to be used in caribou conservation and sustainable land management.

Supplementary Materials: The latest version of code used in this paper and additional processing scripts can be accessed through GitHub: https://github.com/galenrichardson/Lichensen2modelcompare/, accessed on 1 February 2023.

Contributions: Conceptualization, G.R., A.K., W.C., and M.S.; methodology, G.R. and A.K.; software, G.R., A.K., and L.H.; validation, G.R. and A.K.; formal analysis, G.R. and A.K.; investigation, G.R. and A.K.; resources, G.R., A.K., W.C., J.L., L.H., and L.Y.N.; data curation, G.R. and L.Y.N.; writing—original draft preparation, G.R. and A.K.; writing—review and editing, W.C., M.S., J.L., L.H., and L.Y.N.; visualization, G.R.; supervision, A.K., W.C., and M.S.; project administration, A.K. and W.C.; funding acquisition, G.R., A.K., and W.C. All authors have read and agreed to the published version of the manuscript.

Acknowledgments: The authors thank Dr. Sylvain G. Leblanc and Dr. Robert H. Fraser for their data collection and conceptualization contributions. The authors also thank Elisha Richardson and the Shallow-Water Earth Observation Laboratory (SWEOL) at the University of Ottawa for their assistance with editing. This research was funded through the Earth Observation Baseline Data for Cumulative Effects Program (EO4CE) at the Canada Centre for Mapping and Earth Observation (CCMEO); The University of Ottawa; The support of the Natural Sciences and Engineering Research Council of Canada (NSERC); and the Association of Canadian Universities of Northern Studies.

Chapter 3: Addressing misclassified pixels in the regional lichen map

3.1 Introduction

The regional lichen maps in Chapter 2.3.2 is often erroneous in mines, asphalt roads, and urban environments were occasionally misclassified as containing lichen by the dense neural network. Additionally, dark water bodies occasionally returned 'Nodata' values from the dense neural network, which requires additional steps to resolve when creating a regional lichen map. This section of the thesis evaluates whether integrating pixels from outside the original UAV imagery could improve regional lichen prediction.

Train, validation, and test polygons were generated over mine, road, urban, and water land covers found in the Sentinel-2 imagery used to generate the regional lichen map in Chapter 2.3.2. The corresponding Sentinel-2 pixels were sampled, and assigned to the train, validation, and test datasets. The shapefiles used for each dataset were at a minimum of 200 m apart from each other to reduce spatial autocorrelation. A total of 4000 pixels, 1000 each from mine, road, urban, and water land covers, were appended to the training dataset from Chapter 2, resulting in a total of 28,998 training pixels. Due to the smaller size of the validation and test datasets and to maintain a consistent dataset split, 300 pixels from each of these land covers was appended to each dataset, resulting in 8,994 and 5,163 pixels respectively. A dense neural network with identical parameters to the model outlined in Chapter 2.2.8 was trained 100 times using the train and validation dataset containing the additional land covers. The initial weights for these dense neural networks were the previous best weights from the model proposed in Chapter 2. Retraining the model with randomized initial weights was also evaluated 100 times but had a worse MAE on the validation dataset than the models with defined initial weights.

3.2 Model training and performance

The retrained model outperformed the original model proposed in Chapter 2.2.8, with a considerably higher R^2 of 0.79 and lower MAE of 3.8% than the original model (R^2: 0.67, MAE: 5.2%) on the test dataset that included new land covers (Figure 13). The original model substantially misclassified more non-lichen pixels (the true value of zero) as having lichen than the retrained model. To test if including these pixels affected the dense neural network model performance on the test dataset derived from UAV lichen maps, a comparison between the original and the retrained model was conducted on the test dataset used in Chapter 2 (Figure 14). The retrained model had a higher R^2 of 0.78 and a lower MAE of 4.8% than the original model (R^2: 0.76, MAE: 5.1%) proposed in Chapter 2.3.1. The inclusion of additional pixels from other land covers seemed to slightly improve the model performance at predicting lichen in the test dataset from Chapter 2. When comparing between the plots in Figure 14, there are noticeably fewer pixels that were misclassified non-lichen pixels (true value of zero) in the retrained model than in the original model.

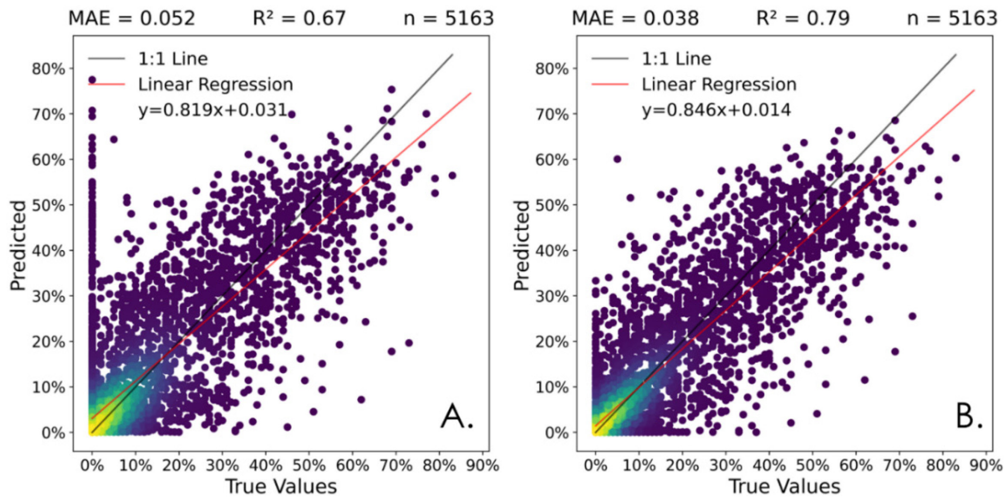

Figure 13. Comparison of the original model (A) and the retrained model (B) on a test dataset including new land covers.

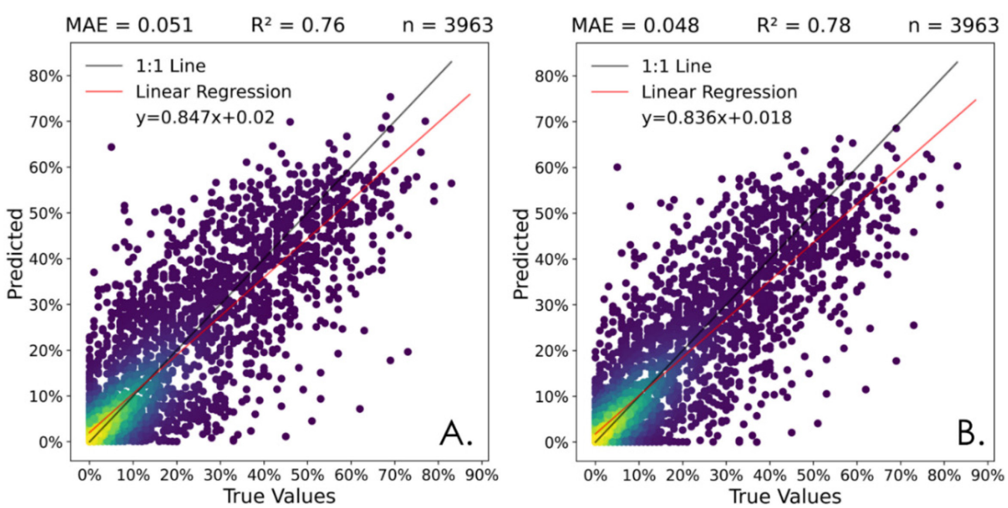

Figure 14. Comparison of the original model (A) and the retrained model (B) on the test dataset used in Chapter 2.

A similar test dataset block comparison to Chapter 2.3.2 was conducted on the original and retrained models using the test dataset from UAV orthomosaics. The performance of the models over larger areas was evaluated by predicting the mean lichen percentage of each block within the test dataset (Figure 15).

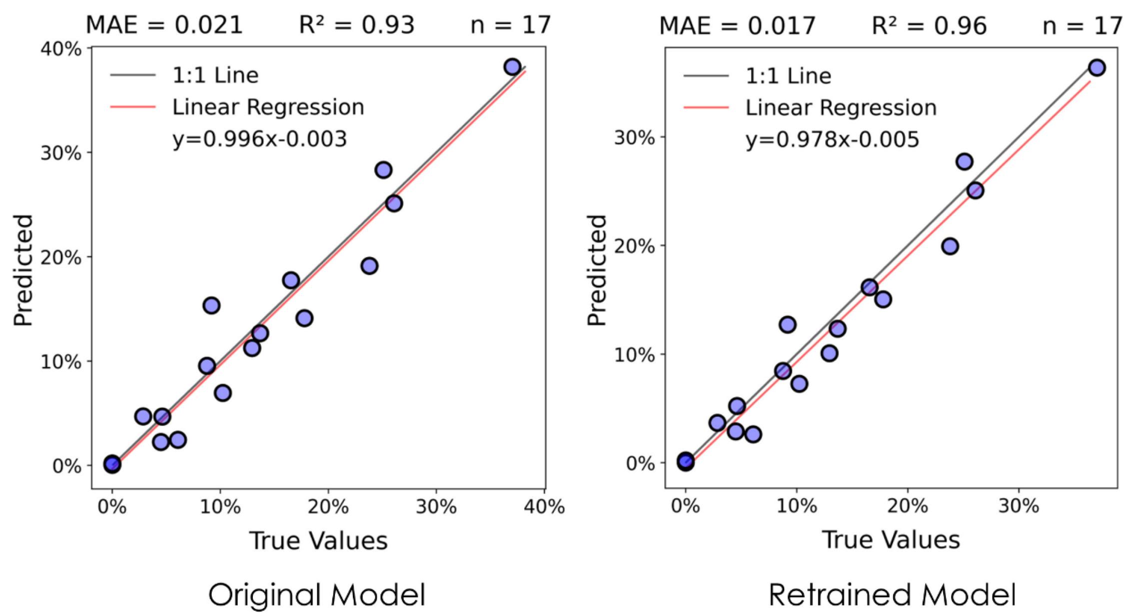

Figure 15. Graphic depicting mean block lichen percent coverage scatterplots, MAE, and R^2 values for the original and retrained dense neural networks.

The retrained model in this comparison had an R^2 of 0.96 and an MAE of 1.7%, both indicating better performance than the original model. The retrained model has a slightly greater systematic tendency to underpredict mean lichen cover (%) in blocks than the original dense neural network model.

A regional lichen map was created using the retrained model (Appendix C) for a visual comparison of the model's performance. Figure 16 contains lichen maps from the original model and retrained model over a smaller area of interest that contains mine, road, and urban land covers.

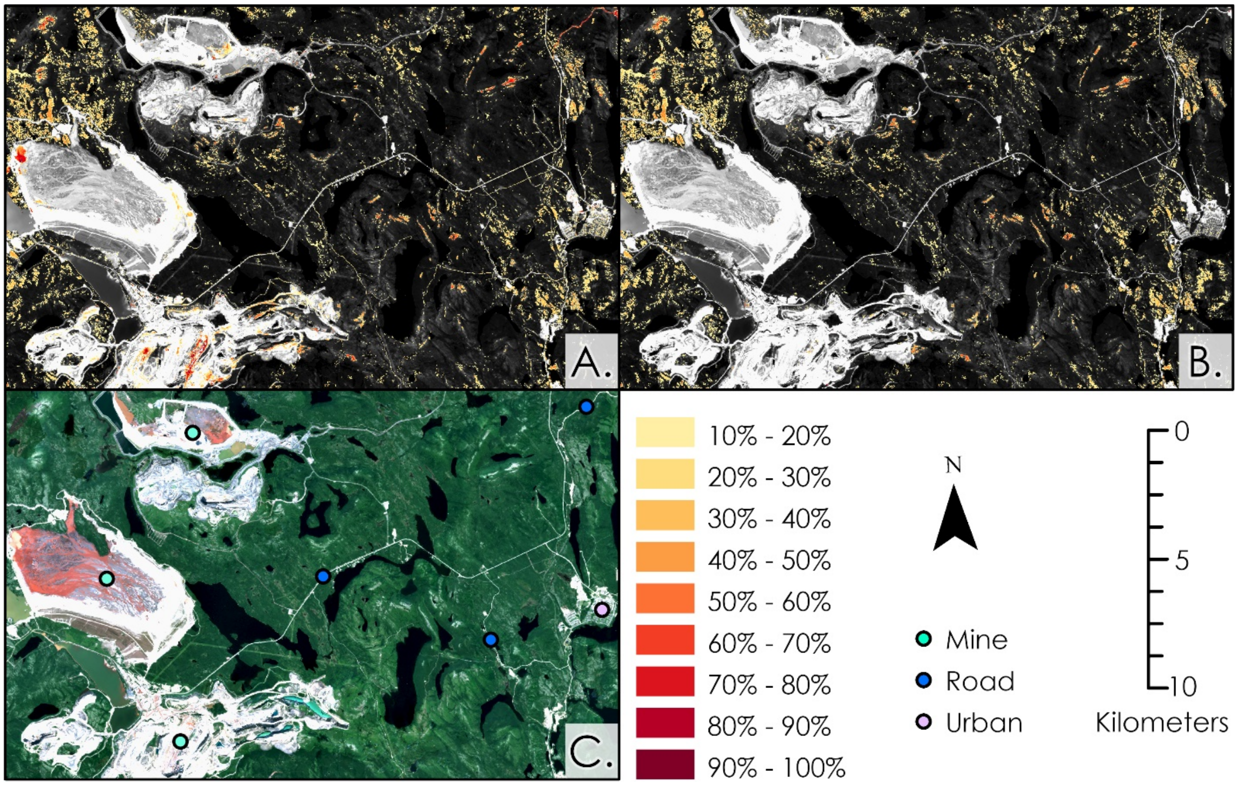

Figure 16. A. and B. lichen maps overlayed on a black and white image of the original and retrained model respectively, C. RGB image with different land covers denoted. Lichen coverage (%) less than 10% is transparent.

The retrained model visually performs better than the original model, avoiding substantial lichen misclassifications of mine, road, and urban land cover (shown as orange/red patches on bright background areas in Figure 16A). The number of 'Nodata' values found in water bodies was also reduced in the retrained model compared to the original model. An intermediate step using the Con tool in ArcGIS Pro was still necessary to set all 'Nodata' values to zero from this output. The road located in the northwest of the map that was classified with a high lichen coverage (%) in the original model seemed to be properly classified as having no lichen in the retrained model. The town of Fermont, Quebec located in the east, which contained a few incorrect low lichen predictions in the original model, had substantially fewer such errors in the retrained model.

3.3 Conclusion

The retrained dense neural network outperformed the original model trained only on UAV lichen maps, with a higher R^2 value and a lower MAE value when assessing pixels found in the original test dataset. The retrained model also had lower MAE and higher R^2 values than the original model during the block comparison and pixel comparison of the test dataset containing additional land covers. The inclusion of pixels not captured in the vegetation surveys increased the lichen coverage (%) predictive performance of dense neural network over Sentinel-2 imagery. A lichen map was created using the retrained model, and many of the previously misclassified land covers visually seemed to be more accurately predicted.

The unique spectral characteristics of mines, roads, and urban land covers can be challenging for machine learning trained on vegetation plots to correctly predict. Including pixels from such land covers can increase the predictive performance of machine learning models that scale-up UAV lichen maps to regional lichen coverage (%) maps.

Chapter 4: Conclusion

4.1 Summary

Chapter 2 determined the optimal machine learning model to create regional lichen coverage (%) maps from UAV datasets using Sentinel-2 imagery. The dense neural network outperformed the Random Forest and CNN models, providing the lowest MAE in both accuracy assessments and with fewer substantial errors in the lichen map comparison (Chapter 2.3.4). Chapter 3 determined that the inclusion of mines, roads, water, and urban land cover pixels in training can improve the results of the dense neural network at predicting lichen in the test dataset used in Chapter 2 and a test dataset containing additional land covers. The resulting lichen map from the retrained model visually performs better than the original model, avoiding substantial lichen misclassifications of mine, road, and urban land cover.

Of the three models assessed in this study, the CNN was the most difficult model to train and optimize, while the Random Forest model was the easiest. Section 2.4 discussed a few of the possibilities why the CNN did not outperform the other models. Here I present a few additional reasons that were not included in the publication. One possible reason is that CNN's and convolutional-2 dimensional layers were originally designed for single-band images and are not optimized to process nine-band imagery. The filters used in convolutional 2-dimensional layers apply to all the bands at the same time and generalize the variabilities found in individual bands. Alternatively, DepthwiseConv2d layers apply filters to each individual band, potentially enhancing additional information that can be used for predictions. Initial tests showed that DepthwiseConv2d layers performed slightly worse than convolutional 2-dimensional layers while substantially increasing the model complexity. However, there could be a possible CNN model that leverages these layers in a different architecture.

Another potential reason for the CNN's results could be the high variability of lichen found in Sentinel-2 imagery tiles from boreal Eastern Canada. The CNN models had a much higher tendency to overfit than the dense neural networks, possibly because the model was learning noise from the training data. The shortcomings of the CNN could be due to the spatial and texture patterns being less meaningful than anticipated to improve results over the dense neural network.

The Random Forest model was outperformed by the dense neural network in this study, but the accuracy this machine learning model provides could still be considered sufficient for regional lichen mapping. There were noticeable misclassifications of roads being classified as lichen from the Random Forest model, but such errors could be rectified by retraining the model with training data of similar pixels. A feature importance bar chart was generated using permutation on the Random Forest model to better understand which bands were most influential in the Random Forest prediction (Figure 17).

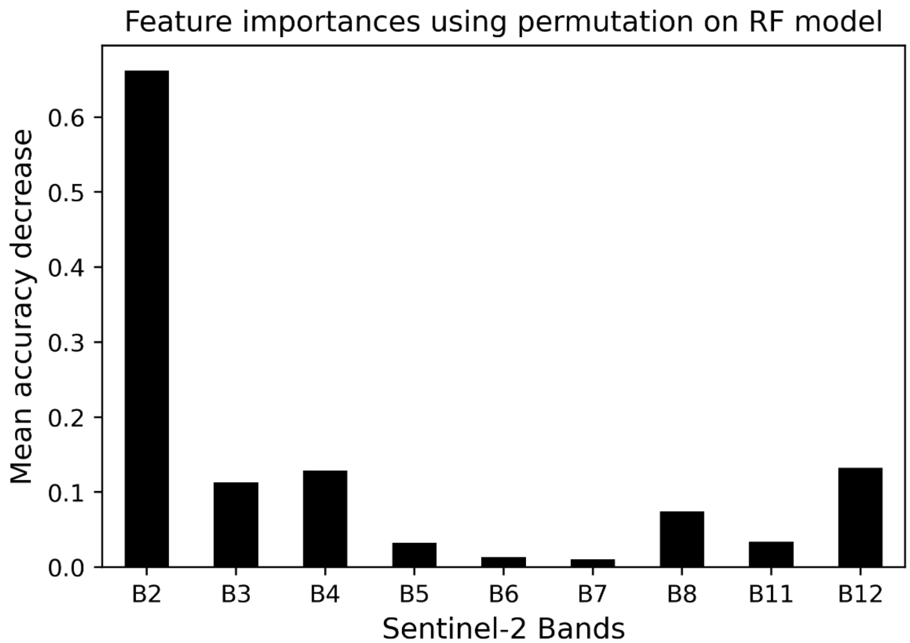

Figure 17. Random Forest feature importance plot.

Band 2 corresponds to the blue band of Sentinel-2, which has the highest feature importance in the prediction of lichen coverage (%) in the Random Forest model. This is consistent with previous research that has reported that lichens such as *Cladonia Stellaris* are separable from other land covers in remotely sensed imagery due to their high blue reflectance (Kuusinen et al., 2020; Nelson et al., 2013). *Cladonia Stellaris* also has relatively higher SWIR reflectance (Bands 11 and 12) than other vegetation, which could account for these bands having some influence on predictions (Kuusinen et al., 2020). Bands 6 and 7 seem to be the least influential in this Random Forest model.

When considering applying such techniques to another region of Canada, the dense neural network and Random Forest models are easier to implement than the CNN and, with a similar or slightly better performance, would be preferable for this application. However, CNNs might perform better than dense neural networks and Random Forest models in other earth observation regression applications with more consistent spatial and texture patterns.

4.2 Further research

Further studies should focus on using regional lichen maps to investigate changes in lichen coverage over time. While Seninel-2 regional lichen maps cannot be used for long-term analyses because the first Sentinel-2 satellite was launched in 2015, but these 10-m lichen maps can be scaled up to the 30 m Landsat resolution and then be used as training data for a lichen mapping model based on Landsat data. GEE has a large archive of Landsat imagery, which could be used to detect changes in lichen over longer time frames. This can further our understanding of the spatiotemporal changes to lichen in Canada.

This thesis provides guidance on methods for scaling-up UAV datasets to Sentinel-2 imagery. Similar applications of this workflow can be applied to other regions of Canada to generate regional lichen maps at 10-m resolution. There are many other Canadian caribou ranges that are impacted by climate change and habitat loss and are vital for healthy ecosystems. Creating regional lichen maps in other caribou ranges can assist sustainable land management across Canada.

Finally, lichen maps need to be integrated with other maps of cumulative effects to better understand threats to caribou. Cumulative effect maps that track changes in landscape due to forest management or forest fires could be used with lichen maps to better track how the caribou habitats have changed over time. Additionally, a greater understanding of other cumulative effects can further our understanding of what threats to lichen abundance exist in Canada. Contextualizing lichen mapping with other landscape changes will allow for researchers to better understand the impacts of climate change on Canadian ecosystems.

References

Al Najar, M., El Bennioui, Y., Thoumyre, G., Almar, R., Bergsma, E. W. J., Benshila, R., Delvit, J. M., & Wilson, D. G. (2022). A COMBINED COLOR AND WAVE-BASED APPROACH TO SATELLITE DERIVED BATHYMETRY USING DEEP LEARNING. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences - ISPRS Archives*, *43*(B3-2022), 9–16. https://doi.org/10.5194/isprs-archives-XLIII-B3-2022-9-2022

Allen, J. L., McMullin, R. T., Tripp, E. A., & Lendemer, J. C. (2019). Lichen conservation in North America: A review of current practices and research in Canada and the United States. *Biodiversity and Conservation*, *28*(12), 3103–3138. https://doi.org/10.1007/s10531-019-01827-3

Alsallakh, B., Kokhlikyan, N., Miglani, V., Yuan, J., & Reblitz-Richardson, O. (2020). *Mind the Pad—CNNs can Develop Blind Spots* (arXiv:2010.02178). arXiv. http://arxiv.org/abs/2010.02178

Alvarez-Taboada, F., Paredes, C., & Julián-Pelaz, J. (2017). Mapping of the invasive species Hakea sericea using Unmanned Aerial Vehicle (UAV) and worldview-2 imagery and an object-oriented approach. *Remote Sensing*, *9*(9). https://doi.org/10.3390/rs9090913

Aptoula, E., & Ariman, S. (2022). Chlorophyll-a Retrieval From Sentinel-2 Images Using Convolutional Neural Network Regression. *IEEE Geoscience and Remote Sensing Letters*, *19*, 1–5. https://doi.org/10.1109/LGRS.2021.3070437

Asplund, J., & Wardle, D. A. (2017). How lichens impact on terrestrial community and ecosystem properties: How lichens impact on communities and ecosystems. *Biological Reviews*, *92*(3), 1720–1738. https://doi.org/10.1111/brv.12305

Astola, H., Seitsonen, L., Halme, E., Molinier, M., & Lönnqvist, A. (2021). Deep Neural Networks with Transfer Learning for Forest Variable Estimation Using Sentinel-2 Imagery in Boreal Forest. *Remote Sensing*, *13*(12), 2392. https://doi.org/10.3390/rs13122392

Bäuerle, A., van Onzenoodt, C., & Ropinski, T. (2021). Net2Vis—A Visual Grammar for Automatically Generating Publication-Tailored CNN Architecture Visualizations. *IEEE Transactions on*

Visualization and Computer Graphics, *27*(6), 2980–2991. https://doi.org/10.1109/TVCG.2021.3057483

Belgiu, M., & Drăguţ, L. (2016). Random forest in remote sensing: A review of applications and future directions. *ISPRS Journal of Photogrammetry and Remote Sensing*, *114*, 24–31. https://doi.org/10.1016/j.isprsjprs.2016.01.011

Bhatnagar, S., Gill, L., & Ghosh, B. (2020). Drone image segmentation using machine and deep learning for mapping raised bog vegetation communities. *Remote Sensing*, *12*(16). https://doi.org/10.3390/RS12162602

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Campbell, M. J., Dennison, P. E., Tune, J. W., Kannenberg, S. A., Kerr, K. L., Codding, B. F., & Anderegg, W. R. L. (2020). A multi-sensor, multi-scale approach to mapping tree mortality in woodland ecosystems. *Remote Sensing of Environment*, *245*(March), 111853. https://doi.org/10.1016/j.rse.2020.111853

Casanovas, P., Black, M., Fretwell, P., & Convey, P. (2015). Mapping lichen distribution on the Antarctic Peninsula using remote sensing, lichen spectra and photographic documentation by citizen scientists. *Polar Research*, *34*(1), 25633. https://doi.org/10.3402/polar.v34.25633

Chen, Y., Qin, S., Qiao, S., Dou, Q., Che, W., Su, G., Yao, J., & Nnanwuba, U. E. (2020). Spatial Predictions of Debris Flow Susceptibility Mapping Using Convolutional Neural Networks in Jilin Province, China. *Water*, *12*(8), 2079. https://doi.org/10.3390/w12082079

Erlandsson, R., Bjerke, J. W., Finne, E. A., Myneni, R. B., Piao, S., Wang, X., Virtanen, T., Räsänen, A., Kumpula, T., Kolari, T. H. M., Tahvanainen, T., & Tømmervik, H. (2022). An artificial intelligence approach to remotely assess pale lichen biomass. *Remote Sensing of Environment*, *280*, 113201. https://doi.org/10.1016/j.rse.2022.113201

Fawagreh, K., Gaber, M. M., & Elyan, E. (2014). Random forests: From early developments to recent advancements. *Systems Science & Control Engineering*, *2*(1), 602–609. https://doi.org/10.1080/21642583.2014.956265

Fernandes, R., Prevost, C., Canisius, F., Leblanc, S. G., Maloley, M., Oakes, S., Holman, K., & Knudby, A. (2018). Monitoring snow depth change across a range of landscapes with ephemeral snowpacks using structure from motion applied to lightweight unmanned aerial vehicle videos. *The Cryosphere*, *12*(11), 3535–3550. https://doi.org/10.5194/tc-12-3535-2018

Fraser, R. H., Pouliot, D., & van der Sluijs, J. (2021). UAV and High Resolution Satellite Mapping of Forage Lichen (Cladonia spp.) in a Rocky Canadian Shield Landscape. *Canadian Journal of Remote Sensing*, *0*(0), 1–21. https://doi.org/10.1080/07038992.2021.1908118

Gunn, A. (2016). *Rangifer tarandus. The IUCN Red List of Threatened Species*. *8235*, e.T29742A22167140.

He, L., Chen, W., Leblanc, S. G., Lovitt, J., Arsenault, A., Schmelzer, I., Fraser, R. H., Latifovic, R., Sun, L., Prévost, C., White, H. P., & Pouliot, D. (2021). Integration of multi-scale remote sensing data for reindeer lichen fractional cover mapping in Eastern Canada. *Remote Sensing of Environment*, *267*(January), 112731. https://doi.org/10.1016/j.rse.2021.112731

Joly, K., Jandt, R. R., & Klein, D. R. (2009). Decrease of lichens in Arctic ecosystems: The role of wildfire, caribou, reindeer, competition and climate in north-western Alaska. *Polar Research*, *28*(3), 433–442. https://doi.org/10.1111/j.1751-8369.2009.00113.x

Jozdani, S., Chen, D., Chen, W., Leblanc, S. G., Prévost, C., Lovitt, J., He, L., & Johnson, B. A. (2021). Leveraging deep neural networks to map caribou lichen in high-resolution satellite images based on a small-scale, noisy UAV-derived map. *Remote Sensing*, *13*(14). https://doi.org/10.3390/rs13142658

Kattenborn, T., Schiefer, F., Frey, J., Feilhauer, H., Mahecha, M. D., & Dormann, C. F. (2022). Spatially autocorrelated training and validation samples inflate performance assessment of convolutional neural networks. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, *5*, 100018. https://doi.org/10.1016/j.ophoto.2022.100018

Kornblith, S., Shlens, J., & Le, Q. V. (2019). Do Better ImageNet Models Transfer Better? *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2656–2666. https://doi.org/10.1109/CVPR.2019.00277

Kuusinen, N., Juola, J., Karki, B., Stenroos, S., & Rautiainen, M. (2020). A spectral analysis of common boreal ground lichen species. *Remote Sensing of Environment*, *247*(February), 111955. https://doi.org/10.1016/j.rse.2020.111955

Leblanc, S. G. (2020). *Off-the-shelf unmanned aerial vehicles for 3D vegetation mapping* (No. 57; p. 57). https://doi.org/10.4095/326065

Lovitt, J., Richardson, G., Rajaratnam, K., Chen, W., Leblanc, S. G., He, L., Nielsen, S. E., Hillman, A., Schmelzer, I., & Arsenault, A. (2022). A New U-Net Based Convolutional Neural Network for Estimating Caribou Lichen Ground Cover from Field-Level RGB Images. *Canadian Journal of Remote Sensing*, 1–24. https://doi.org/10.1080/07038992.2022.2144179

MacAnder, M. J., Palm, E. C., Frost, G. V., Herriges, J. D., Nelson, P. R., Roland, C., Russell, K. L. M., Suitor, M. J., Bentzen, T. W., Joly, K., Goetz, S. J., & Hebblewhite, M. (2020). Lichen cover mapping for caribou ranges in interior Alaska and Yukon. *Environmental Research Letters*, *15*(5). https://doi.org/10.1088/1748-9326/ab6d38

Maggiori, E., Tarabalka, Y., Charpiat, G., & Alliez, P. (2016). Fully convolutional neural networks for remote sensing image classification. *International Geoscience and Remote Sensing Symposium (IGARSS)*, *2016-Novem*, 5071–5074. https://doi.org/10.1109/IGARSS.2016.7730322

Maxwell, A. E., Warner, T. A., & Fang, F. (2018). Implementation of machine-learning classification in remote sensing: An applied review. *International Journal of Remote Sensing*, *39*(9), 2784–2817. https://doi.org/10.1080/01431161.2018.1433343

Myers-Smith, I. H., Forbes, B. C., Wilmking, M., Hallinger, M., Lantz, T., Blok, D., Tape, K. D., MacIas-Fauria, M., Sass-Klaassen, U., Lévesque, E., Boudreau, S., Ropars, P., Hermanutz, L., Trant, A., Collier, L. S., Weijers, S., Rozema, J., Rayback, S. A., Schmidt, N. M., … Hik, D. S. (2011). Shrub expansion in tundra ecosystems: Dynamics, impacts and research priorities. *Environmental Research Letters*, *6*(4). https://doi.org/10.1088/1748-9326/6/4/045509

Myers-Smith, I. H., Kerby, J. T., Phoenix, G. K., Bjerke, J. W., Epstein, H. E., Assmann, J. J., John, C., Andreu-Hayles, L., Angers-Blondin, S., Beck, P. S. A., Berner, L. T., Bhatt, U. S., Bjorkman, A. D., Blok, D., Bryn, A., Christiansen, C. T., Cornelissen, J. H. C., Cunliffe, A. M., Elmendorf, S. C., … Wipf, S. (2020). Complexity revealed in the greening of the Arctic. *Nature Climate Change*, *10*(2), 106–117. https://doi.org/10.1038/s41558-019-0688-1

Nazari, F., & Yan, W. (2021, September 1). *Convolutional versus dense neural networks: Comparing the two neural networks' performance in predicting building operational energy use based on the building shape*. 2021 Building Simulation Conference. https://doi.org/10.26868/25222708.2021.30735

Nelson, P. R., Roland, C., Macander, M. J., & McCune, B. (2013). Detecting continuous lichen abundance for mapping winter caribou forage at landscape spatial scales. *Remote Sensing of Environment*, *137*, 43–54. https://doi.org/10.1016/j.rse.2013.05.026

Nordberg, M. L., & Allard, A. (2002). A remote sensing methodology for monitoring lichen cover. *Canadian Journal of Remote Sensing*, *28*(2), 262–274. https://doi.org/10.5589/m02-026

Olthof, I., & Pouliot, D. (2009). Recent (1986-2006) Vegetation-Specific NDVI Trends in Northern Canada from Satellite Data. *ARCTIC*, *61*(4), 381–394. https://doi.org/10.14430/arctic46

Panahi, M., Sadhasivam, N., Pourghasemi, H. R., Rezaie, F., & Lee, S. (2020). Spatial prediction of groundwater potential mapping based on convolutional neural network (CNN) and support vector regression (SVR). *Journal of Hydrology*, *588*, 125033. https://doi.org/10.1016/j.jhydrol.2020.125033

Petzold, D. E., & Goward, S. N. (1988). Reflectance Spectra of Subarctic Lichens. In *REMOTE SENSING OF ENVIRONMENT* (Vol. 24, pp. 481–492).

Ploton, P., Mortier, F., Réjou-Méchain, M., Barbier, N., Picard, N., Rossi, V., Dormann, C., Cornu, G., Viennois, G., Bayol, N., Lyapustin, A., Gourlet-Fleury, S., & Pélissier, R. (2020). Spatial validation reveals poor predictive performance of large-scale ecological mapping models. *Nature Communications*, *11*(1). https://doi.org/10.1038/s41467-020-18321-y

Richardson, G., Leblanc, S. G., Lovitt, J., Rajaratnam, K., & Chen, W. (2021). Leveraging AI to estimate caribou lichen in UAV orthomosaics from ground photo datasets. *Drones*, *5*(3). https://doi.org/10.3390/drones5030099

Salman, S., & Liu, X. (2019). *Overfitting Mechanism and Avoidance in Deep Neural Networks* (arXiv:1901.06566). arXiv. http://arxiv.org/abs/1901.06566

Shi, Q., Liu, M., Li, S., Liu, X., Wang, F., & Zhang, L. (2021). A Deeply Supervised Attention Metric-Based Network and an Open Aerial Image Dataset for Remote Sensing Change Detection. *IEEE Transactions on Geoscience and Remote Sensing*, 1–16. https://doi.org/10.1109/TGRS.2021.3085870

Stumpf, A., Michéa, D., & Malet, J.-P. (2018). Improved Co-Registration of Sentinel-2 and Landsat-8 Imagery for Earth Surface Motion Measurements. *Remote Sensing*, *10*(2), 160. https://doi.org/10.3390/rs10020160

Théau, J., Peddle, D. R., & Duguay, C. R. (2005). Mapping lichen in a caribou habitat of Northern Quebec, Canada, using an enhancement-classification method and spectral mixture analysis. *Remote Sensing of Environment*, *94*(2), 232–243. https://doi.org/10.1016/j.rse.2004.10.008

Thompson, I. D., Wiebe, P. A., Mallon, E., Rodger, A. R., Fryxell, J. M., Baker, J. A., & Reid, D. (2015). Factors influencing the seasonal diet selection by woodland caribou (rangifer tarandus tarandus) in boreal forests in ontario. *Canadian Journal of Zoology*, *93*(2), 87–98. https://doi.org/10.1139/cjz-2014-0140

Vors, L. S., & Boyce, M. S. (2009). Global declines of caribou and reindeer. *Global Change Biology*, *15*(11), 2626–2633. https://doi.org/10.1111/j.1365-2486.2009.01974.x

Wadoux, A. M. J. C., Heuvelink, G. B. M., de Bruin, S., & Brus, D. J. (2021). Spatial cross-validation is not the right way to evaluate map accuracy. *Ecological Modelling*, *457*. https://doi.org/10.1016/j.ecolmodel.2021.109692

Ying, X. (2019). An Overview of Overfitting and its Solutions. *Journal of Physics: Conference Series*, *1168*, 022022. https://doi.org/10.1088/1742-6596/1168/2/022022

You, K., Long, M., Wang, J., & Jordan, M. I. (2019). *How Does Learning Rate Decay Help Modern Neural Networks?* (arXiv:1908.01878). arXiv. http://arxiv.org/abs/1908.01878

Zhao, T., Yang, Y., Niu, H., Chen, Y., & Wang, D. (2018). *Comparing U-Net convolutional networks with fully convolutional networks in the performances of pomegranate tree canopy segmentation*. *October 2018*, 64. https://doi.org/10.1117/12.2325570

Zimmermann, S., Dubovyk, O., Oldenburg, C., Pape, R., & Löffler, J. (2018). Lichen cover mapping in southern Norway—An analysis with remote sensing and GIS. *Gis.Science - Die Zeitschrift Fur Geoinformatik*, *2*, 60–71.

Appendix

Appendix A

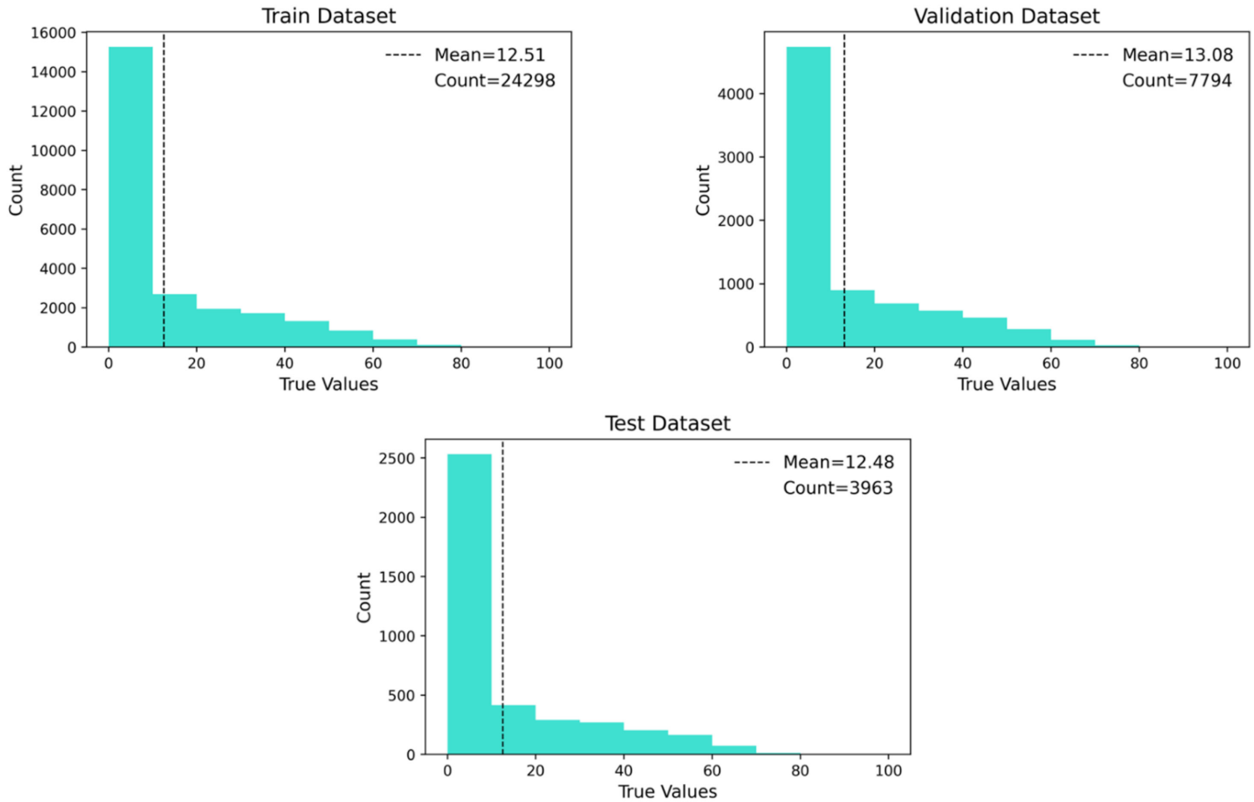

Figure 18. Histograms of the datasets used in Chapter 2.

Appendix B.

We tested rectified linear unit (ReLU, exponential linear unit, linear, tangent, and sigmoid activation functions to see how they affected the dense neural network and CNN. Each activation function was tested with 16 different models before we determined that the ReLU activation function outperformed the other activation functions when used in all layers except for the last layer. A linear activation function was selected for the last layer in order to produce a continuous real output (Aptoula & Ariman, 2022). We then tested the Adadelta, Adagrad, Adam, Adamax, Ftrl, Nadam, RMSprop, and Stochastic Gradient Descent optimizers by evaluating them on 16 different CNNs and dense neural networks. The Adam optimizer consistently outperformed other optimizers when integrated into both CNNs and dense neural networks. We assessed different learning rate decay schedulers, since learning rate decay has been found to improve the learning of complex problems while suppressing the memorization of noisy data (You et al., 2019). We tested constant learning rate, exponentially decreasing, and step decreasing learning rate schedules on over 100 CNNs and 100 dense neural networks. A model consisting of a constant learning rate perform best for dense neural networks and CNNs. While we did not train all the possible model shapes with all possible combinations of activation functions, loss functions, optimizers, and learning rate decay schedulers, we deemed this process sufficient for understanding which kind of model performs well for this application.

Appendix C.

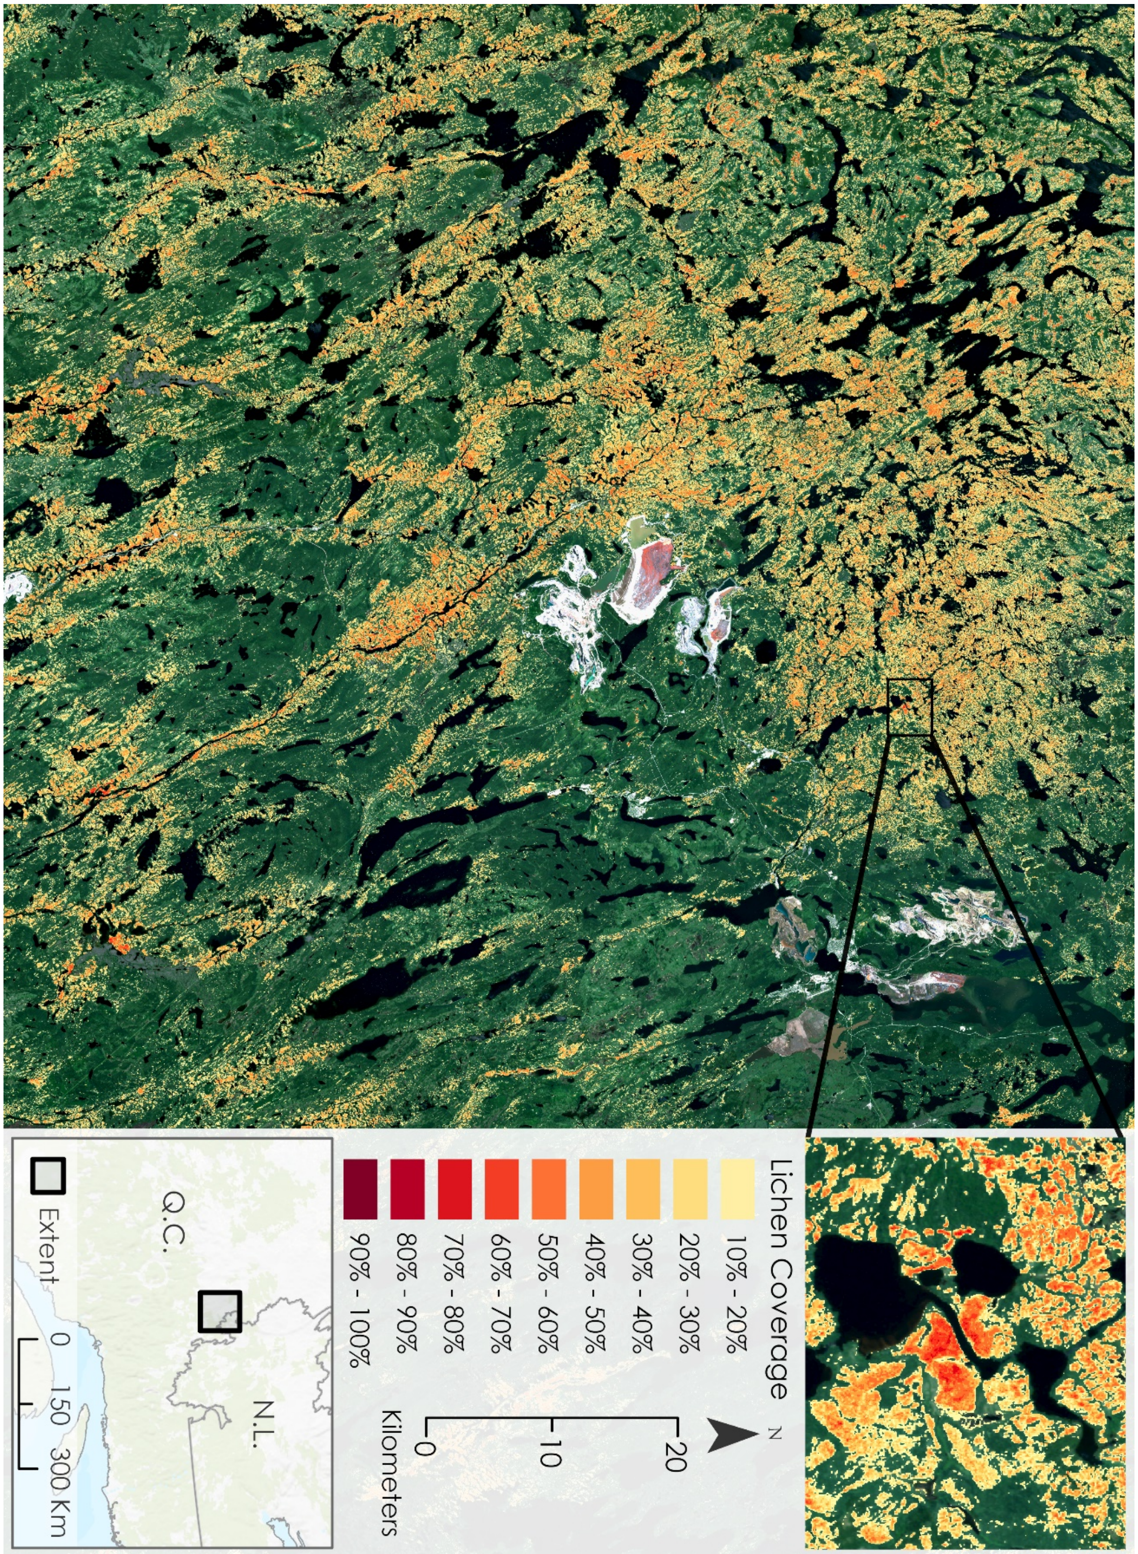

Figure 19. Lichen map of Fermont Quebec with inset of high lichen region using the retrained dense neural network model. A lichen coverage of less than 10% is transparent in this map.