The Use of Machine Learning Algorithms to Predict Crop Yield

Kristiana Anderson

CONTENTS

***Abstract*—By 2050 the global population is expected to increase by more than one billion people. Unless the global crop yield is dramatically increased, food security will continue to decline. Simultaneously, the escalating effects of climate change will continue to disrupt agriculture. To aid in crop production, machine learning algorithms can be used to predict and increase crop yield. This paper will analyze the machine learning algorithms used to predict crop yield as well as compare the practicality of each algorithm. The algorithms discussed include Convolutional Neural Networks, K-Nearest Neighbors, Support Vector Machines, Random Forest Classifiers, and Linear Regression. These machine learning algorithms can improve farmers' understandings of their crops leading to better resource allocation and thus better budgets. Due to the usefulness of these algorithms, machine learning will continue to spread throughout agriculture, aiding farmers in their quest to more efficiently produce crops.**

I. SURVEY OF THE FIELD

A. Introduction

According to the United Nations, the world has an estimated population of 7.9 billion. In the next 30 years, it is expected to increase to over 9 billion people [27]. This population growth will require the food supply to increase by at least 70 percent worldwide to provide an adequate amount of food for all [13]. Simultaneously, farmers will be struggling to adapt to a rapidly changing world due to climate change. Agriculture is especially vulnerable to climate change as weather directly impacts crop yields [7]. One article from the Proceedings of the National Academy of Sciences predicts that climate change will decrease the yields of major crops [29]. This comes at a time when increase crop yield is needed. Combined, these factors signal disaster for global food security. However, farmers can combat this with an improved understanding of their fields by analyzing predicted crop yields.

For many years, farmers have used traditional methods to predict crop yield. These methods begin by partitioning off a section of the field, then individually counting the produce within the selected section. Once this process is complete, the farmers will extrapolate the final crop yield from the collected data [28]. Although this method is widely used, it is time-consuming and can be highly inaccurate. With the use of computer vision, taking a sample would not be needed, and instead a whole field could be quickly and inexpensively surveyed. Thus, the crop yield and data pertaining to a field is more accessible, leading to more educated decisions during farming. For example, if a farmer can view the predicted crop yield multiple times during the growing season, they can alter their approach and focus their resources on the section of the fields with lower predicted yields. The computer vision algorithms used most commonly for crop yield prediction are convolutional neural networks, a subset of artificial neural networks along with several other machine learning algorithms.

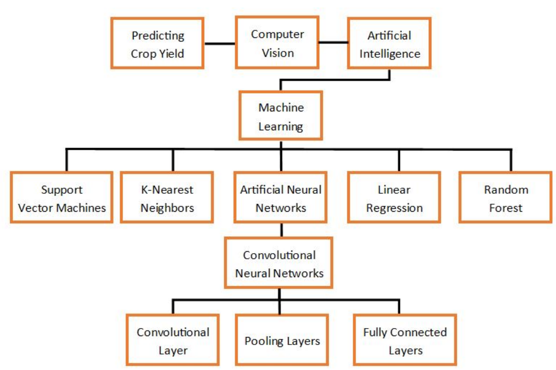

Fig. 1. Diagram depicts a mental model of the algorithms referenced in this article. Crop yield is predicted using computer vision which utilizes artificial intelligence. Machine learning is a subsection of artifical intelligence that includes many useful algorithms. Notably, artificial neural networks evolved into convolutional neural networks which are the machine algorithms most commonly used to predict crop yield. Convolutional neural networks are composed of three main layers.

II. Technical Analysis of Machine Learning Algorithms in Agriculture

A. Artificial Neural Networks

Computer vision is a subset of artificial intelligence (AI) that allows computers to analyze and identify images. Machine Learning, a part of AI, is widely used in computer vision. A machine learning algorithm is given data from a user and then examines that data for patterns that would allow the algorithm to make predictions [21]. In the past, the majority of image recognition research in agriculture has utilized machine learning algorithms, however, with the inception of deep learning - machine learning algorithms with artificial neural networks (ANNs) - an increasing number of studies are abandoning past machine learning algorithms for more complex algorithms. ANNs are algorithms modeled after the human brain and thus are built to solve nonlinear and complex problems [22]. The most basic ANNs are collections of perceptrons and activation functions. Perceptrons are composed of:

- Inputs: The data given to the perceptron. For computer imaging these inputs will be the numerical representation of the image.
- Weights: The weights operate on the inputs passed to them in the weighted sum. The weights are established during training of the ANN.
- Weighted Sum: The sum of all the inputs when multiplied by their corresponding weight.
- Activation or Step Functions: Activation functions dictate when a perceptron should activate.
- Output: The output is the result of the perceptron, which may be passed to succeeding perceptrons.

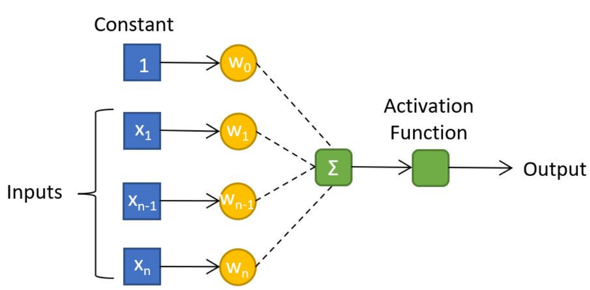

Fig. 2. Diagram depicts a single perceptron in an artificial neural network. In deep learning, there would be many layers of perceptrons, all connected by activation functions.

The weighted sums of the perceptrons are then passed on through activation functions, the simplest activation function's being step functions as shown in the example above. These functions are what makes neural networks non-linear and allows ANNs to solve complex problems. Common activation functions are sigmoid, tanh, and most notably rectified linear unit (RELU). ReLu is a piecewise linear function that outputs a zero if negative and the original input if positive [11]. The simplest neural networks are made up of three layers: the input layer, a hidden layer, and an output layer. The input layer contains nodes that represent each input variable, while the hidden layer processes the data, and the output layer represents the result. If a network has multiple hidden layers, it would be considered a deep neural network [24].

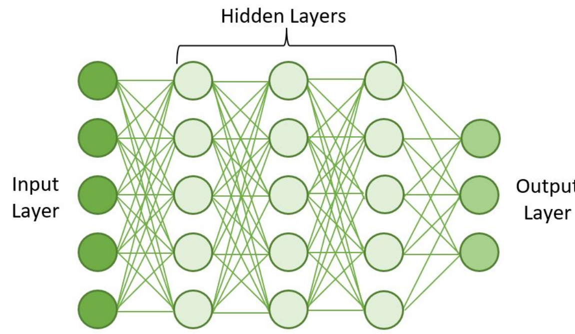

Fig. 3. Pictured above is a diagram of a simplified deep neural network.The circles represent perceptrons, and the lines signify the connection to perceptrons in succeeding layers.

Before a neural network can be used, it must be trained. The training process determines the values of the weights in the perceptrons. When training begins, the weights are set to random values. During training, the error of the neural network is calculated using a loss function, then the weights of the input are adjusted accordingly [22]. There are several different loss functions, each with different use cases. Some of the notable loss functions include hinge loss, softmaz loss, contrastive loss, and triplet loss [11]. Training ends when the loss function can no longer be minimized [22]. Training requires large amounts of data. Most studies reserve at least 75 percent of the dataset for training [12]. Once a neural network has been trained, the weight values of the network are referred to as the model [22].

Because ANNs analyze many different types of data, specialized architectures have been created to solve specific issues within the field. The neural networks that are used most frequently for analyzing images are convolutional neural networks (CNN).

B. Convolutional Neural Networks

Much like ANNs, CNNs are made up of perceptrons and solve complex problems by learning from the data provided to them. However, CNNs are specifically made to process images, meaning there are some key differences. To start, unlike an ANN, the perceptrons within a CNN will only connect to a small area of the following layers. This allows a CNN to analyze images that would be considered too large to use in traditional ANN training. Additionally, the perceptrons within a layer of a CNN are arranged in three dimensions, contrasting with the single dimension in ANNs. CNNs are also composed of three distinct layers: convolutional, pooling, and fully-connected layers [18].

When processing images, the convolutional layer acquires the data from the input layer which contains the pixel values of the image. The convolutional layer aims to identify the distinguishing characteristics of the image. To do this, the convolutional layer utilizes a kernel, an array containing the numeric representation of the desired feature, and walks the kernel over the image [15]. Then as the kernel is moved, the scalar product will be determined for every value in the kernel. This will result in a feature map of the image [18]. Below is an example of a kernel searching for vertical edges in a 4x4 image.

Step 1:

$$\begin{pmatrix} 55 & 208 & 98 & 222 \\ 67 & 236 & 107 & 162 \\ 14 & 214 & 155 & 47 \\ 135 & 197 & 45 & 48 \end{pmatrix} * \begin{pmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} \quad & \quad \\ \quad & \quad \end{pmatrix}$$

Step 2:
$(55 * 1) + (228 * 0) + (98 * -1) + (67 * 1) + (236 * 0) + (107 * -1) + (14 * 1) + (214 * 0) + (155 * -1) = -224$

$$\begin{pmatrix} 55 & 208 & 98 & 222 \\ 67 & 236 & 107 & 162 \\ 14 & 214 & 155 & 47 \\ 135 & 197 & 45 & 48 \end{pmatrix} * \begin{pmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} -224 & \quad \\ \quad & \quad \end{pmatrix}$$

$$\cdots$$

Step n:

$$\begin{pmatrix} 55 & 208 & 98 & 222 \\ 67 & 236 & 107 & 162 \\ 14 & 214 & 155 & 47 \\ 135 & 197 & 45 & 48 \end{pmatrix} * \begin{pmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{pmatrix} = \begin{pmatrix} -224 & 247 \\ -91 & 390 \end{pmatrix}$$

The network then uses the resulting matrix to learn the kernels that react to specific features. Each kernel used in the convolutional layer creates an activation map that corresponds with the depth dimension of the CNN. This is what gives the CNN its three dimensional structure [18].

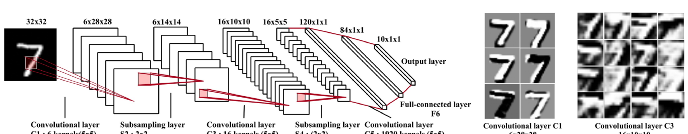

(a) LeNet-5 network (b) Learned features

Fig. 4. Figure depicts the LeNet-5 network, a CNN algorithm. The diagram clearly displays the 3D attribute of the CNN as well as the stacking of the feature maps [11].

The pooling or subsampling layer then reduces the complexity of the model by compressing the dimensions of the activation map. It achieves this by summarizing groups of cells in each of the activation maps using a max, min, or average function [15]. CNNs may be composed of several convolutional and subsampling layers, with some convolutional layers in succession. Nevertheless, the fully-connected layers are always placed at the end of the process. Fully-connected layers calculate class scores for the activations that will be used for classification. Finally, the output layer may feed the result into further algorithms depending on the CNN's task [11]. Some CNNs may continue to pass results to successive CNNs, while others may feed the data to other algorithms [15]. For example, if the CNN is used for object detection, the result may be fed to an ANN which will then predict if the desired object is found in the image.

In the past decade, research into CNNs has grown significantly, leading to improvements in the standard CNN architecture. Some of the most notable advances in image classification CNNs are:

- AlexNet: improved large scale image classification
- Hierarchical Deep CNNs: allowed deep CNNs into a category hierarchy
- R-CNN: system could learn part and whole object detectors [11]

C. Other Machine Learning Classifiers

Although CNNs are the leading machine learning algorithms used to calculate crop yield, several other notable machine learning algorithms are still in use today. These classifiers can be used independently or in conjunction with CNNs.

1) **Support Vector Machines (SVMs)** are a type of supervised learning classifier [13]. Supervised learning means that the system relies on labeled training data. Conversely, unsupervised learning machines use unlabeled data sets. Thus, they find patterns in the data independent of human interaction [8]. To analyze the data, the SVM organizes it into a dimensional space and then performs a nonlinear transformation on the input. This results in a feature space with higher dimensionality. Then the SVM calculates separating hyperplanes. In oversimplified terms, a separating hyperplane can be seen as a plane in three or more dimensions that splits the space in half [20]. Stated plainly, the SVM attempts to separate the different classes of data by creating boundaries around them. Therefore, the SVM can solve complex problems that cannot be solved in a one-dimensional space by transforming them into linear form in a higher dimensional space [12].
2) **K-Nearest Neighbors (KNNs)** are supervised learning classifiers that are often utilized for classification pattern recognition. The goal of KNNs are to determine the "closest training samples in a feature space" [12]. Weighted KKNs achieve this by assigning greater weights to data closer to a value, and lesser weights to those further away [12]. Some of the formulas used to calculate the distance between two points in KNNs are Minkowski, Manhattan, and Euclidean [20].
3) **Random Forest** classifiers are algorithms that depend on decision and regression trees. Decision trees exhibit decisions and their outcomes, forming a tree-like graph. Regression trees are similar to decision trees, however, instead of classifying data, regression trees attempt to make predictions. To make predictions, random forest classifiers use numerous decision trees to generate a large number of regression trees. They then use the resulting regression trees to make a prediction [20].
4) One of the most basic machine learning algorithms today is **linear regression** algorithms. Linear regression makes

predictions based on mathematical variable projections. In its simplest form, linear regression "distinguishes the influence of independent variables from the interaction of dependent variables" [16]. However, multivariable linear regression (MLR) can also be used to predict the result of a variable based on multiple independent variables [16].

$$y = \beta_0 + \beta_1 x + \epsilon \tag{1}$$

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_m x_m + \epsilon \tag{2}$$

D. Advantages and Disadvantages

All of the previously mentioned machine learning algorithms have their use in predicting crop yield. However, some are more accurate than others depending on the situation. For example, the effectiveness of an algorithm may rely on the diversity and size of the data set. The nature of the data can also impact the accuracy of a classifier's prediction. Agricultural datasets often contain outliers, and some classifiers are more vulnerable to outliers than others. Even the capacity of the hardware may dictate which machine learning classifier is used in certain situations [20].

According to a systematic literature review from the Journal of Computers and Electronics in Agriculture, CNNs are the most applied deep learning algorithms, appearing in over thirty percent of the papers analyzed [25]. This is due to the algorithm's ability to learn the important features of the data from the raw input. Regardless of the advantages, some systems cannot handle the high processing costs of CNNs. Additionally, CNNs require a large amount of data to train, data that is not readily available to the masses [20]. Some studies have avoided this issue by rotating the images and thus generating more training data as the model will recognize it as a new image [28]. Depending on the system, images may require pre-processing, preventing some from utilizing CNNs. Thus a CNN may be used in a situation where there is a large collection of images and no constraints on computing power.

Not all machine learning algorithms require high amounts of preprocessing. Due to random forest's low data preprocessing requirements, implementation can be done relatively quickly. Other advantages of random forests are their high-speed operations, and ability to perform on "sparsely annotated data" [20]. The classifier is also resistant to overfitting [3]. Overfitting occurs when the system is trained on too small of a dataset. This leads to the classifier recognizing images in the training data set, but failing to recognize any other images [28]. Overall, random forest classifiers operate efficiently on large datasets and are resistant to outliers [3]. Even so, random forests' efficiency depends on their hyperparameters, the parameters dictating the learning process i.e. decision trees, making optimization of hyperparameters integral to their performance [20].

Similarly to random forests, the training of KNNs is relatively fast. Additionally, the model is resistant to noisy data or data with little meaning. The computational complexity of KNN is one of its largest limitations. KNNs also run slowly because of it being a "lazy learner". "Lazy learning" classifiers do not generalize training data until a query is made. Together, the complexity and slow speed of the KNN can cause the model to be taxing on some systems, especially if the data set is very large. This means that KNN should not be used in situations where there are large datasets and computing restraints [4].

The review from the Journal of Computers and Electronics in Agriculture also notes that the most used machine learning algorithms behind neural networks are linear regressions [25]. Even so, it must be acknowledged that "most widely used" does not translate to "the best performing". Often linear regression models are used as benchmarking algorithms. Unlike CNNs where there is a hidden layer, linear regression allows researchers to easily read the coefficients and say why the model behaved the way it did. Therefore, linear regression may not be the most accurate algorithm, but it still offers value to researchers [20].

While some machine learning algorithms require large datasets, SVMs perform well under small training sets but operate poorly under large data sets [20], [5]. They also have low variance and high stability due to their simple decision rule. SVMs can be more resistant to noisy data found in agricultural data sets than other machine learning because of its generalization of data [20]. Still, the SVMs accuracy decreases if the overall data set is unbalanced [5].

In recent years, scientists have tried to minimize the structural weaknesses in each machine learning algorithm. One strategy is to combine several algorithms. For example, hybrid networks are networks that combine several deep learning algorithms to form a new system [25]. Another study form the Journal of Stored Products Research used a CNN as a generic feature extractor and then used machine learning classifiers such as SVM, KNN, and ANN to classify the images. Researchers found that models trained with "simple features" or features not extracted by a CNN had a lower accuracy than the models trained with CNN-extracted features [12]. Additionally, studies have found that combinations of SVM and other machine learning algorithms perform better than either SVM or the machine learning algorithm alone. This is due to the decreased level of variance when multiple algorithms are combined [20].

All of these technologies allow farmers to receive more accurate crop yields. Today, unmanned aerial vehicles (UAVs) and other motorized vehicles can quickly scan fields [24]. This allows farmers to have a complete map of their crops and ultimately a better understanding of the effects of their farming practices. This greater comprehension of resources can lead to more efficient resource consumption, better educated budgets, and the adoption of eco-friendly policies [3]. Meaning that although computer vision methods do face their own challenges, they have the potential to be much more useful and accurate compared to traditional techniques.

E. Challenges

Even though machine learning algorithms are advancing rapidly, the technology still faces challenges in an agricultural setting. One of the largest issues for image recognition in farming is the irregularity of the natural world. Although crops follow a general shape and size, all are slightly different, meaning if a crop differs too greatly, the computer vision software will not be able to recognize it. In addition, the lighting of a photo can dramatically affect the estimation of the crop yield. This is because over-shadowing and overexposure can cause the crop's features to become less discernible in images. The partial obstruction and clustering of crops in photos can also impact the system's accuracy [19]. Finally, images with visually complex backgrounds may reduce the accuracy of the software [23]. While these are sizable issues, accuracy can be improved by taking more images from distinct angles and by using artificial illumination during image collection. Extensive pre-processing of images may also improve the accuracy of the machine learning algorithms [19]. Even though there are challenges facing computer vision in agriculture, the benefits from more complete data and more accurate crop yield far surpass the negatives

III. Future Trends

Computer vision technology has changed drastically from the first inception of the first multilayer CNN in 198, to present day CNNs [14]. The technology will continue to change as researchers create more complex algorithms and advanced hardware. In the next year, the use of machine learning algorithms to predict crop yield will increase due to their improving accuracy and accessibility.

A. Deep CNNs

In recent years, a large percentage of papers analyzing the use of machine learning to predict crop yield have used deep CNNs [25]. Their popularity is due to their ability to solve complex problems. Unlike CNNs, deep CNNs possess "the ability to learn complex representations at different levels of abstraction" [14]. Studies have found that deep CNNs outperform their less advanced counterparts when used for image recognition [14]. Although deep CNNs are more accurate than other algorithms, they are often avoided as they require massive datasets and advanced hardware [26]. However, within the next year these requirements may no longer limit the algorithm.

In countries like the United States, sizable datasets have become more available as there is a push for large-scale survey data [2]. Datasets such as CropDeep and PlantDoc are open to the public and contain thousands of images of crops [30]. There also has been an increase in commercial datasets such as Crop Tracker and Xyonix. Additionally, equipment that collects data is becoming more accessible. Most unmanned vehicles sold today are fitted with RGB, high-resolution cameras. Therefore, any farmer can utilize these technologies to create their own dataset [17]. The cost of unmanned vehicles, especially drones has rapidly decreased in the past decade [9]. Unmanned vehicles can also use intuitive and automated software to create flight plans allowing even inexperienced owners to efficiently and safely collect their own data [25]. The accuracy of deep CNNs is dependent on the quality and size of the dataset it is trained on, meaning more accessible, sizable datasets can greatly improve the accuracy of deep CNNs' crop yield predictions.

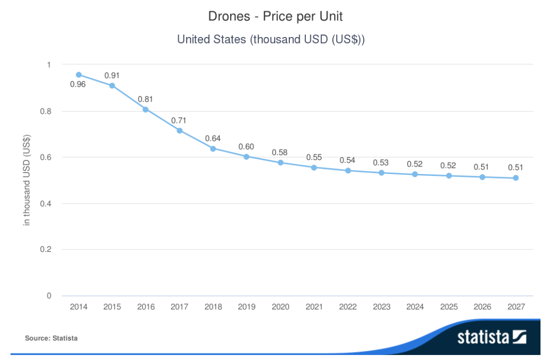

Fig. 5. The graph above depicts the average price per drone sold in the United States in terms of thousands of dollars. From 2014 to 2019, the price of drones rapidly decreased. It is predicted that the price will continue to decrease, however at a much slower pace [9].

If there comes a time when labeled data sets are still not available, researchers can use transfer learning. Transfer learning is a technique that utilizes a pre-trained model to solve a distinct but similar issue. For example, a pre-trained deep CNN model can be used to classify images it was not trained to recognize. Instead of re-training the model, the first few layers are unchanged and the last layers are removed. Therefore, the previously learned weights are transferred to the model, allowing it to continue classifying images. Transfer learning can decrease the computation time and size of the dataset required to train a deep CNN [1]. Both of which improve the accuracy of the models.

Deep CNNs can also pose challenges for resource-limited devices. For instance, a deep CNN can run in reasonable time on a Nvidia DGX-2 supercomputer but will fail to run on the average smartphone. Due to the algorithm's deep and wide architecture, it makes hundreds of multiplications in the convolutional layers which leads to a sizable memory requirement and a high computational cost. This can cause issues for those without specialty equipment. To improve deep CNN's performance, researchers study possible strategies to optimize parameters, so CNNs can operate on resource-constrained hardware. A new field of research is creating features that manipulate feature maps and input representation to alter the size of the network. Other strategies used to lessen the algorithm's requirements are:

1) Knowledge distillation: the method of transferring knowledge from a larger model to a much smaller model which could be used on a less technologically advanced device [14].
2) Training of small networks: researchers have discovered subnetworks in networks that are considerably smaller than the original network, but are still capable of making accurate predictions [6].
3) Squeezing of pre-trained networks: by using methods such as Huffman coding (a way to compress data without losing any details), the size of the network can be reduced [14].

Experts have also sought to convert deep networks from supervised learning to semi-supervised learning, self-taught learning, and generative adversarial networks, which all utilize unsupervised learning to an extent. Unsupervised learning algorithms do not require data to be labeled, easing the coordinator's workload. They can also identify unknown patterns in the dataset as unsupervised learning algorithms are not instructed on how to classify the data [26]. In the next year, new architecture will likely be created that improve current deep CNN's strenuous requirements. This will make the algorithm more accessible and increase the use of AI in agriculture. However, CNNs will not be the only algorithms to be improved upon, other machine learning algorithms will continue to advance in the next year.

B. Improvements for all Algorithms

There have been several other proposals to increase algorithms' efficiency. Studies suggest pipeline parallelism, a method of processing layers of an algorithm simultaneously, to decrease DNN memory and computation requirements. Additionally, cloud computing can be used to process large amounts of data without the cost of high computational efficiency. With the use of technology such as cloud computing, resource-limited devices may run any machine learning algorithm, further increasing machine learning's accessibility. The figure below displays the value of AI in agriculture in 2020 as well as the predicted value each year up to 2026. Overall, the value is predicted to grow to four billion dollars by 2026. Although this value is an aggregate of all AI in agriculture, it signals that the usage of AI to predict crop yield will continue to increase.

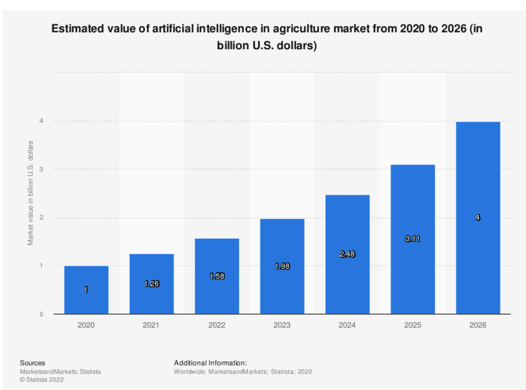

Fig. 6. The graph above displays the predicted growth of AI in agriculture. Within six years, the value of AI in agriculture is estimated to grow by four times [10]

IV. Conclusion

Today, computer vision is aiding the advancement of agriculture by providing more accurate crop yields. Continuation of research into the independent use and combination of machine algorithms and neural networks will create more efficient and powerful systems. Soon, because of image recognition technology, traditional methods of predicting crop yield may become irrelevant. Although there are still challenges facing computer vision, the technology continues to grow in popularity. If computer vision continues to spread throughout agriculture, farmers will have the opportunity to increase global crop yield despite the disastrous effects of climate change.

References

[1] Muhammet Fatih Aslan, Muhammed Fahri Unlersen, Kadir Sabanci, and Akif Durdu. Cnn-based transfer learning–bilstm network: A novel approach for covid-19 infection detection. *Applied Soft Computing*, 98:106912, 2021.

[2] Nishu Bali and Anshu Singla. Emerging trends in machine learning to predict crop yield and study its influential factors: A survey. *Archives of Computational Methods in Engineering*, 29, 03 2021.

[3] Brenon Diennevan Souza Barbosa, Lucas Costa, Yiannis Ampatzidis, Vinay Vijayakumar, Luana Mendes dos Santos, et al. Uav-based coffee yield prediction utilizing feature selection and deep learning. *Smart Agricultural Technology*, 1:100010, 2021.

[4] Nitin Bhatia et al. Survey of nearest neighbor techniques. *arXiv preprint arXiv:1007.0085*, 2010.

[5] Jair Cervantes, Farid Garcia-Lamont, Lisbeth Rodríguez-Mazahua, and Asdrubal Lopez. A comprehensive survey on support vector machine classification: Applications, challenges and trends. *Neurocomputing*, 408:189–215, 2020.

[6] Adam Conner-Simons. Smarter training of neural networks, 2019. https://news.mit.edu/2019/smarter-training-neural-networks-0506.

[7] Andrew Crane-Droesch. Machine learning methods for crop yield prediction and climate change impact assessment in agriculture. *Environmental Research Letters*, 13(11), 2018.

[8] Julianna Delua. Supervised vs. unsupervised learning: What's the difference?, 2021. https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning.

[9] Drones - price per unit, 2022. https://www.statista.com/outlook/cmo/consumer-electronics/drones/united-states.

[10] Estimated value of artificial intelligence in agriculture market from 2020 to 2026, 2022. https://www.statista.com/statistics/1326924/ai-in-agriculture-marketvalue-worldwide/.

[11] Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, Jianfei Cai, and Tsuhan Chen. Recent advances in convolutional neural networks. *Pattern Recognition*, 77:354–377, 2018.

[12] Shima Javanmardi, Seyed-Hassan Miraei Ashtiani, Fons J. Verbeek, and Alex Martynenko. Computer-vision classification of corn seed varieties using deep convolutional neural network. *Journal of Stored Products Research*, 92:101800, 2021.

[13] Vijay Kajani, Basivi Praveen Kumar Van Huan Nguyen, Hakil Kim, and Visweswara Rao Pasupuleti. A critical review on computer vision and artificial intelligence in food industry. *Journal of Agriculture and Food Researchs*, 2, December 2020.

[14] Asifullah Khan, Anabia Sohail, Umme Zahoora, and Aqsa Saeed Qureshi. A survey of the recent architectures of deep convolutional neural networks. *Artif. Intell. Rev.*, 53(8):5455–5516, dec 2020.

[15] Samaya Madhavan. Introduction to convolutional neural networks, 2021. https://developer.ibm.com/articles/introduction-to-convolutional-neural-networks/.

[16] Dastan Maulud and Adnan M Abdulazeez. A review on linear regression comprehensive in machine learning. *Journal of Applied Science and Technology Trends*, 1(4):140–147, 2020.

[17] Petteri Nevavuori, Nathaniel Narra, and Tarmo Lipping. Crop yield prediction with deep convolutional neural networks. *Computers and Electronics in Agriculture*, 163:104859, 2019.

[18] Keiron O'Shea and Ryan Nash. An introduction to convolutional neural networks. *arXiv preprint arXiv:1511.08458*, 2015.

[19] WS Qureshi, Alison Payne, KB Walsh, R Linker, O Cohen, and MN Dailey. Machine vision for counting fruit on mango tree canopies. *Precision Agriculture*, 18(2):224–244, 2017.

[20] Mamunur Rashid, Bifta Sama Bari, Yusri Yusup, Mohamad Anuar Kamaruddin, and Nuzhat Khan. A comprehensive review of crop yield prediction using machine learning approaches with special emphasis on palm oil yield prediction. *IEEE Access*, 9:63406–63439, 2021.

[21] Ranajit Roy, Ankur Paul, Priya Bhimjyani, Nibhash Dey, Debankan Ganguly, Amit Kumar Das, and Suman Saha. A short review on applications of big data analytics. *Emerging Technology in Modelling and Graphics*, pages 265–278, 2020.

[22] Rajalingappaa Shanmugamani, Abdul Ghani Abdul Rahman, Stephen Maurice Moore, and Nishanth Koganti. *Deep Learning for Computer Vision : Expert Techniques to Train Advanced Neural Networks Using TensorFlow and Keras*. Packt Publishing, 2018.
[23] Davinder Singh, Naman Jain, Pranjali Jain, Pratik Kayal, Sudhakar Kumawat, and Nipun Batra. Plantdoc: a dataset for visual plant disease detection. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, pages 249–253. 2020.
[24] Rocío Vallesteros, Diego S. Intrigliolo, José F. Ortega, Juan M. Ramírez-Cuesta, Miguel A. Moreno, and Ignacio Buesa. Vineyard yield estimation by combining remote sensing, computer vision and artificial neural network techniques. *Precision Agriculture*, 21:1242–1262, 2020.
[25] Thomas van Klompenburg, Ayalew Kassahun, and Cagatay Catal. Crop yield prediction using machine learning: A systematic literature review. *Computers and Electronics in Agriculture*, 177:105709, 2020.
[26] Aichen Wang, Wen Zhang, and Xinhua Wei. A review on weed detection using ground-based machine vision and image processing techniques. *Computers and Electronics in Agriculture*, 158:226–240, 2019.
[27] World population day, 2022. https://www.un.org/en/observances/world-population-day.
[28] Huaqing Yu, Shining Song, Shaoxi Ma, and Richard O. Sinnott. Estimating fruit crop yield through deep learning. In *Proceedings of the 6th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, BDCAT '19, page 145–148, New York, NY, USA, 2019. Association for Computing Machinery.
[29] Chuang Zhao, Bing Liu, Shilong Piao, Xuhui Wang, David B. Lobell, Yao Huang, Mengtian Huang, Yitong Yao, Simona Bassu, Philippe Ciais, Jean-Louis Durand, Joshua Elliott, Frank Ewert, Ivan A. Janssens, Tao Li, Erda Lin, Qiang Liu, Pierre Martre, Christoph Müller, Shushi Peng, Josep Peñuelas, Alex C. Ruane, Daniel Wallach, Tao Wang, Donghai Wu, Zhuo Liu, Yan Zhu, Zaichun Zhu, and Senthold Asseng. Temperature increase reduces global yields of major crops in four independent estimates. *Proceedings of the National Academy of Sciences*, 114(35):9326–9331, 2017.
[30] Yang-Yang Zheng, Jian-Lei Kong, Xue-Bo Jin, Xiao-Yi Wang, and Min Zuo. Cropdeep: The crop vision dataset for deep-learning-based classification and detection in precision agriculture. *Sensors*, 19(5):1058, Mar 2019.

Appendix

Due to this project's independent nature, I found myself relying mostly on my past experiences. Although I have never taken any college courses that focused on producing formal papers on scientific and technical topics, I found reviewing previous coursework to be helpful. The courses that I found most beneficial when completing this project were CSCI200, CSCI230, and Software Engineering.

When I first took CSCI200, I assumed that all I would take away from the course was a basic knowledge of data structures. However, during Capstone I found myself looking back on the course and using the knowledge I learned to create the demonstration of my project. In CSCI200, we were tasked with creating our own video game with little aid from the instructors. During this time, I learned how to code independently as well as how to research possible solutions. At the time I did not think it would be that important, but now as I look back I understand the significance of this knowledge.

I also found the CSCI230 project to be very helpful. When I first took the class, I had not taken CSCI200, meaning I had to learn from my peers and resources online. Due to this experience, I have internalized the attitude that I must learn from those around me. In this project, I did this by asking other students for helpful resources and using the shared Github in order to expand my understanding of convolutional neural networks.

Software Engineering was another course that I thought was beneficial. Although I took the course this semester, the independent aspect of the software engineering project showed the importance of continuously reading, listening, and watching. I also learned that you can't cram all the work into one day, it must be spread out in order to create the best product.

Overall, this course felt like a quick review of all the courses I have taken in my college career. The project allowed me to use all of the skills I have learned in classes over the past four years. For example, the CSCI200 and CSCI230 projects seemed tedious at the time, however, now I understand the significance of such assignments. I know that in the future I will be grateful for the things I learned in this course and the opportunities it gave me.

List of Figures