Neural Networks

Lecture 20

Deep Neural Networks (2)

Md. Mijanur Rahman, Prof. Dr.

Dept. of Computer Science and Engineering, Jatiya Kabi Kazi Nazrul Islam University, Bangladesh. Email: mijan@ieee.org

Chapter Contents

- **This chapter covers the following topics:**
  - **Types of Deep Neural Networks**
    - **Multi-Layer Perceptrons (MLP)**
    - **Convolutional Neural Networks (CNN)**
    - **Recurrent Neural Networks (RNN)**
  - **Frameworks for Deep Learning**
  - **Limitations and Challenges**

Multilayer Perceptrons (MLPs)

- A multilayer perceptron (MLP) is a class of a feedforward artificial neural network (ANN).
- MLPs models are the most basic deep neural network, which is composed of a series of fully connected layers.
- Today, MLP machine learning methods can be used to overcome the requirement of high computing power required by modern deep learning architectures.
- Each new layer is a set of nonlinear functions of a weighted sum of all outputs (fully connected) from the prior one.

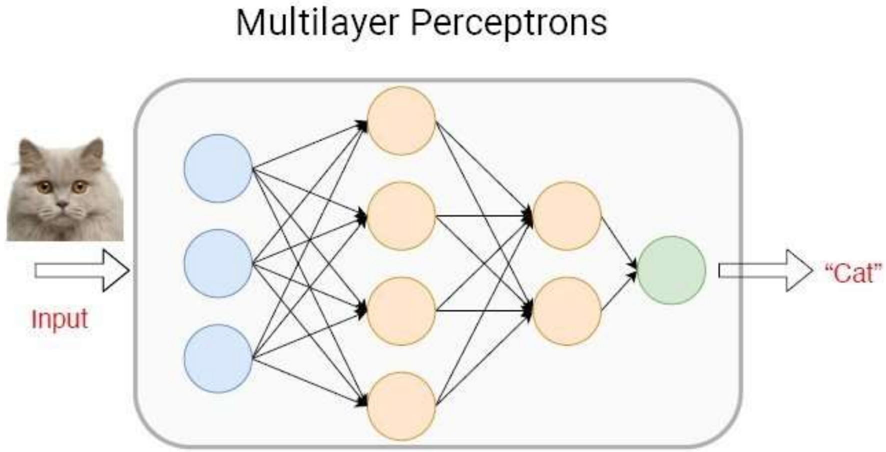

Convolutional Neural Network…

- A convolutional neural network (CNN, or ConvNet) is another class of deep neural networks. CNNs are most commonly employed in computer vision.
- Given a series of images or videos from the real world, with the utilization of CNN, the AI system learns to automatically extract the features of these inputs to complete a specific task, e.g., image classification, face authentication, and image semantic segmentation.

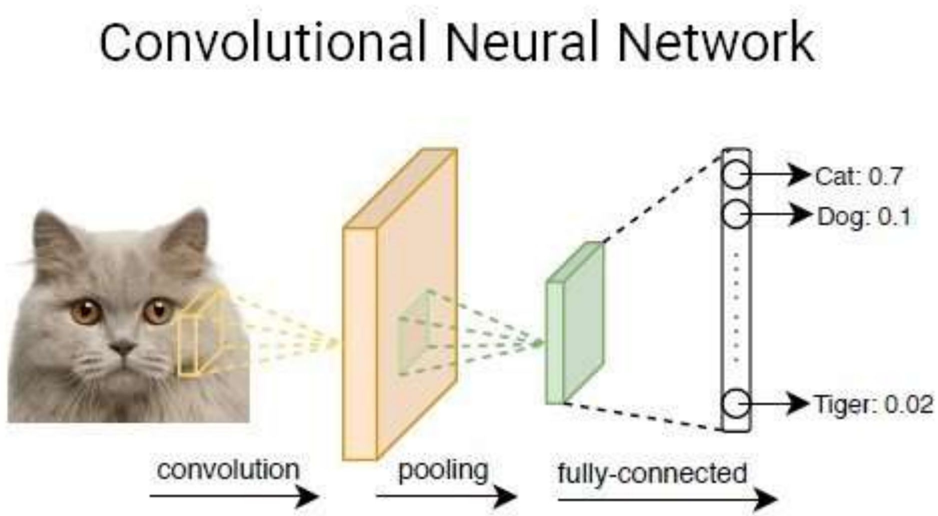

Convolutional Neural Network…

- Different from fully connected layers in MLPs, in CNN models, one or multiple convolution layers extract the simple features from input by executing convolution operations.
- Each layer is a set of nonlinear functions of weighted sums at different coordinates of spatially nearby subsets of outputs from the prior layer, which allows the weights to be reused.

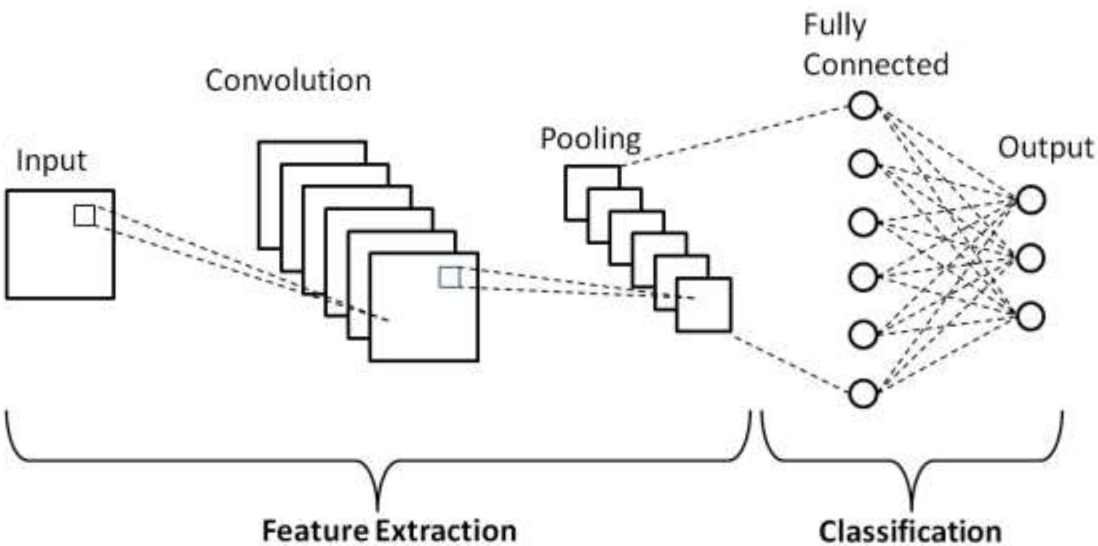

Convolutional Neural Network…

- A typical **convolutional neural network architecture** generally has several components:

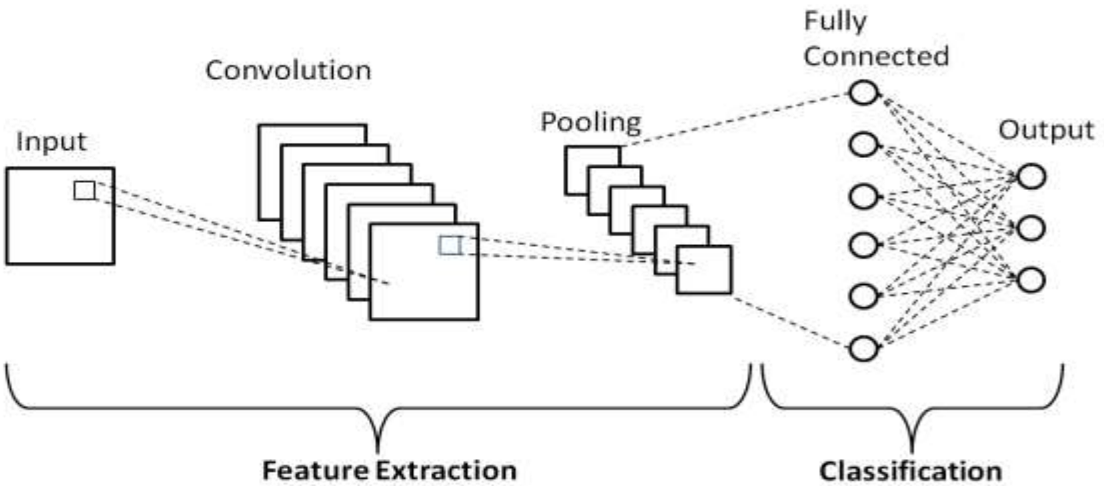

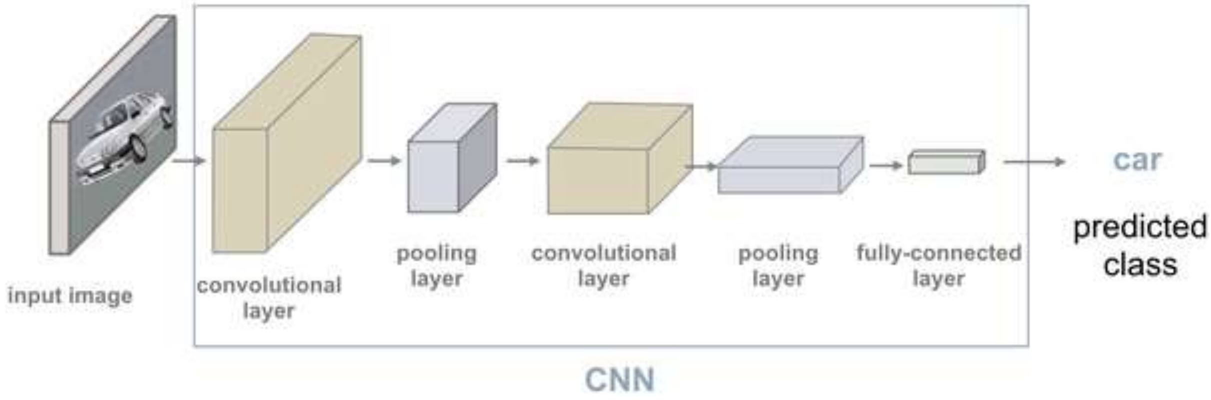

Convolutional Neural Network…

A convolution layer:

- You can think of this layer as **“what relevant features are we picking up in an image?”**
- In a convolutional neural network, we have multiple convolutional layers that extract low to high-level features depending on what specific layer we are focusing on.
- To give an (over-simplified) intuition, earlier convolutional layers pick up lower-level features (i.e. like lines and edges) while later convolutional layers pick up higher-level features based on inputs from lower-level features (i.e. shapes, structures) - analogous to how vision works in the human brain.

Convolutional Neural Network…

A pooling layer:

- Convolutional neural networks are typically used for image classification. However, images are high-dimensional data - so we would prefer to reduce the dimensionality to minimize the possibility of overfitting.
- Pooling essentially reduces the spatial dimensions of the image based on certain mathematical operations such as average or max-pooling.
- We generally incorporate pooling since it (1) generally acts as a noise suppressant (2) makes it invariant to translation movement for image classification and (3) helps capture essential structural features of the represented images without being bogged down by the fine details.

Convolutional Neural Network…

Fully-connected layer:

- You can think of a series of convolution and pooling operations as dimensionality reduction steps prior to passing this information over to the fully connected (Dense) layer.
- Essentially, what the fully connected layer does is that it takes the "compressed" representation of the image and it tries to fit a basic NN (multi-layer perceptron) when doing classification.

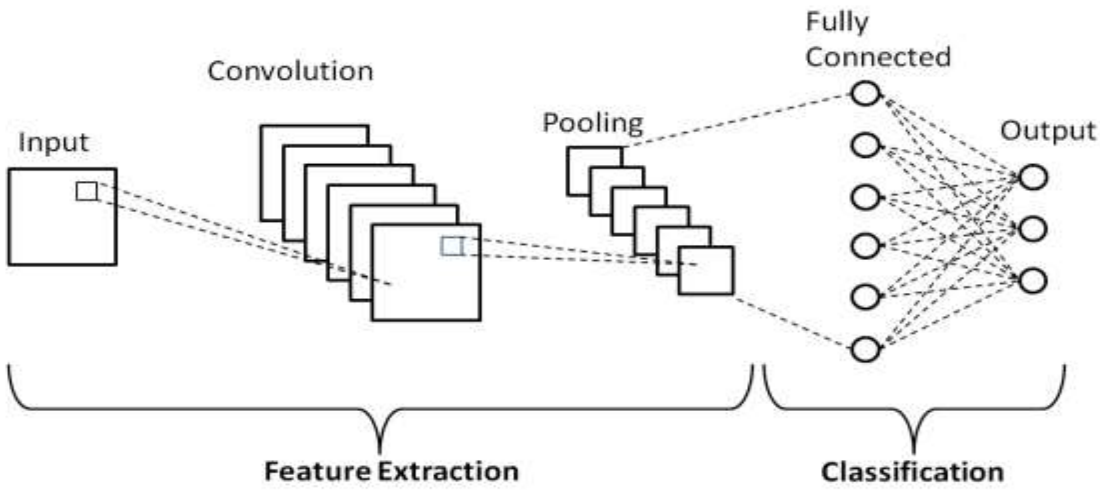

Convolutional Neural Network…

- Applying various convolutional filters, CNN machine learning models can capture the high-level representation of the input data, making it most popular for computer vision tasks, such as-
- Image classification (e.g., AlexNet, VGG network, ResNet, MobileNet), and
- Object detection (e.g., Fast R-CNN, Mask R-CNN, YOLO, SSD).

Convolutional Neural Network…

- **AlexNet:** For image classification, as the first CNN neural network to win the ImageNet Challenge in 2012, AlexNet consists of five convolution layers and three fully connected layers. Thus, AlexNet requires 61 million weights and 724 million MACs (multiply-add computation) to classify the image with a size of 227×227.

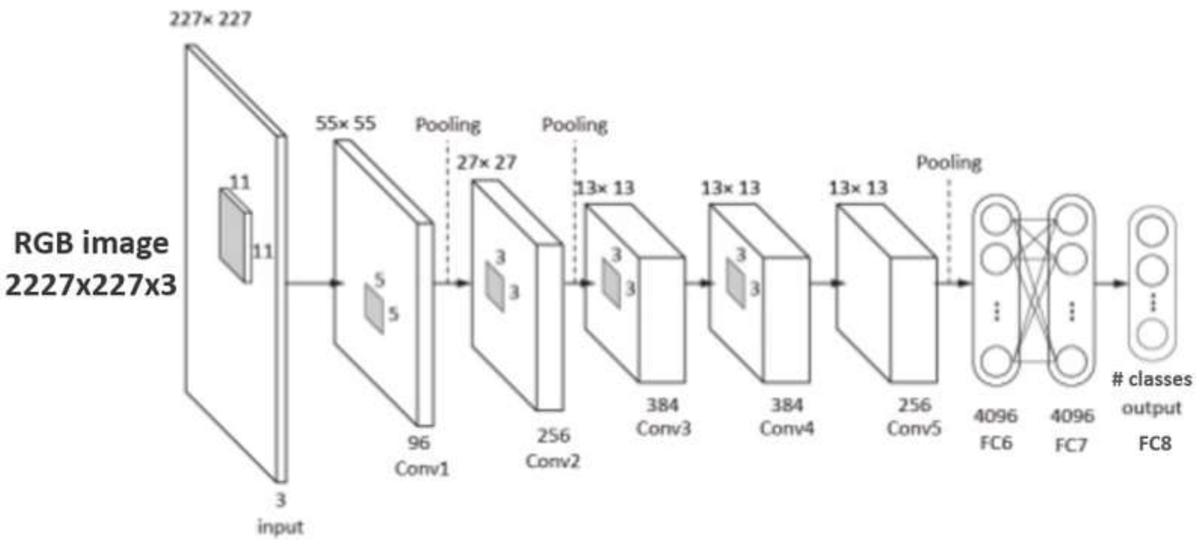

Convolutional Neural Network…

- **VGG-16:** To achieve higher accuracy, VGG-16 is trained to a deeper structure of 16 layers consisting of 13 convolution layers and three fully connected layers, requiring 138 million weights and 15.5G MACs to classify the image with a size of 224×224.

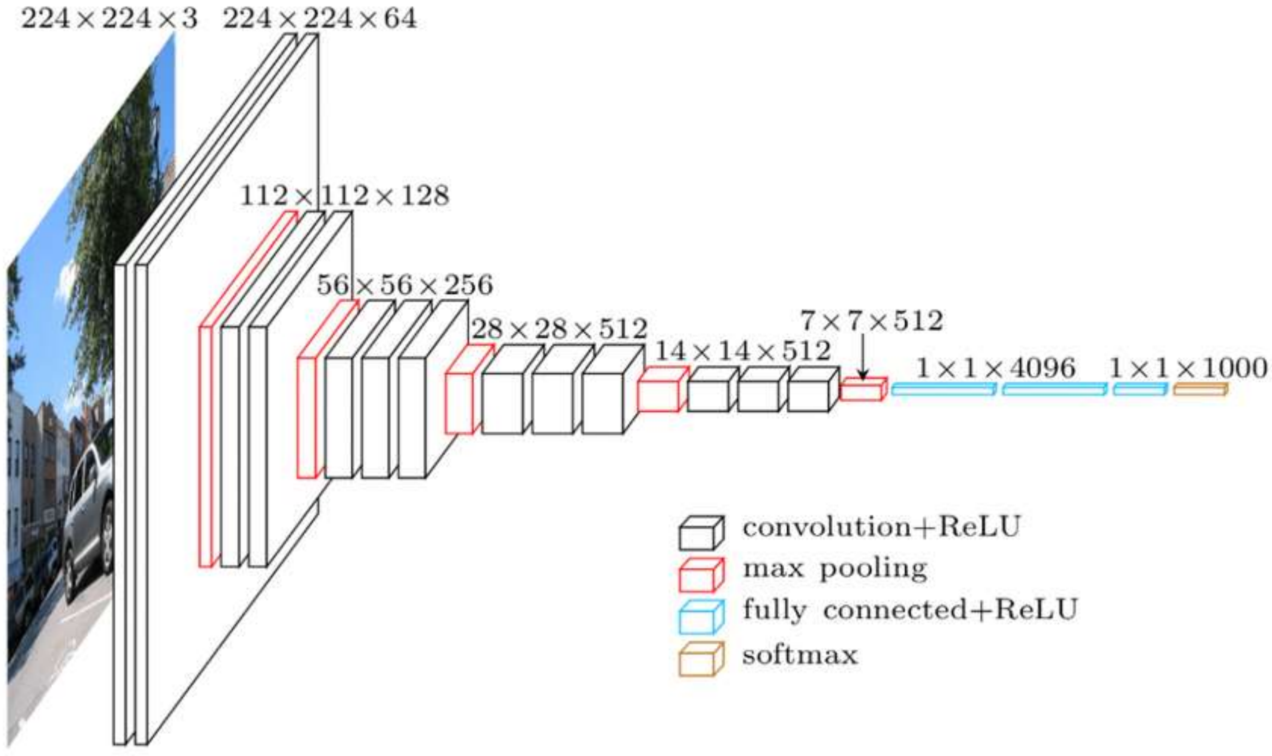

Convolutional Neural Network

- **GoogleNet:** To improve accuracy while reducing the computation of DNN inference, GoogleNet introduces an inception module composed of different sized filters. As a result, GoogleNet achieves a better accuracy performance than VGG-16 while only requiring seven million weights and 1.43G MACs to process the image with the same size.
- **ResNet:** ResNet, the state-of-the-art effort, uses the "shortcut" structure to reach a human-level accuracy with a top-5 error rate below 5%. In addition, the "shortcut" module is used to solve the gradient vanishing problem during the training process, making it possible to train a DNN model with a deeper structure.

Recurrent Neural Networks

Deep Learning – Introduction to Recurrent Neural Networks

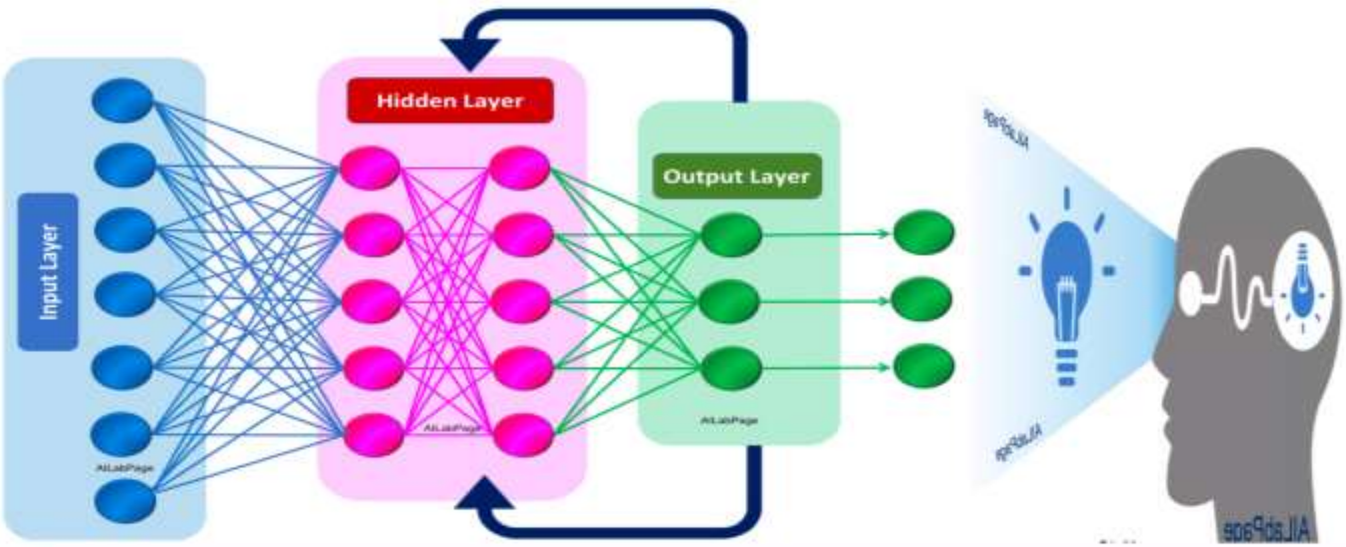

01 Useful
Useful for tasks that are dependent on a sequence of a successive states.

02 Training
The network can be trained by backpropagation.

03 Memory
The network has a form of short-term memory.

04 Simplicity
Simple recurrent network (SRN) has a similar form of short-term memory.

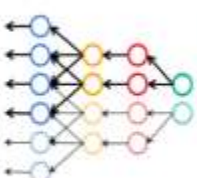

Backpropagation through time

BPTT – An algorithm used for updating weights in the recurrent neural network. to minimize the error of the network outputs. For every recurrent network there is a feedforward network with identical behavior.

Back Activation

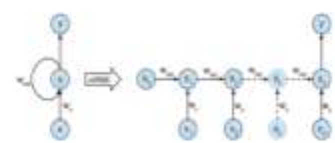

A recurrent connection feeds back activation that will affect the output from the network during subsequent iterations.

Recurrent Neural Network…

- A recurrent neural network (RNN) is another class of artificial neural networks that use sequential data feeding. RNNs have been developed to address the time-series problem of sequential input data.
- RNN works on the principle of saving the output of a particular layer and feeding this back to the input in order to predict the output of the layer.
- Below is how you can convert a Feed-Forward Neural Network into a Recurrent Neural Network.

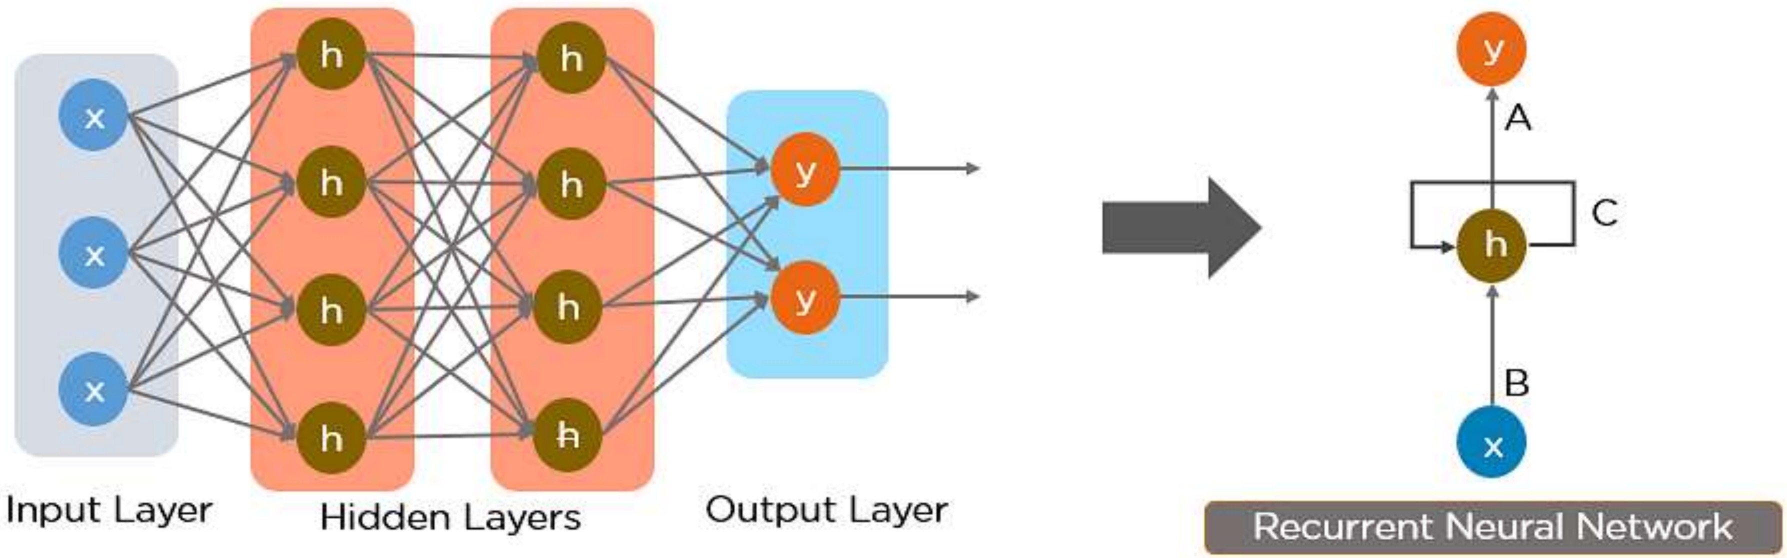

Recurrent Neural Network…

- The input of RNN consists of the current input and the previous samples. Therefore, the connections between nodes form a directed graph along a temporal sequence.
- Furthermore, each neuron in an RNN owns an internal memory that keeps the information of the computation from the previous samples.

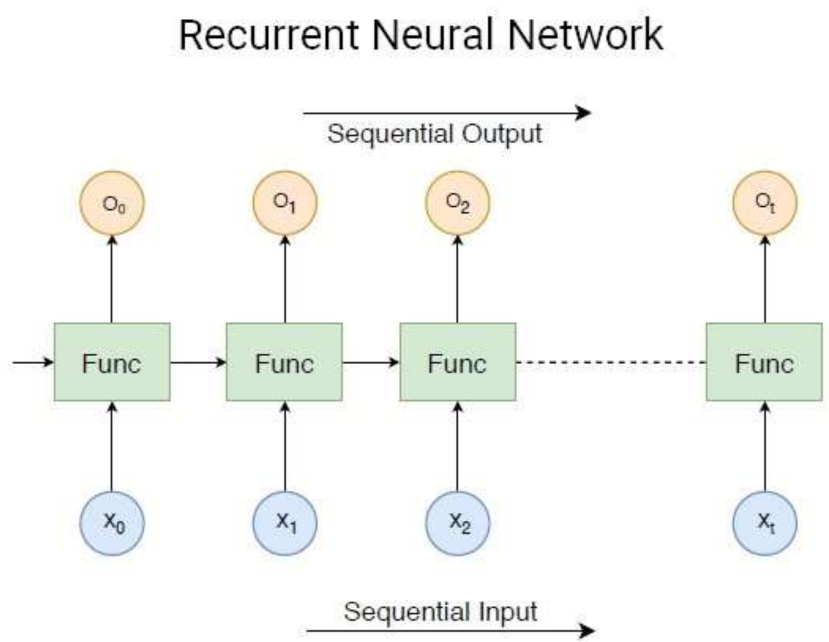

Recurrent Neural Network

- RNN models are widely used in natural language processing (NLP) due to the superiority of processing the data with an input length that is not fixed.
- The task of the AI here is to build a system that can comprehend natural language spoken by humans, e.g., natural language modeling, word embedding, and machine translation.
- In RNNs, each subsequent layer is a collection of nonlinear functions of weighted sums of outputs and the previous state. Thus, the basic unit of RNN is called "cell", and each cell consists of layers and a series of cells that enables the sequential processing of recurrent neural network models.

Frameworks for Deep Learning…

What are some popular frameworks for Deep Learning?

- These deep learning frameworks provide us with reusable code blocks that abstract the logical blocks and also provides several additional handy modules in developing a deep learning model.
- We can classify the available options as a **low level or high-level** deep learning framework.
- **Low-level frameworks** provide a more basic block of abstraction while giving a ton of room for customization and flexibility. **High-level frameworks** further aggregate the abstraction to ease our work while limiting the scope of customization and flexibility.
- **High-level frameworks use a low-level framework as a backend** and typically work by converting the source into the desired low-level framework for execution.

Frameworks for Deep Learning

- Below are a few of the popular choices of frameworks for deep learning:
  - Low-level frameworks
    - TensorFlow
    - MxNet
    - PyTorch
  - High-level frameworks
    - Keras (uses TensorFlow as a backend)
    - Gluon (uses mxnet as a backend)
- At the moment, the most popular choice is **Tensorflow (by Google)**. **Keras** is also popular given the ease it provides in quickly protyping deep learning models. However, **PyTorch (by Facebook)** is another popular framework that is catching up with the race at an astoundingly high speed.

Deep learning vs. Machine learning

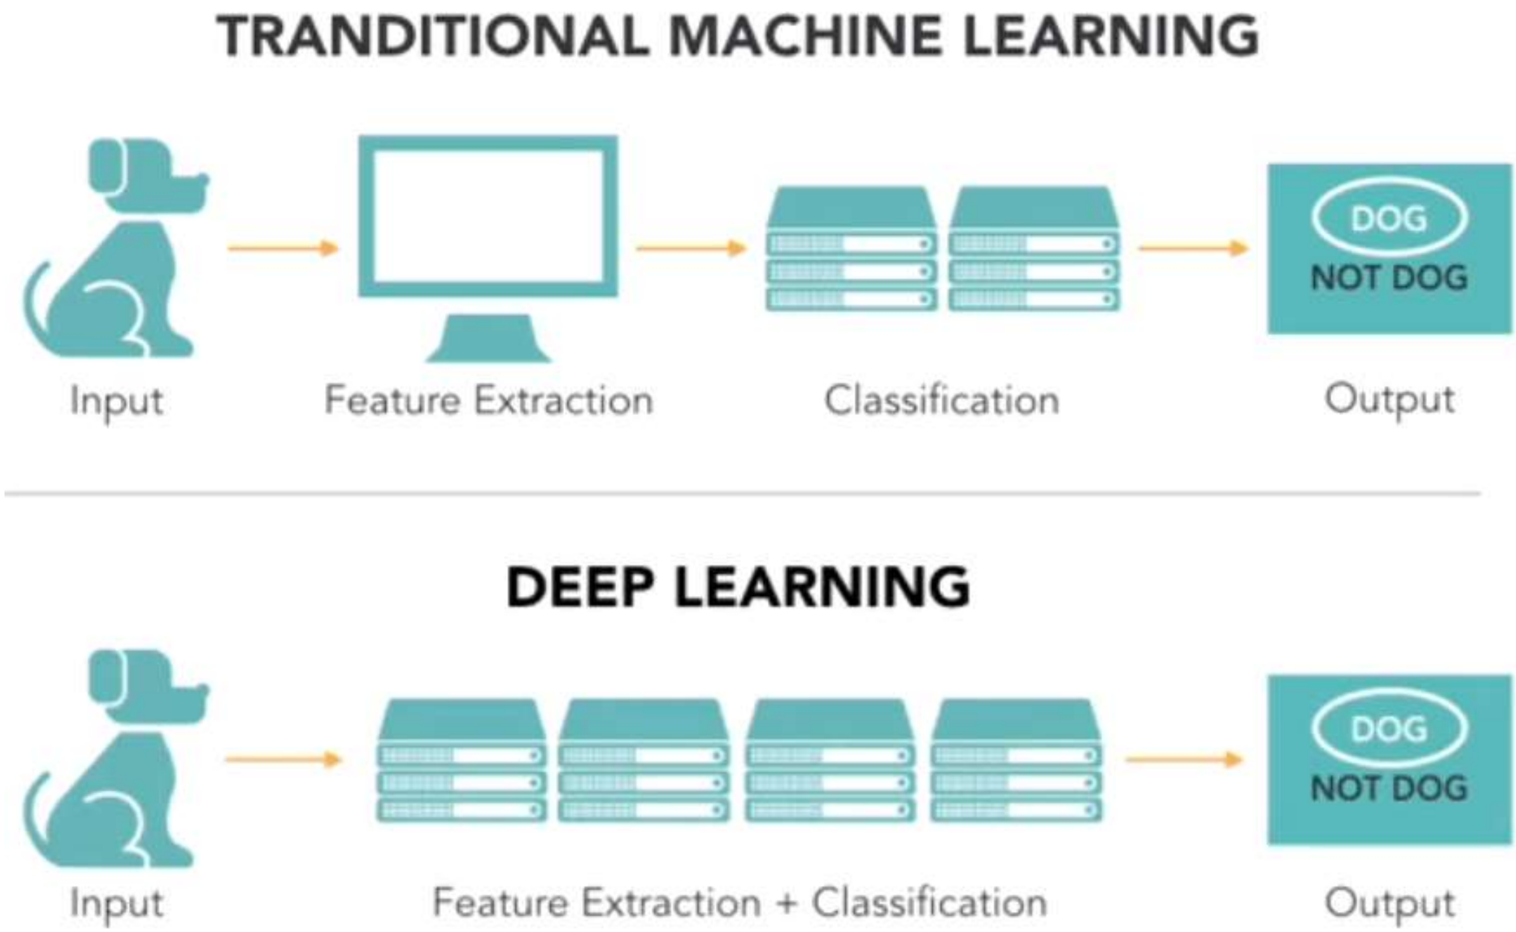

Deep learning vs. Machine learning

- Deep learning is a subset of machine learning that differentiates itself through the way it solves problems. Machine learning requires a domain expert to identify most applied features. On the other hand, deep learning understands features incrementally, thus eliminating the need for domain expertise.
- This makes deep learning algorithms take much longer to train than machine learning algorithms, which only need a few seconds to a few hours. However, the reverse is true during testing. Deep learning algorithms take much less time to run tests than machine learning algorithms, whose test time increases along with the size of the data.
- Furthermore, machine learning does not require the same costly, high-end machines and high-performing GPUs that deep learning does.
- In the end, many data scientists choose traditional machine learning over deep learning due to its superior interpretability, or the ability to make sense of the solutions. Machine learning algorithms are also preferred when the data is small.
- Instances where deep learning becomes preferable include situations where there is a large amount of data, a lack of domain understanding for feature introspection, or complex problems, such as speech recognition and NLP.

Limitations and Challenges

- The biggest limitation of deep learning models is they learn through observations. This means they only know what was in the data on which they trained. If a user has a small amount of data or it comes from one specific source that is not necessarily representative of the broader functional area, the models will not learn in a way that is generalizable.
- The issue of biases is also a major problem for deep learning models. If a model trains on data that contains biases, the model will reproduce those biases in its predictions.
- The learning rate can also become a major challenge to deep learning models. If the rate is too high, then the model will converge too quickly, producing a less-than-optimal solution. If the rate is too low, then the process may get stuck, and it will be even harder to reach a solution.
- The hardware requirements for deep learning models can also create limitations. Multicore high-performing graphics processing units (GPUs) and other similar processing units are required to ensure improved efficiency and decreased time consumption. However, these units are expensive and use large amounts of energy. Other hardware requirements include random access memory and a hard disk drive (HDD) or RAM-based solid-state drive (SSD).

Limitations and Challenges

Other limitations and challenges include the following:

- Deep learning requires large amounts of data. Furthermore, the more powerful and accurate models will need more parameters, which, in turn, require more data.
- Once trained, deep learning models become inflexible and cannot handle multitasking. They can deliver efficient and accurate solutions but only to one specific problem. Even solving a similar problem would require retraining the system.
- Any application that requires reasoning -- such as programming or applying the scientific method -- long-term planning and algorithm like data manipulation is completely beyond what current deep learning techniques can do, even with large data.

Lecture 23: Deep Neural Networks (2)

DEEP NEURAL NETWORKS

The End.