Computers and Electronics in Agriculture 168 (2020) 105044

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Image recognition of Legacy blueberries in a Chilean smart farm through deep learning

Ignacio A. Quiroz[a], Germán H. Alférez[b,*]

[a] Facultad de Ingeniería y Negocios, Universidad Adventista de Chile, Camino a Tanilvoro Km. 12 Las Mariposas Chillán, Biobío, Chile
[b] Facultad de Ingeniería y Tecnología, Universidad de Montemorelos, Av. Libertad 1300 Poniente, Barrio Matamoros, 67530 Montemorelos, N.L., Mexico

ARTICLE INFO

Keywords:
Convolutional neural networks
Smart farms
Image recognition
Legacy blueberry

ABSTRACT

Agriculture is one of the most important pillars of development in Chile. However, it is expected that around the year 2030 there is going to be a gradual decrease in the number of farmers. Therefore, it is necessary to replace this workforce with technology and mechanization. One way to do this is through smart farms to leverage agricultural production. The contribution of this research work is a novel approach for deep-learning image recognition of Legacy blueberries in the rooting stage that can be used in smart farms in Chile. Legacy blueberry is a variety of Southern Highbush blueberry. This species constitutes 80% of the blueberry crops in Chile. Specifically, we propose an image recognition approach based on a convolutional neural network (CNN) to detect the presence of trays with living blueberry plants, the presence of trays without living plants, and the absence of trays. The average results of the evaluation of the predictive model are as follows: accuracy: 86%, precision: 86%, recall: 88%, and F_1 score: 86%.

1. Introduction

Agriculture in Chile has been an important area of development. Until around the mid-seventies it was mainly seen as a means of food production. Today this vision has changed, and agriculture has become one of the cornerstones for the development of the country. However, a gradual abandonment of this area of production by workers is envisioned by the year 2030. Therefore, it is imperative to replace this workforce with technology and mechanization (Carbonell, 2017).

In Chile, 80% of the blueberry crops are of the Southern Highbush species (Morales et al., 2017). The production of Legacy blueberries consists of different stages. The first stage is the in vitro micropropagation (within laboratory specimens) and the second stage is the rooting ex-vitro. Plants arrive in the ex-vitro stage without roots and with a size of not more than five centimeters. These plants are planted in plastic trays.

The rooting stage is a critical part of the production process. Thus, this stage requires a lot of care and attention. We have experienced during a year, in which micropropagation activity has been carried out at the Laboratory of Micropropagation at Universidad Adventista de Chile (Adventist University of Chile), that the highest percentage of plant losses occurs at the rooting stage (in certain seasons, such as winter and summer, the proportion of plant loss is equal to or greater than 50% of loss). It has been observed that this fact is directly related to exposure to extreme temperatures such as sub-zero temperatures during winter and temperatures above 30 °C in summer in addition to the low relative humidity of the environment.

To avoid plant loss, there is an ongoing project at the Laboratory of Micropropagation to automatize the rooting process of Legacy blueberries through an intelligent robotic system that works on plants in trays on a hot bed. In the hot bed, the trays are usually arranged in a fixed position. However, the quantity of the trays may vary when trays are removed. This leaves empty holes in the work area. Also, blueberry plants are randomly planted in the substrate of the trays. Therefore, the robotic system must be able to recognize the presence and absence of trays and even more, to recognize when a tray no longer contains plants of living blueberries. All of this with the need for saving both water and electricity resources.

The main objective of the Laboratory of Micropropagation is to increase profitability. Therefore, it is very important to reduce losses in the production of blueberry plants because these losses are translated into financial losses. Also, empirically we have observed in the rooting stage that the successful generation, development, and survival of roots of Legacy blueberry plants, is directly connected to the precision of care given to them during the rooting stage.

Our contribution in this work is an image recognition solution based on a convolutional neural network (CNN) to detect the presence of trays with living plants of blueberries of the Legacy variety in the rooting

* Corresponding author.
E-mail address: harveyalferez@um.edu.mx (G.H. Alférez).

https://doi.org/10.1016/j.compag.2019.105044
Received 5 July 2019; Received in revised form 2 October 2019; Accepted 5 October 2019
Available online 21 November 2019
0168-1699/ © 2019 Elsevier B.V. All rights reserved.

stage, the presence of trays without living plants, and the absence of trays. The characteristics of blueberry plants in the rooting stage are as follows: (1) *color:* in general, the leaves of blueberry plants vary from green to purple. As the plant withers, its leaves turn brown; (2) *size:* the size of the plants in the rooting stage can vary from 2 cm to 10 cm in height; and (3) *shape:* the shape of the leaves is oval or elliptical, with a pointed apex and a length that varies within 1 cm.

Specifically, this research work has two objectives. First, to create a predictive model by training a CNN using TensorFlow[1] with pictures from the Laboratory of Micropropagation of the Adventist University of Chile to detect through image recognition the presence of trays with living plants of blueberries, the presence of trays without living plants, and the absence of trays. A CNN is employed to achieve image recognition because this approach has proven to be very useful and effective for the recognition of images achieving a high percentage of accuracy in the agriculture domain (Deli et al., 2016; Wang et al., 2018; Durmuş et al., 2017; Milioto et al., 2018; Picon et al., 2019; Habaragamuwa et al., 2018; Steinbrener et al., 2019). Moreover, we propose to use TensorFlow as it is a very powerful library of high-performance numerical computing. It is also open source and has a flexible architecture. Our second objective is to evaluate the results obtained from the model in terms of accuracy and cross-validation with new images (not used during training). We expect that our solution could be used in other smart farms in Chile to increase the production of Legacy Blueberries.

This paper is organized as follows. The second section presents the state of the art. The third section presents the underpinnings of our approach. The fourth section presents the methodology. The fifth section describes the evaluation results. Finally, we present the conclusions and future work.

2. State of the art

One of the most famous projects in the context of smart farms is the Farmbot project that was developed on the concepts of open hardware and software (Aronson, 2014). This is a robotic system that allows to cultivate urban orchards in an autonomous way. This approach is specifically focused on people who want to grow their own vegetables at home. It is also an easy-to-use system for users who do not have exhaustive knowledge in programming, robotics or agriculture.

In (Yalcin, 2017), the authors show how deep learning can be used for image recognition to identify and classify the phenological stages of various types of plants. This study was uniquely based on visual data captured every half hour by cameras mounted on ground stations in Turkey as a part of a neural-network monitoring system. This approach employs a pre-trained CNN to extract the features of the images. The results suggest that the CNN architecture surpasses traditional approaches in the discrimination of the phenological stages.

In another research work (Deli et al., 2016), it is mentioned that, among the different deep learning topologies, CNN offers high performance in image classification. Specifically, the authors use CNNs to classify four classes of farmlands: field of cultivation, field of wood, field of houses, and no field of cultivation. 100 images were used in each class. 80% of the total number of images were used for training and 20% for validation. TensorFlow was used for training the model. After 700 iterations of training, the resulting accuracy was higher than 75% in each class.

In (Durmuş et al., 2017), the author used CNNs to detect diseases affecting the leaves of tomato plants. To this end, AlexNet and SqueezeNet were used for training a model on a Nvidia Jetson TX1. PlantVillage was used as the dataset (only with images of 10 different types of disease classes), as well as healthy plant images. The model was evaluated with images from the Web. The results obtained show that training with the AlexNet architecture was slightly faster than with the SqueezeNet architecture, getting both around 94% of accuracy.

In another research work, CNN-based semantic segmentation of cultivation fields was used to separate beet plants, weeds, and the background by analyzing RGB data (Milioto et al., 2018). The system was fully implemented in a real agricultural robot operating in different fields in Germany and Switzerland.

In (Picon et al., 2019), a CNN topology was proposed for analyzing the performance of the early identification of three relevant European endemic wheat diseases. The analysis was done using different mobiles devices, and more than 8,178 images were captured in two pilot sites in Spain and Germany for three years. The results reveal an overall improvement of the balanced accuracy from 0.78 up to 0.87 under exhaustive testing, and balanced accuracies greater than 0.96 on a pilot test performed in Germany.

In (Habaragamuwa et al., 2018), the authors used a CNN and greenhouse images taken under natural lighting conditions for recognizing two classes: mature and immature strawberries. The classification model achieved an average precision of 88.03% and 77.21%, and a bounding box overlap of 0.7394 and 0.7045 respectively for mature and immature classes.

In another research work (Steinbrener et al., 2019), the authors show an approach to classify hyperspectral images with state-of-the-art CNNs pre-trained for RGB image data. A small, custom dataset of hyperspectral images was recorded from staged but realistic scenes. With this dataset, an ImageNet pre-trained CNN was fine-tuned to create a classifier. To isolate the benefit of increased spectral resolution for the classification, the same analysis was also performed with pseudo-RGB images calculated from the hyperspectral images. The results show that the hyperspectral image data increases the average classification accuracy from 88.15% to 92.23%.

In (Wang et al., 2018), two CNNs, a Residual Network (ResNet) and its improved version named ResNeXt, were used to detect internal mechanical damage of blueberries using hyperspectral transmittance data. The original structure and size of hypercubes are adapted for the CNN training. To ensure that the models are applicable to hypercube, the authors adjusted the number of filters in the convolutional layers. Also, they state that selling damaged berries will lead to greater interest loss than discarding the sound ones. Thus, precision, recall, and F_1 score were also used as the evaluation indicators alongside accuracy to quantify the false positive rate. However, according to the authors the first three indicators are seldom used by researchers in the agricultural engineering domain. The fine-tuned ResNet/ResNeXt achieved average accuracy and F_1 score values of 0.8844/0.8784 and 0.8952/0.8905, respectively. The two generated deep learning models achieved better classification performance than other traditional machine learning methods.

3. Underpinnings of our approach

Computer vision or also known as artificial vision, aims to extract features of the real world by processing images with computers (Vélez et al., 2003). Computer vision tries to describe the world around us in one or more images and rebuild its properties such as shape, lighting, and color distribution, emulating the way humans and animals perform vision (Szeliski, 2010). Computer vision is composed of four stages. First, a digital image is acquired via some sensors. Then, filters and geometric transformations are applied with the aim of removing undesirable information and highlighting the most relevant information. Afterwards, images are segmented to obtain objects of interest. Finally, the image is recognized by distinguishing the segmented elements through the analysis of certain features. It is important to mention that these stages are not always followed in a sequential order, as several times it will be necessary to go backwards (Vélez et al., 2003).

We achieve computer vision through deep learning implemented by means of a CNN. Deep learning is a particular type of machine learning,

[1] https://www.tensorflow.org.

which is subfield of artificial intelligence. Machine learning is defined as the art and science of programming machines that can learn from data (Géron, 2017). These methods have dramatically improved image and audio recognition techniques (Goodfellow et al., 2016).

A CNN is a specialized artificial neural network (Goodfellow et al., 2016). CNNs have proven to be very effective in image recognition and classification. In a CNN four main operations are carried out: convolution, non-linear (rectified linear unit (ReLU) nonlinear function), pooling or sub-sampling, and classification. These operations are applied to images that are represented by numeric-value matrices of pixels (Krizhevsky et al., 2012). Specifically, CNNs depart from regular artificial neural networks in several ways (Davies, 2018): (1) CNN neurons have local connectivity, so they do not have to be connected to all the outputs from the previous layer of neurons; (2) their input fields can overlap; (3) in any layer, neurons have the same weight parameters across the whole layer; (4) CNNs use the ReLU nonlinear function instead of the sigmoidal output function (though each convolutional layer does not have to be fed directly to a ReLU layer); (5) they intersperse convolution layers with subsampling or "pooling" layer; and (6) they may have normalization layers to keep signals from each layer at suitable levels.

The CNN proposed in this research work has been implemented on TensorFlow. TensorFlow is an open source platform for deep learning development that was developed by Google Brain. TensorFlow allows the easy creation of machine-learning models thanks to its simple and flexible architecture.

The robotic system at the Laboratory of Micropropagation of the Adventist University of Chile is implemented via computer numerical control (CNC). CNC is a term that is often used to distinguish a type of technology in which the movements of a machine are controlled by a computer. Commonly, these machines are used in industry to improve efficiency, precision, and reliability (Groover, 2019). This type of automation derives from the concept of Numerical Control. The robot at the Adventist University of Chile has a carriage that runs along a hot bed mounted on a control area of approximately 2 square meters with the container trays of blueberry plants (see Fig. 1).

Legacy blueberry is a variety of Southern Highbush of blueberry, which constitutes 80% of the blueberry crops in Chile. Its fruits have medium to large size. In addition, they are firm and good tasting. It is a variety of high production in Chile. In fact, there are orchards that have reached 18 to 20 tons per hectare at the fourth and fifth years of production (Morales et al., 2017).

4. Methodology

This research work follows the IBM Foundational Methodology for Data Science (Rollins, 2016). This methodology consists of ten iterative stages. The first seven stages are problem understanding, analytic approach, data requirements, data collection, data understanding, data preparation, and modeling. These first stages are described in this section in the context of the problem and our proposed solution. The evaluation, deployment, and feedback stages are presented in the following section in the context of the evaluation of the solution.

4.1. Problem understanding

In this step, we studied the critical stage of rooting of Legacy blueberry plants. The highest percentage of plant losses throughout the year is in this stage. Specifically, at certain times of the year, the proportion of lost plants is equal to or greater than 50% due to the exposure to extreme temperatures, lack of precision in controlling the environment humidity, and inadequate application of irrigation. The robotic system a the Laboratory of Micropropagation has been designed to mitigate plant losses.

Moreover, computer vision is a key aspect to avoid wasting water and electricity when controlling the rooting process with the robot. By means of computer vision, the robotic system is able to control the following: (1) trays with living blueberry plants; (2) trays without living blueberry plants; and (3) absence of trays. These three scenarios are important because even though the trays may have a certain order in the hot bed, it is not guaranteed that there is always a tray in each place or that in one or more trays there are living plants.

4.2. Analytical approach

Deep learning was chosen as the analytical approach to classify the images in three classes: tray with living blueberry plants, tray without living blueberry plants, and no tray. A CNN was chosen to achieve deep learning because this kind of neural networks are ideal for image recognition.

4.3. Data requirements

The data required in this research work consist of images of trays with living blueberry plants, trays without living blueberry plants, and no trays (images of empty holes between trays or the sand surface of the

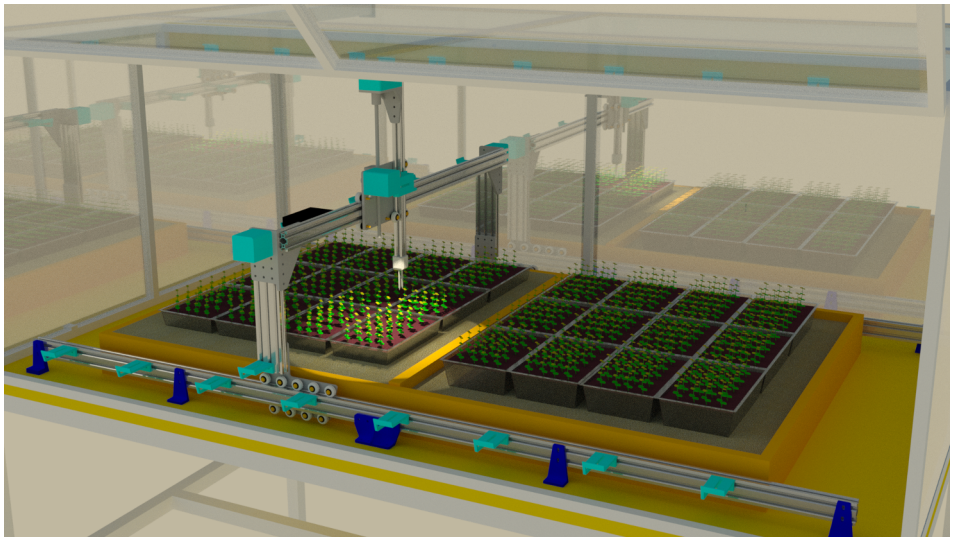

Fig. 1. CNC robotic system.



Fig. 2. Tray with living Legacy blueberry plants.

hot bed). The pictures were taken with a Microsoft Lifecam Studio digital camera with a resolution of 1,920 × 1,080 pixels. The pictures had to fulfill several requirements. First, the pictures had to depict three specific situations: tray with living blueberry plants, tray without living blueberry plants, and no tray. Also, the pictures had to be taken with good light conditions (illuminance of 10,000–25,000 lux – full daylight or 32,000–100,000 lux – direct sun), not blurred or without distracting objects in the background (such as wires, stones, rubbish, etc.). Finally, all the images had to be taken from the same point of view of the robotic arm (from the top of the hot bed).

4.4. Data collection

The images that were used to train the classification model are as follows: (1) Fig. 2 shows a picture of a tray with living Legacy blueberry plants; (2) Fig. 3 shows a picture of a tray without living Legacy blueberry plants; and (3) Fig. 4 shows a picture without any tray. This picture only shows the sand surface that covers the hot bed.



Fig. 3. Tray without living Legacy blueberry plants.



Fig. 4. No tray.

Table 1
Distribution of images used in the experiments (total values are in bold).

Stages	Classes		
	Tray with living Legacy blueberry plants	*Tray without living Legacy blueberry plants*	*No tray*
Training	56	56	56
Validation	18	18	18
Prediction	12	12	12
Total number of images per class	**86**	**86**	**86**
Total number of images		**258**	



Fig. 5. Tray in orange with living Legacy blueberry plants.

Table 2
Topology of the CNN.

Layer (Type)	*Output Shape*	*Parameters*
conv2d (Conv2D)	(None, 150, 150, 32)	896
max_pooling2d (MaxPooling2D)	(None, 150, 150, 32)	0
conv2d_1 (Conv2D)	(None, 75, 75, 64)	8,256
max_pooling2d_1 (MaxPooling2)	(None, 37, 37, 64)	0
flatten (Flatten)	(None, 87,616)	0
dense (Dense)	(None, 256)	22,429,952
dropout (Dropout)	(None, 256)	0
dense_1 (Dense)	(None, 3)	771

In total, 258 images were used in the experiments. These images were divided into 3 classes as shown in Table 1. The classification model was generated in the *training stage* with 168 pictures, 56 pictures per class. Initially, we started with a low number of images for training. Then, the generated model was validated in the *validation stage*. We repeated these steps until we reached an average accuracy, precision, and recall values higher than 85%, which is an adequate percentage in our case. Finally, the classification of a specific image was made in the *prediction stage*. The images used in the validation and prediction stages were images different from those used in the training stage to avoid overfitting.

4.5. Data understanding

The original trays are made of transparent plastic as seen in Figure 2 and Fig. 3. However, we found out in our first experiments that transparency could make it hard to locate the trays by the robotic arm. Therefore, to facilitate the training of the underlying CNN, we decided to paint the edges of the trays in orange. For instance, Fig. 5 shows a tray with living plants of Legacy blueberries with a painted tray.

4.6. Modeling

In the modeling stage, the predictive model was generated through a CNN with the topology shown in Table 2. This topology is composed of eight layers. The first layer applies the convolution operation. In this layer, 32 filters are applied to the input signals. The output is a 150 × 150 pixels image. The next layer is max pooling. This combination is common in a CNN. In the max pooling layer, the signals of outputs (map of features) obtained from the convolution layer are reduced in terms of their spatial dimension, where the size of the pool is 2 × 2 pixels. Next, there is another convolution of 64 additional filters. The convolutional layer extracts features of higher abstraction level, such as borders and lines from the reduced feature maps in the previous max pooling layer. Next, there is another max pooling layer. This layer allows to keep decreasing the spatial dimension of the features while keeping the most important ones. Then comes a flatten layer. In this stage, the images are flattened in a single dimension. The next layer is the dense layer, which has 256 fully-connected neurons, followed by a dropout layer that controls the activation of the neurons of the previous layer to avoid overfitting of the network and improve the adaptation of the network to new information. Finally, there is a dense layer with only 3 neurons. This is because each one of these neurons represents a class to be recognized (tray with living blueberries plants, tray without living plants, and no tray).

A program written in Python uses TensorFlow version 1.13.1 for CPU to train and evaluate the classification model and make predictions. Also, we used Keras[2] because it provides high level instructions for the development of deep learning.

[2] https://keras.io.

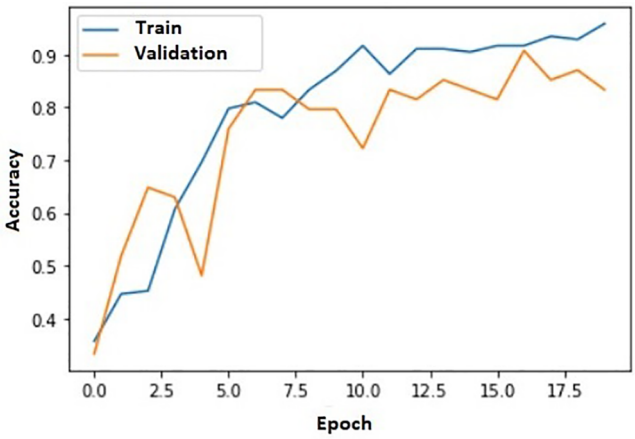

Fig. 6. Model's accuracy per epoch.

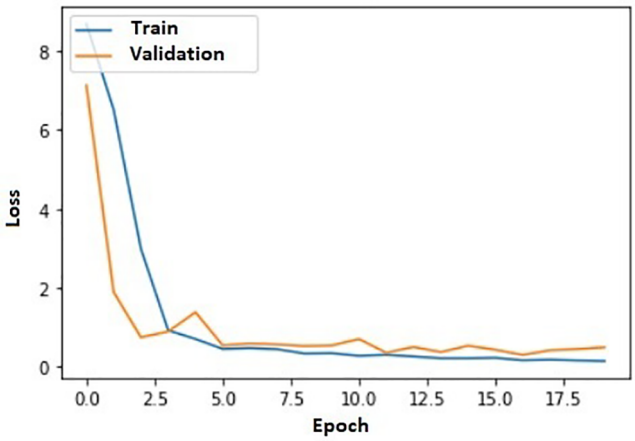

Fig. 7. Model's loss per epoch.

Table 3
Multi-class confusion matrix of the predictive model.

Actual	Prediction		
	Tray with living plants	*Tray without living plants*	*No tray*
Tray with living plants	9	1	0
Tray without living plants	3	11	1
No Tray	0	0	11

5. Evaluation

This section presents the results obtained with the classification model. The computer used to create the classification model, to evaluate it, and to make predictions was a laptop Acer Nitro 5 AN515-5 with the following characteristics: AMD Ryzen 5 2500U processor with Radeon Vega Mobile Gfx, RAM memory of 12 GB, and with Windows 10 of 64 bits. Python version 3.7 and TensorFlow version 1.13.1 for CPU were used in the experiments.

The robotic arm, controlled by a Raspberry Pi 3 B+, moves over the hot bed. The hot bed is abstracted as an $m \times n$ matrix. The size of each cell *a*equals the size of a tray. In every cell, a program written in Python takes a picture of the hot bed half-hourly, starting from the first cell $a_{0,0}$. After the robotic arm reaches the last cell, it returns to the first cell and

keeps taking pictures during the day (in order to take pictures under good light conditions). The Python program is deployed on an Nvidia Jetson Nano. As soon as the picture is taken, the program loads the classification model and uses the new picture as a parameter to make a new classification of the current location. Currently, the robotic system is in a testing phase. Feedback will be received from the staff of the Micropropagation Lab for extending the model to classify other possible scenarios as part of our future work.

The levels of accuracy and loss were measured to validate the model. Figs. 6 and 7 present the results. These figures show that the accuracy increases, and the loss decreases as the number of iterations increases until reaching epoch number 20.

In the prediction stage, we used 36 images not used during training to make predictions according to the classification model, 12 images for each one of the three classes. These images were also taken at the Laboratory of Micropropagation. Table 3 shows that the classes with the highest number of right classifications were the classes of trays without living Legacy blueberry plants and no tray. The metrics that were used for evaluation are accuracy, precision, recall, and F_1 score.

The accuracy metric represents the fraction of predictions the model got right, and is formally defined as:

$$Accuracy = \frac{Number\ of\ correct\ predictions}{Total\ number\ of\ predictions}$$

The precision metric represents the ratio of the real positives among the examples that are expected to be positive in relation to the total quantity of true positives plus false positives. The precision metric is defined as:

$$Precision = \frac{TP}{TP + FP},$$

where TP is the number of true positives and FP the number of false positives. The precision is intuitively the ability of the classifier not to label as positive a sample that is negative.[3]

The recall metric is defined as:

$$Recall = \frac{TP}{TP + FN},$$

where TP is the number of true positives and FN the number of false negatives. The recall is intuitively the ability of the classifier to find all the positive samples.[4]

The F_1 score is the weighted harmonic mean of the precision and recall. It is defined as:

$$F_1\ Score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

The resulting accuracy classification score of the model is 0.86 (i.e, 86 correct predictions out of 100 total examples). Moreover, the outcomes in Table 4 show the quality of the classification model in terms of precision, recall, and F_1 score. In all the cases, the average results are close to 1, which is adequate in our case. The metrics were measured with Scikit-learn.[5]

6. Conclusions and future work

This work presented a deep-learning solution for image recognition of Legacy blueberries in the rooting stage. Specifically, the classification model was created by training a CNN with pictures of trays with living Legacy blueberry plants in the process of rooting, trays without living Legacy blueberry plants, and no trays. The results obtained from the evaluation of the model when making predictions with new images from the Laboratory of Micropropagation at the Adventist University of Chile reached promising averages: accuracy: 86%, precision: 86%, recall: 88%, and F_1 score: 86%. This results show the technical feasibility of using deep learning in smart farms in Chile to try to prevent blueberry plants losses in the rooting stage.

As future work, the classification model of the computer vision module of the robotic arm will be extended to recognize other situations, such as the presence of pests. Also, filters for noise reduction will be applied to recognize objects in pictures taken under extreme conditions, for instance, in a foggy environment due to humidity. Moreover, we will explore approaches for real-time multi-object detection to detect more than one situation per picture. To this end we will work with YOLO (Redmon and Farhadi, 2018). Furthermore, we expect to use generative adversarial networks (GANs) to generate synthetic images that closely resemble real ones. GANs are a type of artificial neural network that generates new data from scratch. The new synthetic images could be used to train and validate the classification model, minimizing the need for accessing real data.

Table 4
Results obtained from the evaluation of the model (average values are in bold).

Class	Metric		
	Precision	*Recall*	*F_1 Score*
Tray with living Legacy blueberry plants	0.75	0.90	0.82
Tray without living Legacy blueberry plants	0.92	0.73	0.81
No tray	0.92	1.00	0.96
Average	**0.86**	**0.88**	**0.86**

Declaration of Competing Interest

There are no conflict of interest in this work.

Acknowledgements

The authors gratefully acknowledge the support of Universidad Adventista de Chile, Universidad de Montemorelos, Red Latinoamericana de Carreras de Informática (RELACI), Banco Santander, and Rodrigo Leiva for making this collaboration a reality. Also, we thank Roberto Rojas for taking the pictures used in this research work.

References

Aronson, R., 2014. Humanity's open-source automated precision farming machine, Tech. Rep., Farmbot, URL <https://bit.ly/2XpvktC>.
Carbonell, C., 2017. Agricultura Chilena, Reflexiones y Desafíos al 2030, Tech. Rep., Oficina de Estudios y Políticas Agrarias (ODEPA).
Davies, E., 2018. Chapter 15 - Deep-learning networks. In: Davies, E. (Ed.), Computer Vision, 5 ed. Academic Press, pp. 453–493 ISBN 978-0-12-809284-2. doi: 10.1016/B978-0-12-809284-2.00015-0.
Deli, Z., Bingqi, C., Yunong, Y., 2016. Farmland Scene Classification Based on Convolutional Neural Network. In: 2016 International Conference on Cyberworlds (CW), pp. 159–162, https://doi.org/10.1109/CW.2016.33.
Durmuş, H., Güneş, E.O., Kırcı, M., 2017. Disease detection on the leaves of the tomato plants by using deep learning. In: 2017 6th International Conference on Agro-Geoinformatics, 1–5, https://doi.org/10.1109/Agro-Geoinformatics.2017.8047016.
Géron, A., 2017. Hands-On Machine Learning with Scikit-Learn and TensorFlow. O'Reilly.
Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press.
Groover, M.P., 2019. Fundamentals of Modern Manufacturing: Materials, Processes and Systems, 7 ed. Wiley ISBN 978-1-119-47521-7.
Habaragamuwa, H., Ogawa, Y., Suzuki, T., Shiigi, T., Ono, M., Kondo, N., 2018. Detecting greenhouse strawberries (mature and immature), using deep convolutional neural network. Eng. Agric., Environ. Food 11 (3), 127–138 ISSN 1881–8366.
Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet Classification with Deep Convolutional Neural Networks. In: Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1, NIPS'12, Curran Associates Inc., USA, 1097–1105, 2012.
Milioto, A., Lottes, P., Stachniss, C., 2018. Real-time semantic segmentation of crop and weed for precision agriculture robots leveraging background knowledge in CNNs.

[3] https://scikit-learn.org/stable/modules/generated/sklearn.metrics.precision_recall_fscore_support.html.
[4] https://scikit-learn.org/stable/modules/generated/sklearn.metrics.precision_recall_fscore_support.html.
[5] https://scikit-learn.org/stable/.

Proc. IEEE Int. Conf. Robot. Autom. (ICRA).
Morales, C., González, A., Riquelme, J., Hirzel, J., France, A., Pedreros, A., Uribe, H., Defilippi, B., Robledo, P., Becerra, C., 2017. Manual de Manejo Agronómico del Arándano, Instituto de Investigaciones Agropecuarias (INIA), URL <https://bit.ly/2WebOQw>.
Picon, A., Alvarez-Gila, A., Seitz, M., Ortiz-Barredo, A., Echazarra, J., Johannes, A., 2019. Deep convolutional neural networks for mobile capture device-based crop disease classification in the wild, Comput. Electron. Agric. 161, 280–290, ISSN 0168-1699, https://doi.org/10.1016/j.compag.2018.04.002, bigData and DSS in Agriculture.
Redmon, J., Farhadi, A., 2018. YOLOv3: An Incremental Improvement, arXiv https://arxiv.org/abs/1804.02767.
Rollins, J.B., 2016. Foundational Methodology for Data Science, Tech. Rep., IBM, URL <https://www.ibm.com/downloads/cas/4QAOJKDN>.
Steinbrener, J., Posch, K., Leitner, R., 2019. Hyperspectral fruit and vegetable classification using convolutional neural networks. Comput. Electron. Agric. 162, 364–372. https://doi.org/10.1016/j.compag.2019.04.019. ISSN 0168-1699.
Szeliski, R., 2010. Computer Vision: Algorithms and Applications. Springer.
Vélez, J., Moreno, A., Sánchez, Angel, Sánchez, J., 2003. Visión por computador, Dykinson, ISBN 84-9772-069-5.
Wang, Z., Hu, M., Zhai, G., 2018. Application of deep learning architectures for accurate and rapid detection of internal mechanical damage of blueberry using hyperspectral transmittance data. Sensors 18 (4), ISSN 1424–8220, https://doi.org/10.3390/s18041126, <https://www.mdpi.com/1424-8220/18/4/1126>.
Yalcin, H., 2017. Plant phenology recognition using deep learning: Deep-Pheno. In: 2017 6th International Conference on Agro-Geoinformatics, 1–5, https://doi.org/10.1109/Agro-Geoinformatics.2017.8046996.