The International Society of Precision Agriculture presents the
15th International Conference on
Precision Agriculture
26–29 JUNE 2022
Minneapolis Marriott City Center | Minneapolis, Minnesota USA

Meta Deep Learning using Minimal Training Images for Weed Classification in Wild Blueberry

Patrick J. Hennessy[1], Travis J. Esau[1], Arnold W. Schumann[2], Aitazaz A. Farooque[3], Qamar U. Zaman[1], Scott N. White[4]

[1]Department of Engineering, Faculty of Agriculture, Dalhousie University, Canada

[2]Citrus Research and Education Center, University of Florida, United States

[3]School of Sustainable Design Engineering, University of Prince Edward Island, Canada

[4]Department of Plant, Food, and Environmental Sciences, Faculty of Agriculture, Dalhousie University, Canada

A paper from the Proceedings of the
15th International Conference on Precision Agriculture
June 26-29, 2022
Minneapolis, Minnesota, United States

Abstract.

Deep learning convolutional neural networks (CNNs) have gained popularity in recent years for their ability to classify images with high levels of accuracy. In agriculture, they have been applied for disease identification, crop growth monitoring, animal behavior tracking, and weed classification. Datasets traditionally consisting of thousands of images of each desired target are required to train CNNs. A recent survey of Nova Scotia wild blueberry (Vaccinium angustifolium Ait.) fields, however, determined that there are more than 200 unique species of weeds present. Collecting an image dataset containing thousands of images of each weed species to train a CNN would therefore be time-consuming and impractical. Meta deep learning allows for classification of images using a small number of labelled training examples, typically one or five images per class. To achieve this, the CNN is pre-trained using a standard dataset containing thousands of generic images. A support dataset containing a small number of images per class is provided for additional training of the specific target identities. A Siamese Neural Network (SNN) then uses the features learned by the CNN to differentiate between the classes in the support dataset. In this study, an SNN was trained to identify six species of weeds using the Keras-TensorFlow deep learning framework. Four different feature embedding sizes were tested for the SNN. The CNN training dataset contained three weed classes with 800 images per class collected in April through June during the 2019 and 2020 field seasons. Support datasets containing 1, 5, 10, 15, and 20 images per species were collected in April through July

The authors are solely responsible for the content of this paper, which is not a refereed publication. Citation of this work should state that it is from the Proceedings of the 15th International Conference on Precision Agriculture. EXAMPLE: Hennessy, P.J., Esau, T.J., Schumann A.W., Farooque, A.A., Zaman, Q.U., & White, S.N. (2022). Meta Deep Learning using Minimal Training Images for Weed Classification in Wild Blueberry. In Proceedings of the 15th International Conference on Precision Agriculture (unpaginated, online). Minneapolis, MN: International Society of Precision Agriculture.

2021 to train the SNN. The SNN achieved accuracies of 88.3% and 88.0% on the original validation and testing datasets with an embedding size of 512 neurons. Fine-tuning with a support set of 5 images yielded accuracies of 70.1% and 70.0% on the query validation and testing sets. Future work will involve using meta deep learning to identify common diseases in the wild blueberry crop including Monilinia blight (Monilinia vaccinii-corymbosi) and Botrytis blight (Botrytis cinerea). The trained SNNs will be deployed in a downloadable smartphone application and an online web-based application to facilitate streamlined delivery of pest identification and management information to wild blueberry growers.

Keywords.
Deep learning, few-shot learning, machine vision, Siamese neural networks, weed identification

Introduction

Wild Blueberry Cropping System

Wild blueberries (*Vaccinium angustifolium* Ait.) are a perennial crop which grows in eastern Canada and Maine, USA. Commercial fields are often developed on deforested areas after the removal of trees and other vegetation or abandoned farmland (Hall et al., 1979). The wild blueberry plants spread over these areas through rhizomes. Despite their perennial nature, wild blueberries are typically managed in a two-year cycle (Hall et al., 1979). During the first (sprout) year, plant growth begins and flower buds begin to grow in August. Harvesting occurs during August of the following (crop) year (Farooque et al., 2014). After harvest, the plants are pruned through mechanical flail mowing or burning, restarting the growth cycle.

Growers and industry professionals have indicated a desire for smart tools to improve field management. One such tool is a smartphone application which can identify visual field features such was weeds and crop diseases. A recent survey found that there are more than 200 unique species of weeds present in wild blueberry fields in Nova Scotia, Canada (Lyu et al., 2021).

Convolutional Neural Networks

Deep learning convolutional neural networks (CNNs) are image processing algorithms which can automatically classify images (Krizhevsky et al., 2012; LeCun et al., 1998) or objects within them (Girshick et al., 2014; Redmon et al., 2015). They consist of computational neurons, based on the perceptron (Rosenblatt, 1958), which are organized in layers to process data (Goodfellow et al., 2016; LeCun et al., 2015). The final layer contains one neuron for each possible class. Originally designed to recognize handwritten digits (LeCun et al., 1998), and later adapted for broad-scale image recognition (Krizhevsky et al., 2012), the adoption of CNNs has greatly increased over the past decade. Large datasets containing thousands of images are used to train CNNs through backpropagation of errors (Rumelhart et al., 1986) and iterative optimization algorithms based on gradient decent (Cauchy, 1847). They have been used in agriculture for crop growth monitoring (MacEachern et al., 2020; Tian et al., 2019), weed recognition (Sharpe et al., 2020; Yu et al., 2019), disease recognition (Fuentes et al., 2017), and monitoring livestock behavior (Wu et al., 2020; Yang et al., 2018).

Images are often downscaled from their original sizes to smaller resolutions between 28x28 pixels (LeCun et al., 1998) to 608x608 pixels (Redmon & Farhadi, 2018). This reduces the computational cost at the expense of accuracy. Studies performed on weeds in wild blueberry fields found that higher resolutions such as 1280x736 were necessary to optimize accuracy (Hennessy, Esau, Farooque, et al., 2021; Hennessy, Esau, Schumann, et al., 2021). Graphics processing units (GPUs) are typically required to train CNNs (Raina et al., 2009), but lightweight models can be deployed on common Central Processing Units (CPUs) found in personal computers and smartphones.

Siamese Neural Networks

Meta deep learning techniques, such as prototypical networks, Siamese neural networks, and model agnostic meta learning, adapt learned information from a trained neural network to new tasks using minimal amount of new training data. Siamese neural networks (SNNs) are a pair of identical neural networks used to compare the similarity of two pieces of data (Bromley et al., 1994; Dey et al., 2017; Koch et al., 2015). They were originally designed to compare and verify the similarity of handwritten signatures (Bromley et al., 1994). The SNN produces a vector embedding for each piece of data, then calculates the distance between the two.

Koch et al. (2015) used an SNN to recognize handwritten characters in the Omniglot dataset (Lake et al., 2015) using only one example, far fewer than the thousands of examples required by CNNs. This was achieved by pre-training the network using images from the MNIST dataset created by (LeCun et al., 1998). Support datasets containing one example of each character were then used to tune the network. The SNN created by Koch et al. (2015) used four layers from a CNN to create a 4096-neuron embedding of 105x105 resolution grayscale images containing handwritten character. An improved signature verification SNN, SigNet, creates 128-neuron embeddings from 155x220 grayscale images of signatures (Dey et al., 2017). Plant leaf diseases were identified using an SNN created by Argüeso et al. (2020) which used an Inception CNN architecture (Szegedy et al., 2015). Leaves were removed from the plants and placed on a single colour background for image capture. Li & Yang (2021) used an SNN with four convolutional layers to identify plant diseases and crop pests. They found that increasing the support set size improved accuracy on query datasets, but the rate of improvement declined when more than five images per class were used.

In this study, an SNN with four convolutional layers was trained to identify select weed species in wild blueberry fields. The accuracy of the SNN was evaluated using four feature embedding sizes ranging from 128 to 512 neurons. Accuracy of the SNN on query images was tested before and after fine-tuning with support datasets of 1, 5, 10, 15, and 20 images. Using an SNN, rather than a CNN, to create a smartphone app for weed recognition in wild blueberry fields would greatly reduce the size of the image dataset needed.

Materials and Methods

Computer Hardware and Software Environment

A Dell[1] Alienware Aurora R11 desktop computer with an Intel[2] Core i9-10900K CPU, 128 GB of system RAM, and an Nvidia[3] GeForce RTX 3090 24GB GPU was used for training and validating the SNNs. The computer used the Windows[4] 11 Pro operating system. The Python[5] programming language (v3.9.12) was installed using Anaconda[6]. The TensorFlow[7] machine learning platform (v.2.8.0) was installed in a virtual environment. The Nvidia graphics driver for the RTX 3090 GPU (v.512.13), the Compute Unified Device Architecture toolkit (CUDA, v.11.6), and the CUDA Deep Neural Network library (cuDNN, v.8.2.4.15) were installed to allow TensorFlow access to the GPU for processing.

Dataset Acquisition and Preprocessing

Images of nine weed species were captured in wild blueberry fields using rear-facing smartphone cameras with resolutions greater than 3000x2000 pixels (Figure 1). Images were captured from April to August 2019, 2020, and 2021 in sprout-year and crop-year fields. This resulted in an image dataset which was highly varied within each class because it encompassed various growth

---

[1] Dell Technologies Inc., Round Rock, TX, USA
[2] Intel Corporation, Santa Clara, CA, USA
[3] Nvidia Corporation, Santa Clara, CA, USA
[4] Microsoft Corporation, Redmond, WA, USA
[5] Python Software Foundation, DE, USA
[6] Anaconda Inc., Austin, TX, USA
[7] Google LLC, Mountain View, CA, USA

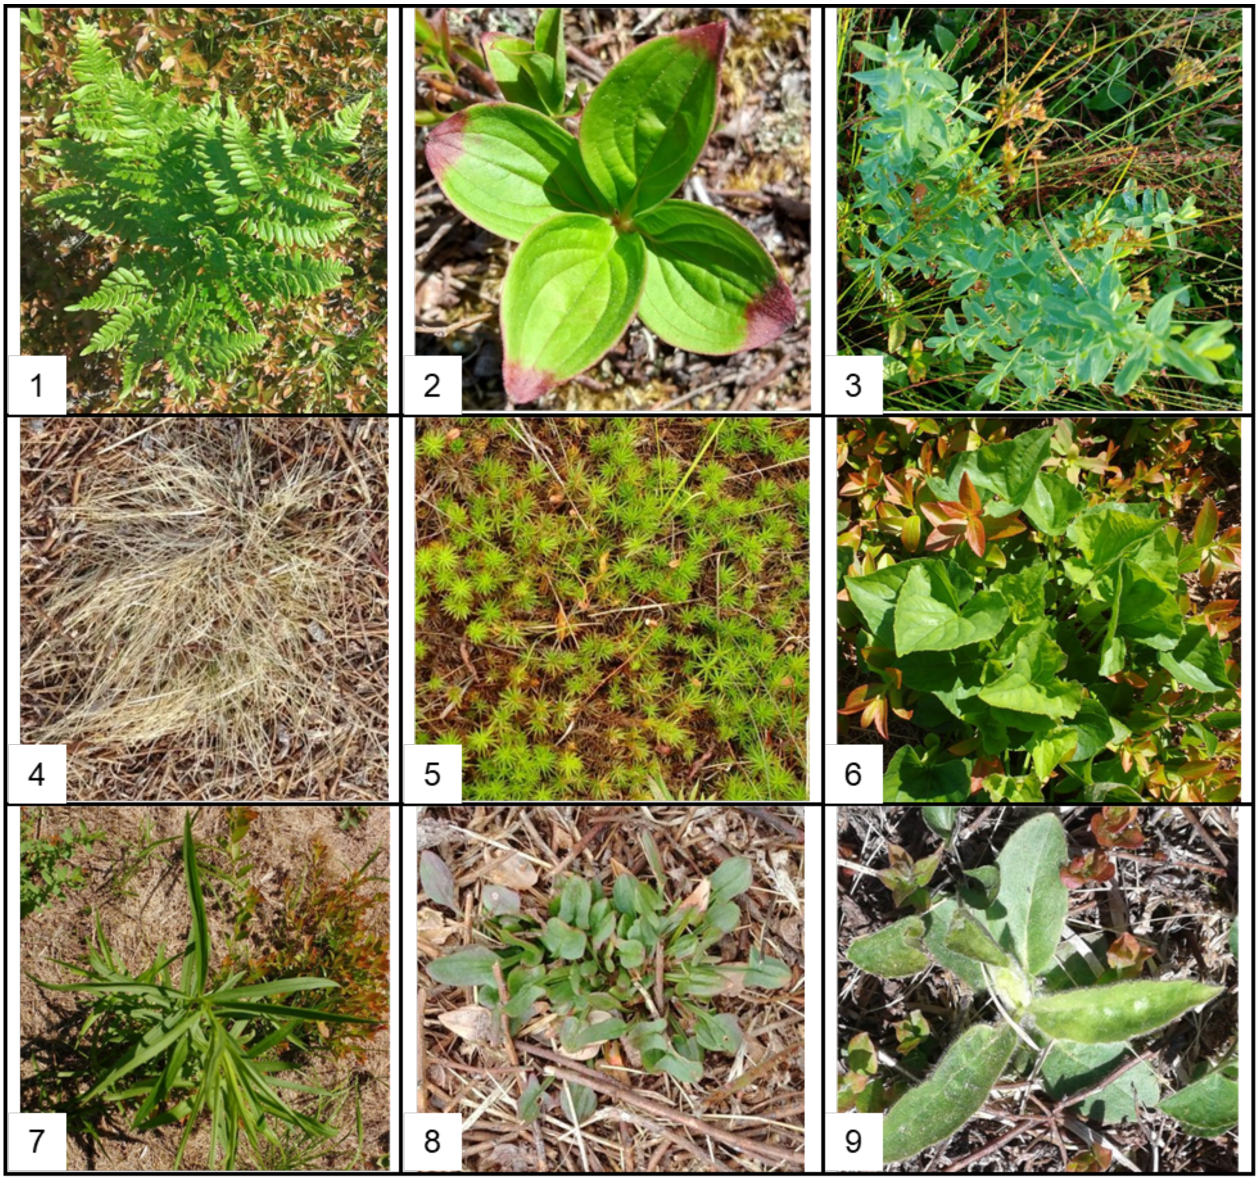

Figure 1: Examples of the nine weed species used in this study: (1) Bracken Fern, (2) Bunchberry, (3) Common St. John's Wort, (4) Hair Fescue, (5) Haircap Moss, (6) Marsh Violet, (7) Narrow-Leaved Goldenrod, (8) Sheep Sorrel, (9) Yellow Hawkweed.

stages of the weeds (Figure). For initial network training, validation, and testing, images of hair fescue (*Festuca filiformis* Pourr.), sheep sorrel (*Rumex acetosella* L.), and bunchberry (*Cornus canadensis* L.) were used. The training dataset consisted of 800 images of each weed, while the validation and independent testing datasets both contained 100 images of each weed.

Six additional weed species were used in the support and query datasets: bracken fern (*Pteridum acquilinum* (L.) Kuhn.), haircap moss (*Polytrichum commune* Hedw.), marsh violet (*Viola obliqua* Hill.), narrow-leaved goldenrod (*Euthamia graminifolia* (L.) Nutt.), common St. John's wort (*Hypericum perforatum* L.), and yellow hawkweed (*Hieracium caespitosum* Dumort). Query validation and testing datasets both contained 40 images of each weed species. Support datasets were created which contained 1, 5, 10, 15, and 20 example images ("shots") of each species. Query validation and testing datasets contained 40 examples of each weed species. The images

Figure 2: Examples of variability within the dataset. Each column contains the same weed species.

were down sampled using the IrfanView[8] batch processing tool (v.4.60). The original testing datasets for hair fescue, sheep sorrel, and bunchberry were divided in the same manner to create additional support and query datasets. The images were first cropped from their original aspect ratios to 16:9. The cropped images were then scaled down to 224x244 resolution for processing.

Network Architecture

A modified version of the SNN defined by Koch et al. (2015) was written and trained in this study using TensorFlow's Keras API (Figure 3). The input resolution of the network was increased from 105x105 to 224x224 to allow finer image details to be retained. Furthermore, the input layer was increased from one channel to three for processing of color images rather than grayscale. Directly after the input layer, a randomized image augmentation layer was added to make the SNN more robust (Goodfellow et al., 2016). Augmentations included rotation (+/- 180 degrees), translation (+/- 2% vertically and horizontally), zooming (+0%, -25%), and flipping (horizontal and vertical). The convolutional filters and max pooling layers defined by Koch et al. (2015) were not modified. Four convolutional layers are used in the SNN, each separated by a max pooling layer. The final convolutional layer in the original SNN is flattened and densely connected to the embedding layer. The SNN used in this study adds a dropout layer with a factor of 20% between the flattened layer

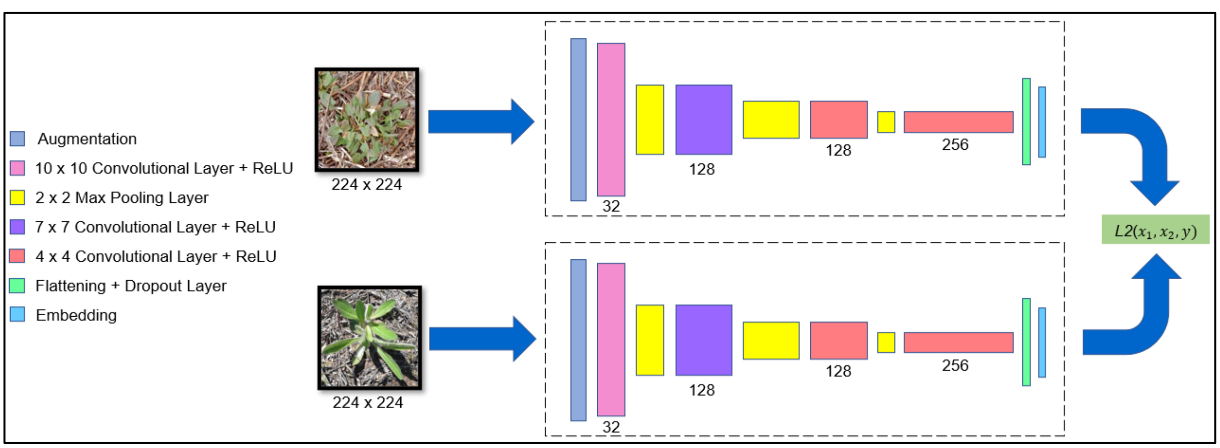

Figure 3: Diagram of the SNN used in this study. A pair of identical CNNs with shared weights each produce an embedding from input images. The L2 distance between the embeddings is calculated to determine if they are of the same class or different classes.

and the embedding layer to help prevent overfitting to the training data (Srivastava et al., 2014).

With the increased image resolution, the SNN could not be trained with an embedding layer of 4,096, as the memory required exceeded the 24 GB available in the RTX 3090 GPU. Four SNNs, each with a different embedding size, were trained and evaluated. The smallest embedding was 128, as used by SigNet (Dey et al., 2017). The other embedding sizes tested were 256, 384, and 512 neurons. The embeddings produced from two given images are compared using an L2 (Euclidian) distance calculation (Goodfellow et al., 2016) rather than the L1 (Manhattan) distance used by Koch et al. (2015).

Network Training

Images in the training, validation, and testing datasets were randomly grouped into positive and negative pairs. Positive pairs contain two images of the same weed species, while negative pairs contain two images of different species. A total of 4800 image pairs were created for training, while 600 pairs were created for both the validation and testing datasets.

The four SNNs were trained using a binary cross entropy loss function (i.e.: the SNN had to determine if an image pair was positive or negative). Network weights were updated using the Adam optimizer (Kingma & Ba, 2015) and an initial learning rate of 0.00001. Larger learning rates would not converge on a solution. The largest batch size possible within the GPU memory was used for training (16, 8, 6, and 4 image pairs for SNN embeddings of 128, 256, 384, and 512 neurons, respectively). Training epochs used every image pair in the training dataset. The SNNs were trained until the accuracy on the validation dataset did not improve for 20 epochs. The weights achieving the highest accuracy on the validation dataset were used for evaluation with the support and query datasets.

Few-Shot Testing

The SNN embedding and weights achieving the highest validation accuracy were further evaluated using the support datasets. The accuracy query validation and testing sets were evaluated with the SNN using positive and negative image pairs. Then, the SNN was fine-tuned using the support datasets containing 1, 5, 10, 15, and 20 images. The fine-tuned SNN was then re-evaluated using the query validation and testing datasets to determine the optimal number of images for fine-tuning. The network weights achieving the best accuracy on the query validation dataset were recorded.

Results and Discussion

Effect of Embedding Size

In general, increasing the embedding size improved the accuracy of the SNN on the training, validation, and testing datasets (Figure 4). The only exception was the decrease in training accuracy from 93.7% to 91.3% when the embedding was increased from 384 to 512 neurons. However, the accuracy on the validation and testing datasets increased by 0.8% and 0.7%, respectively. The peak accuracies achieved on the validation and testing datasets, achieved with the 512-neuron embedding, were 88.3% and 88.0%, respectively. The SNNs overfit to the training data at all embedding sizes other than 128 neurons. The accuracy reduced by 5.7%, 6.3% and 3.3%, respectively between the training and testing datasets at the 256, 384, and 512-neuron embedding sizes.

For smartphone deployment, the reduced accuracy of the 384-neuron embedding may be acceptable because of the increase in processing speed and decrease in memory required compared to the 512-neuron embedding. The results with the 256-neuron embedding compared to the 128-neuron embedding indicate that SNNs for use in this weed classification task require larger embeddings than the signature verification SNN created by Dey et al. (2017). However, the much larger 4096-neuron embedding used by Koch et al. (2015) is likely unnecessary, as the

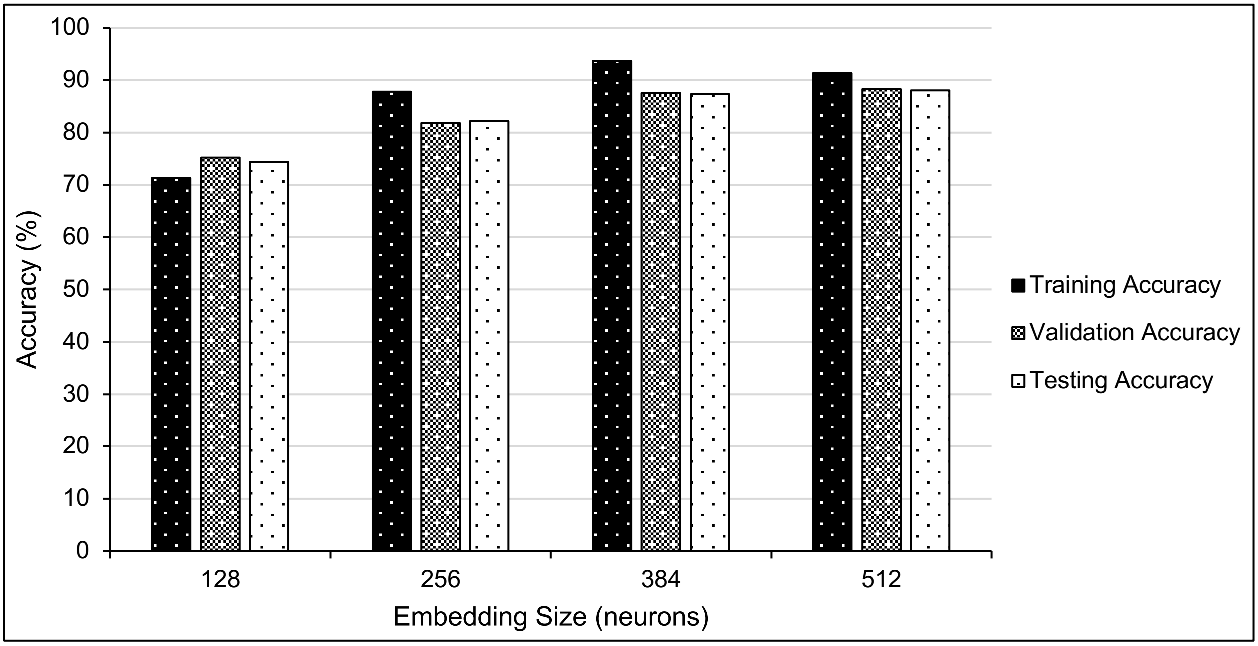

Figure 4: Comparison of training, validation, and testing dataset accuracy at four feature embedding sizes.

increase in accuracy between the 384 and 512-neuron embeddings in this study was minimal. Accuracy may be further improved through processing images at a higher resolution, as seen in other neural network weed classification studies in wild blueberry (Hennessy, Esau, Farooque, et al., 2021; Hennessy, Esau, Schumann, et al., 2021).

Fine Tuning with Few-Shot Datasets

Before fine-tuning, the SNN achieved accuracies of 67.5% and 66.6% on the query validation and testing datasets, respectively (Figure 5). Fine-tuning with a support set of 5 images improved the results on the query validation and testing datasets to 70.1% and 70.0%. As the number of support images increased, the query validation and testing accuracies diverged, eventually reaching

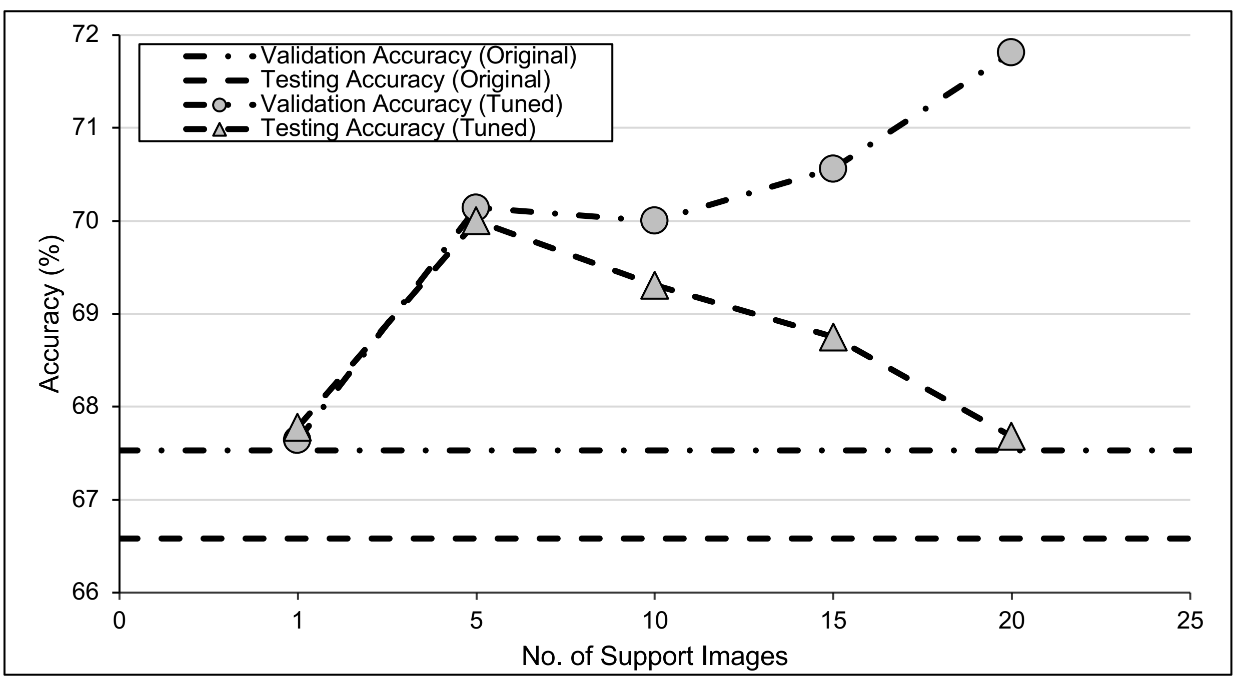

Figure 5: Improvements in SNN accuracy on the query validation and testing datasets after fine-tuning with different sizes of support datasets.

71.8% on the validation set and 67.7% on the testing dataset. Fine-tuning with a single image improved the accuracy of the SNN to 67.6% and 67.8% on the query validation and testing datasets.

The improvement in accuracy resulting from increasing the fine-tuning dataset from 1-shot to 5-shots is consistent with the results from Li & Yang (2021). The divergence of accuracy on the query validation and testing datasets with 10-shot and larger fine-tuning may indicate that the SNN weights overfitted to the validation dataset, as the accuracy on the validation dataset was used as the acceptance criteria. The reduced accuracy on the query datasets compared to the original datasets may be the result of the small number of original training classes. Training with more classes may result in more weed image features being represented in the network weights, which could improve accuracy. Furthermore, other CNN architectures such as ResNets (He et al., 2016) and Inception (Szegedy et al., 2015) have been effective for other SNNs (Argüeso et al., 2020). These architectures should be evaluated for use in weed classification in wild blueberry.

Conclusion

A Siamese neural network (SNN) based on four convolutional layers was trained to classify positive and negative pairs of 224x224 resolution weed images captured in wild blueberry fields. Four feature embedding sizes were tested, for the SNN architecture. An embedding size of 512 neurons achieved accuracies of 88.3% and 88.0%, respectively on the original validation and testing datasets. The accuracy of the SNN was substantially lower on the query validation and testing datasets, 67.5% and 66.6% respectively before fine-tuning. Fine-tuning with 5-shots improved the accuracy of the SNN to 70.1% and 70.0% on the query validation and testing datasets. To improve the accuracy of the SNN on the query datasets, other CNN architectures and an increased number of training classes should be investigated. If the accuracy of the SNN can be improved, it could serve as the backbone of a weed identification smartphone application for the wild blueberry industry. Further work will investigate possible improvements to the SNN, and the application of SNNs in other machine vision tasks in the wild blueberry industry such as plant disease identification. A field scouting smartphone application would help wild blueberry growers ensure they are using the most current management practices for weeds, diseases, and other yield-limiting factors.

Acknowledgements

This work was supported by Mitacs through the Mitacs Accelerate Program (IT23150), in partnership with the Wild Blueberry Producers Association of Nova Scotia. Additional funding for this research was provided by 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). The authors would like to thank the wild blueberry growers in Nova Scotia for use of their fields during image collection. Also, the authors acknowledge the efforts of the Agricultural Mechanized Systems (AMS) and Precision Agriculture research teams at Dalhousie University's Faculty of Agriculture. In particular, the authors thank Spencer Hauser and Chloe Toombs of the AMS research team for helping create the image database and generate the figures in this paper.

References

Argüeso, D., Picon, A., Irusta, U., Medela, A., San-Emeterio, M. G., Bereciartua, A., & Alvarez-Gila, A. (2020). Few-Shot Learning approach for plant disease classification using images taken in the field. *Computers and Electronics in Agriculture*, *175*(May). https://doi.org/10.1016/j.compag.2020.105542

Bromley, J., Guyon, I., LeCun, Y., Sackinger, E., & Shah, R. (1994). Signature Verification using a "Siamese" Time Delay Neural Network. *Advances in Neural Information Processing*

Systems 6 (NIPS 1993).

Cauchy, A.-L. (1847). Methode generale pour la resolution des systemes d'equations simultanees. *Compte Rendu Des Seances de L'Acad'emie Des Sciences*, *25*(2), 536–538.

Dey, S., Dutta, A., Toledo, J. I., Ghosh, S. K., Llados, J., & Pal, U. (2017). SigNet: Convolutional Siamese Network for Writer Independent Offline Signature Verification. *Pattern Recognition Letters*, 1–7.

Farooque, A. A., Zaman, Q. U., Groulx, D., Schumann, A. W., Yarborough, D. E., & Nguyen-Quang, T. (2014). Effect of ground speed and header revolutions on the picking efficiency of a commercial wild blueberry harvester. *Applied Engineering in Agriculture*, *30*(4), 535–546. https://doi.org/10.13031/aea.30.10415

Fuentes, A., Yoon, S., Kim, S. C., & Park, D. S. (2017). A robust deep-learning-based detector for real-time tomato plant diseases and pests recognition. *Sensors*, *17*(9). https://doi.org/10.3390/s17092022

Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 580–587. https://doi.org/10.1109/CVPR.2014.81

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* (1st ed.). The MIT Press.

Hall, I. v., Aalders, L. E., Nickerson, N. L., & vander Kloet, S. P. (1979). The biological flora of Canada 1. Vaccinium angustifolium Ait., sweet lowbush blueberry. *The Canadian Field-Naturalist*, *93*, 415–430.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CPVR)*. https://doi.org/10.1109/CVPR.2016.90

Hennessy, P. J., Esau, T. J., Farooque, A. A., Schumann, A. W., Zaman, Q. U., & Corscadden, K. W. (2021). Hair Fescue and Sheep Sorrel Identification Using Deep Learning in Wild Blueberry Production. *Remote Sensing*, *13*(5), 943. https://doi.org/10.3390/rs13050943

Hennessy, P. J., Esau, T. J., Schumann, A. W., Zaman, Q. U., Corscadden, K. W., & Farooque, A. A. (2021). Evaluation of Cameras and Image Distance for CNN-Based Weed Detection in Wild Blueberry. *Smart Agricultural Technology*, *2*(100030). https://doi.org/10.1016/j.atech.2021.100030

Kingma, D. P., & Ba, J. L. (2015). Adam: A method for stochastic optimization. *ArXiv*, 1–15. https://arxiv.org/pdf/1412.6980.pdf

Koch, G., Zemel, R., & Salakhutdinov, R. (2015). Siamese Neural Networks for One-Shot Image Recognition. *Proceedings of the 32nd International Conference on Machine Learning*.

Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). *ImageNet classification with deep convolutional neural networks*.

Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338. https://doi.org/10.1126/science.aab3050

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. https://doi.org/10.1038/nature14539

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Biochemical and Biophysical Research Communications*, *86*(11), 2278–2324.

Li, Y., & Yang, J. (2021). Meta-learning baselines and database for few-shot classification in agriculture. *Computers and Electronics in Agriculture*, *182*(October 2020), 106055.

https://doi.org/10.1016/j.compag.2021.106055

Lyu, H., McLean, N., McKenzie-Gopsill, A., & White, S. N. (2021). Weed Survey of Nova Scotia Lowbush Blueberry (Vaccinium Angustifolium Ait.) Fields. *International Journal of Fruit Science*, *21*(1), 359–378. https://doi.org/10.1080/15538362.2021.1890674

MacEachern, C., Esau, T., Schumann, A., Hennessy, P., & Zaman, Q. (2020). Deep Learning Artificial Neural Networks for Detection of Fruit Maturity Stage and Yield Determination in Wild Blueberries. *Manuscript in Progress*.

Raina, R., Madhavan, A., & Ng, A. Y. (2009). Large-scale deep unsupervised learning using graphics processors. *Proceedings of the 26th International Conference on Machine Learning*.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2015). You Only Look Once: Unified, real-time object detection. *ArXiv*.

Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *ArXiv*.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, *65*(6), 386–408. https://doi.org/10.1037/h0042519

Rumelhart, D., Hinton, G., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.

Sharpe, S. M., Schumann, A. W., Yu, J., & Boyd, N. S. (2020). Vegetation detection and discrimination within vegetable plasticulture row-middles using a convolutional neural network. *Precision Agriculture*, *21*, 264–277. https://doi.org/10.1007/s11119-019-09666-6

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, *15*, 1929–1958.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CPVR)*. https://doi.org/10.1109/CVPR.2015.7298594

Tian, Y., Yang, G., Wang, Z., Wang, H., Li, E., & Liang, Z. (2019). Apple detection during different growth stages in orchards using the improved YOLO-V3 model. *Computers and Electronics in Agriculture*, *157*(January), 417–426. https://doi.org/10.1016/j.compag.2019.01.012

Wu, D., Wu, Q., Yin, X., Jiang, B., Wang, H., He, D., & Song, H. (2020). Lameness detection of dairy cows based on the YOLOv3 deep learning algorithm and a relative step size characteristic vector. *Biosystems Engineering*, *189*, 150–163. https://doi.org/10.1016/j.biosystemseng.2019.11.017

Yang, Q., Xiao, D., & Lin, S. (2018). Feeding behavior recognition for group-housed pigs with the Faster R-CNN. *Computers and Electronics in Agriculture*, *155*(October), 453–460. https://doi.org/10.1016/j.compag.2018.11.002

Yu, J., Sharpe, S. M., Schumann, A. W., & Boyd, N. S. (2019). Detection of broadleaf weeds growing in turfgrass with convolutional neural networks. *Pest Management Science*, *75*(8), 2211–2218. https://doi.org/10.1002/ps.5349