Towards practical object detection for weed spraying in precision agriculture*

Adrian Salazar Gomez[1]   Madeleine Darbyshire[2]   Junfeng Gao[1]   Elizabeth I Sklar[1]   Simon Parsons[2]

Abstract—The evolution of smaller, faster processors and cheaper digital storage mechanisms across the last 4-5 decades has vastly increased the opportunity to integrate intelligent technologies in a wide range of practical environments to address a broad spectrum of tasks. One exciting application domain for such technologies is *precision agriculture*, where the ability to integrate on-board machine vision with data-driven actuation means that farmers can make decisions about crop care and harvesting at the level of the individual plant rather than the whole field. This makes sense both economically and environmentally. However, the key driver for this capability is fast and robust machine vision—typically driven by machine learning (ML) solutions and dependent on accurate modelling. One critical challenge is that the bulk of ML-based vision research considers only metrics that evaluate the accuracy of object detection and do not assess practical factors. This paper introduces three metrics that highlight different aspects relevant for real-world deployment of precision weeding and demonstrates their utility through experimental results.

Index Terms—precision agriculture, automated weeding, computer vision, object detection

I. INTRODUCTION

The current agricultural approach to weeding in arable crops is to spray an entire field with a selective herbicide that kills the weeds, but does not harm the crops. Such a *broadcast spraying* approach is easy to deliver—requiring only a sprayer to dispense the herbicide—but wasteful, since much of the area sprayed does not contain weeds, being either bare earth or crops (see Figure 3). *Precision agriculture* aims to use ideas from *artificial intelligence (AI)* and robotics to create agricultural solutions that can be delivered at the level of individual plants, instead of entire fields. An important application within precision agriculture is *automated weeding*, which aims to detect and target individual weeds, resulting in the precise delivery of herbicide [1] on the weeds while avoiding wastage, using a laser [2], or a mechanical tool [1]. In the work presented here, we are concerned with precise delivery of herbicide in a real-world farm setting.

Now, it is clear that any approach to automated weeding needs to identify the weeds, and there have been numerous attempts to use computer vision to do this (see Section II), many using AI methods based on *machine learning (ML)*. However, most of this work has treated automated weeding as purely a computer vision problem: datasets are collected and annotated, object identifiers are trained, and the resulting models are optimised with respect to *accuracy* and/or *mAP* for classifying crops and weeds. We argue that while these metrics are important, there are additional measures that must be considered in assessing the feasibility of ML models for use in precision spraying. Here we focus on three of these.

First, the *weed coverage rate (WCR)* is more important than conventional ML-based metrics. WCR identifies the proportion of the weeds that could be targeted by a sprayer triggered by the model. WCR depends on the accuracy of the model, but also takes into account the resolution of a spray. Second, we need to understand the *area sprayed*, that is the area covered in herbicide by the precision sprayer in order to understand the saving in herbicide compared with current practice. Third, it is important to know whether the approach is *practical*—because you can't just pile more GPUs on-board a tractor operating in an open field, where issues like compute power and energy consumption come into play. Real-world automated weeding (see Figure 1) will need to process many images very rapidly, so here we use *inference speed* as a proxy for comparing the practicality of different ML approaches.

The primary contribution of this paper is to assess the feasibility of precision spraying by comparing a selection of standard ML-based vision methods applied to weed detection using the additional metrics we have introduced (WCR, area sprayed and inference speed). Section II briefly describes the state-of-the-art in ML-based vision approaches applied to agriculture. Section III explains our methodology and Section IV details the experiments we conducted which form the basis of our comparison. Section V analyses these results and Section VI closes with conclusions.

II. RELATED WORK

Initial approaches to weed detection used machine learning algorithms with handcrafted features based on the differences in colour, shape or texture. [3] extracted local binary patterns with support vector machines for plant discrimination. This method generally requires a relatively small dataset for model development. However, it might fail to generalise under different field conditions. The deep learning-based methods for computer vision increasingly gain more popularity, offering an end-to-end weed detection solution that deals with the issue of generalisation. An adjusted YOLOv3
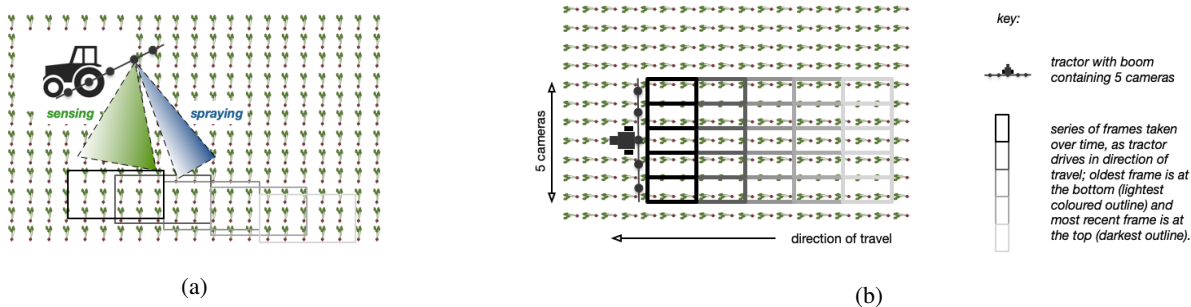
Fig. 1: A diagrammatic explanation of a small precision sprayer. (a) A boom is mounted on the back of a vehicle, carrying 5 cameras and a set of spray nozzles. Images are captured in real-time, as the vehicle moves. The system processes images, identifies weeds and their location within images, and targets spraying to hit the weeds. Sequences of images captured may provide overlapping views of the field, or there may be gaps between images. (b) A top-down view of the sprayer. The direction of travel for the is indicated in the drawing, from the bottom to the top of the illustration. Each camera captures its own image stream over time. Each row of images represents the set captured by the cameras at the same point in time; each column represents the passage of time as the tractor drives. The images here do not overlap.

model was developed for bindweed detection in sugar beet fields [4]. This work provides good detection accuracy (mAP=0.829), but it was based on the assumption of only one weed species in crop fields. [5] trained a CentreNet model only for crop detection and treated other remaining vegetation objects as weed species. In this case, the model focuses on only crop detection without consideration the specific weed species in fields. Besides, [6] developed a fully convolutional network integrating sequential information for robust weed detection in fields. The results show a better generalizability to unseen fields compared to other approaches in [7], [8]. In order to remedy the scarcity of labelled samples for training the current deep supervised networks, [9] proposed a pipeline to generate annotated weed and crop synthetic data. Similar work based on the cut-and-paste approach for synthetic image generation can be seen in [4] for weed detection. Additionally, generative adversarial networks (GANs) were exploited in weed identification with transfer learning [10].

Conventional CNNs require millions of floating point parameters and consume a lot of compute resources. Therefore, there have been many attempts to optimise networks for fast inference. An approach that has been extremely successful is quantization. This aims to retain the performance of CNNs while decreasing the precision of weights and/or activations. Binarisation, a 1-bit quantization first used in [11], is the most extreme approach. BinaryNet [12] and XNOR-Nets [13] binarize both the weights and activations to achieve a further speed increase and a further reduction in memory usage. Bi-Real Net [14] aimed to improve the accuracy of these initial approaches by retaining the real-valued outputs of 1-bit convolution using a shortcut. BiDet [15] uses a binary network implementation based on Bi-Real Net for object detection using both SSD [16] and Faster R-CNN [17]. To our knowledge, there has been no other study investigating the use of binary neural networks for identifying weeds. In [18], nine varieties of weed in the DeepWeeds dataset were classified using a binary neural network. The

authors found that most of the accuracy could be retained after downsampling the images to just 32x32x3 so running on the Intel DE1-SoC FPGA, an impressive inference time of 1.059ms could be achieved using just 6.04 watts of power. However, in this study we will focus on object detection instead of image classification and only look at how these networks perform on GPUs.

After weed detection, following work on robotic weed control system can be studied for precision weed management. Spot spraying or selective spraying, the application case in this paper, is referred as switching the individual nozzle ON/OFF to deliver chemicals only for weed species. Hussain et al. [19], [20] described a smart sprayer based on a deep learning detection model and evaluated the entire system under different weather conditions. Other than spraying, robotic mechanical precision weeding [1] is an alternative, suitable for weed management in organic farming, but typically sacrificing the operation efficiency compared with spraying approach. Our study estimates the most suitable object detector based on deep learning by exploiting multiple datasets at different locations. This detector is determined with an emphasis on the balance between detection accuracy and inference speed for robotic spraying weeding development in field environments.

III. METHODOLOGY

The overall scenario we are considering is shown in Figure 1. A boom, mounted on the back of a farm vehicle, carries both cameras and nozzles for dispensing herbicide. To be practical, images captured by the camera must be processed, and weeds detected, in time that they can be sprayed as the nozzles pass over them. Here we describe how we evaluate the effectiveness of different approaches to this detection problem.

A conventional metric to evaluate the performance of object detectors is *mean Average Precision* (*m*AP). To calculate mAP, first, a *threshold* is defined for the *IoU* value and this is used to distinguish between true positive (TP), false positive

(*FP*) and false negative (*FN*) detections. The *IoU* value, or *Intersection over Union*, measures the accuracy of the predicted *bounding box* circumscribing the object identified and is equal to $(A \cap B)/(A \cup B)$, where A is the area of the ground truth (labelled) bounding box and B is the area of the bounding box predicted by the model.

Precision (*P*) is the proportion of correctly identified objects over all objects identified, $TP/(TP+FP)$. *Recall* (*R*) is the proportion of correctly identified objects over all identified objects (correct or incorrect), $TP/(TP+FN)$. These metrics are used to compute *average precision* (*AP*):

$$AP = \sum_n (R_n - R_{n-1}).P_n$$

where n is the IoU threshold rank. Because *AP* is dependent on the IoU threshold value, mAP takes into account different threshold values and corresponding variations in the relationship between precision and recall. Typically, $n = 10$ threshold values are chosen, where the threesholds range from 0.5 to 0.95, i.e. $\{0.5, 0.55, 0.6, \ldots, 0.95\}$.

While *m*AP is a good measure of the performance of an object detector, it focuses on identifying the object very precisely. In practice, for a sprayer, the level of precision for spraying is limited by the spray nozzle, and, as we will see, it is possible for an object detector that has a relatively low *m*AP to still be good enough to be effective in ensuring that weeds are covered with herbicide.

To get a better idea of how the precision of the detector impacts the precision of spraying, we have devised a new metric which we call *weed coverage rate (WCR)*. This estimates how many of the weeds in the test data would be sprayed, rather than how many are detected. We start by deciding how many sprayer nozzles n will operate on the area within an image. In practice it is possible to spray at a resolution of around 10cm which (see below) means that we could have 3 or 4 nozzles spraying independent sections of the ground covered by one image. Each image i is then split, along it length, into n stripes. If a weed is detected in a given stripe (that is the bounding box intersects with the stripe), an area that is the height of the stripe, and stretches to either side of the bounding box, is added to the spray area S_i. This is illustrated in Figure 2.

The spray area is larger than just the bounding boxes because the size is partly determined by the width of the spray (compare Figure 2a and Figure 2c). WCR differs from *m*AP because that additional area, while increasing the herbicide used, will sometimes "hit" weeds that have evaded detection. WCR is defined as follows. A weed is counted as having been sprayed if it is wholly contained in the spray area:

$$Sprayed(A_i) = \begin{cases} 1, \text{if } A \subseteq S \\ 0, \text{otherwise} \end{cases}$$

Then WCR is the fraction of the weeds that are counted as having been sprayed:

$$WeedCoverageRate = \frac{\sum_i^n Sprayed(A_i)}{n} \times 100$$

In addition, we compare the area that has been sprayed with the total area of all the images:

$$AreaSprayed = \frac{\sum_i^n S_i}{\sum_i^m I_i} \times 100$$

Since area sprayed is a proportion of the total area that would be covered by broadcast spraying, The volume of herbicide saved is proportional to $100 - AreaSprayed$.

In order to establish the requirements in terms of frame-rate, we need to make some assumptions about the design of a sprayer. The boom on a typical sprayer is 24m, and the recommended height to operate the boom above the crop canopy is 50cm [21]. We established empirically that, at this height, a typical camera with a 1.8:1 aspect ratio (see below), can cover 550mm by 305mm. The maximum speed of a sprayer is widely taken to be around 15mph, or 6.7m/s (see, for example, [22]) to prevent spray drift. With the long edge of the image aligned along the spray boom, we need 44 cameras, and each camera can cover 305mm in the direction of travel. At 6.7m/s, we will need to process 22 frames per second (assuming no overlap), and across all 44 cameras the required frame-rate will be 968. With the short edge of the image aligned along the boom we would need 79 cameras and process 13 frames a second per camera, giving a required frame-rate of 1027.

IV. EXPERIMENTS

A. Datasets

We used two weed/crop datasets: the *Lincoln beet (LB)* dataset[1] which we collected and annotated as part of this work and the *Belgium beet (BB)* dataset from [4]. Both datasets contain pictures from fields in which sugar beet was grown commercially and the images contain pictures of beet and malicious weeds with their respective bounding box locations. The BB dataset contains 2506 images of 1800×1200 pixels. The LB dataset consists of 4405 images of 1920×1080 pixels. The LB images were extracted from videos[2] recorded at different points in time and three sugar beet fields. These data collection dates range from May to June 2021, where each field was scanned, at minimum, on four different dates with a week of separation to record weeds at different stages of growth. For all the scanning sessions, the distance from the camera to the ground was approximately 50cm. Two cameras were used; one with 12 megapixels, 26mm focal length and an f1.6 aperture and a camera with 64 megapixels, 29mm focal length and an aperture of f2.0. The original size of pictures from both cameras was 2160×3840 pixels. The fields used for the LB dataset are in Lincoln, UK, at different locations with conditions as the type of soil, distribution of the plants, and weed varieties. Fig. 3 shows examples of the BB dataset and three fields used to create the LB dataset.

---

[1]The dataset will be available at the following link: **https://github.com/LAR/lincolnbeet_dataset**.

[2]The video frames converted to images were separated with enough frame distance to avoid repeated images in the dataset.

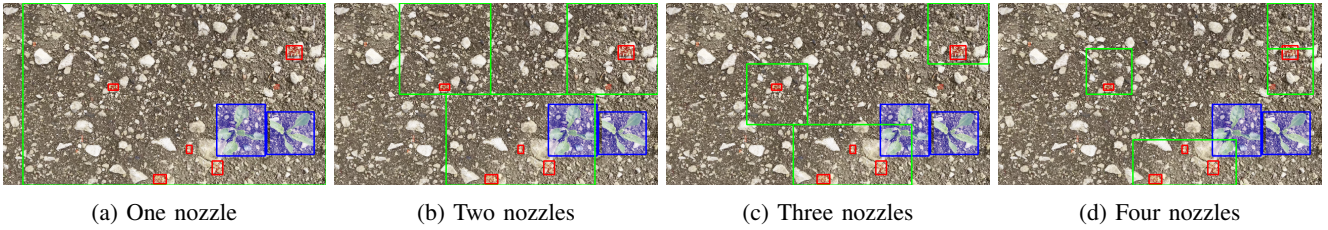| (a) One nozzle | (b) Two nozzles | (c) Three nozzles | (d) Four nozzles |

Fig. 2: WCR. The spray area (green), weeds (red) and beet (blue) for the same image and one to four nozzles. When a nozzle is triggered, it sprays an area that starts before and finishes after (in the *x* direction) the relevant weed bounding boxes and fills (in the *y* direction) the relevant fraction of the image for the number of nozzles. Note how the area sprayed decreases as the number of nozzles increases, and how the rightmost weed triggers two "sprays" for 4 nozzles.

Fig. 3: Examples of the different fields in our datasets. Clockwise from top left: BB dataset, LB location 1, LB location 2, and LB location 3

Both datasets present different item distribution and item visibility. The BB dataset has a lower number of items per picture than the LB dataset. In terms of visibility, the items in the LB dataset are proportionally smaller than the items in the BB dataset, and the items in the BB dataset have higher levels of inter-item occlusion. Table I shows the visibility and distribution characteristics of both datasets, and the characteristics of each of the items in the dataset.

TABLE I: Characteristics of the BB and LB datasets at dataset level(top) and at item level(bottom)

	BB dataset	LB dataset
Number of images	2506	4402
Number of items	5578	39246
Average items per picture	2.22	8.915
Average percentage of the bounding box that is occluded	0.02187	0.0176
Average area of the image occupied by bounding boxes	0.09883	0.0717
Average area of a bounding box compared with the image area	0.0662	0.0272
Average items per pixel	2.5199e-06	4.2795e-06

	BB dataset	LB dataset
Number of sugar beets	2654	16399
Average sugar beet plants per picture	1.059	3.725
Number of weeds	2924	22847
Average weed plants per picture	1.166	5.190
Average area occluded in sugar beet bounding boxes	0.0239	0.0324
Average area occluded in weed bounding boxes	0.0159	0.001
Average image area occupied by a sugar beet bounding box	0.0899	0.033
Average image area occupied by a weed bounding box	0.0360	0.002

B. Data preparation

For the experiments, each dataset was randomly split into training, test and validation sets with 70%, 20%, and 10% of the dataset images, respectively. To evaluate the approaches under different image resolutions, the images and the annotations were resized into three different sizes that maintain the width/height ratio; 960×540 pixels, 640×360 pixels, and 320×180 pixels.

C. Models and Parameters

For the identification of weeds, we implemented one-stage detectors, two-stage detectors and a binary one-stage detector. The one-stage models are Yolov5m [23], Yolov3 [24] and Yolov4 [25], where all three models use Darknet-53 [24] (DN-53) as a backbone. The two-stage detectors are based on Faster R-CNN [17] with three different backbones: one with a ResNet-50 backbone [26] and a Feature Pyramid Network [27] (FPN) neck; namely, ResNet-50-FPN (R-50-FPN), one with a ResNet-101-FPN [26] (ResNet-101-FPN), and another detector with a ResNeXt-101-FPN [28] (Rx-101-FPN). During training, for the one-stage and two-stage models, the optimiser was stochastic gradient descent (SGD), the learning rate was 1×10^{-2}, the momentum was 0.937, and the learning decay was 5×10^{-4}. In both model varieties, the networks were pre-trained on the COCO dataset [29]. The binary detector uses BiDet [15] and is based on the single shot detector (SSD) [16] with a Bi-Real Net [14] backbone based on VGG-16 [30]. The optimiser used for the binary model was ADAM [31]; the learning rate was 1×10^{-3}, momentum 0.9, learning decay 0 and the network was pre-trained on Imagenet [32].

Training and testing were run on a GTX2080Ti processor with 11 GB VRAM. For all detectors, the batch size,

TABLE II: Performance for object detectors with the BB (top) and LB (bottom) datasets. The fastest speeds (FPS) and highest *m*AP scores are highlighted.

Model	Backbone	FPS	*m*AP Score	*m*AP Sugar Beet	*m*AP Score Weed
yolov5m	DN-53	68.96	**87.5**	**93.3**	**81.7**
Yolov4	DN-53	46.29	63.2	74.8	51.5
Yolov3	DN-53	**75.18**	85.6	92.1	79.1
Faster R-CNN	R-50-FPN	41.75	71.4	83.3	59.5
Faster R-CNN	ResNet-101-FPN	33	72.0	83.9	60.2
Faster R-CNN	Rx-101-FPN	18.2	75.7	85.8	65.6
SSD	VGG-16 (Bi-Real)	58.01	63.9	75.6	52.2

Model	Backbone	FPS	*m*AP Score	*m*AP Sugar Beet	*m*AP Score Weed
yolov5m	DN-53	49.26	**51.0**	**67.5**	**34.6**
Yolov4	DN-53	24.75	41.2	59.5	23.0
Yolov3	DN-53	**82.64**	50.4	66.3	34.6
Faster R-CNN	R-50-FPN	40.13	42.4	62.2	22.6
Faster R-CNN	R-101-FPN	32.7	42.2	62.2	22.3
Faster R-CNN	Rx-101-FPN	25.97	43.2	62.8	23.6
SSD	VGG-16 (Bi-Real)	52.57	37.4	54	20.8

TABLE III: Speed in frames per second (FPS) for object detectors with different batch sizes for the BB dataset (top) and LB (bottom).

Model	Backbone	44 batch	22 batch	15 batch	11 batch	8 batch	4 batch
yolov5m	DN-53	277	294.11	285	285	250	250
Yolov4	DN-53	133	123	114	120	113	106
Yolov3	DN-53	153	166	166	178	166	144
Faster R-CNN	R-50-FPN	58	55	55	55	50	50
Faster R-CNN	R-101-FPN	41	45.54	50.1	45.5	45.5	38
Faster R-CNN	Rx-101-FPN	41	41	41	41	38	35
SSD	VGG-16 (Bi-Real)	62.8	66.4	70.3	67.7	63.9	61.2

Model	Backbone	44 batch	22 batch	15 batch	11 batch	8 batch	4 batch
yolov5m	DN-53	333	333	344	312	312	208
Yolov4	DN-53	40	39	40	38	39	38
Yolov3	DN-53	188	175	175	181	188	169
Faster R-CNN	R-50-FPN	44	46	49	49	49	46
Faster R-CNN	R-101-FPN	38	41	44	41	41	40
Faster R-CNN	Rx-101-FPN	35	36	36	36	35	33
SSD	VGG-16 (Bi-Real)	25.9	25.8	26.1	25.1	23.5	21.5

that is the number of images that are fed to the models simultaneously, used for training was the maximum number of images that can fit in the GPU along with the detector, and the number of epochs for training were 300 for the one-stage and two-stage detectors and 350 for the binary detector. For each model, the model weights used for testing were the ones with the highest *m*AP on the validation set during the training process. For testing and inference, the non-maximum-suppression threshold was 0.45.

D. Results

We evaluated the trained models in several ways across both the BB and LB datasets and report both traditional metrics (*m*AP) as well as our additional metrics: inference speed in *frames per second (FPS)*, weed coverage rate (WCR) and area sprayed. Analysis of results and discussion are given in the next section (V). Table II provides a conventional evaluation giving *m*AP and speed of inference with batch size= 1. The approaches with the best performance on these metrics are highlighted. Table III also provides inference speed, but exploits the ability of the models and the hardware they are run on to operate in parallel. Results are given for a range of batch sizes. Tables IV reports weed coverage rate and area sprayed for one to four nozzles. Figure 4 shows how WCR and spray area vary with the number of sprayer nozzles across the set of detectors. While weed coverage decreases with the number of nozzles, WCR decreases more slowly than area sprayed. We also plotted the WCR and area sprayed against *m*AP, Figure 5. Note that while we have results for all image sizes, space limitations mean we only show results for 640×360 images.

V. DISCUSSION

Table II shows that YoloV5m has the highest *m*AP values for both sugarbeet and weeds, on both the LB and BB datasets—a traditional ML evaluation would stop there. But recalling our practical motivation to devise a strategy for assessing the feasibility of ML methods for precision spraying applications, we need to examine the results in more depth. Table II shows that YoloV3 has the fastest inference
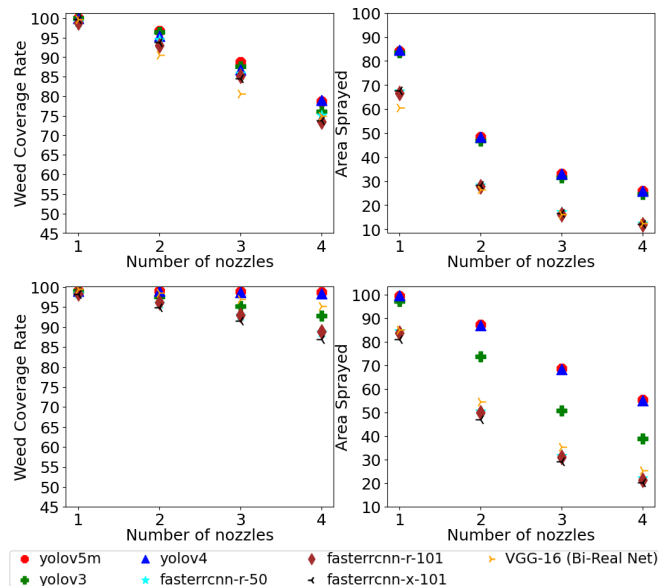

Fig. 4: The number of nozzles plotted against WCR (left) and area sprayed (right) for BB (top row) and LB (bottom).

speed. The trade-off here is to select YoloV5m and process 68.96 frames per second very accurately or select YoloV3 and process 75.18 fps, but less accurately.

Next, we look at the trade-offs that come with different batch sizes. Table II only reports results where batch size= 1, whereas table III compares results for higher batch sizes. If images from 44 cameras were processed in parallel (i.e. batchsize= 44) a frame rate of 277 or 333 (BB or LB, respectively) could be achieved with YoloV5m on a single GPU. While this is below the speed required to spray at our target pace of 15mph[3], a total of 3–4 GPUs would make in-the-field spraying with YoloV5m feasible.

Further, we investigate investigate the trade-offs when considering our WCR and area spray measures. Figure 4 and table IV contrast results when varying the number of sprayer nozzles per image from 1 to 4. For both datasets, using

[3]As per Section III, that would be a frame rate of 968 or 1027, depending on the number of cameras.

TABLE IV: WCR and area sprayed for 1–4 nozzles, for BB (top) and LB (bottom). Highlighted values explained in text.

Model	Backbone	1 nozzle		2 nozzles		3 nozzles		4 nozzles	
		Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed
yolov5m	DN-53	100	84.21	96.75	48.46	**88.82**	**33.03**	78.73	25.97
yolov4	DN-53	100	83.52	96.43	46.88	87.68	31.51	76.05	24.67
yolov3	DN-53	100	84.51	95.45	48.34	86.89	33.03	79.06	25.9
Faster R-CNN	R-50-FPN	99.03	67.31	94.52	28.34	**85.78**	**17.08**	75.08	12.51
Faster R-CNN	R-101-FPN	98.86	66.53	93.06	27.77	85.21	16.28	73.47	12.04
Faster R-CNN	Rx-101-FPN	99.68	67.61	93.87	28.18	84.57	16.67	73.79	12.11
SSD	VGG-16 (Bi-Real)	99.68	60.63	90.58	26.55	80.58	16.18	74.88	12.57

| Model	Backbone	1 nozzle		2 nozzles		3 nozzles		4 nozzles	
		Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed	Weed coverage rate	Area sprayed
yolov5m	DN-53	99.1	99.4	99	87.2	**98.9**	**68.6**	98.7	55.3
yolov4	DN-53	99.1	97.3	97.6	73.6	95.2	50.8	92.9	39
yolov3	DN-53	99.1	99.5	99	87	98.8	68.3	98.5	55.2
Faster R-CNN	R-50-FPN	98.5	84.6	95.9	50.4	92.8	31.7	88.7	22.2
Faster R-CNN	R-101-FPN	98.4	83.7	96.2	49.9	93	31.2	88.8	21.5
Faster R-CNN	Rx-101-FPN	98.2	81.1	94.9	46.9	91.6	29.2	86.9	20.2
SSD	VGG-16 (Bi-Real)	99.58	84.97	98.51	54.52	**96.94**	**35.36**	95.31	25.5
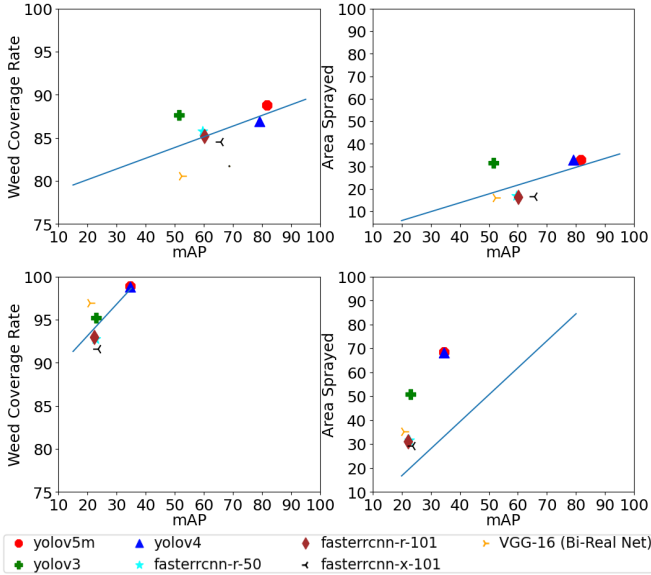
Fig. 5: *m*AP plotted against WCR (left) and area sprayed (right) with 3 nozzles for BB (top row) and LB (bottom).

positively correlated with *m*AP, these results suggest that the detector which exhibits the ideal balance between weed coverage and herbicide usage may not be the detector with the highest *m*AP.

One limitation of our results is that they were computed on non-contiguous images. This means the entire spray resolution needs to fit within one image. However, if our images were contiguous we could calculate the impact on WCR and area by allowing the spray to run over into the next image frame. Other limitations are that our model does not account for the time taken for the sprayer to turn on and off, and it assumes rectangular spray zones, whereas these are typically circular. Spray drift due to wind force and vehicle movement is also an influencing factor for precision targeting. Modelling these aspects will provide more accurate estimates of sprayer performance, and they are all elements that we will address in future work.

3 nozzles seems to produce a significant reduction in area sprayed while retaining a good weed coverage rate—which is what we want: high WCR combined with low area sprayed means that we are hitting large numbers of weeds while wasting less herbicide. While the highest WCR on the BB dataset is achieved by YoloV5m at 88.2 for 3 nozzles, Faster-RCNN (50-FPN) only exhibits a small reduction in WCR with 85.78 but only sprays 17.08% of the area compared to 33.03%. For the LB dataset with 3 nozzles, YoloV5m achieves a WCR of 98.9, but 68.6% of the area is sprayed, whereas SSD achieves WCR of 96.94 but only needs to spray 35.36% of the area.

Figure 5 illustrates the relationships between the traditional *m*AP results (from table II) and the new WCR and area sprayed metrics. While WCR and area sprayed are both

VI. CONCLUSION

This paper has demonstrated that selective spraying is feasible with state-of-the-art object detection. Our experiments show that the models considered are fast enough to detect and spray weeds on-the-fly, while our WCR metric helps to demonstrate, more clearly than conventional ML metrics, that the accuracy of our weed detection would be adequate to achieve a similar hit rate to broadcast spraying. At the same time, the area sprayed metric highlight options that produce clear reductions in area sprayed and hence herbicide required. As with many multi-criteria optimisation problems, there is no single clear winner. Our new metrics help highlight the advantages and drawbacks of different approaches, demonstrating that when it comes to practical deployment, it's not just about *m*AP.

In future, we will look to improve the frame rate achievable on a single GPU. Additionally, we will improve our WCR and area sprayed estimations by using contiguous images and take into account more properties of the spray, like area shape and sprayer response times.

REFERENCES

[1] X. Wu, S. Aravecchia, P. Lottes, C. Stachniss, and C. Pradalier, "Robotic weed control using automated weed and crop classification," *Journal of Field Robotics*, vol. 37, no. 2, pp. 322–340, 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21938

[2] S. K. Mathiassen, T. Bak, S. Christensen, and P. Kudsk, "The effect of laser treatment as a weed control method," *Biosystems Engineering*, vol. 95, no. 4, pp. 497–505, 2006. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1537511006002984

[3] V. Nguyen Thanh Le, B. Apopei, and K. Alameh, "Effective plant discrimination based on the combination of local binary pattern operators and multiclass support vector machine methods," *Information Processing in Agriculture*, vol. 6, no. 1, pp. 116–131, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214317318300258

[4] J. Gao, A. P. French, M. P. Pound, Y. He, T. P. Pridmore, and J. G. Pieters, "Deep convolutional neural networks for image-based convolvulus sepium detection in sugar beet fields," *Plant Methods*, vol. 16, no. 1, pp. 1–12, 2020.

[5] X. Jin, J. Che, and Y. Chen, "Weed identification using deep learning and image processing in vegetable plantation," *IEEE Access*, vol. 9, pp. 10 940–10 950, 2021.

[6] P. Lottes, J. Behley, A. Milioto, and C. Stachniss, "Fully convolutional networks with sequential information for robust crop and weed detection in precision farming," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, p. 2870–2877, Oct 2018. [Online]. Available: http://dx.doi.org/10.1109/LRA.2018.2846289

[7] A. Milioto, P. Lottes, and C. Stachniss, "Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in CNNs," *arXiv preprint arXiv:1709.06764*, 2018.

[8] P. Lottes and C. Stachniss, "Semi-supervised online visual crop and weed classification in precision farming exploiting plant arrangement," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 5155–5161.

[9] M. Di Cicco, C. Potena, G. Grisetti, and A. Pretto, "Automatic model based dataset generation for fast and accurate crop and weeds detection," *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep 2017. [Online]. Available: http://dx.doi.org/10.1109/IROS.2017.8206408

[10] B. Espejo-Garcia, N. Mylonas, L. Athanasakos, E. Vali, and S. Fountas, "Combining generative adversarial networks and agricultural transfer learning for weeds identification," *Biosystems Engineering*, vol. 204, pp. 79–89, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1537511021000155

[11] M. Courbariaux, Y. Bengio, and J.-P. David, "BinaryConnect: Training deep neural networks with binary weights during propagations," *arXiv preprint arXiv:1511.00363*, 2016.

[12] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1," *arXiv preprint arXiv:1602.02830*, 2016.

[13] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: ImageNet classification using binary convolutional neural networks," *arXiv preprint arXiv:1603.05279*, 2016.

[14] Z. Liu, B. Wu, W. Luo, X. Yang, W. Liu, and K.-T. Cheng, "Bi-Real Net: Enhancing the performance of 1-bit CNNs with improved representational capability and advanced training algorithm," in *Proceedings of the European Conference on Computer Vision*, September 2018.

[19] N. Hussain, A. A. Farooque, A. W. Schumann, F. Abbas, B. Acharya, A. McKenzie-Gopsill, R. Barrett, H. Afzaal, Q. U. Zaman, and M. J. Cheema, "Application of deep learning to detect Lamb's quarters (Chenopodium album L.) in potato fields of Atlantic Canada,"

[15] Z. Wang, Z. Wu, J. Lu, and J. Zhou, "BiDet: An efficient binarized object detector," *arXiv preprint arXiv:2003.03961*, 2020.

[16] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single shot multibox detector," in *European Conference on Computer Vision*. Springer, 2016, pp. 21–37.

[17] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *arXiv preprint arXiv:1506.01497*, 2015.

[18] C. Lammie, A. Olsen, T. Carrick, and M. Rahimi Azghadi, "Low-power and high-speed deep FPGA inference engines for weed classification at the edge," *IEEE Access*, vol. 7, pp. 51 171–51 184, 2019. *Computers and Electronics in Agriculture*, vol. 182, p. 106040, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169921000582

[20] N. Hussain, A. A. Farooque, A. W. Schumann, A. McKenzie-Gopsill, T. Esau, F. Abbas, B. Acharya, and Q. Zaman, "Design and development of a smart variable rate sprayer using deep learning," *Remote Sensing*, vol. 12, no. 24, 2020. [Online]. Available: https://www.mdpi.com/2072-4292/12/24/4091

[21] O. Hill, "How to properly set up a crop sprayer," Farmers Weekly, https://www.fwi.co.uk/arable/crop-management/video-how-to-properly-set-crop-sprayer, 16th February 2017, (accessed 13th September 2021).

[22] "Speed + Spraying: What growers need to know," https://www.etsprayers.com/blog/2017/03/20/speed-spraying-growers-need-know/, March 2017, (accessed 13th September 2021).

[23] G. Jocher, A. Stoken, J. Borovec, NanoCode012, A. Chaurasia, TaoXie, L. Changyu, A. V, Laughing, tkianai, yxNONG, A. Hogan, lorenzomammana, AlexWang1900, J. Hajek, L. Diaconu, Marc, Y. Kwon, oleg, wanghaoyang0106, Y. Defretin, A. Lohia, ml5ah, B. Milanko, B. Fineran, D. Khromov, D. Yiwei, Doug, Durgesh, and F. Ingham, "ultralytics/yolov5: v5.0 - YOLOv5-P6 1280 models, AWS, Supervise.ly and YouTube integrations," Apr. 2021. [Online]. Available: https://doi.org/10.5281/zenodo.4679653

[24] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[25] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," *arXiv preprint arXiv:2004.10934*, 2020.

[26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[27] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2117–2125.

[28] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1492–1500.

[29] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.

[30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[32] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.