remote sensing

MDPI

Article

Hair Fescue and Sheep Sorrel Identification Using Deep Learning in Wild Blueberry Production

Patrick J. Hennessy [1], Travis J. Esau [1,*], Aitazaz A. Farooque [2], Arnold W. Schumann [3], Qamar U. Zaman [1] and Kenny W. Corscadden [4]

[1] Department of Engineering, Faculty of Agriculture, Dalhousie University, Truro, NS B2N 5E3, Canada; Patrick.Hennessy@Dal.Ca (P.J.H.); qzaman@dal.ca (Q.U.Z.)
[2] School of Sustainable Design Engineering, University of Prince Edward Island, Charlottetown, PE C1A 4P3, Canada; afarooque@upei.ca
[3] Citrus Research and Education Center, University of Florida, Lake Alfred, FL 33850, USA; schumaw@ufl.edu
[4] Centre for Technology, Environment & Design, Lethbridge College, Lethbridge, AB T1K 1L6, Canada; kenny.corscadden@lethbridgecollege.ca
* Correspondence: tesau@dal.ca

Abstract: Deep learning convolutional neural networks (CNNs) are an emerging technology that provide an opportunity to increase agricultural efficiency through remote sensing and automatic inferencing of field conditions. This paper examined the novel use of CNNs to identify two weeds, hair fescue and sheep sorrel, in images of wild blueberry fields. Commercial herbicide sprayers provide a uniform application of agrochemicals to manage patches of these weeds. Three object-detection and three image-classification CNNs were trained to identify hair fescue and sheep sorrel using images from 58 wild blueberry fields. The CNNs were trained using 1280x720 images and were tested at four different internal resolutions. The CNNs were retrained with progressively smaller training datasets ranging from 3780 to 472 images to determine the effect of dataset size on accuracy. YOLOv3-Tiny was the best object-detection CNN, detecting at least one target weed per image with F_1-scores of 0.97 for hair fescue and 0.90 for sheep sorrel at 1280 × 736 resolution. Darknet Reference was the most accurate image-classification CNN, classifying images containing hair fescue and sheep sorrel with F_1-scores of 0.96 and 0.95, respectively at 1280 × 736. MobileNetV2 achieved comparable results at the lowest resolution, 864 × 480, with F_1-scores of 0.95 for both weeds. Training dataset size had minimal effect on accuracy for all CNNs except Darknet Reference. This technology can be used in a smart sprayer to control target specific spray applications, reducing herbicide use. Future work will involve testing the CNNs for use on a smart sprayer and the development of an application to provide growers with field-specific information. Using CNNs to improve agricultural efficiency will create major cost-savings for wild blueberry producers.

Keywords: convolutional neural networks; artificial intelligence; machine vision; precision agriculture; weed detection; wild blueberries; *Vaccinium angustifolium* Ait.; *Festuca filiformis* Pourr.; *Rumex Acetosella* L.

check for updates

Citation: Hennessy, P.J.; Esau, T.J.; Farooque, A.A.; Schumann, A.W.; Zaman, Q.U.; Corscadden, K.W. Hair Fescue and Sheep Sorrel Identification Using Deep Learning in Wild Blueberry Production. *Remote Sens.* **2021**, *13*, 943. https://doi.org/10.3390/rs13050943

Academic Editor: Asa Gholizadeh

Received: 20 January 2021
Accepted: 25 February 2021
Published: 3 March 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Background

Wild blueberries (*Vaccinium angustifolium* Ait.) are an economically important crop native to northeastern North America. Wild blueberry plants grow through naturally occurring rhizomes in the soils. Commercial fields are typically developed on abandoned farmland or deforested areas after the removal of trees and other vegetation [1]. In 2016, there were more than 86,000 ha of fields in production in North America, yielding approximately 119 million kg of fruit [2]. Wild blueberries contributed over $100 million to Nova Scotia's economy in 2017, including $65.9 million in exports [3]. The crop is desired for

its health benefits ranging from anti-aging and anti-inflammatory properties [4] to high antioxidant content [5] which helps reduce the risk of cardiovascular disease and cancer [6].

Infestations of hair fescue (*Festuca filiformis* Pourr.), sheep sorrel (*Rumex acetosella* L.), and over 100 other weed species [7] limit yield by competing with the wild blueberry plants for nutrients and interfering with harvesting equipment. In 2019, sheep sorrel and hair fescue were the first and fourth most common weeds in Nova Scotia wild blueberry fields, respectively [8]. Weeds are typically managed by applying a uniform application of liquid herbicide using a boom sprayer fixed to an agricultural tractor. Hexazinone was used to manage a broad array of weeds in wild blueberry fields beginning in 1982 but is no longer used for hair fescue as it has developed resistance from repeated use [9]. Hexazinone was shown to reduce instances of sheep sorrel in Nova Scotia wild blueberry fields [9,10], but was ineffective in Maine [9]. Pronamide costing more than twice that of hexazinone [11], is currently used to manage hair fescue [12]. Other options for hair fescue management include glufosinate, sulfentrazone, and terbacil [13]. A study by [14] resulted in pronamide reducing the number of sheep sorrel plants in three of four test sites, but only reducing the biomass in one of four sites. Sulfentrazone is currently being studied as a management option for sheep sorrel, with promising initial results [8].

1.2. Related Work

Smart sprayers such as See & Spray (Blue River Technology, Sunnyvale, CA, USA), GreenSeeker and WeedSeeker (Trimble Inc., Sunnyvale, CA, USA), when used in other cropping systems such as cotton, can detect and spray specific areas of fields that require application, reducing the volume of agrochemical needed. Machine vision systems relying on imaging data from red-green-blue (RGB) cameras have previously been researched for controlling smart sprayers in wild blueberry production [15–19]. A green colour segmentation algorithm was used to detect areas with weed cover, which resulted in herbicide savings of up to 78.5% [17]. However, this system did not distinguish different species of green-coloured weeds. Colour co-occurrence matrices were used to detect and spray goldenrod (*Solidago* spp.) in wild blueberry fields [18,19]. The algorithm developed by [18] was accurate, but inconvenient due to long processing times and because it had to be purpose-built for goldenrod. Reference [20] used an edge-detection approach to identify weeds in RGB images of lawns, resulting in accuracies from 67 to 83%. Light Distance and Ranging (LIDAR) sensing and hyperspectral imaging provide alternative methods to RGB cameras for identifying weeds [21,22]. A weed detection system relying on LIDAR correctly discriminated weeds from soil and maize 72.2% of the time [23]. Weed detection systems relying on hyperspectral imaging produced classification accuracies from 88 to 95% [24] and 85 to 90% [25]. Salazar-Vazquez and Mendez-Vazquez [26] developed an open-source hyperspectral camera as an alternative to cost-prohibitive commercial options; however, the imaging resolution was limited to 116 × 110 pixels. Considering the lower level of accuracy produced by LIDAR sensors and cost of high-resolution hyperspectral cameras, RBG imagining is currently a promising option for weed identification systems.

Convolutional neural networks (CNNs) are a recent processing technique which can classify entire images or objects within an image [27]. Image-classification CNNs provide an inference about an entire image, while object-detection CNNs identify objects within images and label them with bounding boxes [28]. CNNs intelligently identify visual features and find patterns associated with the target with minimal input from the user, making them easily adaptable for new targets. They are trained to detect new targets through backpropagation, which involves repeatedly showing a computer many labelled pictures of the desired target [29]. This technology has been used in wild blueberry production for detecting fruit ripeness stages and estimating potential fruit yield [30]. In other cropping systems, CNNs have been effective for detecting weeds in strawberry fields [31], potato fields [32], turfgrasses [33,34], ryegrasses [35] and Florida vegetables [36]. CNNs have also been used for detecting diseases on tomato [37,38], apple, strawberry and various other plants [38]. A web-based application for aiding citrus nutrient management

relies on a CNN [39]. Smartphone apps in other research [40–42] and agricultural cropping systems such as Pocket Agronomist (Perceptual Labs, https://www.perceptuallabs.com/ [accessed on 19 October 2020]), SenseAgro (Shenzhen SenseAgro Technology Co. Ltd., https://en.senseagro.com/ [accessed on 19 October 2020]), Agrobase (Farmis, https://farmis.com/ [accessed on 19 October 2020]), and xarvio Scouting (BASF Digital Farming GmbH, https://www.xarvio.com/ [accessed on 19 October 2020]) use CNNs to provide growers with instantaneous analysis of images from their fields. The wild blueberry industry has a free economic management tool for growers [11], but lacks intelligent software tools which can automatically analyze field information. CNNs which process images from low-cost RGB cameras with high levels of accuracy provide a promising solution for remote sensing of weeds in agricultural fields.

The aforementioned CNNs used between 1472 [31] and 40,800 [35] labelled images for training and validation. Given that there are more than 100 unique weed species in Nova Scotia wild blueberry fields, it would be best to train CNNs to identify more than just fescue and sheep sorrel. It is impractical to collect 40,800 images for weed detection in wild blueberry during the short herbicide spraying window without significant resources. Spring applications of herbicide must occur before new blueberry plant growth, and autumn applications of herbicide must occur after plant defoliation but before the soil is frozen [43]. This leaves only a few weeks during each spray timing interval for image collection.

1.3. Convolutional Neural Networks Used In This Study

This study evaluated the effectiveness of three object-detection CNNs and three image-classification CNNs for identifying hair fescue and sheep sorrel in images of wild blueberry fields. The object-detection CNNs used were YOLOv3, YOLOv3-Tiny [44], and YOLOv3-Tiny-PRN [45]. YOLOv3 and YOLOv3-Tiny are leading object-detection CNNs designed on the Darknet framework [46] which processed 416x416 images at 35 and 220 frames per second (FPS) respectively on an Nvidia Titan X graphics processing unit (GPU, Nvidia Corporation, Santa Clara, CA, USA), while achieving mean average precision (mAP) [47,48] scores of 55.3% and 31.0% on the COCO dataset (Common Objects in Context, https://cocodataset.org/ [accessed on 15 November 2020]) [44,49]. YOLOv3-Tiny-PRN added elements from Partial Residual Networks to improve the inference speed of YOLOv3-Tiny by 18.2% and maintain its accuracy on the COCO dataset [45]. YOLOv3 produced F_1-scores from 0.73 to 0.93 for single-class weed detection in Florida vegetable crops [36]. Images used in this study were captured 1.30 m from the target and scaled to 1280 × 720 resolution [36]. Hussain et al. [32] produced mAP scores of 93.2% and 78.2% with YOLOv3 and YOLOv3-Tiny, respectively for detection of lamb's quarters (*Chenopodum album* L.) in potato fields. The image-classification CNNs used in this study were Darknet Reference [46], EfficientNet-B0 [50], and MobileNetV2 [51]. Darknet Reference is an image classifier which achieved a Top-1 accuracy [48] of 61.1% on the ImageNet dataset (http://www.image-net.org/ [accessed on 15 November 2020]) and processed 224 × 224 images at 345 FPS on a Titan X GPU [46]. EfficientNet-B0 and MobileNetV2 achieved higher Top-1 scores (77.1% [50] and 74.7% [51], respectively) than Darknet Reference at the same resolution. The internal resolution of the CNNs can be increased for identifying small visual features, as was performed by [30], but large increases can result in a limiting return in accuracy improvements [50]. This technology can be used to discriminate between weed species in real-time, providing a level of control not present in previous wild blueberry smart sprayers [15–19]. Using a CNN to control herbicide spray has the potential to greatly increase application efficiency, leading to significant cost-savings for wild blueberry growers.

1.4. Outline

Section 2 of this paper will detail the creation, organization, and labelling of the image datasets used for training, validation and testing of the CNNs. Additionally, it will provide

information regarding the training procedures and performance measures used to evaluate the trained CNNs. Section 3 will present results regarding identification accuracy with the object-detection and image-classification CNNs, respectively. It will also present the results of training the CNNs with varying training dataset sizes. Section 4 will discuss the results of identifying hair fescue and sheep sorrel with the CNNs. Conclusions will be discussed in Section 5.

2. Materials and Methods

2.1. Image Datasets

Images used for this study were collected from 58 fields in northern and central Nova Scotia during the 2019 field season (Figure 1). Eight digital cameras with resolutions ranging from 4000 × 3000 to 6000 × 4000 pixels captured colour pictures of wild blueberry fields containing fescue and sheep sorrel. All pictures were captured with the camera lens pointed downward and without zoom. The photographers were instructed to capture approximately 70% of images from waist height (0.99 ± 0.09 m), 15% from knee height (0.52 ± 0.04 m) and 15% from chest height (1.35 ± 0.07 m) to allow the CNNs to recognize targets at a range of distances. A total of 4200 images with hair fescue, 2365 images with sheep sorrel, 2337 with both weeds and 2041 with neither weed were collected from 24 April to 17 May 2019 (spring), during the herbicide application timing interval. An additional 2442 images with hair fescue and 1506 images with sheep sorrel were collected from 13 November 13 to 17 December 2019 (autumn), during the herbicide application timing interval.



Figure 1. Map of image collection sites in Nova Scotia, Canada. Images were collected in 58 different fields to ensure variability of the dataset (adapted from [52]).

The images were scaled and cropped to 1280 × 720 pixels (720p) using IrfanView (v4.52, IrfanView, Wiener Neustadt, Austria) for target labeling, training, validation, and testing processes (Figure 2). The hair fescue and sheep sorrel datasets were randomly subdivided into three datasets for training, validation, and independent testing of the object-detection CNNs in the Darknet framework [46]. A script was written in the Python programming language (v3.8.2, Python Software Foundation, Wilmington, DE, USA) to randomly select a user-defined number of images for each dataset using the Mersenne Twister pseudo-random number generator [53]. For hair fescue, 3,780 images were used for training, 420 images were used for validation and 250 images were used for testing. Training and validation images were randomly selected from the spring dataset, while the testing images were randomly selected from the autumn dataset. The training,

validation, and testing datasets for sheep sorrel consisted of 960, 120 and 120 images from spring. Instances of hair fescue and sheep sorrel in the cropped images were labeled using custom-built software developed by [30]. One bounding box was created for each hair fescue tuft or sheep sorrel plant whenever possible. Weeds were densely packed in some images, making it impossible to determine where one instance ended, and another started. Hair fescue tufts would overlap one another, and it was unclear which sheep sorrel leaves belonged to a common plant. In these cases, a box was drawn around the entire region encompassing the densely packed weeds. An average of 4.90 instances of hair fescue and 14.10 instances of sheep sorrel were labelled per image.

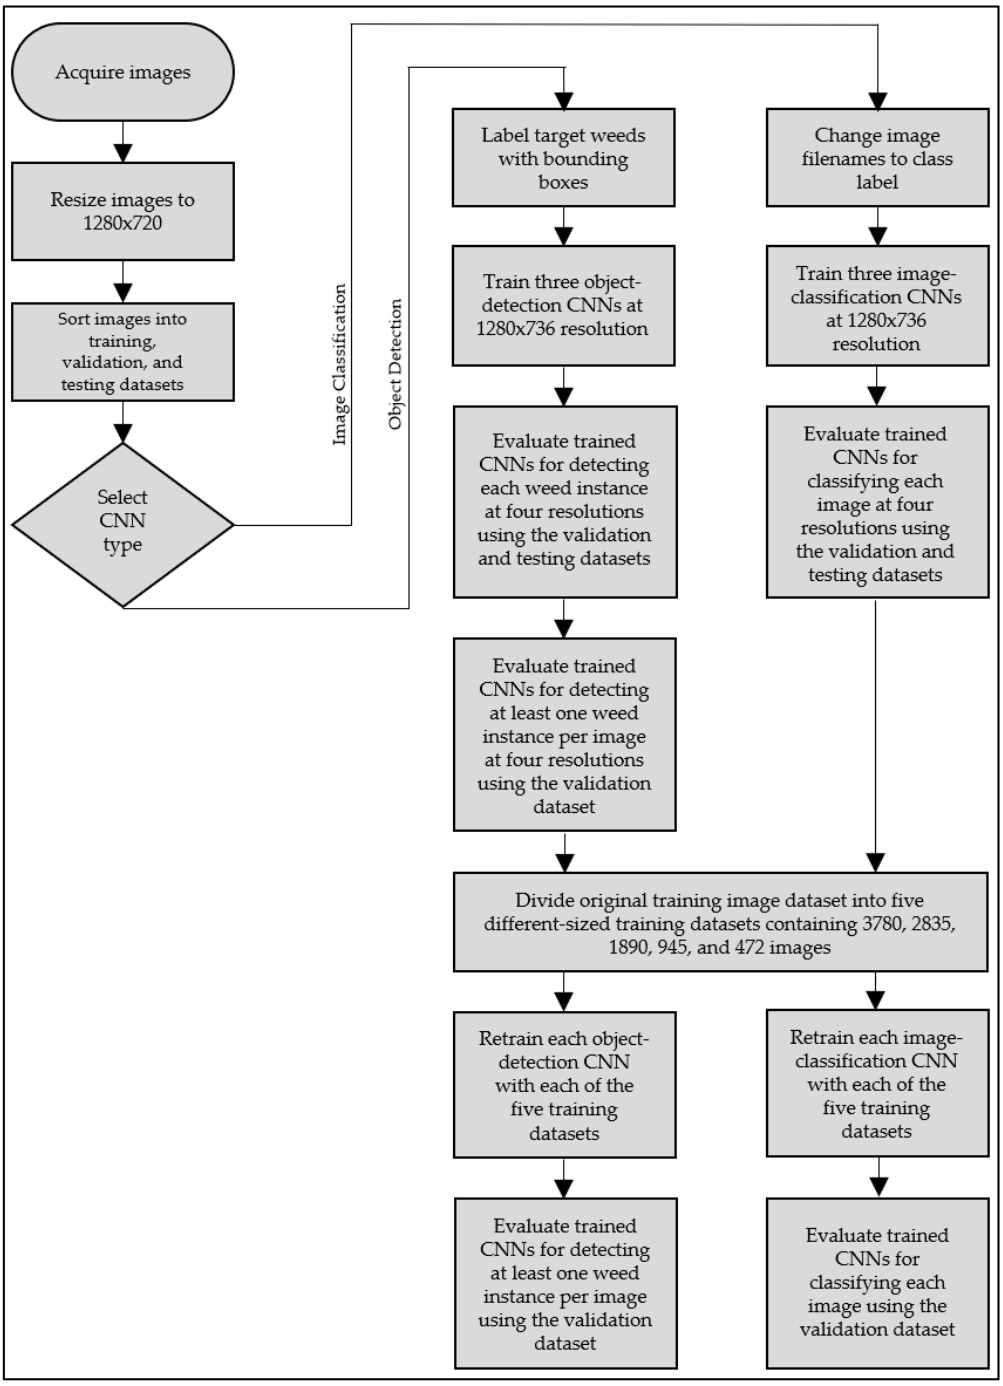

Figure 2. Flowchart of data sorting, training and evaluation procedures used in this study.

For image classification CNNs, 2000 images were randomly selected from each image set (hair fescue, sheep sorrel, both weeds, neither weed) from spring. Image-classification CNNs for hair fescue identification were trained to put images into one of two classes: "hair fescue" and "not hair fescue". The datasets containing images with hair fescue and images with both weeds were labelled as "hair fescue", while the datasets containing images with sheep sorrel and neither weed were labelled as "not hair fescue". The same methodology was used for classifying images as "sheep sorrel" or "not sheep sorrel", with the hair fescue dataset replaced by the sheep sorrel dataset and vice-versa. Randomly selected subsets containing 10% of the images were used for validation, while the remaining images were used for training. The testing datasets consisted of 700 randomly selected images from autumn.

2.2. Convolutional Neural Network Training, Validation, and Testing

The networks were trained and validated on a custom-built Ubuntu 16.04 (Canonical Ltd., London, UK) computer with an Nvidia RTX 2080 Ti GPU. The object-detection networks were trained for 5000 iterations each, with trained weights being saved every 50 iterations. The initial learning rate was set at 0.001, the default learning rate for each network, and was decreased by a factor of 10 after 2000, 3000 and 4000 iterations. A general rule-of-thumb states that object detectors require 2000 iterations per class to train [30,54], so the learning rate was decreased after 2000 iterations to avoid overfitting. The default batch size, 64, for all networks was used for training. The batch was subdivided into 16 for YOLOv3-Tiny and YOLOv3-Tiny-PRN, and 64 for YOLOv3 so the CNNs would fit within the GPU memory limits while training. All other hyperparameters were left identical to the settings used by [44,45] in each network. After training was complete, the weight file achieving the highest average precision (*AP*) score [47] was used for analysis. The image-classification networks were trained for 10,000 iterations each at the default learning rates of 0.1 for Darknet Reference and MobileNetV2, and 0.256 for EfficientNet-B0. The batch size of 64 was subdivided by 16 for Darknet Reference, 16 for MobileNetV2, and 64 for EfficientNet-B0. All other parameters for each CNN were left as defined by their respective authors. The weight file achieving the highest Top-1 score was used for analysis. The network configuration files used for EfficientNet-B0 and MobileNetV2 were converted to the Darknet framework by [55]. The Darknet framework requires image resolutions to be a multiple of 32, so the internal network resolutions were set to 1280 × 736 pixels for training and four different resolutions (1280 × 736, 1024 × 576, 960 × 544 and 864 × 480) in the network configuration file for testing.

All networks were evaluated on three additional performance metrics with the detection threshold of object-detectors set at 0.25: precision, recall and F_1-score [56]. Precision, recall and F_1-score are functions of true positive detections (*tp*), false positive detections (*fp*), and false negative detections (*fn*) of targets. Precision is ratio of true positives to all detections:

$$\mathrm{Pr} = \frac{tp}{tp + fp} \tag{1}$$

Recall is the ratio of true positives to all actual targets:

$$Re = \frac{tp}{tp + fn} \tag{2}$$

F_1-score is the harmonic mean of precision and recall:

$$F_1 = 2 \cdot \frac{\mathrm{Pr} \cdot Re}{\mathrm{Pr} + Re} \tag{3}$$

Decreasing the detection threshold increases the number of true and false positive detections, and vice-versa. This results in an increase in recall at the expense of precision. The *AP* score is calculated using the maximum precision at eleven levels of recall [47]:

$$AP = \frac{1}{11} \sum_{Re \in \{0, 0.1, \ldots, 1\}} \max_{\widetilde{Re}: \widetilde{Re} \geq Re} Pr\left(\widetilde{Re}\right) \quad (4)$$

Object-detection networks were first tested on their ability to detect all instances of hair fescue or sheep sorrel in each validation image. A detection was counted as a true positive if the Intersection-over-Union (IoU) [57], the overlap between the ground-truth box and the detection box, was greater than 50%. However, a machine vision system on a smart sprayer would require one or more detections in an image to trigger a spray event. The area sprayed due to a single detection would be wide enough to cover all target weeds in an image, even if only one is detected. A second set of tests were performed on object-detection networks to measure their ability to detect one or more instances of the target weed in a picture, regardless of how many total targets were in each image. A result was considered a true positive if the network detected at least one instance of the target weed in an image containing the desired target. The second set of tests were performed at two detection thresholds, 0.15 and 0.25, and the same four resolutions. The networks were evaluated on precision, recall, and F_1-score.

Finally, all networks for hair fescue detection were retrained with progressively smaller datasets to examine the effect that the number of training images had on detection accuracy. Five training datasets containing 3780, 2835, 1890, 945, and 472 images and the original 420-image validation dataset were used for this test. All networks were trained for the same number of iterations, at the same learning rate, and with the same batch sizes as in the initial tests. The F_1-scores of all networks at 1280 × 736 resolution were compared to determine a preliminary baseline for the number of images needed to train a CNN for weed detection.

3. Results

3.1. Hair Fescue and Sheep Sorrel Targeting with Object Detection Networks

YOLOv3, YOLOv3-Tiny, and YOLOv3-Tiny-PRN were successfully trained to detect hair fescue and sheep sorrel (Figure 3). The largest difference between validation and testing *AP* scores on ground-truth images was 2.99 percentage points with YOLOv3 at 1280 × 736 resolution, indicating that the trained networks did not overfit to the validation images (Table 1). The highest validation *AP* score for hair fescue detection (75.83%) was achieved with YOLOv3 at 1280 × 736 resolution, although the difference in *AP* score was within 1% for YOLOv3 and YOLOv3-Tiny networks with resolutions from 960 × 544 to 1280 × 736. The highest F_1-score (0.62) was achieved by YOLOv3 at 1280 × 736 resolution and YOLOv3-Tiny at 1024 × 576 resolution. YOLOv3-Tiny-PRN yielded the lowest *AP* and F_1-scores at each resolution. The precision of all networks at all resolutions was 0.90 or greater, indicating minimal false-positive detections. The recall varied from 0.34 to 0.46, indicating that more than half of the hair fescue plants were not detected at a threshold of 0.25 (Figure 4).

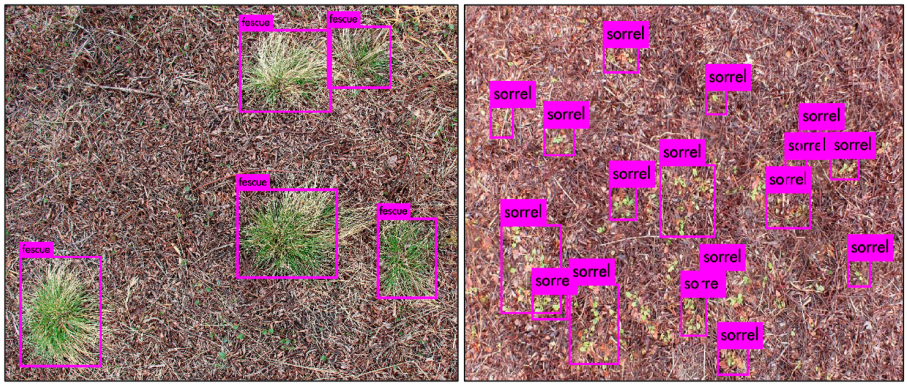

Figure 3. Examples of hair fescue (**left**) and sheep sorrel (**right**) detections on images captured in wild blueberry fields using YOLOv3-Tiny at 1280 × 736 resolution. The base hair fescue picture was captured on the morning of 24 April 2019 in Murray Siding, NS (45.3674° N, 63.2124° W), while the base sheep sorrel picture was captured on the morning of 25 April 2019 in Kemptown, NS (45.4980° N, 63.1023° W). Both pictures contain the respective target weeds with wild blueberry plants after flail mowing. The hair fescue detection picture also contains some sheep sorrel.

Table 1. Object-detection results for hair fescue plants in images of wild blueberry fields. Training and validation images were captured in spring 2019, while the testing images were captured in autumn 2019.

Network	Resolution	Validation				Testing
		Precision	Recall	F_1-Score	AP_{50} (%)	AP_{50} (%)
YOLOv3	1280 × 736	0.94	0.46	0.62	75.83	72.84
	1024 × 576	0.94	0.44	0.60	75.56	73.87
	960 × 544	0.94	0.43	0.59	75.13	72.81
	864 × 480	0.95	0.39	0.55	71.68	70.61
YOLOv3-Tiny	1280 × 736	0.91	0.45	0.60	74.83	75.18
	1024 × 576	0.92	0.46	0.62	75.46	76.57
	960 × 544	0.91	0.46	0.61	75.27	75.63
	864 × 480	0.90	0.45	0.60	71.40	74.06
YOLOv3-Tiny-PRN	1280 × 736	0.92	0.44	0.59	71.40	73.26
	1024 × 576	0.94	0.39	0.56	66.40	63.91
	960 × 544	0.94	0.38	0.54	61.38	61.04
	864 × 480	0.94	0.34	0.50	55.29	54.89

Detection accuracy for sheep sorrel was considerably lower than the accuracy for hair fescue. The peak validation *AP* score (60.81%) and F_1-score (0.51) were achieved by YOLOv3 at 1280 × 736 resolution (Table 2). The largest difference between validation and testing *AP* scores was 4.85 percentage points, indicating that the trained networks did not overfit to the validation images. The precision and recall metrics were lower for sheep sorrel than fescue for every test, with a mean change of 0.15 and 0.22, respectively. The peak recall (0.37) was achieved by YOLOv3 at 1280 × 736 resolution, with no other resolutions achieving above 0.29 for any network. YOLOv3-Tiny-PRN had lower scores than the other two networks, with the *AP* score decreasing by an average of 35.82 percentage points compared to YOLOv3 and 31.64 percentage points compared to YOLOv3-Tiny.

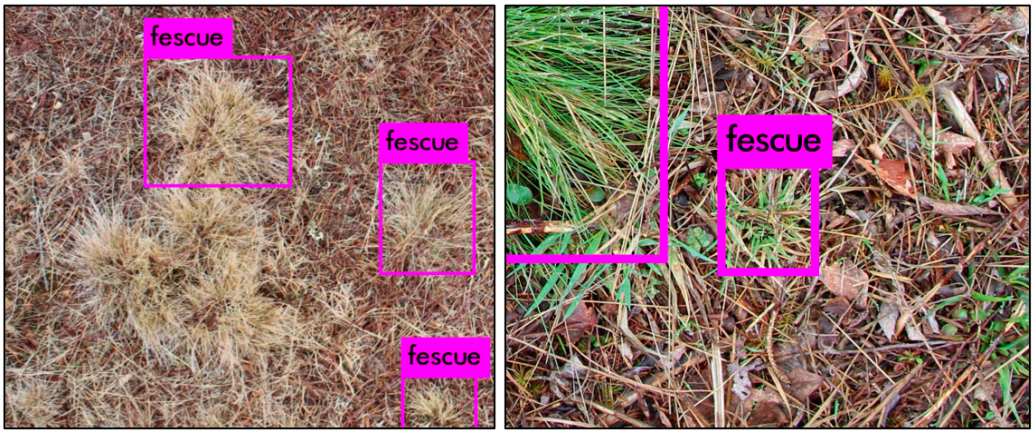

Figure 4. Examples of false negative (**left**) and false positive (**right**) hair fescue detections from spring 2019 with YOLOv3-Tiny at 1280 × 736 resolution. In the left image, several hair fescue tufts are clumped together and are not detected. In the right image, a small, bladed weed is misidentified as hair fescue. Part of a hair fescue tuft is shown directly left of the false positive detection.

Table 2. Object-detection results for sheep sorrel plants in images of wild blueberry fields. Training, validation and testing images were captured in spring 2019.

Network	Resolution	Validation				Testing
		Precision	Recall	F_1-Score	AP_{50} (%)	AP_{50} (%)
YOLOv3	1280 × 736	0.79	0.37	0.51	60.18	62.98
	1024 × 576	0.82	0.28	0.42	52.20	57.05
	960 × 544	0.83	0.26	0.40	49.15	53.04
	864 × 480	0.83	0.21	0.34	42.06	46.63
YOLOv3-Tiny	1280 × 736	0.79	0.29	0.42	54.02	57.00
	1024 × 576	0.81	0.24	0.37	48.29	52.32
	960 × 544	0.80	0.23	0.35	45.16	49.46
	864 × 480	0.78	0.19	0.31	39.40	42.91
YOLOv3-Tiny-PRN	1280 × 736	0.80	0.13	0.22	26.80	29.41
	1024 × 576	0.74	0.07	0.13	14.26	15.28
	960 × 544	0.72	0.06	0.11	11.54	12.26
	864 × 480	0.69	0.04	0.08	7.73	7.98

Given the high precision and low recall scores for detection of each weed, the networks' ability to detect at least one target weed per image was tested at a threshold of 0.15 in addition to the original threshold of 0.25. Through lowering the threshold, average recall and F_1-scores in fescue detection were improved by 0.08 and 0.03 on the validation dataset, respectively, at the expense of 0.02 from precision. In hair fescue detection, the peak F_1-score (0.98) was achieved by YOLOv3-Tiny-PRN at 1280 × 736 and 1024 × 736 and a threshold of 0.15, although no network at any resolution or threshold had an F_1-score of less than 0.90. (Table 3). Precision varied between 0.93 and 0.99, with the latter being achieved by YOLOv3-Tiny-PRN at the three lower resolutions. YOLOv3-Tiny at 864x480 produced a recall score of 1.00, and all other networks produced a recall of 0.93 or greater at a threshold of 0.15 and 0.83 or greater at a threshold of 0.25. All three object-detection networks at the tested resolutions are highly effective at detection of at least one hair fescue plant per image.

Table 3. Results for correct detection of at least one hair fescue plant per image in the validation dataset.

Network	Resolution	Threshold = 0.15			Threshold = 0.25		
		Precision	Recall	F_1-Score	Precision	Recall	F_1-Score
YOLOv3	1280 × 736	0.97	0.93	0.95	0.98	0.88	0.93
	1024 × 576	0.96	0.94	0.95	0.98	0.88	0.93
	960 × 544	0.96	0.95	0.95	0.98	0.88	0.93
	864 × 480	0.96	0.96	0.96	0.98	0.86	0.91
YOLOv3-Tiny	1280 × 736	0.97	0.97	0.97	0.97	0.91	0.94
	1024 × 576	0.96	0.97	0.96	0.97	0.93	0.95
	960 × 544	0.95	0.97	0.96	0.97	0.92	0.94
	864 × 480	0.93	1.00	0.96	0.96	0.94	0.95
YOLOv3-Tiny-PRN	1280 × 736	0.97	0.99	0.98	0.98	0.88	0.93
	1024 × 576	0.97	0.98	0.98	0.99	0.87	0.93
	960 × 544	0.97	0.96	0.97	0.99	0.86	0.92
	864 × 480	0.97	0.95	0.96	0.99	0.83	0.90

Similar to hair fescue, detection of at least one sheep sorrel plant produced greatly improved results compared to detection of all sheep sorrel plants (Table 4). Lowering the detection threshold yielded variable results on the F_1-scores. YOLOv3-Tiny-PRN improved at all four resolutions by an average of 0.06 while YOLOv3-Tiny performed better by an average of 0.01 at the original 0.25 threshold. The F_1-score for YOLOv3 showed no change at the two middle resolutions, increased by 0.01 at 864 × 480, and decreased by 0.02 at 1280 × 736. The peak F_1-score (0.91) was achieved with YOLOv3 at 1280 × 736 and a threshold of 0.25, while the minimum F_1-score (0.68) was produced with YOLOv3-Tiny-PRN at 860 × 480 and a threshold of 0.25. All three networks except YOLOv3-Tiny-PRN at the three lowest resolutions had F_1-scores of 0.87 or higher at both thresholds.

Table 4. Results for correct detection of at least one sheep sorrel plant per image in the validation dataset.

Network	Resolution	Threshold = 0.15			Threshold = 0.25		
		Precision	Recall	F_1-Score	Precision	Recall	F_1-Score
YOLOv3	1280 × 736	0.83	0.96	0.89	0.88	0.94	0.91
	1024 × 576	0.86	0.94	0.90	0.91	0.89	0.90
	960 × 544	0.87	0.91	0.89	0.91	0.87	0.89
	864 × 480	0.89	0.88	0.88	0.92	0.83	0.87
YOLOv3-Tiny	1280 × 736	0.81	0.97	0.88	0.87	0.93	0.90
	1024 × 576	0.82	0.94	0.88	0.88	0.91	0.89
	960 × 544	0.82	0.94	0.88	0.89	0.90	0.89
	864 × 480	0.82	0.93	0.87	0.89	0.87	0.88
YOLOv3-Tiny-PRN	1280 × 736	0.84	0.89	0.86	0.88	0.78	0.82
	1024 × 576	0.82	0.79	0.80	0.88	0.65	0.75
	960 × 544	0.82	0.76	0.79	0.88	0.64	0.74
	864 × 480	0.82	0.69	0.75	0.86	0.56	0.68

3.2. Classification of Images Containing Hair Fescue and Sheep Sorrel

The Darknet Reference network at 1280 × 736 resolution achieved the highest Top-1 accuracy for hair fescue classification with 96.25% on the validation dataset (Table 5). An example of a false negative detection is shown in Figure 5. All networks at all resolutions scored above 90% except EfficientNet-B0 at 864 × 480 on the validation dataset. MobileNetV2 was the most consistent across different resolutions, with only a 1.13 percentage point difference between its best (95.63%) and worst (94.50%) Top-1 scores. Recall was higher than precision for each network at every resolution except MobileNetV2 and EfficientNet-B0 at 1280 × 736, while the recall of Darknet-Reference at this resolution was only slightly (0.01) higher than the precision.

Table 5. Classification results for images with or without hair fescue.

Network	Resolution	Validation				Testing
		Precision	Recall	F_1-Score	Top-1 (%)	Top-1 (%)
Darknet Reference	1280 × 736	0.96	0.97	0.96	96.25	96.29
	1024 × 576	0.91	0.98	0.94	93.88	95.29
	960 × 544	0.93	0.98	0.95	95.25	96.14
	864 × 480	0.86	0.99	0.92	91.38	93.00
EfficientNet-B0	1280 × 736	0.96	0.95	0.96	95.50	95.71
	1024 × 576	0.90	0.96	0.93	92.75	97.29
	960 × 544	0.86	0.96	0.91	90.25	96.14
	864 × 480	0.83	0.95	0.88	87.38	95.57
MobileNetV2	1280 × 736	0.98	0.93	0.96	95.63	95.43
	1024 × 576	0.94	0.98	0.96	95.88	97.29
	960 × 544	0.95	0.97	0.96	95.75	98.14
	864 × 480	0.91	0.99	0.95	94.50	97.28



Figure 5. An image from spring 2019 with hair fescue that was misclassified as "not hair fescue" by Darknet Reference at 1280 × 736 resolution. The hair fescue tufts in this image are less prominent in this image than in Figures 3 and 4, which may have caused the classification error.

The best Top-1 validation accuracy for sheep sorrel classification (95.25%) was achieved by Darknet Reference at 1024 × 576 and 864 × 480 and by MobileNetV2 at 864 × 480 (Table 6). Darknet-Reference was the most consistent on the validation dataset, producing an F_1-score of 0.95 at all resolutions. The lowest Top-1 score and F_1-score were produced by EfficientNet-B0 at 1024 × 576 resolution. False negative misclassifications may have been the result of images where sheep sorrel did not contrast as prominently from the wild blueberry branches (Figure 6). Another observation is that EfficientNet-B0 consistently performed better on the testing dataset than the validation dataset by an average of 10.34 percentage points. Darknet Reference and MobileNetV2 performed worse on the testing dataset by 1.78 and 2.94 percentage points, respectively.

Table 6. Classification results for images with or without sheep sorrel.

Network	Resolution	Validation				Testing
		Precision	Recall	F_1-Score	Top-1 (%)	Top-1 (%)
Darknet Reference	1280 × 736	0.96	0.93	0.95	94.75	94.14
	1024 × 576	0.93	0.98	0.95	95.25	92.42
	960 × 544	0.93	0.97	0.95	95.00	92.42
	864 × 480	0.94	0.96	0.95	95.25	94.14
EfficientNet-B0	1280 × 736	0.97	0.87	0.92	92.13	96.57
	1024 × 576	0.98	0.74	0.84	86.25	98.00
	960 × 544	0.97	0.76	0.85	87.38	96.57
	864 × 480	0.97	0.63	0.76	80.88	96.86
MobileNetV2	1280 × 736	1.00	0.75	0.86	87.63	82.29
	1024 × 576	0.98	0.86	0.91	92.13	89.71
	960 × 544	0.96	0.91	0.94	93.88	92.71
	864 × 480	0.97	0.93	0.95	95.25	92.43



Figure 6. An image from spring 2019 with sheep sorrel that was misclassified as "not sheep sorrel" by Darknet Reference at 1280 × 736 resolution. Many of the sheep sorrel leaves in this image contain red hues, resulting in less contrast from the wild blueberry branches than the sheep sorrel leaves in Figure 3.

3.3. Effect of Training Dataset Size on Detection Accuracy

The CNNs were retrained for hair fescue identification with progressively smaller datasets and their F_1-scores were compared to understand the minimum requirements for training dataset size (Figure 7). For object-detection networks, the F_1-score for detecting at least one weed per image at a threshold of 0.15 was used. All networks except Darknet Reference consistently produced F_1-scores above 0.90 for all training dataset sizes from 472 to 3780 images. F_1-scores for Darknet Reference varied from 0.92 to 0.77.

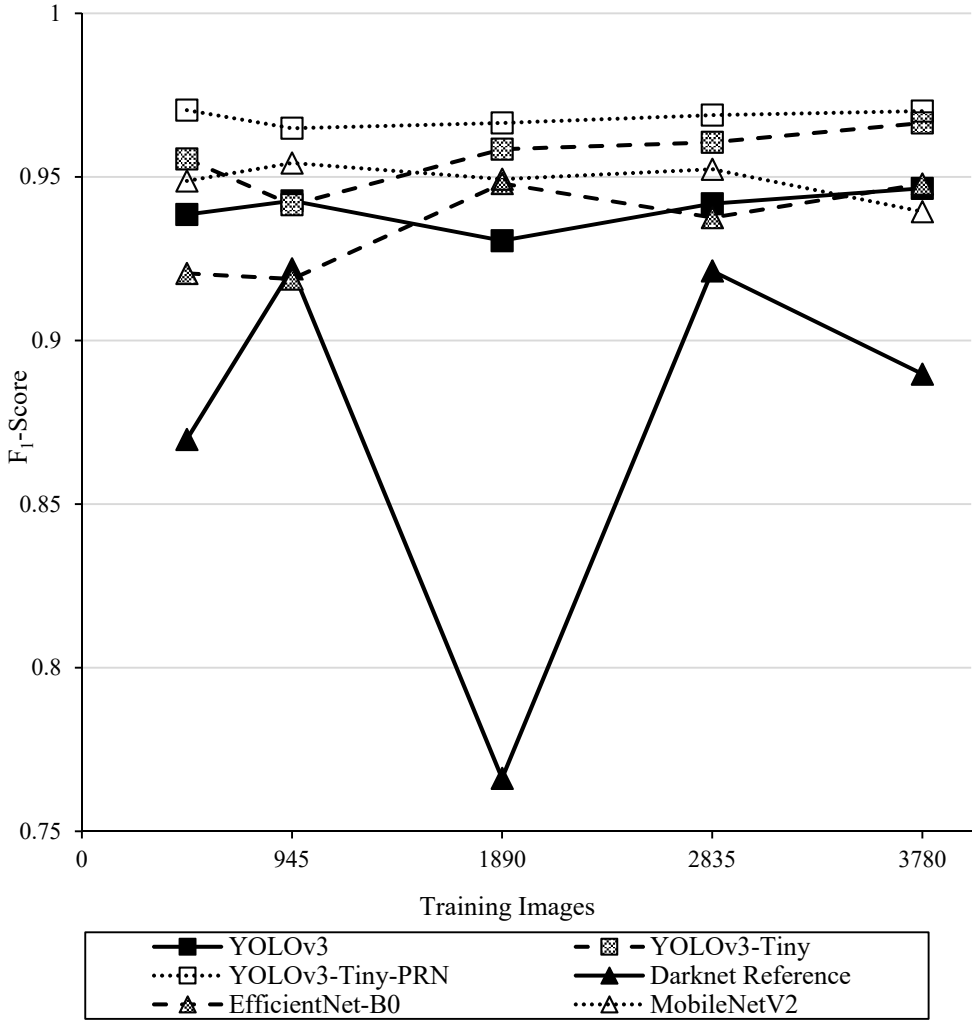

Figure 7. F_1-scores for object-detection CNNs at 1280 × 736 resolution when trained with different dataset sizes.

4. Discussion

4.1. Accuracy of Object-Detection Networks

The recall values produced by object-detection networks for hair fescue detection ranged from 0.34 to 0.36 (Table 1), indicating that many instances of hair fescue were not detected. This was lower than in other agricultural applications, including [36], which saw recall values ranging from 0.59 to 0.99 for detection of vegetation in vegetable row middles. This may be a result of the light-coloured soil in the images used by [36] creating a larger contrast between the targets and the background compared to the images used in this study. Preprocessing the images to accentuate green hues in the weeds may improve the results; however, it may also reduce processing speed. Higher resolutions were particularly important for YOLOv3 and YOLOv3-Tiny-PRN, which saw improvements in F_1-score by 0.07 and 0.09, respectively, when the resolution was increased from 864 × 480 to 1280 × 736. Resolution was less important for YOLOv3-Tiny in this test, with the F_1-score changing by a maximum of 0.02. Given that a machine vision system on a smart sprayer would require one or more detections in an image to trigger a spray event, the current level of detection accuracy may be acceptable. Alternatively, the detection threshold could be lowered to decrease the overall number of true and false positive detections, thereby increasing the recall at the expense of precision.

For sheep sorrel detection, the recall values ranged from 0.04 to 0.37 (Table 2). YOLOv3-Tiny-PRN produced particularly low results, with a peak recall of 0.13, indicating that it was not effective for detecting sheep sorrel. Higher processing resolutions were also important

in this test, with the F_1-scores produced by YOLOv3, YOLOv3-Tiny and YOLOv3-Tiny-PRN improving by 0.17, 0.11 and 0.14, respectively. The labelling strategy may have affected the sheep sorrel detection accuracy. The training bounding boxes were drawn to encompass each entire sheep sorrel plant, but this resulted in background pixels encompassing more than 50% of the bounding box area in some cases. Additionally, the network may be drawing a box around a group of sheep sorrel leaves that is inconsistent with the box drawn in the labelling software. If the IoU between the ground-truth and detected boxes is less than 50%, the result would be considered a false negative and a false positive instead of a true positive. Labelling each individual leaf instead of each plant may improve the results. Alternatively, semantic image segmentation methods like the one used by [58] could be effective as well. However, only one instance of sheep sorrel needs to be detected per image to trigger a spray application, so the current CNN may be viable for this application.

For detection of at least one hair fescue plant per image, all three object-detection networks produced an F_1-score of at least 0.90 at every resolution (Table 3). For sheep sorrel, the minimum F_1-score was 0.87 with YOLOv3 and YOLOv3-Tiny, and 0.68 for YOLOv3-Tiny-PRN (Table 4). Changing the resolution resulted in the F_1-score for hair fescue detection changing by a maximum of 0.03, indicating that most images had at least one hair fescue tuft that was detectable at the lowest resolution. Similar results are seen with YOLOv3 and YOLOv3-Tiny for detection of at least one sheep sorrel plant, with the F_1-scores changing with resolution by a maximum of 0.04 and 0.02, respectively. YOLOv3-Tiny-PRN was more sensitive to changes in resolution, with the F_1-score changing by 0.14 between 1280 $\times$ 736 and 864 $\times$ 480 resolution at the 0.25 threshold. Given the consistently high F_1-scores across all four resolutions for both weeds, YOLOv3 and YOLOv3-Tiny are more ideal CNNs for this task than YOLOv3-Tiny-PRN for detecting individual weed instances. The F_1-scores produced by YOLOv3-Tiny-PRN improved when the condition for a true positive was changed to detecting at least one weed in an image. Considering the much lower accuracy for detecting sheep sorrel than hair fescue, and the difference in results when the true-positive benchmark was changed, the lower results may be an extension of the IoU problem seen in the sheep sorrel dataset. Further experiments with the IoU threshold reduced from 50% to greater than 0%, as was done by Sharpe et al. [36], may help confirm this.

Considering the effect that resolution had on detection results, future work should involve testing the CNNs at different height and resolution combinations. Although images were captured at a variety of heights to train robust models, a limitation of this study was that the effectiveness of the CNNs at different image heights was not tracked. Due to the variable topography of wild blueberry fields, capturing images at various heights was a necessary design choice for future implementation on a smart sprayer. Future testing should be done to determine the required resolution at each height to adequately capture the visual features of the weeds for CNN processing.

4.2. Processing Speed Considerations for Sprayer Implementation

Inference speed is also an important consideration for using CNNs to control spray applications. YOLOv3, YOLOv3-Tiny, YOLOv3-Tiny-PRN process 416 $\times$ 416 images at 46, 330 and 400 FPS respectively, using a GTX 1080 Ti GPU [55]. Sprayers used in the wild blueberry industry typically range up to 36.6 m in width, meaning that multiple cameras will be required to automate spray applications. The prototype research sprayer developed by [17] used nine cameras for capturing images along the spray boom. Assuming the single-image inference time scales proportionally with the number of cameras and one camera is required per nozzle, using YOLOv3 at 416 $\times$ 416 resolution would result in an inference speed of 5.11 FPS per camera in a 9-camera system. The higher resolution images used in fescue and sheep sorrel detection will also take longer to process. Additional testing will be required to confirm, but a framerate of 5.11 FPS on a GTX 1080 Ti GPU is not ideal for scaling to real-time spray applications. Using this GPU in a sprayer would be inconvenient because it is only available for desktop computers and has a 600 W system power require-

ment. Mobile processing hardware will take longer to process images, so using YOLOv3 for real-time spray may not be feasible. YOLOv3-Tiny and YOLOv3-Tiny-PRN process images 7.2 [49] and 8.7 [45] times faster respectively than YOLOv3 [44]. Considering their superior processing speed, YOLOv3-Tiny and YOLOv3-Tiny-PRN are more feasible options than YOLOv3 for real-time spray applications.

4.3. Accuracy of Image-Classification Networks

All three image-classification networks produced F_1-scores of 0.96 at 1280 × 736 resolution on the hair fescue validation dataset (Table 5), indicating they are all effective for classifying images containing hair fescue. Precision decreased with resolution for all three networks, indicating that lower resolutions yield more false positive hair fescue classifications than false negatives. Decreasing the resolution resulting in the reduction of the F_1-scores for Darknet Reference, EfficientNet-B0 and MobileNetV2 by 0.04, 0.08 and 0.01, respectively. Given the high accuracy of these networks at lower resolutions, future work should involve further reducing the resolution to determine the minimum image size for accurate classification.

An interesting observation is that the validation Top-1 scores of MobileNetV2 on the sheep sorrel dataset increased as the resolution was decreased. Decreasing the resolution results in the finer details of the sheep sorrel plants being removed, indicating that MobileNetV2 may be classifying the images on the colour and round overall shape. The results show that this works well when the network must choose between two very different classes but MobileNetV2 may struggle if it must differentiate between sheep sorrel and another weed with similar features. Like with hair fescue, the accuracy of the CNNs at lower resolutions indicates that further testing should involve a determination of the minimum image size for accurate classification. Additionally, multiclass CNNs should be trained and tested to determine if they can distinguish between sheep sorrel and other similarly shaped weeds.

4.4. Identification Accuracy with Different Training Dataset Sizes

Interestingly, Darknet Reference was consistently on the worst performing networks in this test, despite being one of the best when trained with 7200 images. The results from the other networks indicate that datasets of approximately 500 images can be used to train all tested CNNs except Darknet Reference for hair fescue detection in wild blueberry fields. The number of images needed to effectively train these CNNs is much lower than the 1472 [31] to 40,800 [35] images used in other agricultural applications of CNNs. This test should be performed with other datasets to confirm whether this is always the case for other weeds. If the results are consistent across other targets, this can greatly decrease the amount of time spent collecting images for training.

5. Conclusions

Deep learning convolutional neural networks provide an opportunity to improve agricultural practices in the wild blueberry industry. Limitations with previous machine vision systems for real-time spraying were their inability to discriminate between weed species of the same colour or were not practical for scaling to other targets. The wild blueberry industry lacks a software tool which provides growers with information about weeds in their fields. Convolutional neural networks provide a promising solution to these problems, as they can successfully discriminate between different weeds and can be repurposed for other targets without manual feature extraction. The results of this study indicate that CNNs provide a high level of accuracy for weed identification in wild blueberry. This burgeoning technology should be developed into innovative solutions to improve wild blueberry field management practices. Processing images at a highest resolution, 1280 × 736, typically produced the highest level of accuracy. Furthermore, fewer training images may be required for training CNNs than have been used in similar studies. YOLOv3 produced comparable accuracy to YOLOv3-Tiny for detecting hair fescue and

sheep sorrel; however, the increased processing speed of YOLOv3-Tiny over YOLOv3 makes it more ideal for a real-time spraying application. YOLOv3-Tiny-PRN produced lower *AP* and F_1-scores than YOLOv3 and YOLOv3-Tiny for detection of these weeds, particularly for sheep sorrel, but showed greatly improved results for detecting at least one weed instance per image. Further testing should be done at lower IoU thresholds to determine how the labelling strategy may have affected per-weed detection accuracy with YOLOv3-Tiny-PRN. For detection of at least one target weed per image, all three networks produced F_1-scores of at least 0.95 for hair fescue at 1280 $\times$ 736 resolution after lowering the detection threshold from the default 0.25 to 0.15. The results of changing the threshold were variable for sheep sorrel detection, with only small changes in the F_1-score. Object-detection results for sheep sorrel may be improved by labelling each individual sheep sorrel leaf or by selection regions of weed cover using semantic segmentation methods. Among image-classification networks, Darknet Reference produced the best overall results with F_1-scores of 0.92 or greater at all resolutions for classification of both weeds when trained with 7200 images. MobileNetV2 and EfficientNet-B0 were more consistent than Darknet Reference with smaller dataset sizes, with F_1-scores above 0.92 with all training dataset sizes. Future work will involve testing the processing speed of all networks on a mobile GPU to determine if they can process images quickly enough for use in a real-time herbicide smart sprayer. The CNNs will also be tested on a moving sprayer, as the motion blur may limit their effectiveness. Using a CNN as part of a machine vision system in a smart sprayer could reduce herbicide use, resulting in major cost-savings for wild blueberry growers. Field scouting software relying on CNNs to detect weeds should also be developed to aid growers with management decisions. To maximize the impact of the smart sprayer and field scouting software, CNNs should be trained to detect other problematic weeds found in wild blueberry fields. Future target weeds should include bunchberry (*Cornus canadensis* L.), goldenrod (*Solidago* spp.), haircap moss (*Polytrichum commune* Hedw.) and hawkweed (*Heiracium* spp.).

Future Work: Additional CNNs will be trained at lower resolutions and for additional target weeds. The trained CNNs will be tested on a mobile GPU to compare their processing speed and one will be selected for use on a smart sprayer to selectively apply herbicides on wild blueberry fields. Additionally, the selected CNN will be used in an application accessible on smartphones to aid growers in field management decisions.

Author Contributions: Conceptualization, T.J.E. and P.J.H.; Methodology, P.J.H., A.W.S., T.J.E., and A.A.F.; Software, P.J.H. and A.W.S.; Validation, P.J.H., T.J.E., A.A.F., A.W.S., K.W.C., and Q.U.Z.; Formal Analysis, P.J.H.; Investigation, P.J.H.; Resources, T.J.E., Q.U.Z., A.A.F., and P.J.H.; Data Curation, P.J.H.; Writing—Original Draft Preparation, P.J.H.; Writing—Review & Editing, P.J.H., T.J.E., A.A.F., A.W.S., Q.U.Z., and K.W.C.; Visualization, P.J.H., T.J.E.; Supervision, T.J.E., Q.U.Z., K.W.C., A.W.S., and A.A.F.; Project Administration, T.J.E., P.J.H.; Funding Acquisition, K.W.C., T.J.E.. All authors have read and agreed to the published version of the manuscript.

Funding: Funding for this research was generously provided by the Natural Sciences and Engineering Research Council of Canada Discovery Grants Program (RGPIN-06295-2019), Doug Bragg Enterprises, Ltd., and New Brunswick Canadian Agricultural Partnership (CAP).

Acknowledgments: The authors would like to thank Scott White for helping with field selection for development of the image dataset and the wild blueberry growers for use of their fields during image collection. Also, the authors would like to acknowledge the efforts of the advanced mechanized systems and precision agriculture research teams at Dalhousie's Faculty of Agriculture.

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

References

1. Hall, I.V.; Aalders, L.E.; Nickerson, N.L.; Vander Kloet, S.P. The biological flora of Canada 1. Vaccinium angustifolium Ait., sweet lowbush blueberry. *Can. Field-Nat.* **1979**, *93*, 415–430.
2. Yarborough, D.E. Blueberry crop trends 1996–2017. In Proceedings of the WBPANS Annual General Meeting, Truro, NS, Canada, 17–18 November 2017.

3. Dinshaw, F. Frost batters Pictou County wild blueberry plants. Truro News 2018. Available online: https://www.saltwire.com/news/frost-batters-pictou-county-wild-blueberry-plants-216587/ (accessed on 1 March 2021).
4. Beattie, J.; Crozier, A.; Duthie, G. Potential health benefits of berries. *Curr. Nutr. Food Sci.* **2005**, *1*, 71–86. [CrossRef]
5. Kay, C.D.; Holub, B.J. The effect of wild blueberry (Vaccinium angustifolium) consumption on postprandial serum antioxidant status in human subjects. *Br. J. Nutr.* **2002**, *88*, 389–397. [CrossRef]
6. Lobo, V.; Patil, A.; Phatak, A.; Chandra, N. Free radicals, antioxidants and functional foods: Impact on human health. *Pharmacogn. Rev.* **2010**, *4*, 118–126. [CrossRef]
7. McCully, K.V.; Sampson, M.G.; Watson, A.K. Weed survey of Nova Scotia lowbush blueberry (Vaccinium angustifolium) fields. *Weed Sci.* **1991**, *39*, 180–185. [CrossRef]
8. White, S.N. Final weed survey update and research progress on priority weed species in wild blueberry. In Proceedings of the Wild Blueberry Producers Association of Nova Scotia Annual General Meeting, Truro, NS, Canada, 14–15 November 2019.
9. Jensen, K.I.N.; Yarborough, D.E. An overview of weed management in the wild lowbush blueberry—Past and present. *Small Fruits Rev.* **2004**, *3*, 229–255. [CrossRef]
10. Kennedy, K.J.; Boyd, N.S.; Nams, V.O. Hexazinone and fertilizer impacts on sheep sorrel (Rumex acetosella) in wild blueberry. *Weed Sci.* **2010**, *58*, 317–322. [CrossRef]
11. Esau, T.; Zaman, Q.U.; MacEachern, C.; Yiridoe, E.K.; Farooque, A.A. Economic and management tool for assessing wild blueberry production costs and financial feasibility. *Appl. Eng. Agric.* **2019**, *35*, 687–696. [CrossRef]
12. White, S.N.; Kumar, S.K. Potential role of sequential glufosinate and foramsulfuron applications for management of fescues (*Festuca spp*.) in wild blueberry. *Weed Technol.* **2017**, *31*, 100–110. [CrossRef]
13. White, S.N. Determination of Festuca filiformis seedbank characteristics, seedling emergence and herbicide susceptibility to aid management in lowbush blueberry (*Vaccinium angustifolium*). *Weed Res.* **2018**, *58*, 112–120. [CrossRef]
14. Hughes, A.; White, S.N.; Boyd, N.S.; Hildebrand, P.; Cutler, C.G. Red sorrel management and potential effect of red sorrel pollen on Botrytis cinerea spore germination and infection of lowbush blueberry (*Vaccinium angustifolium Ait.*) flowers. *Can. J. Plant. Sci.* **2016**, *96*, 590–596. [CrossRef]
15. Esau, T.; Zaman, Q.U.; Chang, Y.K.; Groulx, D.; Schumann, A.W.; Farooque, A.A. Prototype variable rate sprayer for spot-application of agrochemicals in wild blueberry. *Appl. Eng. Agric.* **2014**, *30*, 717–725. [CrossRef]
16. Esau, T.; Zaman, Q.; Groulx, D.; Corscadden, K.; Chang, Y.K.; Schumann, A.; Havard, P. Economic analysis for smart sprayer application in wild blueberry fields. *Precis. Agric.* **2016**, *17*, 753–765. [CrossRef]
17. Esau, T.; Zaman, Q.; Groulx, D.; Farooque, A.; Schumann, A.; Chang, Y. Machine vision smart sprayer for spot-application of agrochemical in wild blueberry fields. *Precis. Agric.* **2018**, *19*, 770–788. [CrossRef]
18. Rehman, T.U.; Zaman, Q.U.; Chang, Y.K.; Schumann, A.W.; Corscadden, K.W.; Esau, T. Optimising the parameters influencing performance and weed (goldenrod) identification accuracy of colour co-occurrence matrices. *Biosyst. Eng.* **2018**, *170*, 85–95. [CrossRef]
19. Rehman, T.U.; Zaman, Q.U.; Chang, Y.K.; Schumann, A.W.; Corscadden, K.W. Development and field evaluation of a machine vision based in-season weed detection system for wild blueberry. *Comput. Electron. Agric.* **2019**, *162*, 1–13. [CrossRef]
20. Parra, L.; Marin, J.; Yousfi, S.; Rincón, G.; Mauri, P.V.; Lloret, J. Edge detection for weed recognition in lawns. *Comput. Electron. Agric.* **2020**, *176*, 105684. [CrossRef]
21. Lu, B.; Dao, P.D.; Liu, J.; He, Y.; Shang, J. Recent advances of hyperspectral imaging technology and applications in agriculture. *Remote Sens.* **2020**, *12*, 2659. [CrossRef]
22. Peteinatos, G.G.; Weis, M.; Andújar, D.; Rueda Ayala, V.; Gerhards, R. Potential use of ground-based sensor technologies for weed detection. *Pest. Manag. Sci.* **2014**, *70*, 190–199. [CrossRef] [PubMed]
23. Andújar, D.; Rueda-Ayala, V.; Moreno, H.; Rosell-Polo, J.R.; Escolà, A.; Valero, C.; Gerhards, R.; Fernández-Quintanilla, C.; Dorado, J.; Griepentrog, H.W. Discriminating crop, weeds and soil surface with a terrestrial LIDAR sensor. *Sensors* **2013**, *13*, 14662–14675. [CrossRef]
24. Eddy, P.R.; Smith, A.M.; Hill, B.D.; Peddle, D.R.; Coburn, C.A.; Blackshaw, R.E. Weed and crop discrimination using hyperspectral image data and reduced bandsets. *Can. J. Remote Sens.* **2014**, *39*, 481–490. [CrossRef]
25. Liu, B.; Li, R.; Li, H.; You, G.; Yan, S.; Tong, Q. Crop/weed discrimination using a field imaging spectrometer system. *Sensors* **2019**, *19*, 5154. [CrossRef]
26. Salazar-Vazquez, J.; Mendez-Vazquez, A. A plug-and-play Hyperspectral Imaging Sensor using low-cost equipment. *HardwareX* **2020**, *7*, e00087. [CrossRef]
27. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef]
28. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*, 1st ed.; MIT Press: Cambridge, MA, USA, 2016.
29. Rumelhart, D.; Hinton, G.; Williams, R.J. Learning representations by back-propagating errors. *Nature* **1986**, *323*, 533–536. [CrossRef]
30. Schumann, A.W.; Mood, N.S.; Mungofa, P.D.K.; MacEachern, C.; Zaman, Q.U.; Esau, T. Detection of three fruit maturity stages in wild blueberry fields using deep learning artificial neural networks. In Proceedings of the 2019 ASABE Annual International Meeting, Boston, MA, USA, 7–10 July 2019.
31. Sharpe, S.M.; Schumann, A.W.; Boyd, N.S. Detection of Carolina geranium (Geranium carolinianum) growing in competition with strawberry using convolutional neural networks. *Weed Sci.* **2019**, *67*, 239–245. [CrossRef]

32. Hussain, N.; Farooque, A.A.; Schumann, A.W.; McKenzie-Gopsill, A.; Esau, T.; Abbas, F.; Acharya, B.; Zaman, Q. Design and development of a smart variable rate sprayer using deep learning. *Remote Sens.* **2020**, *12*, 4091. [CrossRef]
33. Yu, J.; Sharpe, S.M.; Schumann, A.W.; Boyd, N.S. Detection of broadleaf weeds growing in turfgrass with convolutional neural networks. *Pest. Manag. Sci.* **2019**, *75*, 2211–2218. [CrossRef] [PubMed]
34. Yu, J.; Sharpe, S.M.; Schumann, A.W.; Boyd, N.S. Deep learning for image-based weed detection in turfgrass. *Eur. J. Agron.* **2019**, *104*, 78–84. [CrossRef]
35. Yu, J.; Schumann, A.W.; Cao, Z.; Sharpe, S.M.; Boyd, N.S. Weed detection in perennial ryegrass with deep learning convolutional neural network. *Front. Plant. Sci.* **2019**, *10*, 1–9. [CrossRef]
36. Sharpe, S.M.; Schumann, A.W.; Yu, J.; Boyd, N.S. Vegetation detection and discrimination within vegetable plasticulture row-middles using a convolutional neural network. *Precis. Agric.* **2020**, *21*, 264–277. [CrossRef]
37. Fuentes, A.; Yoon, S.; Kim, S.C.; Park, D.S. A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors* **2017**, *17*, 2022. [CrossRef]
38. Venkataramanan, A.; Honakeri, D.K.P.; Agarwal, P. Plant disease detection and classification using deep neural networks. *Int. J. Comput. Sci. Eng.* **2019**, *11*, 40–46.
39. Schumann, A.; Waldo, L.; Mungofa, P.; Oswalt, C. Computer Tools for Diagnosing Citrus Leaf Symptoms (Part 2): Smartphone Apps for Expert Diagnosis of Citrus Leaf Symptoms. *EDIS* **2020**, *2020*, 1–2. [CrossRef]
40. Qi, W.; Su, H.; Yang, C.; Ferrigno, G.; De Momi, E.; Aliverti, A. A fast and robust deep convolutional neural networks for complex human activity recognition using smartphone. *Sensors* **2019**, *19*, 3731. [CrossRef]
41. Pertusa, A.; Gallego, A.J.; Bernabeu, M. MirBot: A collaborative object recognition system for smartphones using convolutional neural networks. *Neurocomputing* **2018**, *293*, 87–99. [CrossRef]
42. Maeda, H.; Sekimoto, Y.; Seto, T.; Kashiyama, T.; Omata, H. Road Damage Detection and Classification Using Deep Neural Networks with Smartphone Images. *Comput. Civ. Infrastruct. Eng.* **2018**, *33*, 1127–1141. [CrossRef]
43. Government of New Brunswick. 2020 Wild Blueberry Weed Control Selection Guide—Fact Sheet C1.6. Available online: https://www2.gnb.ca/content/dam/gnb/Departments/10/pdf/Agriculture/WildBlueberries-BleuetsSauvages/C162-E.pdf (accessed on 5 June 2020).
44. Redmon, J.; Farhadi, A. YOLOv3: An incremental improvement. *arXiv* **2018**, arXiv:1804.02767. Available online: https://arxiv.org/abs/1804.02767 (accessed on 11 March 2019).
45. Wang, C.-Y.; Liao, H.-Y.M.; Chen, P.-Y.; Hsieh, J.-W. Enriching variety of layer-wise learning information by gradient combination. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Seoul, South Korea, 27 October–2 November 2019.
46. Redmon, J. Darknet: Open Source Neural Networks in C. Available online: http://pjreddie.com/darknet/ (accessed on 8 June 2020).
47. Everingham, M.; Van Gool, L.; Williams, C.K.I.; Winn, J.; Zisserman, A. The Pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.* **2010**, *88*, 303–338. [CrossRef]
48. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going deeper with convolutions. In Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CPVR), Boston, MA, USA, 7–12 June 2015.
49. Redmon, J. YOLO: Real-Time Object Detection. Available online: https://pjreddie.com/darknet/yolo/ (accessed on 23 October 2020).
50. Tan, M.; Le, Q.V. EfficientNet: Rethinking model scaling for convolutional neural networks. In Proceedings of the International Conference on Machine Learning (ICML), Long Beach, CA, USA, 10–15 June 2019.
51. Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; Chen, L.C. MobileNetV2: Inverted residuals and linear bottlenecks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 4510–4520. [CrossRef]
52. Google LLC.; Hennessy, P. Map of Wild Blueberry Sites Used for Image Collection in Spring 2019 – Google Maps™. Available online: https://www.google.com/maps/@45.3747723,-63.691178,68373m/data=!3m1!1e3 (accessed on 23 February 2021).
53. Matsumoto, M.; Nishimura, T. Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator. *ACM Trans. Model. Comput. Simul.* **1998**, *8*, 3–30. [CrossRef]
54. Redmon, J.; Bochkovskiy, A.; Sinigardi, S. Darknet: YOLOv3—Neural Network for Object Detection. Available online: https://github.com/AlexeyAB/darknet (accessed on 20 January 2020).
55. Wang, C.-Y.; Liao, H.-Y.M.; Yeh, I.-H.; Wu, Y.-H.; Chen, P.-Y.; Hsieh, J.-W. CSPNet: A new backbone that can enhance learning capability of CNN. In Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, Online, 16–18 June 2020; pp. 390–391.
56. Sokolova, M.; Lapalme, G. A systematic analysis of performance measures for classification tasks. *Inf. Process. Manag.* **2009**, *45*, 427–437. [CrossRef]
57. Rahman, M.A.; Wang, Y. Optimizing intersection-over-union in deep neural networks for image segmentation. In Proceedings of the International Symposium on Visual Computing, Las Vegas, NV, USA, 12–14 December 2016; pp. 234–244.
58. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CPVR), Las Vegas, NV, USA, 26 June–1 July 2016.