

Introduction to Artificial Intelligence in Agriculture¹

Young Gu Her, Nikolay Bliznyuk, Yiannis Ampatzidis, Ziwen Yu, and Haimanote Bayabil²

Introduction

The Fourth Industrial Revolution, also known as Industry 4.0, represents a fundamental change in the way we live, work, and relate to one another, driven by advancements in technology. ML and AI are central to this transformation. We now see more self-driving cars on the roads, chatbots that quickly answer customer questions, web browsers that recommend goods, and even automatic sentence completion before punctuation. AI has also changed the way we grow crops and raise animals, improving farm decision-making and automating various practices. For instance, unmanned aircraft vehicles (UAVs) or drones monitor the conditions of fields and crops. Data collected by a drone sensor is quickly transmitted and processed in a cloud to provide actionable information. Based on these monitored conditions, AI determines the schedules and rates for applying water, fertilizers, and herbicides. Robots can also find and harvest ripe fruit, while tractor-pulled implements scan vegetables, detect weeds, and apply agrochemicals precisely where needed. AI technology is making crop and livestock farming more efficient. AI evolves to become smarter and more autonomous as computing and sensing technologies continue to advance. In recent years, AI spread has shown a drastic increase in applications in agricultural production. This article is part of a series on AI applications, and its purpose is to outline basic concepts and principles of AI. The use of AI technologies is expected

to become more affordable and common in agriculture in the future. However, it is often difficult to understand terms and ideas introduced in communications, including news and research articles and conversations regarding AI applications in agriculture. Such difficulty may bring miscommunication or misperceptions about AI and its use. This article is prepared to help Extension agents, farmers, farm managers, researchers, graduate students, and the general public better understand AI, associated terms, and technologies. This effort is part of the UF/IFAS Department of Agricultural and Biological Engineering's Ask IFAS series on AI that discusses the broad aspects of AI in agriculture and natural resources management and facilitates communications between AI scientists and stakeholders including growers, ranchers, decision-makers at local and state governments, 4-H youth, and concerned residents.

What are machine learning and AI?

There are multiple definitions of AI depending on its application field. In general, AI can be defined as “the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings” (Cope-land 2023). Artificial intelligence (AI) is a broad field that includes machine learning (ML), as well as other subfields such as natural language processing, computer vision, and robotics. ML is broadly a branch of computer science that provides theoretical and algorithmic foundations for AI.

1. This document is AE605, one of a series of the Department of Agricultural and Biological Engineering, UF/IFAS Extension. Original publication date September 2024. Visit the EDIS website at <https://edis.ifas.ufl.edu> for the currently supported version of this publication. © 2024 UF/IFAS. This publication is licensed under [CC BY-NC-ND 4.0](https://creativecommons.org/licenses/by-nc-nd/4.0/).
2. Young Gu Her, associate professor, hydrology and water quality, Department of Agricultural and Biological Engineering, UF/IFAS Tropical Research and Education Center; Nikolay Bliznyuk, associate professor, statistics and machine learning, Department of Agricultural and Biological Engineering; Yiannis Ampatzidis, associate professor, precision agriculture engineering, UF/IFAS Southwest Florida Research and Education Center; Ziwen Yu, assistant professor, big data analytics, Department of Agricultural and Biological Engineering; and Haimanote Bayabil, associate professor, water resources, Department of Agricultural and Biological Engineering, UF/IFAS TREC; UF/IFAS Extension, Gainesville, FL 32611.

Loosely, ML is concerned with the development of algorithms by which a computer model “learns” about patterns and rules in data. Here, an algorithm is a set of instructions arranged to solve a particular problem such as delineating a face in a photo and identifying weeds among crops.

Running ML models may be viewed as a “black box” where its internal workings are not visible or known to the user. Users only see the input that goes into the box and the output it generates; the detailed functioning is either too complex to understand or is deliberately obscured. Users rely on these models to make predictions or decisions but may not fully understand how those decisions are made. Thus, the quality of ML model predictions heavily depends on how well the tuning parameters are calibrated or trained with the input data. For this reason, ML models are often criticized for their lack of transparency and interpretability.

AI usually includes processes for sensing and data acquisition and coupling of multiple ML models trained to perform individual tasks required to achieve its overall goal (Figure 1). For example, one model is trained to identify areas where crops do not grow well, from images taken by a drone overlooking the field. Then, another model predicts how much water or fertilizer should be applied to the areas to maintain the crop health based on the information obtained from the image. An AI system that integrates the two trained models can help improve crop productivity by providing information and implementing actions required to manage crop growth and yield.

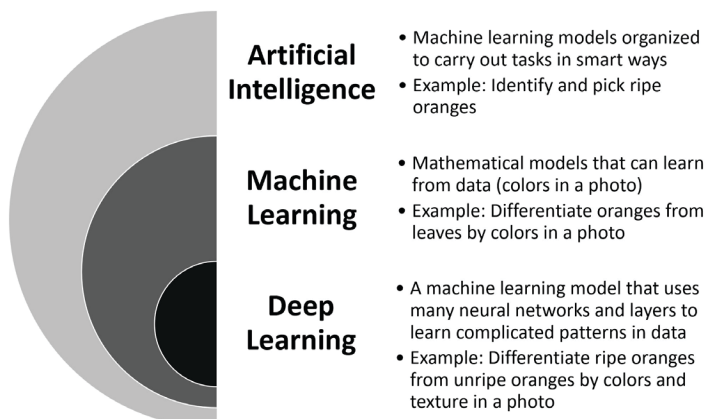


Figure 1. Relationship between artificial intelligence, machine learning, and deep learning. Credits: Young Gu Her, UF/IFAS

What are the kinds of machine learning?

Supervised Learning

There are two major types of ML methods largely determined by the types of available training data and/or data acquisition: supervised and unsupervised learning (Figure 2). In supervised learning, individual data points in a training dataset are labeled and divided into different groups based on the users’ understanding and knowledge of the dataset. Labeling means that each data point in the dataset is given a specific tag or category. For example, if we are classifying oranges, each orange and leaf (data point) might be labeled as “ripe orange,” “unripe orange,” and “leaf.” These labels help the machine learning model understand what each data point represents (Figure 3). Then, an ML algorithm tries to find the unique characteristics of the multiple groups. In this example, a ripe orange may have a bright and vibrant orange color, and others including its leaves and unripe oranges may have more green colors in a photo. A scientist delineates areas (or pixels) covered by the ripe orange and the others in the photo and feeds the delineated parts of the photo to an ML algorithm for its training. The ML algorithm learns the spectral (color) characteristics, such as the percentages of red, green, and blue colors, of the areas (or pixels) covered by the ripe oranges and the others in the photo (Figure 3). Based on the learned or trained spectral characteristics, the ML algorithm classifies all pixels of the entire photo.




Figure 2. Machine learning methods. Credits: Young Gu Her, UF/IFAS

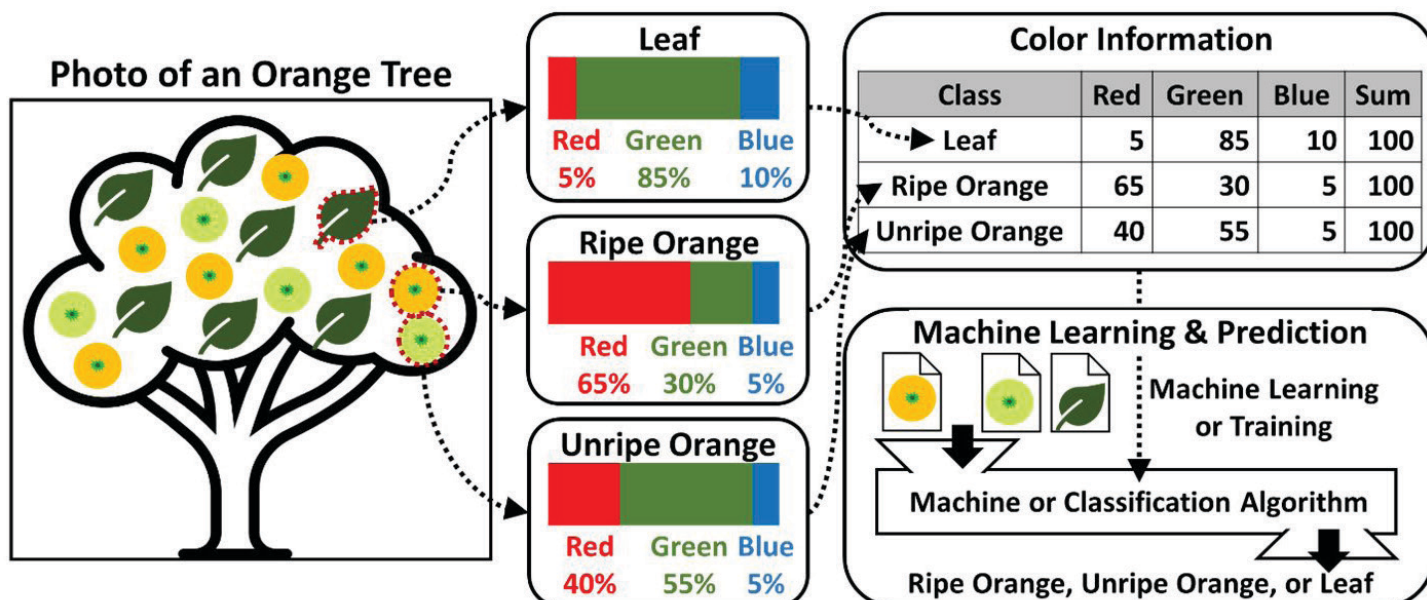


Figure 3. An example of supervised learning for identifying ripe oranges in a photo.
Credits: Young Gu Her, UF/IFAS

Unsupervised Learning

In unsupervised learning, on the other hand, the intervention of a user, such as the delineation of ripe and unripe oranges in a photo, is minimized. An ML algorithm trained by an unsupervised learning method clusters all pixels of a photo based on their spectral characteristics (Figure 2). The number of clusters can be one or as many as the number of pixels, depending on the spectral characteristics, which may be chosen by the user based on experience or understanding, or learned from preferences about within- and between-cluster similarity. In the example of the orange classification, unsupervised learning identifies all potential groups of pixels based on their color combinations, and the number of groups may be at least three (e.g., ripe oranges, leaves, and unripe oranges), and can be as many as five, ten, or thirty. An unsupervised learning process that does not require the intervention of a user is still desirable, and cumulative datasets and training assessments are expected to help make self-supervised training possible in the near future.

Reinforcement Learning

Reinforcement learning is a special type of supervised learning that works with labeled data. Reinforcement learning is also often considered a distinct category separate from supervised and unsupervised learning, as it does not require labeled data or a training set. In reinforcement learning, new data can be added to the training dataset and labeled in the middle of training, unlike in other supervised learning methods. Reinforcement learning learns from the training datasets obtained or updated from interacting with the environment; it receives feedback such as penalties

(negative feedback) and rewards (positive feedback) from the interaction. The interaction and feedback help ML more efficiently learn a system of interest. Reinforcement learning is mostly implemented in video games and robotics. Exposing the fundamental concept of reinforcement learning used in farming applications, many agriculture robots learn from mistakes such as colliding with obstacles or failing to pick fruit through penalty scores. At the same time, reinforcement learning helps robots figure out the shortest path to bypass obstacles or grab a fruit with the minimum number of motions, through rewarding optimum practices.

What are artificial neural networks and deep learning?

Artificial Neural Networks (ANNs)

An ANN is a way of learning structures or patterns in data by mimicking how brain neurons signal to each other to process information. An ANN is a kind of supervised learning method, requiring training datasets (or reference input and output data) and user intervention. Specifically, an artificial neuron (perceptron) takes (many) input datasets (e.g., signals to the human brain), weighs them based on their contributions to getting the right answers, and passes them to other neurons in the network. An ANN can modify the artificial networks by updating or changing the weights as it learns more about a system of interest from a series of runs (or comparing the previous outputs with the reference data of a training dataset). Such a process is often called “backpropagation” as errors in the outputs are propagated (or reapportioned) back to the weights in the neural networks. In data-rich scenarios, ANN learning is often

capable of solving more complex problems (e.g., image recognition) compared to other ML learning methods. The number of nodes and the complexity of connections are dependent on the characteristics of a problem of interest, computation capacity, and training efficiency. ANNs have been widely used for computer vision, speech recognition, and natural language processing. Given its scalability in computation and complexity, the evolution of ANN in solving various problems has inspired the development of deep learning.

Deep Learning

In a simple ANN, there is a hidden layer between the input layer and the output layer. The input layer contains neurons (units, nodes, or perceptions) that take information from a training dataset, while the output layer provides the results after processing the data. Neurons in the hidden layer receive weighted input data from the neurons in the input layer. These hidden neurons then decide whether and to what extent to consider the input neurons based on their influence on the outputs. The mathematical function used to determine the importance or influence of the input neurons is called an activation function. Input neurons with importance above a certain threshold will be considered in calculating the output.

The number of hidden layers required to accurately represent the connections between input and output neurons increases with the complexity of the connections. A complex relationship between inputs and outputs requires many hidden layers. Deep learning methods were proposed to employ several hidden layers and thus handle more complex relationships. The deep learning method is a special type of ANN learning method with many hidden layers so that it can learn much more complex data and map its networks. On many data-rich problems, this technique offers outstanding performance in terms of feature extraction and prediction accuracy. Deep learning is widely used for image processing, with promising results for addressing farming problems such as plant disease diagnosis and pesticide recommendations. There are several types of neural network structures for deep learning methods, including convolutional neural networks (CNN) and recurrent neural networks (RNN) including its variants such as long short-term memory (LSTM) networks. CNN are used for image classification and object detection, such as ripe orange identification (Figure 4). CNN handles two-dimensional data such as a photo, a two-dimensional array of pixels that have spectral information like the depths of red, green, and blue colors. On the other hand, RNN and LSTM were developed to work with one-dimensional (or

sequential) data such as words, audio signal, and daily air temperature time series. LSTM is an advanced version of RNN with the capacity of handling long-term dependencies such as context and storyline in sequential data, series of words in a paragraph, or even a book. More details of the neural networks can be found in the following sources: neural network (<https://rb.gy/u465qi>), ANN (<https://rb.gy/3yg9ak>), and CNN (<https://rb.gy/pqzq92>).

Recent Advances in Deep Learning

Other ML methods stemming from deep learning such as transfer learning are also playing an important role in disease detection and are expected to advance decision-making by multiple factors such as environmental conditions, harvesting practices, financial needs, soil characteristics, or water availability. Transfer learning is a type of deep learning, and it employs an ML model trained for a task with a set of data to develop a new ML model for another task. For example, an ML model trained with a set of photos of orange trees (e.g., Figure 3) can be a good starting point for the development of another ML model for identifying ripe grapefruits and tangerines, even though the orange ML model should be adjusted or fine-tuned to the new task. Usually, overcoming the challenge of training from scratch of an image classifier (or ML algorithm) for distinguishing healthy and infected plant leaves requires a massive amount of data, substantial computing resources, and considerable human effort. Instead, researchers build models, trained on large image datasets like ImageNet, COCO, and Open Images, and share these models with the general public for reuse. Transfer learning is effective for image classification problems (Espejo-Garcia et al. 2020).

Another technique that is expected to be relevant in the near future is the generative adversarial networks (GANs). GANs are a type of ML model that allows neural networks to generate synthetic data similar to the original one. This type of network can address the problem of data scarcity within agricultural computer vision. The main advantage of GANs is that they can create pictures with synthetic “real” crops instead of just rotating or adding noise to existing ones. This synthetic data could be used to improve the generalization ability of a deep learning model that obtains poor results at disease detection tasks, due to the constraints in the size of the original dataset and the limitations of traditional techniques.

With a similar architecture but a different algorithm, a diffusion model is a different type of generative AI model that generates data such as images. Inspired by physical diffusion processes, diffusion models are used to generate data

by gradually adding noise to and removing it from a sample to structure the process of generating synthetic data to mimic real-world conditions. Potential applications include generating synthetic images of crops, pests, diseases, and other agricultural elements to augment existing training datasets for AI models. Diffusion models can enhance and generate high-resolution images from lower-quality inputs to provide farmers with clearer and more detailed visual information. The models also can simulate realistic images of crop growth stages impacted by different variables such as weather, soil nutrients, and irrigation patterns to help farmers understand the potential impacts of extreme weather events and different management practices.

Examples of ML Applications

Two examples are provided here to help the reader appreciate how ML and AI are applied to solve agricultural problems. There are two types of data: quantitative and categorical data. Quantitative data is expressed with numbers and directly allows algebraic operations on them; these are typically a result of a numerical measurement, such as the height, width, and weight of a plant. Categorical data, also known as qualitative, is a type of data that can be labeled with natural language and used to describe categories, classes, clusters, and groups. The types of lawn grass, such as bahiagrass, bermudagrass, fescue, St. Augustinegrass, and carpetgrass, are an example of categorical data. Depending on the data types, supervised ML can be subdivided into regression and classification.

The first example is to determine (or predict) the presence or absence of apple scab disease (categorical data) from a photo taken above an apple tree (Figure 4). Detecting apple scab disease is crucial for maintaining healthy apple orchards, ensuring economic viability for growers, and providing high-quality fruit to consumers. First, an ML modeler selects groups of pixels covering parts of healthy and infested leaves and labels them with 0 (healthy) or 1 (infested). Then, an ML model investigates the spectral characteristics such as the combinations of blue, green, and red colors of the labeled pixel groups. Using statistical techniques, the model learns the common spectral features shared among each group of pixels, allowing it to identify patterns and classify similar groups in the dataset. The spectral characteristics of pixels are not the only information ML models can use but also shape (of a group of pixels), texture (the spatial arrangement of colors and/or intensities), and the spatial relationships with the surrounding pixels. Considering that apple scabs appear as visible color anomalies on the leaves, the major differences

will be in color features during the classification process in the example.

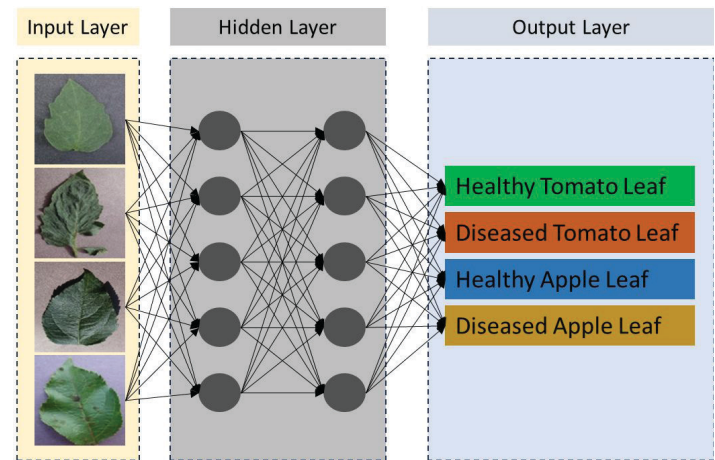


Figure 4. An application of deep learning to the image classification of apple and tomato leaves.

Credits: Young Gu Her and Yiannis Ampatzidis, UF/IFAS

The results of ML modeling in identifying diseased plants, such as those with apple scab, provide farmers with precise, actionable information that can enhance their ability to manage crops effectively, reduce costs, improve yield and quality, and support sustainable agricultural practices. By accurately identifying diseased plants, for example, farmers can apply pesticides specifically to the affected areas rather than treating the entire orchard. This targeted approach reduces the overall use of chemicals, lowering costs and minimizing environmental impact. In addition, early detection through ML modeling allows farmers to take timely actions, such as removing infected leaves or fruits, applying fungicides, or implementing other control measures before the disease spreads. This helps in managing the disease more effectively and prevents large-scale outbreaks.

Another example of an ML application is predicting monthly streamflow based on past observations (Figure 5). Monthly streamflow forecasts are crucial for making efficient environmental and economic decisions, such as managing water supplies and planning for floods or droughts. To make these predictions, all available data up to the current month is used to predict the streamflow for the next month. If predictions further into the future are needed, this one-month-ahead prediction process can be repeated. By leveraging historical streamflow data and climate indicators (like the El Niño-Southern Oscillation indices), statistical and ML approaches can build data-driven models to forecast monthly streamflow. For instance, a simple yet effective model might predict the current month's streamflow as a linear function of the previous month's streamflow. The accuracy of this prediction depends on the correlation between past and present streamflow measurements—the

stronger the correlation, the more accurate the prediction. Due to the significant variability in streamflow over time and across different locations, a promising strategy is to develop station-specific ML models trained with data from each particular station. Training involves estimating model coefficients and fine-tuning parameters using a validation set to avoid overfitting. Overfitting occurs when a model performs exceptionally well on training data but fails to generalize to new, unseen data. Prediction accuracy is typically measured using the root-mean-squared-error between predictions and actual observations. Many statistical ML models provide not only point predictions (such as the average streamflow in cubic feet per second) but also uncertainty estimates, allowing for confidence or prediction intervals to be derived.

Concluding Remarks

ML and AI technologies are bringing the Fourth Industrial Revolution into agriculture and providing opportunities to further improve agricultural productivity and sustainability. In the era of digital agriculture, ML and AI are the definitive tools for smart farming and precision agriculture. The ML and AI tools are allowing us to do things that we could never do before. A variety of ML and AI technologies will be created, used, and communicated. ML and AI technologies heavily rely on advances in computer hardware, network technologies, and computational techniques, and thus it is often challenging even for scientists to clearly understand the mechanisms and features and to follow the technology trends. The adoption and use of the new technologies may not require knowing all the details of ML and AI technologies, but a good understanding of the benefits and risks ML and AI may bring to one's farm business. The use of ML and AI in agriculture is expected to be similar to buying and driving a car in the future; one can buy and drive a car without knowing what is under the hood. Further details of ML and AI can be found in the UF/IFAS Department of Agricultural and Biological Engineering AI series.

References

Copeland, B. 2023. "artificial intelligence." *Encyclopedia Britannica*. <https://www.britannica.com/technology/artificial-intelligence>

Espejo-Garcia, B., N. Mylonas, L. Athanasakos, S. Fountas, and I. Vasilakoglou. 2020. "Towards Weeds Identification Assistance through Transfer Learning." *Computers and Electronics in Agriculture* 171:105306. <https://doi.org/10.1016/j.compag.2020.105306>

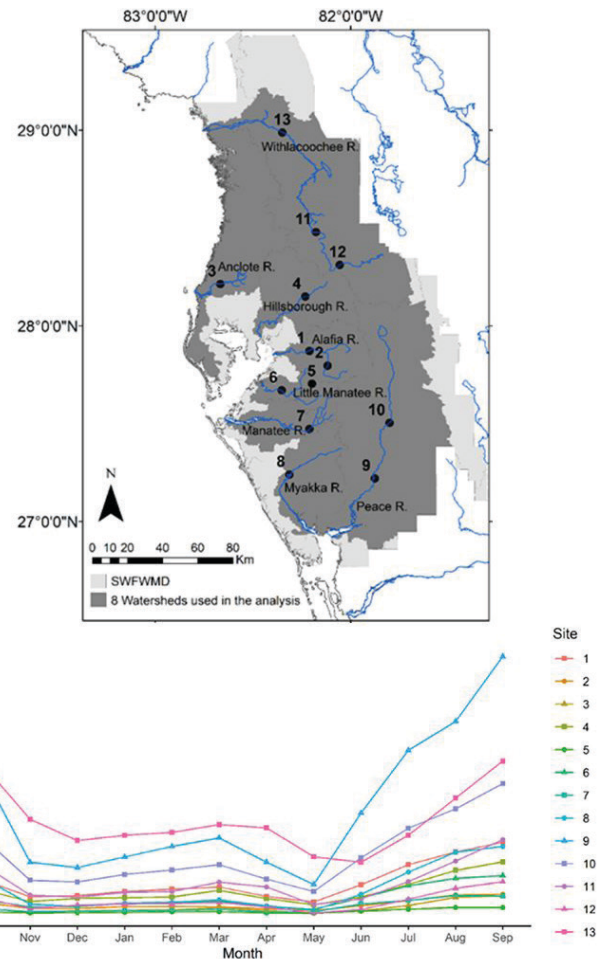


Figure 5. Top: The locations of United States Geological Survey (USGS) streamflow monitoring stations in west-central Florida. Bottom: Average monthly streamflow predicted using statistical and ML approaches at the USGS monitoring stations. Credits: Jia-Yi Ling and Nikolay Bliznyuk, UF/IFAS