

Land-Use Classification via Transfer Learning with a Deep Convolutional Neural Network

Chu-Yin Weng

Jericho High School, New York, USA

Email: cyweng03@gmail.com

How to cite this paper: Weng, C.-Y. (2022) Land-Use Classification via Transfer Learning with a Deep Convolutional Neural Network. *Journal of Intelligent Learning Systems and Applications*, 14, 15-23. <https://doi.org/10.4236/jilsa.2022.142002>

Received: February 9, 2022

Accepted: May 24, 2022

Published: May 31, 2022

Copyright © 2022 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Land cover classification provides efficient and accurate information regarding human land-use, which is crucial for monitoring urban development patterns, management of water and other natural resources, and land-use planning and regulation. However, land-use classification requires highly trained, complex learning algorithms for accurate classification. Current machine learning techniques already exist to provide accurate image recognition. This research paper develops an image-based land-use classifier using transfer learning with a pre-trained ResNet-18 convolutional neural network. Variations of the resulting approach were compared to show a direct relationship between training dataset size and epoch length to accuracy. Experiment results show that transfer learning is an effective way to create models to classify satellite images of land-use with a predictive performance. This approach would be beneficial to the monitoring and predicting of urban development patterns, management of water and other natural resources, and land-use planning.

Keywords

Land-Use Classification, Machine Learning, Transfer Learning, Convolutional Neural Network

1. Introduction

In recent decades, the human population has more than doubled and is projected to continue to grow at the same rate. The expansion of infrastructure and agriculture required by this population growth has increased the pace of land degradation. This current rate of land transformation is unsustainable. To control this issue, land-use classification can regulate the monitoring and predict of urban development patterns, management of water and other natural resources,

and land-use planning and regulation. Therefore, the creation of an efficient and accurate means of classifying land use is important. While it is possible to classify manually, it is not a plausible means of classification due to its time-consuming and costly nature. This creates a necessity for an inexpensive and efficient means of land-use classification. The application of machine learning to this issue would provide a good solution as it trains computers to make decisions without explicit programming. Specifically, since land-use classification is done through satellite images, we should utilize deep learning and convolutional neural networks, which are designed for working with images. An issue with deep learning is that the training of such an algorithm requires a large amount of data. In this experiment, the data used are satellite images, which are rarely released for public use. The images readily available are often either outdated or of low quality. These types of data are not ideal for the use of image classification because of their negative effect on the performance of the models. An efficient way to solve this issue is with the use of transfer learning, which is a machine learning method where a pre-trained model is borrowed as the starting point for modeling a new task. The use of transfer learning would not only decrease the amount of data required for us to train the model, but also allow us to create an image classification model within a short time period. In this paper, we explore the possibility of using transfer learning to create an accurate model that classifies types of land use. Specifically, we utilized the ResNet-18 Convolutional Neural Network and tested the effects of different training dataset sizes as well as the number of training periods. As supported by our findings, transfer learning can be used to create an accurate land-use classification model.

2. Background

Our approach employs techniques from machine learning, a subfield of artificial intelligence. In this section, we briefly review relevant background material on these topics before describing our application of these techniques to land use classification.

Fundamental Principles:

Artificial intelligence (AI) is a broad field that refers to the ability of computers to complete tasks that typically require human intelligence. The goal of AI is to create machines that are ultimately able to mimic human behavior. An important subfield of artificial intelligence used within this experiment is machine learning, which further specifies deep learning. Machine learning is a field of AI that deals with machines that are able to learn and make predictions based on data processed [1]. Therefore, machine learning allows for machines to make decisions with higher accuracy by processing large amounts of data. One technique that has been shown to obtain high performance is the use of neural networks. Neural networks are composed of nodes that mimic the functions of human brains. Each node is connected to another and has an associated weight. Deep Learning algorithms consist of neural networks with multiple layers. The

layered network enables the computer to process and gathers progressively more information from the inputs as data is passed through the layers.

There are different types of neural networks, and each is used for different purposes. The two most common types are the multilayer perceptron (MLP), and the convolutional neural networks (CNN). The MLP has a feedforward architecture and is one of the most commonly used neural network architectures.

When it comes to image classification, CNNs are more commonly used. This is a favored algorithm over others due to its property of spatial invariance, which refers to its ability to recognize features in an image even if the feature does not look exactly the same as the images used during training [2]. Currently, convolutional neural network models have been applied to multiple fields, including image classification, expression recognition and more. CNN's typically consist of three types of layers, convolutional layer, pooling layer, and fully-connected layer. This is demonstrated in **Figure 1**. With each progressing layer, the CNN increases in complexity and starts to identify different levels of features from lines and corners to objects [1] [2] [3]. The convolutional layer is where the majority of calculations take place within the network [4]. When an image is processed, its RGB values (which represent the intensity of the colors red, green, and blue) are translated and combined into tensors for faster and more efficient calculations. Then, as shown in **Figure 2**, the image first gets padded by blank values around it to ensure that no pixels are passed over less than the others [2].

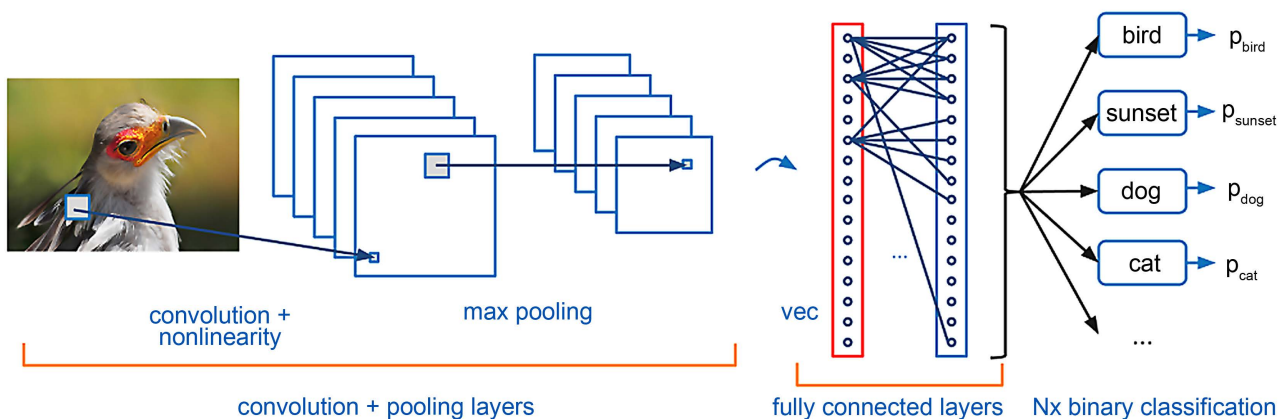


Figure 1. A representation of the layers within a CNN.




Figure 2. A visual representation of padding.

A convolution is when a filter is passed through the image, checking if the feature is present. In different algorithms, there can be different amounts of filters used, and strides of the filters. After each convolution operation, a CNN applies a Rectified Linear Unit (ReLU), which is a piecewise linear function that will output the input directly if it is positive; otherwise, it will output zero [2] [4]. The ReLU function is a popularly used activation function for many types of neural networks because a model that uses it is easier to train and often achieves better performance. After the convolutional layers, the CNN has pooling layers that reduce the number of parameters from the input [2]. This is to decrease the computational power required to process the data while also forcing the neural net to the abstract critical information needed by future layers [3]. Finally, there is a fully-connected layer where each node is connected to a node in the previous layer [5]. This is the layer where classification is done through the features extracted from the previous layers [4]. A full pass of the dataset through the algorithm is called an epoch, and is sometimes referred to as an iteration. Within each epoch, the weights of the algorithm are changed to better fit the dataset [5]. **Figure 3** demonstrates the layers of an epoch.

As mentioned before, the training of a deep learning model requires a large amount of data. Transfer learning addresses this issue because it utilizes pre-trained models with a large dataset, thereby reducing the time and resources needed to train a new model. This method speeds up the process of creating a model for the second task and increases the precision of said model because the pre-trained model allows for the reuse of learned features, giving it a better starting point for the search

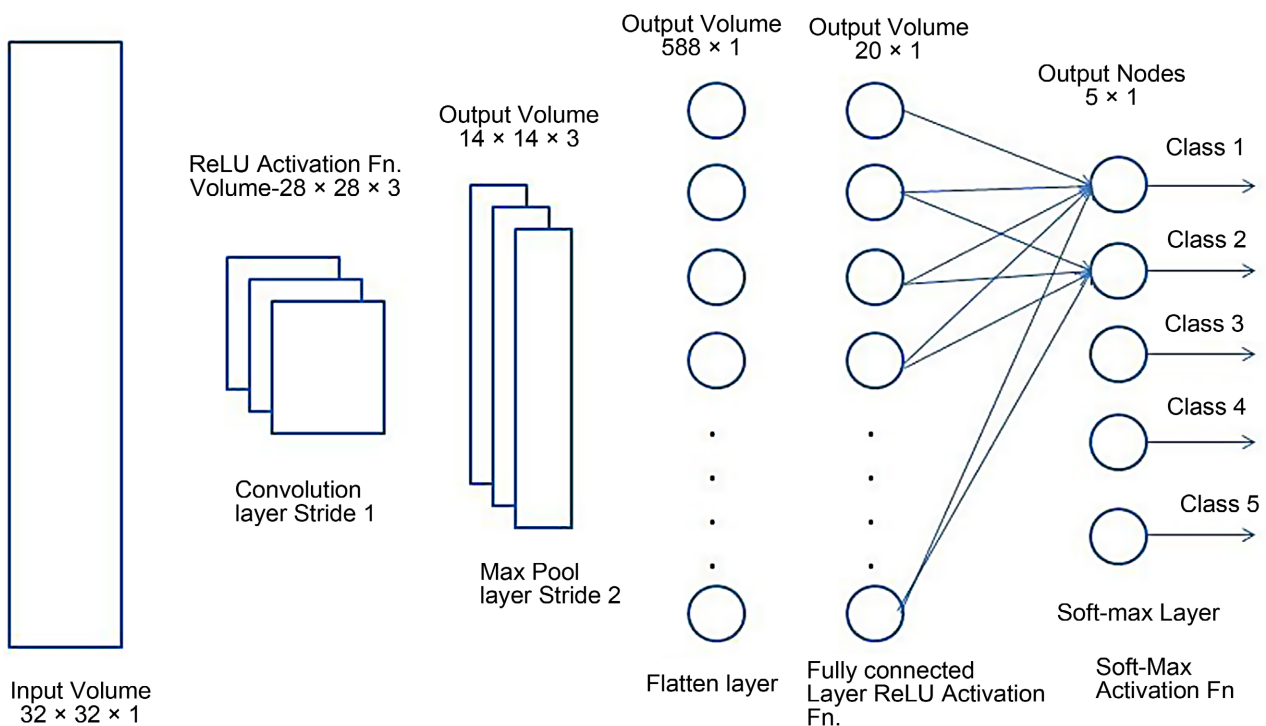


Figure 3. A visual representation of fully connected layers.

process. This is especially useful when there isn't enough data to train a new model. The key to transfer learning is the common features used within image recognition models. The ability to reuse these features means that the trained network can in some form be repurposed for a new problem. It is very common for this method to be used in models trained to take image inputs. In the case of image recognition, models with high precision require the model to have processed a large amount of data to effectively learn the importance of certain features.

3. Approach

This experiment used a pre-existing model, ResNet-18, which is an 18-layer deep model pre-trained with more than a million images from the ImageNet database. The previous training done with these images allows it to identify general features and classify general images, decreasing the amount of time required for us to create a relatively accurate model. This ability is applied in this experiment as we modify the weights to fit our specific satellite images. Our dataset composed of 21 classes of images is acquired from UC Merced, created by professor Shawn D. Newsam [6].

After loading the dataset, we split the data randomly into three: training dataset, validation dataset, and testing dataset [7]. From the pre-trained model, we modified it by removing its final classification layer, which was then replaced with a 21-element new classification layer and with all of the other weights kept as initialized. We then used backpropagation on the training data in order to train up the final classification layer, as well as fine tune all of the other weights. Originally, we trained the model with 75% of the data for 25 epochs, and during these epochs, we tracked the accuracy of each model and reported the model that provided the best accuracy out of all of the epochs. The purpose of splitting into the validation dataset and the testing dataset is to make sure that when we choose the best model from the validation dataset, the final model is tested by a fresh set of images set aside in the validation dataset beforehand; therefore yielding the true accuracy of the model.

We first normalized the images, which shifts and rescales the pixel values across images to make them easier to learn, thus simplifying the calculation process [7]. The images are augmented, randomly resized and flipped before normalization and manipulation. Then we created a function to train the model for the classification of the data. As the model is trained, the best learning rate and the most accurate model are saved. For each epoch, we implemented a training and validation phase that iterates through the datasets for each of the phases. Before continuing onto the next iteration of data, we accumulate the gradients over each minibatch and then perform a step of gradient descent to reduce the predictive error [7]. Next, the forward pass of the model is started to calculate the values of the layers from the inputs, and the loss of the model during this iteration is calculated. After each pass completes, we check if it has the highest accuracy thus far, and if it does, backward descent takes place and opti-

mizes the model. After each epoch, we calculated the loss and accuracy.

4. Experiments and Results

We obtained and examined the accuracy of the data processed by the machine learning model we created. The dataset used consists of 21 classes, each with 100 images. Examples of these classes include agriculture, forest, and river. Example images are shown in **Figure 4**. The experiment also tested the effects of training data size and epoch length on the accuracy of the model by training the model with different numbers of epochs. We collected the accuracy of the models from each of the epochs trained, which show us the trends as values were modified. For a more accurate result, we ran each model five times, and the accuracies of each epoch were averaged to show the general trend between accuracy throughout the iterations.

Within each of the experiments, the training, validation, and testing accuracies increase significantly in the first few epochs. After this initial increase, the accuracies increase by smaller margins.

We first tested the effects of changing the training data size of the model. As shown in **Figure 5**, the experiment showed a positive correlation of accuracy and the size of the training data. For example, the accuracy increases from the range between 0.75 and 0.80 to above 0.90 (Shown in **Figure 5**). This can be explained because a larger training dataset exposes the model to a greater variety of features needed to identify the objects within each image. With each of the models, the training accuracy starts below the validation and testing accuracies.

In a second experiment, we evaluated the effects of different numbers of epochs on the accuracies of the model. As shown in **Figure 6**, the experiment showed a positive correlation between accuracy and the number of epochs used to train the model. This result is because through the use of more epochs, the model is able to best fine tune the weights of the original image recognition model to



Figure 4. Examples of classes in the data set used for this experiment.

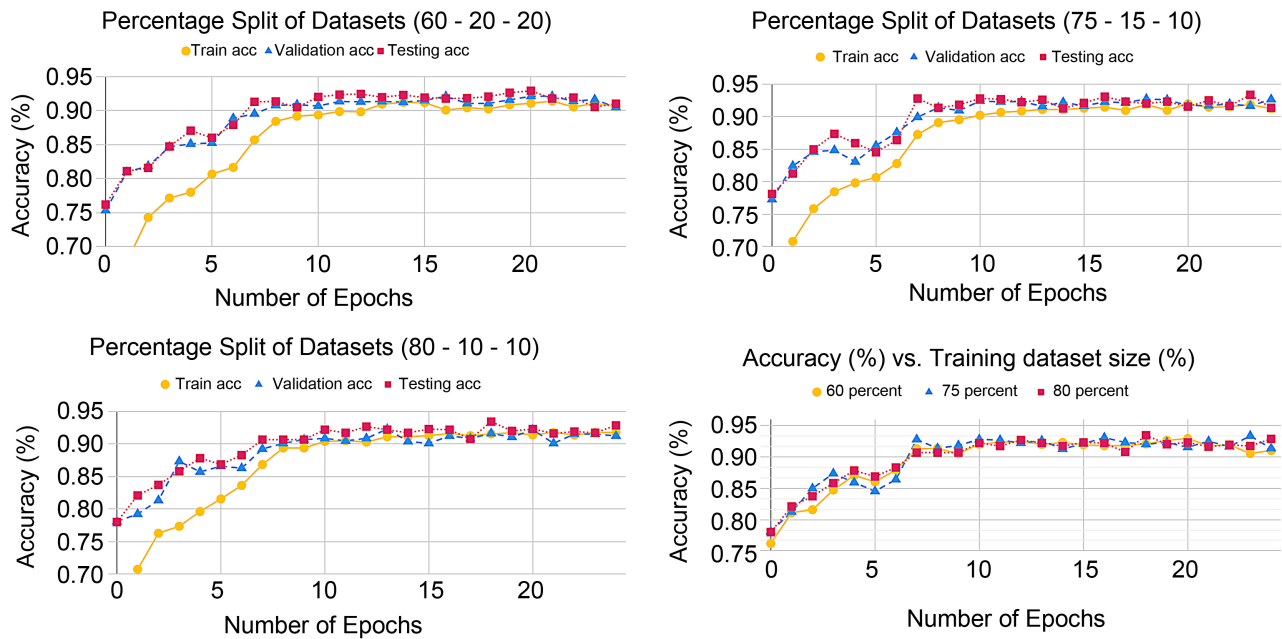


Figure 5. Accuracies of the model with different training data sizes. Models trained with larger training data sizes are more accurate than those that are trained with a smaller training data size. As shown in the graphs, the numbers within the parenthesis represent the percentage of data used for training, validation, and testing.

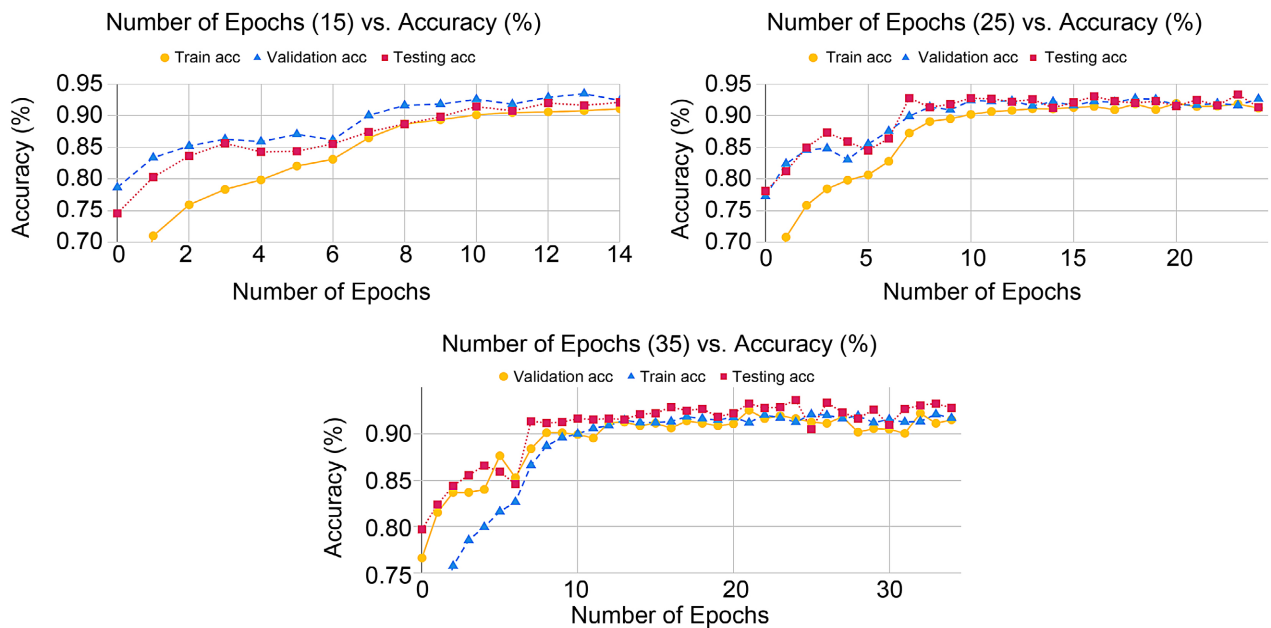


Figure 6. Accuracies of the model with different epochs. Models trained with a greater number of epochs are more accurate than those that are trained with a lower number of epochs.

fit the land-use dataset.

However, it is important to note that too many epochs used can result in overfitting the training dataset, and too few epochs can result in an underfit model. We can measure if a model is overfitted by checking if the validation loss is decreasing or increasing. If it is decreasing, the model is underfitting, and if it

is increasing, the model is overfitting.

5. Conclusions

The increasing land use in recent years has been proved to have a significant effect on the environment. In order to monitor and mitigate this issue, human land-use needs to be classified. Machine learning proves to be an efficient way of classifying this increasing human land use. While the accuracy of the model still isn't at 100%, the upward trend of the learning curve shown by the graph demonstrates that the accuracy can be improved through larger training datasets and a greater number of epochs. We learn that transfer learning is a greatly efficient way of creating an accurate image classification algorithm within a short period of time. With this information, this method can be applied in the future with larger databases covering more categories of human land use.

We have created a machine learning model through transfer learning that is able to classify satellite images of land use with relatively high accuracy. This would be beneficial to the monitoring and predicting of urban development patterns, management of water and other natural resources, and land-use planning and regulation. As a result of companies and governments keeping satellite data private, there is currently a limited amount of satellite data available to the public. The accuracy of this model can be further improved through larger datasets and longer training periods. With this model, we as a society can better understand the change in land use over the years and the increase in human land use on this planet. It will not only be useful for urban planning purposes but also for the study of how human activity affects landscape and more.

Acknowledgements

Thanks to Professor Eric Eaton for his guidance and feedback on this work, and the contributions of William Chen. Most importantly, thank Professor Shawn D. Newsam from UC Merced for providing the dataset used in this experiment.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Sara, B. (2021) Machine Learning, Explained.
<https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>
- [2] IBM Cloud Education (2020) Convolutional Neural Networks.
<https://www.ibm.com/cloud/learn/convolutional-neural-networks>
- [3] SuperDataScience (2018) The Ultimate Guide to Convolutional Neural Networks (CNN).
<https://www.superdatascience.com/blogs/the-ultimate-guide-to-convolutional-neural-networks%20cnn>
- [4] Shahid, D. (2019) Convolutional Neural Network.

-
- <https://towardsdatascience.com/covolutional-neural-network-cb0883dd6529>
- [5] Desai, J.V., Jayanth, K. and Suddpalli, K.C. (2017) Fully Connected Layer. 2017. Embedded Systems Lab, 6 May.
http://shukra.cedt.iisc.ernet.in/edwiki/Neural_Network_based_classification_of_traffic_signs
- [6] Yang, Y., and Newsam, S. (2010) Bag-of-Visual-Words and Spatial Extensions for Land-Use Classification. *18th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS 10)*, New York, November 2010, 270-279.
<https://doi.org/10.1145/1869790.1869829>
- [7] Chilamkurthy, S. (2021) Transfer Learning for Computer Vision Tutorial.
https://pytorch.org/%20tutorials/beginner/transfer_learning_tutorial.html