

# Avocado Classification Using Deep Learning

Mohammed N. Abu Alqumboz and Samy S. Abu-Naser

Department of Information Technology, Faculty of Engineering and Information Technology,  
Al-Azhar University - Gaza, Palestine  
[m.naji1512@gmail.com](mailto:m.naji1512@gmail.com)

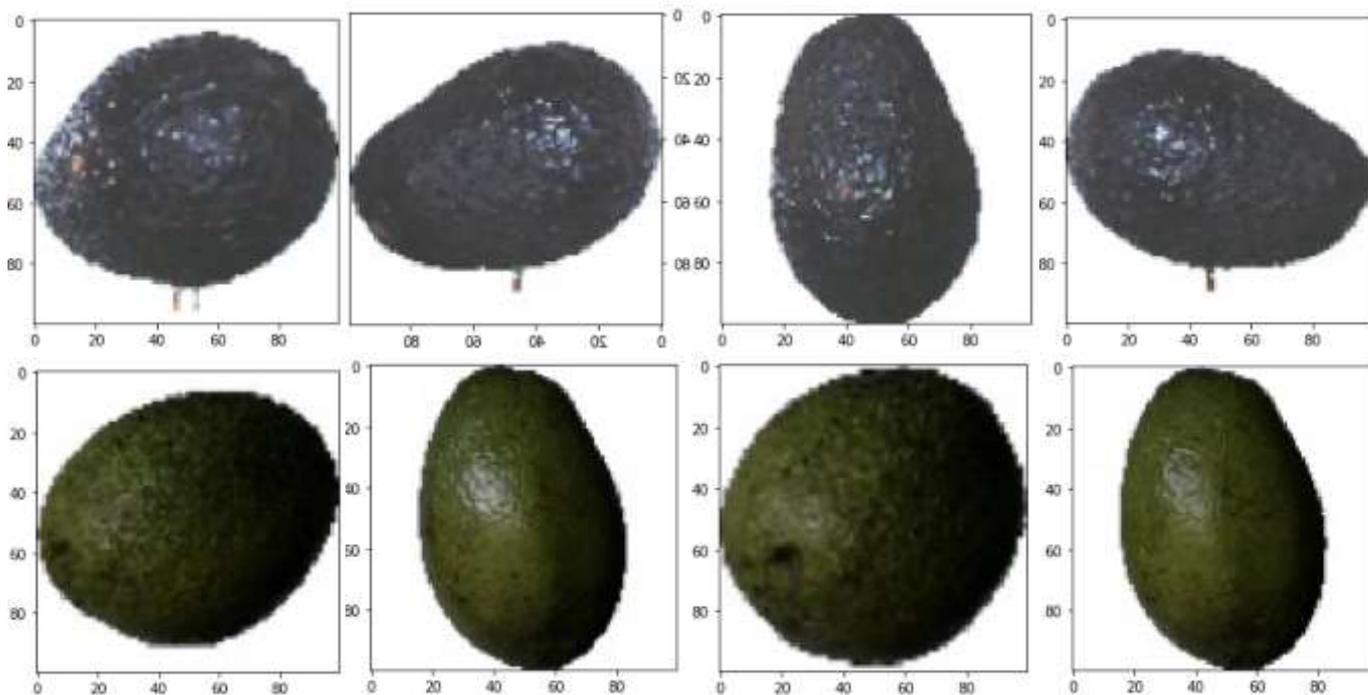
**Abstract:** Avocado is the fruit of the avocado tree, scientifically known as *Persia Americana*. This fruit is prized for its high nutrient value and is added to various dishes due to its good flavor and rich texture. It is the main ingredient in guacamole. These days, the avocado has become an incredibly popular food among health-conscious individuals. It's often referred to as a superfood, which is not surprising given its health properties. Using a public dataset of 1,234 images of Avocado collected under controlled conditions, we trained a deep convolutional neural network to identify two types of avocado. The trained model achieved an accuracy of 99.84% on a held-out test set, demonstrating the feasibility of this approach. Overall, the approach of training deep learning models on increasingly large and publicly available image datasets presents a clear path toward types of avocado.

**Keywords:** Avocado, Classification, Deep Learning

## INTRODUCTION

Deep Learning has grown hand-in-hand with the digital era, which has conveyed about an explosion of data in all forms and from every area of the world. This data, recognized as Big Data, is pinched from sources like social media, search engines, e-commerce platforms and more. This huge amount of data is freely accessible and can be shared through fintech applications like cloud computing. Though, the data, which normally is unstructured, is so massive that it could take years for humans to understand it and extract pertinent information. Companies understand the unbelievable potential that can result from disentangling this wealth of information, and are progressively adapting to Artificial Intelligence systems for automated support. Deep learning, a division of machine learning, uses a hierarchical level of artificial neural networks to perform the process of machine learning. The artificial neural networks are constructed like the human brain, with neuron nodes linked together like a web. While traditional programs build to do analysis with data in a linear way, the hierarchical task of deep learning systems allows machines to process data with a nonlinear approach. A traditional approach to detecting fraud or money laundering might depend on the amount of transaction that precedes, while a deep learning nonlinear technique would include geographic, IP address, time, location, type of retailer and any other feature that is likely to indicate a fraudulent activity. The first layer of the neural network processes a raw data input like the amount of Deep learning is used across all industries for a number of different tasks. Commercial apps that use image recognition, open source platforms with consumer recommendation apps and medical research tools that explore the possibility of reusing drugs for new ailments are a few of the examples of deep learning incorporation [1,2,3].

## DATASET



The dataset used, provided by ImageNet, contains a set of 1,234 images of approximately 960 unique plants belonging to tow type of avocado. The images were resized into 100×100 for faster computations but without compromising the quality of the data.

## CONVOLUTIONAL NEURAL NETWORKS

In machine learning, a Convolutional Neural Network (CNN, or ConvNet) is a class of deep, feed-forward artificial neural networks, most commonly applied to analyzing visual imagery.

CNNs use a variation of multilayer perceptron's designed to require minimal preprocessing. They are also known as shift invariant or Space Invariant Artificial Neural Networks (SIANN), based on their shared-weights architecture and translation invariance characteristics.

Convolutional networks were inspired by biological processes in that the connectivity pattern between neurons resembles the organization of the animal visual cortex. Individual cortical neurons respond to stimuli only in a restricted region of the visual field known as the receptive field. The receptive fields of different neurons partially overlap such that they cover the entire visual field.

CNNs use relatively little pre-processing compared to other image classification algorithms. This means that the network learns the filters that in traditional algorithms were hand engineered. This independence from prior knowledge and human effort in feature design is a major advantage.

They have applications in image and video recognition, recommender systems and natural language processing[4].

### Design

A CNN consists of an input and an output layer, as well as multiple hidden layers. The hidden layers of a CNN typically consist of convolutional layers, pooling layers, fully connected layers and normalization layers [1-5].

Description of the process as a convolution in neural networks is by convention. Mathematically it is a cross- correlation rather than a convolution. This only has significance for the indices in the matrix, and thus which weights are placed at which index.

### Convolutional

Convolutional layers apply a convolution operation to the input, passing the result to the next layer. The convolution emulates the response of an individual neuron to visual stimuli.

Each convolutional neuron processes data only for its receptive field.

Although fully connected feedforward neural networks can be used to learn features as well as classify data, it is not practical to apply this architecture to images. A very high number of neurons would be necessary, even in shallow (opposite of deep) architecture, due to the very large input sizes associated with images, where each pixel is a relevant variable. For instance, a fully connected layer for a (small) image of size 100 x 100 has 10000 weights for each neuron in the second layer. The convolution operation brings a solution to this problem as it reduces the number of free parameters, allowing the network to be deeper with fewer parameters. For instance, regardless of image size, tiling regions of size 5 x 5, each with the same shared weights, requires only 25 learnable parameters. In this way, it resolves the vanishing or exploding gradients problem in training traditional multi-layer neural networks with many layers by using backpropagation[6-7].

### Pooling

Convolutional networks may include local or global pooling layers[7], which combine the outputs of neuron clusters at one layer into a single neuron in the next layer. For example, max pooling uses the maximum value from each of a cluster of neurons at the prior layer. Another example is average pooling, which uses the average value from each of a cluster of neurons at the prior layer[11-15].

### Fully Connected

Fully connected layers connect every neuron in one layer to every neuron in another layer. It is in principle the same as the traditional multi-layer perceptron neural network (MLP)

### Receptive Field

In neural networks, each neuron receives input from some number of locations in the previous layer. In a fully connected layer, each neuron receives input from every element of the previous layer. In a convolutional layer, neurons receive input from only a restricted subarea of the previous layer. Typically the subarea is of a square shape (e.g., size 5 by 5). The input area of a neuron is called its receptive field. So, in a fully connected layer, the receptive field is the entire previous layer. In a convolutional layer, the receptive area is smaller than the entire previous layer.

## Weights

Each neuron in a neural network computes an output value by applying some function to the input values coming from the receptive field in the previous layer. The function that is applied to the input values is specified by a vector of weights and a bias (typically real numbers). Learning in a neural network progresses by making incremental adjustments to the biases and weights. The vector of weights and the bias are called a filter and represents some feature of the input (e.g., a particular shape). A distinguishing feature of CNNs is that many neurons share the same filter. This reduces memory footprint because a single bias and a single vector of weights is used across all receptive fields sharing that filter, rather than each receptive field having its own bias and vector of weights.

## METHODS

We experimented with two types of images to see how the model work and what exactly it learns, first we take the image as it is with 3 color channels, and then we experimented with 1 color channel images (Gray-Scale). And as expected the model learns different patterns in each approach.

## Data augmentation

In order to make the most of our few training examples and increase the accuracy of the model, we augmented the data via a number of random transformations. The selected data augmentation techniques were: size re-scaling, rotations of 40, horizontal shift, image zooming, and horizontal flipping. Furthermore, it is expected that data augmentation should also help prevent overfitting (a common problem with small datasets, when the model, exposed to too few examples, learns patterns that do not generalize to new data) and, for this reason, improving the models ability to generalize.

## MODEL

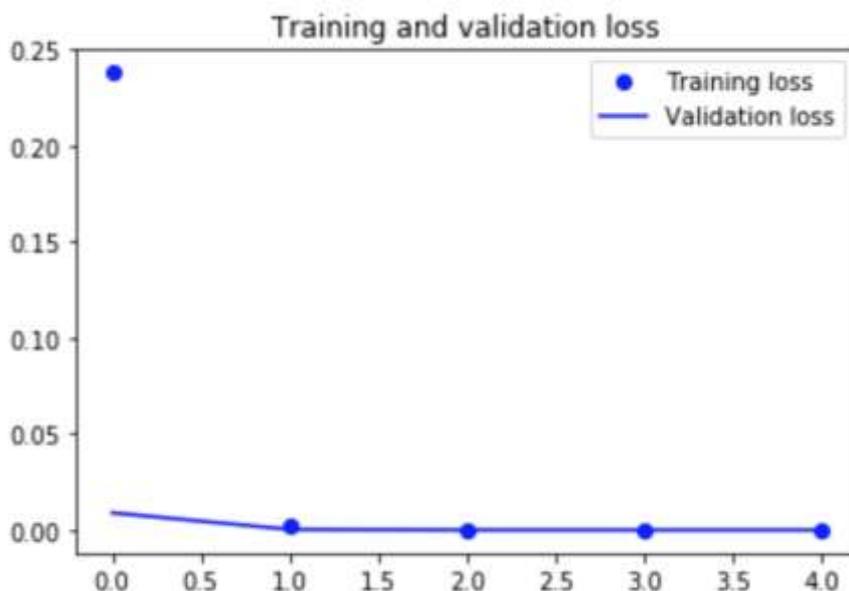
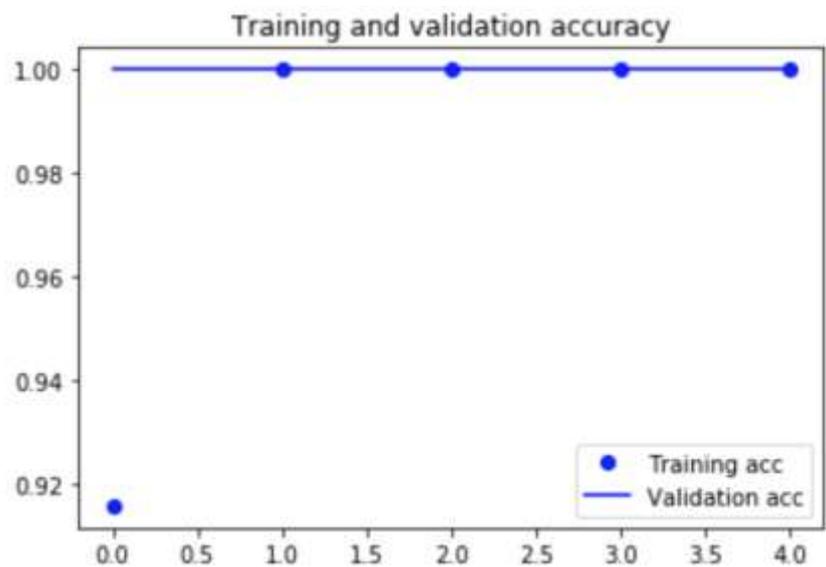
Our model takes raw images as an input, so we used Convolutional Neural Networks (CNNs) to extract features.

Model: "sequential\_1"

Layer (type)	Output Shape	Param #
conv2d_1 (Conv2D)	(None, 148, 148, 32)	896
max_pooling2d_1 (MaxPooling2D)	(None, 74, 74, 32)	0
conv2d_2 (Conv2D)	(None, 72, 72, 64)	18496
max_pooling2d_2 (MaxPooling2D)	(None, 36, 36, 64)	0
conv2d_3 (Conv2D)	(None, 34, 34, 128)	73856
max_pooling2d_3 (MaxPooling2D)	(None, 17, 17, 128)	0
conv2d_4 (Conv2D)	(None, 15, 15, 128)	147584
max_pooling2d_4 (MaxPooling2D)	(None, 7, 7, 128)	0
flatten_1 (Flatten)	(None, 6272)	0
dense_1 (Dense)	(None, 512)	3211776
dense_2 (Dense)	(None, 1)	513
Total params: 3,453,121		
Trainable params: 3,453,121		
Non-trainable params: 0		

**DATA VISUALISATION**

To see how the model works and what exactly learns we choose to visualize intermediate activations that consists of displaying the feature maps that are output by various convolution and pooling layers in a network, given a certain input (the output of a layer is often called its activation, the output of the activation function). This gives a view into how an input is decomposed into the different filters learned by the network



**CONCLUSION**

we have designed and implemented a two-class classifier that takes avocado images with 2 different species as input, builds a model using deep learning convolutional neural networks, and uses this model to predict the type of (previously unseen) images of avocado.

The proposed approach achieves promising results – most notably, validation accuracy of 100% .

## References

1. Werbos, P.J. (1975). Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences.
2. <https://towardsdatascience.com/types-of-machine-learning-algorithms-you-should-know-953a08248861>
3. <https://www.sciencedirect.com/topics/computer-science/supervised-learning>
4. <https://whatis.techtarget.com/definition/unsupervised-learning>
5. Deng, Li, and Dong Yu. "Deep learning: methods and applications." *Foundations and Trends® in Signal Processing* 7.3–4 (2014): 197-387.
6. Liu, Ziwei, Ping Luo, Xiaogang Wang, and Xiaoou Tang. "Deep learning face attributes in the wild." In *Proceedings of the IEEE international conference on computer vision*, pp. 3730-3738. 2015.
7. Bengio, Yoshua. "Deep learning of representations for unsupervised and transfer learning." In *Proceedings of ICML workshop on unsupervised and transfer learning*, pp. 17-36. 2012.