

Using Semi-Supervised Learning to Predict Weed Density and Distribution for Precision Farming

Mr. Abhilash Sam Paulstin K.C.

Research Scholar

Department of Computer Science

AJK College of Arts and Science

Coimbatore, Tamilnadu, India.

e-mail: abhi.sam83@gmail.com

Dr. Angeline Prasanna G.

Former Associate Professor and Head

Department of Computer Science

AJK College of Arts and Science

Coimbatore, Tamilnadu, India.

e-mail: drganelineprasanna@gmail.com

Abstract— If weed growth is not controlled, it can have a devastating effect on the size and quality of a harvest. Unrestrained pesticide use for weed management can have severe consequences for ecosystem health and contribute to environmental degradation. However, if you can identify problem spots, you can more precisely treat those areas with insecticide. As a result of recent advances in the analysis of farm pictures, techniques have been developed for reliably identifying weed plants. . On the other hand, these methods mostly use supervised learning strategies, which require a huge set of pictures that have been labelled by hand. Therefore, these monitored systems are not practicable for the individual farmer because of the vast variety of plant species being cultivated. In this paper, we propose a semi-supervised deep learning method that uses a small number of colour photos taken by unmanned aerial vehicles to accurately predict the number and location of weeds in farmlands. Knowing the number and location of weeds is helpful for a site-specific weed management system in which only afflicted areas are treated by autonomous robots. In this research, the foreground vegetation pixels (including crops and weeds) are first identified using an unsupervised segmentation method based on a Convolutional Neural Network (CNN). There is then no need for manually constructed features since a trained CNN is used to pinpoint polluted locations. Carrot plants from the (1) Crop Weed Field Image Dataset (CWFID) and sugar beet plants from the (2) Sugar Beets dataset are used to test the approach. The proposed method has a maximum recall of 0.9 and an accuracy of 85%, making it ideal for locating weed hotspots. So, it is shown that the proposed strategy may be used for too many kinds of plants without having to collect a huge quantity of labelled data.

Keywords- Artificial intelligence, convolutional neural networks, machine learning, precision agriculture, semi-supervised learning.

I. INTRODUCTION

When considering the importance of industries for human existence, agriculture remains at the top of the list. There has been significant progress in farming equipment in recent years. In farming, weeding refers to the process of removing or treating undesired vegetation. Weeds can compete with your crops for resources like water, fertiliser, and natural light, so it's essential to keep weeds under control. For this reason, their deliberate removal is required to guarantee a high-quality harvest [1, 2]. But the usual practise of using agrochemicals to treat all farms the same way to kill weeds is ineffective and can hurt soil biodiversity, the quality of fresh water, and human

health. Rather than using substances to kill the weeds could try hand-weeding. Even though this way of carrying out things gets the job done, it takes a lot of time and work. By definition, "precision agriculture" is a "management approach that takes into account temporal and geographical variability to improve agricultural productivity over the long term." [3]. Precision agriculture is often used to find weeds, check the health of crops and soil, control operations like tillage, sawing, mechanical weeding, and fertiliser distribution, estimate crop output, find fruits and vegetables, and pick them[4]. It has been demonstrated that autonomous robots can be used for chemical weeding of weed plant patches [5,6]. These robots utilise

machine vision and other detection and localization methods to focus in on and eradicate undesirable plant life.

There are four primary steps make up a traditional image processing-based weed detection method: pre-processing, segmentation, feature extraction, and classification. In pre-processing, the input image is prepared for segmentation by applying various image enhancement techniques, such as a change to the colour space. Afterward, a segmentation technique is used to separate the enhanced image into a foreground and a background. The two most common kinds of this procedure are index-based segmentation and learning-based segmentation. The index-based technique differentiates between plant and background by comparing the intensity value of each pixel to a threshold setting. This approach is notoriously unreliable due to factors such as overlapping crop and weed plants and fluctuating lighting conditions [7,8]. In order to accurately identify the vegetation, it has been shown that learning-based techniques are better [2]. In segmentation, a plant mask is made, and its parts could be crop pixels or weed pixels. This is why a feature vector is made by hand using the biological shape, spectral properties, visual textures, and geographical settings of food plants and weed plants. The feature vectors are then sent into a classifier, which decides whether or not the segmented plant in question is in fact a weed. Conventional weeding methods have a number of drawbacks, including the necessity for a lot of manual labour and the expense and inconvenience of chemical spraying. Because they rely on human-created features, however, these techniques can only be employed with a limited number of plant species or invasive weeds. Deep learning-based techniques [10] have been proposed in recent years as a means of eliminating the need for such characteristics.

Most of these methods, however, are supervised, which means that they need a lot of training data. This means that they can only be used on a small group of crops and weeds. It is hard to create a strong, scalable vision system for the independent robots as a result of factors such as 1) various kinds of lighting, 2) weed and crop plants that overlap and hide each other, 3) different weed densities, and 4) different types of crops and weed plants. For the supervised learning method to work, labelled data is also very important. By looking at the sorts of weeds included in the image, a species name may be assigned to the complete picture [11]. This approach may identify weeds in the field, but it can't tell how many there are. This study comes up with a way to figure out how many weeds there are and where they are without having to label each pixel. Our study looks at a semi-supervised method for weed localization and density estimates with the goal of reducing the amount of human annotation needed to train deep networks. By using less data-intensive segmentation networks, we may be

able to speed up adoption for a wider range of crop/weed types and settings.

Our main goal is to make a semi-supervised decision-support system that can successfully predict where and how many weeds are from a single colour picture taken by a self-driving robot. Our main focus is not on pixel-level segmentation, but on the more basic question of whether or not pesticides deserve to be used specifically in specific regions. Either the weed's expected spread or its location and density can be used to figure this out. By applying the proposed approach, sources of weed damage can be found accurately. Find out how many weeds there are in the affected areas.

Because it doesn't depend on pixel-by-pixel annotations like standard end-to-end deep learning segmentation networks, it is more scalable and can be used in more places.

To determine if a picture is background or vegetation, the suggested technique uses an unsupervised Convolutional Neural Network. Neither only that, but the proposed approach identifies as background every pixel that is not ground or a plant. The vegetation mask is applied on top of the tiled parts of the input colour image. Next, the algorithm labels each tile that is covered by plants as either a weed or a crop. Unlike prior image-based approaches to weed classification, the proposed method does not rely on manually created attributes. Also, the proposed answer does not need a lot of segmentation tagging of crop and weed plant pixels, like the methods in [10], [12].

II. RELATED WORKS

Here, we'll take a quick look back at both traditional and modern approaches to weed categorization in photos using deep learning. The latest advancements in deep learning have been applied to precision agriculture, allowing farmers to avoid the pitfalls of older methods. In [13], the most recent applications of deep learning to agricultural problems such weed identification, land cover classification, and fruit counting are outlined.

A. Supervised Technique

Recent years have seen state-of-the-art outcomes for applications such as autonomous driving achieved through the application of deep learning techniques to challenging datasets. Yet, they are generic in that they may be used with many different kinds of things. Class management is minimal while doing weed identification and mapping. Several research [10] propose an end-to-end semantic segmentation network based on earlier efforts like SegNet to distinguish crop plants from weed plants. Networks are trained using 465 multispectral images, and impressive F1 scores (>0.95) are produced in another research paper. Although not a large number of training images are utilized, high-priced multispectral sensors are required for reliable results. Training networks on the same set of 10,000 RGB images as the authors of [9], [10] resulted in

comparable performance (F1 score > 0.90). These results support the feasibility of using deep learning models to distinguish between crop plants and weeds in a training environment. However, a big, manually annotated dataset is required to train a supervised learning model's network. This issue is less pressing in contexts where models can adequately generalise to numerous environments without suffering a performance hit (such as object detection for common items such as chairs, humans, etc). The authors of [9] further demonstrate the need for adaptable networks by testing the trained network on a new set of plants and achieving success. Vegetation in the background seems pretty similar in both datasets, though. Instead, our study offers an alternative to pixel-wise segmentation algorithms. Using a scatter transform to build feature vectors, the authors of [14] classify cultivated plants. When applied to domesticated vegetation, the accuracy of this method—which is taught using a made-up data set—is around 85%. Another example of a supervised learning technique from [15] is presented. Using artificial IDs for planted crops, this strategy achieves 99.7 percentile results in computer vision detection accuracy. Instead, then using photographs that have been edited in any manner or marked with physical markers, our solution relies on raw RGB images as input.

Object detection is advocated as a means of identifying weeds. A deep neural network is trained on the data, and then it generates coverage maps and bounding boxes to find the locations of plants and weeds. Since accurate findings from this approach need manual annotation of covering maps and bounding boxes, it is highly data-intensive. Making a multispectral orthomosaic map [16] involves projecting a 3D point cloud onto a 2D plane. They provide a possible answer to the challenge of scanning a large area without losing fine-grained information on plant distribution. These maps are then used as input for a modified SegNet model, which extracts the weeds from the background noise. Data-intensive and requiring sensors able to generate point clouds and the training of an end-to-end segmentation model (the study employed a dataset with more than 10,000 pictures). As shown in [17], a binary vegetation mask is produced initially by employing an end-to-end segmentation network. The generated landscape mosaic is then sent into a sophisticated VGG-16 network for labelling. Two-stage pipelines are effective, but both networks must be trained on the relevant agricultural domains. Our study applies unsupervised learning for vegetation segmentation based on the idea of a two-step method for weed identification (which is the first stage). Making tile labels is now all that's needed to begin training a classifier. These modules may be used to reduce dependency on data, and they can be easily adapted to fit other crop/weed scenarios.

B. Semi-Supervised Techniques

Semi-supervised and unsupervised methods of learning have also been explored for their potential use in weed identification. As an illustration, in [1] we see a comparison of the deep unsupervised learning algorithms JULE and DeepCluster [18], in conjunction with a deep network like VGG-16 or ResNet-50; they are used for weed categorization and automated labelling. K-Means pre-training is used to fine-tune the network weights before a LeNet-5 model is used for weed classification. These algorithms, in contrast to others, do not produce a precise map of weed or weed pixels but rather a general forecast regarding the image's categorization. They need to know the weed density in order to use variable herbicide spraying, which increases application efficiency and minimises environmental damage, but they have no method of doing so. An unsupervised technique for classifying plants is provided. They achieve competitive performance if there is no overlap between weeds and crops. When dealing with such a wide variety of plant species, that is not a reasonable assumption to make. Another challenging aspect of the unsupervised technique is determining how many clusters to use when dividing the image. The suggested approach solves this problem by relying solely on an unsupervised method to partition the vegetation mask, resulting in a maximum of two distinct groups.

The approach described in [17] is the one that most closely resembles the one that is being proposed. The authors employ a deep learning approach to weed identification. A two-stage network was implemented, with a convolutional neural network (CNN) doing the initial mask extraction to differentiate weeds from crops. On the other hand, there is a significant difference between these parts and the ones used in the proposed job. The proposed method requires significantly less data for training than the supervised learning networks utilized in [17]. (Vegetation segmentation is unsupervised while the classifier is trained with a small number of region labels). Comparing the 2000 pictures needed to train a network in [17] with the 90 and 500 images used to test the proposed method (including the upgraded versions) reveals a significant reduction in the number of images required for testing. In contrast to the networks in [17], which require pixel-wise annotations for training, the proposed pipeline does not require them. Using instances of previously undiscovered plant species and their binary labels, only the classifier has to be adjusted in the proposed study. They also point out that the overlap between plants is a common source of errors. It has been shown that the proposed approach is robust under low-light settings, occlusions, and high-plant densities, and that it is also adaptable enough to work with a wide variety of plant species. The proposed technique provides more accurate estimates of

weed density and dispersion from RGB images than previous semi-supervised algorithms.

C. Estimation of Weed Density

The number of weeds is a good sign of which places need to be treated with chemicals. There are ways to figure out how many weeds are in a row crop in [19] and [20]. The number of weeds per unit of land and the number of weeds per unit of crop are used to measure the quantity of weeds. Positional histograms, which are mentioned in [20], are used to figure out where the weeds are. By counting the number of white pixels in a binary vegetation mask down each column, we can get the horizontal pixel distribution, which is then shown as a histogram. It involves finding out the weed density (the number of weed pixels at a certain interval as a percentage of the size of the whole picture) for a set number of time intervals. This strategy only works for weeds growing in the spaces between rows, and it doesn't account for the potential of weeds and crops growing in the same space. When trying to estimate weed densities, it is also necessary to make assumptions regarding the positioning of the crop rows. As a result, it can only be employed in a limited range of agricultural contexts.

III. PROPOSED METHODOLOGY

To ensure that only severely infested regions are treated, the proposed technique seeks for and evaluates weed density. All that is needed to feed into the pipeline is a single RGB picture.

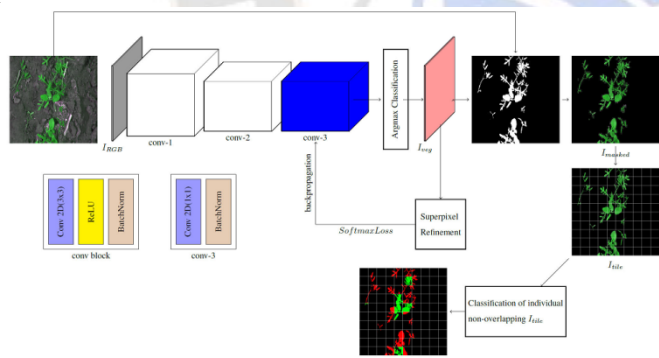


Figure 1. Proposed Method Overview

Figure 1 describes about getting a better grasp on the intended pipeline layout. Using an unsupervised deep learning-based segmentation system, each pixel in a picture is initially classified into one of two groups: vegetation or background. Two masks are made: one for the plants in the foreground and other is the background. Imasked defines the area of interest (ROI) by applying the vegetation mask, Iveg, on top of the original RGB picture. Itile then cuts this into square tiles that are smaller and smaller. Each tile has a feature vector (Itile) that describes the plants shown by the pixels on that tile. These vectors are used by a binary classifier to decide

whether Itile is a food plant or a weed. We also look at how well a trained CNN (ResNet50) can sort things on Itile. From where "weed" Itile zones are, you can tell where weeds are a problem. We can use the number of plants per square metre as a proxy for the number of weeds. By dividing the number of pixels in an area by the total number of pixels in that area, we can get an idea of how many crops and weeds are in that area. The proposed method is easily scalable and applicable to a broad variety of weed and agricultural plant kinds since just a tiny percentage of it is trained under supervision. The steps are described in further depth below.

A. Segmentation of Vegetation

First, the bicubic interpolation method from the OpenCV package is used to stretch the IRGB picture to 500x500 square pixels. The image's pixels need to be split into two different groups: the background and the centre. Here, we use the unsupervised segmentation technique with convolutional neural networks (CNNs). However, the piece stands on its own, thus just the most vital information about the work is provided. This iterative approach consists of two parts: the forward pass of the network (label prediction) and the back-propagation (learning network parameters while assuming the labels are fixed). The following restrictions are suggested by the approach for identifying the most likely cluster or class to which each pixel belongs: The first limitation is in the necessity of feature similarity. Similar pixels tend to group together in clusters. Each pixel requires its own response map, which is generated for this purpose. Each pixel is assigned to a cluster with its neighbours based on the response map. Second, there can be no gaps in the continuity between distinct sites. The authors use an image's superpixels and randomly give each of those labels the same cluster label. The term "superpixel" refers to a group of pixels that are all physically similar in some way, such as their proximity to one another or their brightness levels. The network employs the Simple Linear Iterative Clustering (SLIC) method to extract the superpixels from a five-dimensional space (three channels of CieLabcolorspace and two-dimensional image coordinates (x, y)). The number of individual clusters into which the image is divided is then constrained. If a maximum of q clusters is employed, under segmentation can be prevented even with a large number of classes. This necessitates doing intra-axis normalization on the response map prior to applying cluster labels.

It is based our methodological choice on the restrictions imposed by pixel-wise segmentation. This technique prioritizes spatially continuous pixels and allows us to set the bare minimum number of clusters at two, which is useful for identifying weeds and agricultural plants thanks to their closed-loop structures (background and vegetation). Pixel-by-pixel segmentation is refined in this way until (1) most pixels can be

separated into two groups, or (2) the maximum number of classification rounds has been reached. This limits the time it takes for the segmentation to converge, preventing either under- or over-segmentation. The cluster with the fewest pixels is utilized as a mask for the plants when the image is divided. This is so because there will be more pixels in the background than there will be plants.

To improve the performance of the unsupervised segmentation, we randomly choose 30% of the data from each dataset and use that to fine-tune the network's parameters. The ideal values for the network's parameters, including the number of superpixels, their density, and their rate of learning, are determined by a thorough examination of the data. By adjusting one setting at a time, we can determine which is optimal. Furthermore, during this time, all other variables are ignored. Mean intersection over union (mIOU) values are selected as the optimal choices for each parameter. The following experimental parameters were determined: The parameters are as follows: (1) learning rate = 0.1, (2) number of superpixels = 2500, and (3) superpixel compactness = 25. Photos in the experimental group have had vegetation masks applied to them using the ideal parameters (I_{veg}). To compare the effectiveness of the supervised and unsupervised methods, we additionally trained U-Net on the training subset of the datasets. U-net efficacy has been shown to be a supervised learning solution for pixel-wise segmentation in several uses, such as medical image segmentation and autonomous driving.

Using an encoder-decoder architecture, the network first downsamples the picture to get the prediction, and then upsamples the image to get the real data. At each step in the "upsampling" process, the feature map from the most recent "downsampling" is added to the original. Even though traits are lost when downsampling, the network may still be able to learn something from them. The network was taught to put binary class labels on the white pixels in the centre to show that there were plants there.

B. Tile Classification

The masked image I_{masked} is the result of applying the vegetation mask I_{veg} on the input picture I_{RGB} . This masked image consists of just the RGB pixels for the vegetation, ensuring that only plant-related attributes are used in the classification process (crops and weeds). The masked image (I_{masked}) is then divided into even smaller pieces (called I_{tile}) of 50x50 pixels. Quite a few places could have few or no greenery pixels. So, in I_{tiles} , if the percentage of land covered by plants (measured in terms of the number of vegetation pixels) is less than 10%, weeds are not considered to be an issue (in pixels). Figure 2 depicts image I_{masked} , region rejection owing to a lack of vegetation pixels, and area selection for classifier training.




Figure 2. I_{masked} is divided into smaller tiles (I_{tile})

Various machine learning techniques, including SVMs, RFs, and MLPs, have been developed to help with this problem. A few examples of well-known machine learning methods have been used to classification issues. In this study, we evaluate the classifiers' capacity to accurately categorize I_{tile} as either a weed or a crop. In this part, we first go through feature vector-based classifiers for determining if I_{tile} is a weed or a crop. Furthermore, we discuss an alternate image-based classifier for I_{tile} , one that use a trained convolutional neural network rather than calculating the feature vector directly.

C. Weed Density Estimation

After identifying weed-infested regions (I_{tiles} with the weed attribute), weed density may be determined by measuring the total area covered by plants in those locations. Here, a cluster rate (CR) from is used to quantify and estimate weed density ([19]). Estimates of weed densities are crucial for site-specific weed management [2]. It's possible that this population density estimate will help when choosing where to spray herbicides in the field. This choice would be influenced by factors such as the sorts of crops and weeds to be grown and the distance between plants.

$CR = \text{Weed plant coverage in the region (in pixels)} / \text{Total land area of the region (in pixels)}$

Algorithm: Weed Distribution and Density Estimation

Input: Color image (I_{RGB}) of the field acquired from an autonomous robot;

Output: Weed density and distribution;

Given (I_{RGB}), Generate the vegetation mask (I_{veg}) using CNN based unsupervised segmentation;

Overlay I_{RGB} with I_{veg} to get I_{masked} ;

Divide the image I_{masked} into smaller regions $I_{tile}(\text{squaretiles})$;

For (I_{tile} in I_{masked})do

Classify I_{tile} into crop, weed or background;

If I_{tile} is weed then

Estimate weed density

end

end

IV. RESULTS AND DISCUSSION

The evaluations are carried out through qualitative analysis and quantitative analysis.

A. Evaluation by Qualitative Analysis

Figure 3 displays the segmentation outcomes for some sample instances in both datasets. The unsupervised segmentation network outperforms the supervised segmentation network in distinguishing the vegetative pixel data from the background, according to the inferred performance (U-Net). Figure 3 shows how the unsupervised segmentation method may be used to identify and distinguish between various plant structures (3rd row of the image). It's important to find U-Net can classify vegetation even whether it's represented by a single pixel or a very sparse collection of them. This is not a typical behaviour for the unsupervised approach. That's because our approach gives greater weight to the spatial continuity of the vegetation clusters, while U-Net focuses more on the neighbourhood of a single pixel (down sampling using max pooling). This pattern was considerably more apparent in the Sugar Beets dataset (which has less contrast than the CWFID).




Figure 3(a). Crop/Weed Field Image Dataset




Figure 3(b). Sugar Beet Dataset

B. Evaluation by Quantitative Analysis

The test splits of both datasets are used to get the mean intersection over union (mIoU) number, which is then used to compare the performance of the two networks. Table 1 shows how things turned out. The unsupervised network did much better than U-Net on the Sugar Beets dataset, and it also did better than U-Net by a small amount on the CWFID dataset.

This might be because, unlike U-Net, the unsupervised method doesn't depend on learning mapping from a limited set of features to tell the difference between pixels in the foreground and pixels in the background. Both supervised and uncontrolled algorithms do better on the CWFID dataset than on the Sugar Beets dataset. This is because the pictures in the Sugar Beets collection have bad lighting and, as a result, not much contrast. The results of this study support the idea that an unsupervised network can be used to pull vegetation pixels from shots of different plants in different places.

TABLE I. QUANTITATIVE EVALUATION OF VEGETATION

Model	Dataset	mIoU
Unsupervised Segmentation	CWFID (Crop/Weed Field Image Dataset)	0.9
Unsupervised Segmentation	Sugar Beet Dataset	0.8
UNet	CWFID	0.9
UNet	Sugar Beet Dataset	0.7

C. Estimation of Weed Density

Once the weedy spots have been found, the cluster rate per tile can be found by using the weeded vegetation pixels. In Table 2, the rate of clusters that was seen in the weedy areas is compared to the rate that was expected. These results show that for both types of data, it is possible to make accurate estimates of the number of weeds. There are four main reasons for the loss of weed density pixels: 1) ignoring tiles or areas where plants cover less than 10% of the total area, 2) incorrect vegetation segmentation, 3) mislabeling weed-infested areas as crop plants, and 4) plants in a given tile that overlap. Possible mistakes could start with the lack of plants in some places. For the goals of this study, a 10% threshold was chosen because it can be changed enough to fit a wide range of crop plants and weed plants. When used on the CWFID and Sugar Beets datasets, the proposed method gets a mean absolute error of 5% for vegetation segmentation and 1% for weed spread. This backs up the idea that the proposed method might be a good way to fix mistakes that come from the above sources. Since the RMSE between datasets for two different crop/weed species is less than 8%, it is clear that the proposed method can be used for any crop/weed species. Before choosing where to carefully apply agrochemicals, it is important to find out where and how many weeds there are.

TABLE II. WEED DENSITY ESTIMATION ACCURACY

Dataset	Mean Accuracy (%)	MAE	RMSE
CFWID	75	5	7.5
SugarBeets	85	1	3

V. CONCLUSION

Once the weedy spots have been found, the weeded vegetation pixels can be used to figure out the rate of clusters per tile. In Table 2, the rate of clusters that was seen in the weedy areas is compared to the rate that was expected. These results show that for both types of data, it is possible to make accurate estimates of the number of weeds. With the help of a computer vision system, we might be able to treat only certain places and cut down on the amount of chemicals we use by a lot. This study suggests a semi-supervised method that could be used to help precision agriculture get more accurate estimates of weed densities and locations.

For the suggested method to work, it needs colour pictures as input. A binary flora mask can be made after the background pixels have been taken out. Using an unsupervised neural network, the bits that make up the background and the plants are put together. Second, we use the mask to divide the original colour picture into smaller parts (50-pixel-square tiles). Then, each piece is labelled as either a crop or a weed. In this study, a fine-tuned ResNet50 is compared to several other classifiers, such as SVM, Gaussian Naive Bayes, Neural Network, and Random Forest, which use a pre-trained ResNet50 as a feature generator. The suggested method is tried on two sets of images called Crop/Weed Field Image and Sugar Beets, which show a wide range of crops and weeds. With a maximum memory of 0.90, weedy areas can be found, and the number of weeds in those areas can be estimated with an accuracy of 85%. One of our goals is to reduce the need for elaborately annotated data sets. In order to identify weeds without the need for new features, the proposed work employs unsupervised segmentation and a pre-trained ResNet50. A pixel-wise segmentation network is shown to be unnecessary for estimating weed distribution and density, in contrast to previous techniques. The recommended pipeline is adaptable enough to process low-contrast images, images with overlapping plants, and images of distinct plant species. This strategy could be advantageous for agricultural organizations looking for cost-effective implementations due to the low data needs for training and tuning. A standard RGB camera is sufficient so long as a platform is in place from which to photograph the plants from above.

REFERENCES

- [1] A. dos Santos Ferreira, D. M. Freitas, G. G. da Silva, H. Pistori, and M. T. Folhes, "Unsupervised deep learning and semi-automatic data labeling in weed discrimination," *Comput. Electron. Agricult.*, vol. 165, Oct. 2019, pp. 104963.
- [2] A. Wang, W. Zhang, and X. Wei, "A review on weed detection using ground-based machine vision and image processing techniques," *Comput. Electron. Agricult.*, vol. 158, pp. 226–240, Mar. 2019.
- [3] International Society of Precision Agriculture, Monticello, IL, USA. Precision Ag Definition. Accessed: Feb 9, 2021. [Online]. Available: <https://www.ispag.org/about/definition>
- [4] Y. Xiong, C. Peng, L. Grimstad, P. J. From, and V. Isler, "Development and field evaluation of a strawberry harvesting robot with a cable-driven gripper," *Comput. Electron. Agricult.*, vol. 157, pp. 392–402, Feb. 2019.
- [5] W. McAllister, D. Osipychev, A. Davis, and G. Chowdhary, "Agbots: Weeding a field with a team of autonomous robots," *Comput. Electron. Agricult.*, vol. 163, Aug. 2019, Art. no. 104827.
- [6] M. N. Abd. Kharim, A. Wayayok, A. R. M. Shariff, A. F. Abdullah, and E. M. Husin, "Droplet deposition density of organic liquid fertilizer at low altitude UAV aerial spraying in rice cultivation," *Comput. Electron. Agricult.*, vol. 167, Dec. 2019, Art. no. 105045.
- [7] M. Hassanein, Z. Lari, and N. El-Sheimy, "A new vegetation segmentation approach for cropped fields based on threshold detection from hue histograms," *Sensors*, vol. 18, no. 4, p. 1253, Apr. 2018.
- [8] S. Sabzi, Y. Abbaspour-Gilandeh, and G. García-Mateos, "A fast and accurate expert system for weed identification in potato crops using metaheuristic algorithms," *Comput. Ind.*, vol. 98, pp. 80–89, Jun. 2018.
- [9] 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 Robot. Autom. Lett.*, vol. 3, no. 4, pp. 2870–2877, Oct. 2018.
- [10] A. Milioto, P. Lottes, and C. Stachniss, "Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in CNNs," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 2229–2235.
- [11] A. Olsen, D. A. Konovalov, B. Philippa, P. Ridd, J. C. Wood, J. Johns, W. Banks, B. Girgenti, O. Kenny, J. Whinney, B. Calvert, M. R. Azghadi, and R. D. White, "DeepWeeds: A multiclass weed species image dataset for deep learning," *Sci. Rep.*, vol. 9, no. 1, Dec. 2019, Art. no. 2058.
- [12] P. Lottes, J. Behley, N. Chebrolu, A. Milioto, and C. Stachniss, "Robust joint stem detection and crop-weed classification using image sequences for plant-specific treatment in precision farming," *J. Field Robot.*, vol. 37, no. 1, pp. 20–34, Jan. 2020.
- [13] A. Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Comput. Electron. Agricult.*, vol. 147, pp. 70–90, Apr. 2018.
- [14] P. Rasti, A. Ahmad, S. Samiei, E. Belin, and D. Rousseau, "Supervised image classification by scattering transform with application to weed detection in culture crops of high density," *Remote Sens.*, vol. 11, no. 3, p. 249, Jan. 2019.
- [15] R. Raja, D. C. Slaughter, S. A. Fennimore, T. T. Nguyen, V. L. Vuong, N. Sinha, L. Tourte, R. F. Smith, and M. C. Siemens, "Crop signalling: A novel crop recognition technique for robotic weed control," *Biosyst. Eng.*, vol. 187, pp. 278–291, Nov. 2019.
- [16] I. Sa, M. Popović, R. Khanna, Z. Chen, P. Lottes, F. Liebisch, J. Nieto, C. Stachniss, A. Walter, and R. Siegwart, "WeedMap: A

large-scale semantic weed mapping framework using aerial multispectral imaging and deep neural network for precision farming,” *Remote Sens.*, vol. 10, no. 9, p. 1423, Sep. 2018.

- [17] M. Fawakherji, A. Youssef, D. Bloisi, A. Pretto, and D. Nardi, “Crop and weeds classification for precision agriculture using context-independent pixel-wise segmentation,” in *Proc. 3rd IEEE Int. Conf. Robotic Comput. (IRC)*, Feb. 2019, pp. 146–152.
- [18] M. Caron, P. Bojanowski, A. Joulin, and M. Douze, “Deep clustering for unsupervised learning of visual features,” in *Proc. ECCV*, 2018, pp. 132–149.
- [19] Y. Xu, R. He, Z. Gao, C. Li, Y. Zhai, and Y. Jiao, “Weed density detection method based on absolute feature corner points in field,” *Agronomy*, vol. 10, no. 1, p. 113, Jan. 2020.
- [20] Y. Xu, Z. Gao, L. Khot, X. Meng, and Q. Zhang, “A real-time weed mapping and precision herbicide spraying system for row crops,” *Sensors*, vol. 18, no. 12, p. 4245, Dec. 2018.

