

TensorFlow in 1 Day: Make your own Neural Network

By Krishna Rungta

Copyright 2019 - All Rights Reserved – Krishna Rungta

ALL RIGHTS RESERVED. No part of this publication may be reproduced or transmitted in any form whatsoever, electronic, or mechanical, including photocopying, recording, or by any informational storage or retrieval system without express written, dated and signed permission from the author.

Table Of Content

Chapter 1: What is Deep learning?

1. [What is Deep learning?](#)
2. [Deep learning Process](#)
3. [Classification of Neural Networks](#)
4. [Types of Deep Learning Networks](#)
5. [Feed-forward neural networks](#)
6. [Recurrent neural networks \(RNNs\)](#)
7. [Convolutional neural networks \(CNN\)](#)

Chapter 2: Machine Learning vs Deep Learning

1. [What is AI?](#)
2. [What is ML?](#)
3. [What is Deep Learning?](#)
4. [Machine Learning Process](#)
5. [Deep Learning Process](#)
6. [Automate Feature Extraction using DL](#)
7. [Difference between Machine Learning and Deep Learning](#)
8. [When to use ML or DL?](#)

Chapter 3: What is TensorFlow?

1. [What is TensorFlow?](#)
2. [History of TensorFlow](#)
3. [TensorFlow Architecture](#)
4. [Where can Tensorflow run?](#)
5. [Introduction to Components of TensorFlow](#)
6. [Why is TensorFlow popular?](#)
7. [List of Prominent Algorithms supported by TensorFlow](#)

Chapter 4: Comparison of Deep Learning Libraries

1. [8 Best Deep learning Libraries /Framework](#)
2. [MICROSOFT COGNITIVE TOOLKIT\(CNTK\)](#)
3. [TensorFlow Vs Theano Vs Torch Vs Keras Vs infer.net Vs CNTK Vs MXNet Vs Caffe: Key Differences](#)

Chapter 5: How to Download and Install TensorFlow Windows and Mac

1. [TensorFlow Versions](#)
2. [Install Anaconda](#)
3. [Create .yml file to install Tensorflow and dependencies](#)
4. [Launch Jupyter Notebook](#)
5. [Jupyter with the main conda environment](#)

Chapter 6: Jupyter Notebook Tutorial

1. [What is Jupyter Notebook?](#)
2. [Jupyter Notebook App](#)
3. [How to use Jupyter](#)

Chapter 7: Tensorflow on AWS

1. [PART 1: Set up a key pair](#)
2. [PART 2: Set up a security group](#)
3. [Launch your instance \(Windows users\)](#)
4. [Part 4: Install Docker](#)
5. [Part 5: Install Jupyter](#)
6. [Part 6: Close connection](#)

Chapter 8: TensorFlow Basics: Tensor, Shape, Type, Graph, Sessions & Operators

1. [What is a Tensor?](#)

2. [Representation of a Tensor](#)
3. [Types of Tensor](#)
4. [Shape of tensor](#)
5. [Type of data](#)
6. [Creating operator](#)
7. [Variables](#)

Chapter 9: Tensorboard: Graph Visualization with Example

Chapter 10: NumPy

1. [What is NumPy?](#)
2. [Why use NumPy?](#)
3. [How to install NumPy?](#)
4. [Mathematical Operations on an Array](#)
5. [Shape of Array](#)
6. [np.zeros and np.ones](#)
7. [Reshape and Flatten Data](#)
8. [hstack and vstack](#)

Chapter 11: Pandas

1. [What is Pandas?](#)
2. [Why use Pandas?](#)
3. [How to install Pandas?](#)
4. [What is a data frame?](#)
5. [What is a Series?](#)
6. [Concatenation](#)

Chapter 12: Scikit-Learn

1. [What is Scikit-learn?](#)
2. [Download and Install scikit-learn](#)
3. [Machine learning with scikit-learn](#)
4. [Step 1\) Import the data](#)

5. [Step 2\) Create the train/test set](#)
6. [Step 3\) Build the pipeline](#)
7. [Step 4\) Using our pipeline in a grid search](#)

Chapter 13: Linear Regression

1. [Linear regression](#)
2. [How to train a linear regression model](#)
3. [How to train a Linear Regression with TensorFlow](#)
4. [Pandas](#)
5. [Numpy Solution](#)
6. [Tensorflow solution](#)

Chapter 14: Linear Regression Case Study

1. [Summary statistics](#)
2. [Facets Overview](#)
3. [Facets Deep Dive](#)
4. [Install Facet](#)
5. [Overview](#)
6. [Graph](#)
7. [Facets Deep Dive](#)

Chapter 15: Linear Classifier in TensorFlow

1. [What is Linear Classifier?](#)
2. [How Binary classifier works?](#)
3. [How to Measure the performance of Linear Classifier?](#)
4. [Linear Classifier with TensorFlow](#)

Chapter 16: Kernel Methods

1. [Why do you need Kernel Methods?](#)
2. [What is a Kernel in machine learning?](#)

3. [Type of Kernel Methods](#)
4. [Train Gaussian Kernel classifier with TensorFlow](#)

Chapter 17: TensorFlow ANN (Artificial Neural Network)

1. [What is Artificial Neural Network?](#)
2. [Neural Network Architecture](#)
3. [Limitations of Neural Network](#)
4. [Example Neural Network in TensorFlow](#)
5. [Train a neural network with TensorFlow](#)

Chapter 18: ConvNet(Convolutional Neural Network): TensorFlow Image Classification

1. [What is Convolutional Neural Network?](#)
2. [Architecture of a Convolutional Neural Network](#)
3. [Components of Convnets](#)
4. [Train CNN with TensorFlow](#)

Chapter 19: Autoencoder with TensorFlow

1. [What is an Autoencoder?](#)
2. [How does Autoencoder work?](#)
3. [Stacked Autoencoder Example](#)
4. [Build an Autoencoder with TensorFlow](#)

Chapter 20: RNN(Recurrent Neural Network) TensorFlow

1. [What do we need an RNN?](#)
2. [What is RNN?](#)
3. [Build an RNN to predict Time Series in TensorFlow](#)


Chapter 1: What is Deep learning?

What is Deep learning?

Deep learning is a computer software that **mimics the network of neurons in a brain**. It is a subset of machine learning and is called deep learning because it makes use of deep **neural networks**.

Deep learning algorithms are constructed with connected layers.

- The first layer is called the Input Layer The last
- layer is called the Output Layer
- All layers in between are called Hidden Layers. The word deep means the network join neurons in more than two layers.



Each Hidden layer is composed of neurons. The neurons are connected to each other. The neuron will process and then propagate the input signal it receives the layer above it. The strength of the signal given the neuron in the next layer depends on the weight, bias and activation function.

The network consumes large amounts of input data and operates them through multiple layers; the network can learn increasingly complex

features of the data at each layer.

Deep learning Process

A deep neural network provides state-of-the-art accuracy in many tasks, from object detection to speech recognition. They can learn automatically, without predefined knowledge explicitly coded by the programmers.



To grasp the idea of deep learning, imagine a family, with an infant and parents. The toddler points objects with his little finger and always says the word 'cat.' As its parents are concerned about his education, they keep telling him 'Yes, that is a cat' or 'No, that is not a cat.' The infant persists in pointing objects but becomes more accurate with 'cats.' The little kid, deep down, does not know why he can say it is a cat or not. He has just learned how to hierarchies complex features coming up with a cat by looking at the pet overall and continue to focus on details such as the tails or the nose before to make up his mind.

A neural network works quite the same. Each layer represents a deeper level of knowledge, i.e., the hierarchy of knowledge. A neural network with four layers will learn more complex feature than with that with two layers.


The learning occurs in two phases.

- The first phase consists of applying a nonlinear transformation of the input and create a statistical model as output.
- The second phase aims at improving the model with a mathematical method known as derivative.

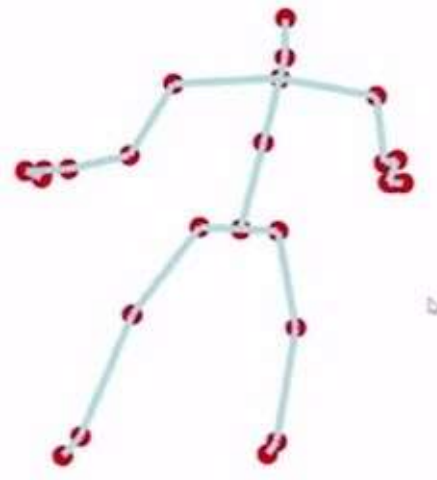
The neural network repeats these two phases hundreds to thousands of time until it has reached a tolerable level of accuracy. The repeat of this

two-phase is called an iteration.

To give an example, take a look at the motion below, the model is trying to learn how to dance. After 10 minutes of training, the model does not know how to dance, and it looks like a scribble.



After 48 hours of learning, the computer masters the art of dancing.



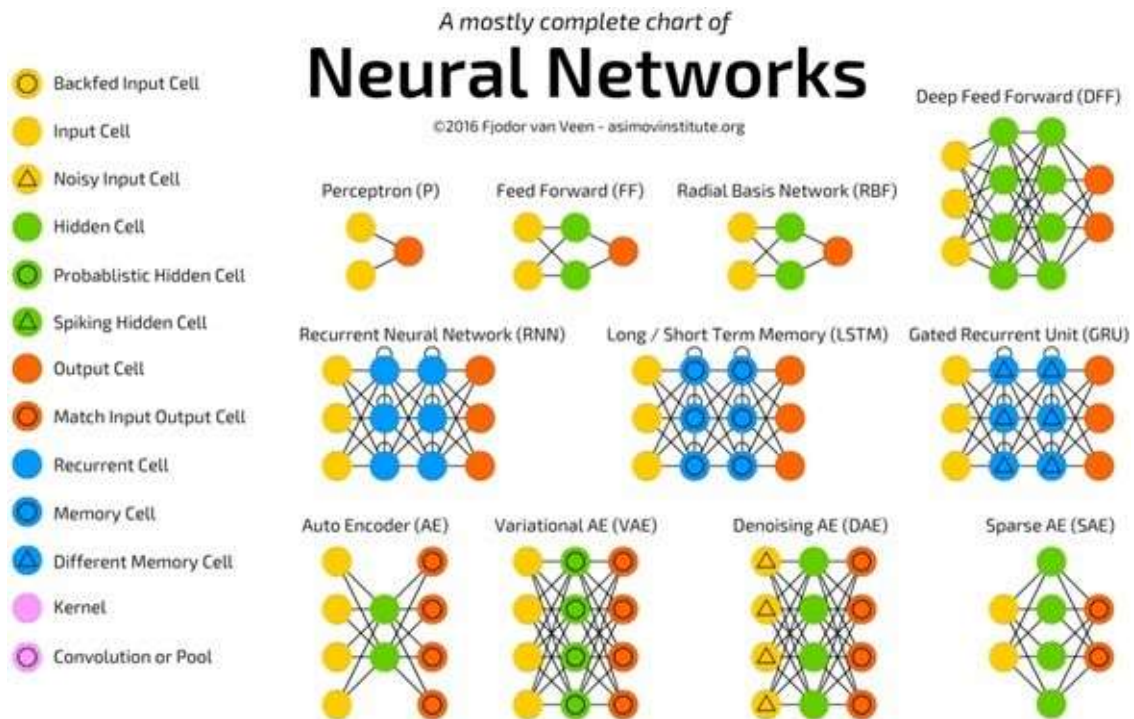
Classification of Neural Networks

Shallow neural network: The Shallow neural network has only one hidden layer between the input and output.

Deep neural network: Deep neural networks have more than one layer. For instance, Google LeNet model for image recognition counts 22 layers.

Nowadays, deep learning is used in many ways like a driverless car, mobile phone, Google Search Engine, Fraud detection, TV, and so on.

Types of Deep Learning Networks

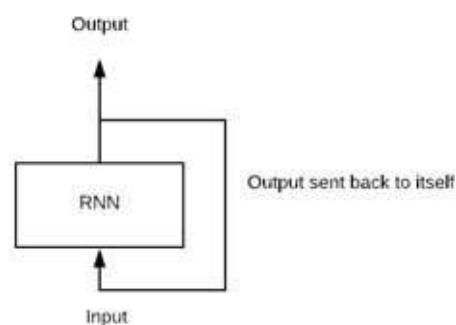


Feed-forward neural networks

The simplest type of artificial neural network. With this type of architecture, information flows in only one direction, forward. It means, the information's flows starts at the input layer, goes to the "hidden" layers, and end at the output layer. The network does not have a loop. Information stops at the output layers.

Recurrent neural networks (RNNs)

RNN is a multi-layered neural network that can store information in context nodes, allowing it to learn data sequences and output a number or another sequence. In simple words it an Artificial neural networks whose connections between neurons include loops. RNNs are well suited for processing sequences of inputs.



Example, if the task is to predict the next word in the sentence “Do you want a ?

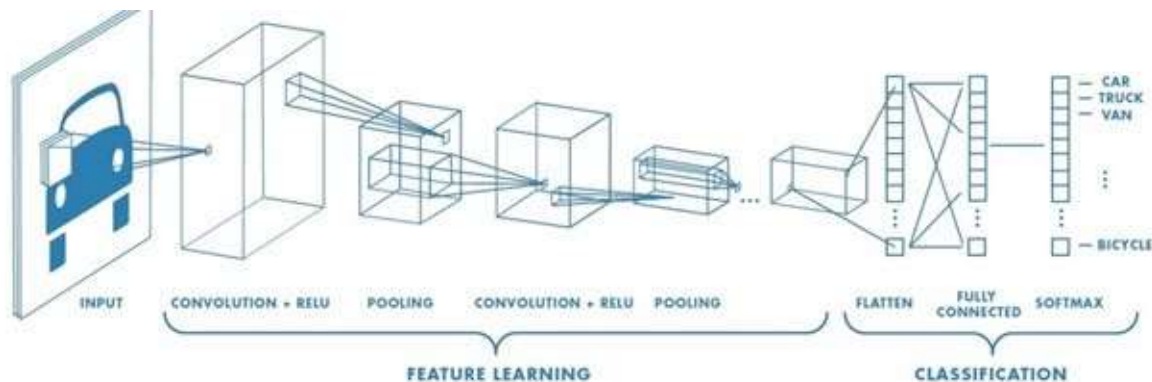
- The RNN neurons will receive a signal that point to the start of the sentence.
- The network receives the word “Do” as an input and produces a vector of the number. This vector is fed back to the neuron to provide a memory to the network. This stage helps the network to remember it received “Do” and it received it in the first position. The network will similarly proceed
- to the next words. It takes the word “you” and “want.” The state of the neurons is updated upon receiving each word.
- The final stage occurs after receiving the word “a.” The neural network will provide a probability for each English word that can be used to complete the sentence. A well-trained RNN probably assigns a high probability to “café,” “drink,” “burger,” etc.

Common uses of RNN

- Help securities traders to generate analytic reports
- Detect abnormalities in the contract of financial statement
- Detect fraudulent credit-card transaction
- Provide a caption for images Power
- chatbots
- The standard uses of RNN occur when the practitioners are working with time-series data or sequences (e.g., audio recordings or text).

Convolutional neural networks (CNN)

CNN is a multi-layered neural network with a unique architecture designed to extract increasingly complex features of the data at each layer to determine the output. CNN's are well suited for perceptual tasks.



CNN is mostly used when there is an unstructured data set (e.g., images) and the practitioners need to extract information from it

For instance, if the task is to predict an image caption:

- The CNN receives an image of let's say a cat, this image, in computer term, is a collection of the pixel. Generally, one layer for the greyscale picture and three layers for a color picture.
- During the feature learning (i.e., hidden layers), the network will identify unique features, for instance, the tail of the cat, the ear, etc.
- When the network thoroughly learned how to recognize a picture, it can provide a probability for each image it knows. The label with the highest probability will become the prediction of the network.

Reinforcement Learning

Reinforcement learning is a subfield of machine learning in which systems are trained by receiving virtual "rewards" or "punishments,"

essentially learning by trial and error. Google's DeepMind has used reinforcement learning to beat a human champion in the Go games. Reinforcement learning is also used in video games to improve the gaming experience by providing smarter bot.

One of the most famous algorithms are:

- Q-learning Deep
- Q network
- State-Action-Reward-State-Action (SARSA) Deep
- Deterministic Policy Gradient (DDPG)

Applications/ Examples of deep learning applications

AI in Finance: The financial technology sector has already started using AI to save time, reduce costs, and add value. Deep learning is changing the lending industry by using more robust credit scoring. Credit decision-makers can use AI for robust credit lending applications to achieve faster, more accurate risk assessment, using machine intelligence to factor in the character and capacity of applicants.

Underwrite is a Fintech company providing an AI solution for credit makers company. underwrite.ai uses AI to detect which applicant is more likely to pay back a loan. Their approach radically outperforms traditional methods.

AI in HR: Under Armour, a sportswear company revolutionizes hiring and modernizes the candidate experience with the help of AI. In fact, Under Armour Reduces hiring time for its retail stores by 35%.

Under Armour faced a growing popularity interest back in 2012. They had, on average, 30000 resumes a month. Reading all of those applications and begin to start the screening and interview process was taking too long. The lengthy process to get people hired and on- boarded impacted Under Armour's ability to have their retail stores

fully staffed, ramped and ready to operate.

At that time, Under Armour had all of the 'must have' HR technology in place such as transactional solutions for sourcing, applying, tracking and onboarding but those tools weren't useful enough. Under armour choose **HireVue**, an AI provider for HR solution, for both on-demand and live interviews. The results were bluffing; they managed to decrease by 35% the time to fill. In return, the hired higher quality staffs.

AI in Marketing: AI is a valuable tool for customer service management and personalization challenges. Improved speech recognition in call-center management and call routing as a result of the application of AI techniques allows a more seamless experience for customers.

For example, deep-learning analysis of audio allows systems to assess a customer's emotional tone. If the customer is responding poorly to the AI chatbot, the system can be rerouted the conversation to real, human operators that take over the issue.

Apart from the three examples above, AI is widely used in other sectors/industries.

Why is Deep Learning Important?

Deep learning is a powerful tool to make prediction an actionable result. Deep learning excels in pattern discovery (unsupervised learning) and knowledge-based prediction. Big data is the fuel for deep learning. When both are combined, an organization can reap unprecedented results in term of productivity, sales, management, and innovation.

Deep learning can outperform traditional method. For instance, deep learning algorithms are 41% more accurate than machine learning algorithm in image classification, 27 % more accurate in facial recognition and 25% in voice recognition.

Limitations of deep learning

Data labeling

Most current AI models are trained through “supervised learning.” It means that humans must label and categorize the underlying data, which can be a sizable and error-prone chore. For example, companies developing self-driving-car technologies are hiring hundreds of people to manually annotate hours of video feeds from prototype vehicles to help train these systems.

Obtain huge training datasets

It has been shown that simple deep learning techniques like CNN can, in some cases, imitate the knowledge of experts in medicine and other fields. The current wave of machine learning, however, requires training data sets that are not only labeled but also sufficiently broad and universal.

Deep-learning methods required thousands of observation for models to become relatively good at classification tasks and, in some cases, millions for them to perform at the level of humans. Without surprise, deep learning is famous in giant tech companies; they are using big data to accumulate petabytes of data. It allows them to create an impressive and highly accurate deep learning model.

Explain a problem

Large and complex models can be hard to explain, in human terms. For instance, why a particular decision was obtained. It is one reason that acceptance of some AI tools are slow in application areas where interpretability is useful or indeed required.

Furthermore, as the application of AI expands, regulatory requirements could also drive the need for more explainable AI models.

Summary

Deep learning is the new state-of-the-art for artificial intelligence. Deep learning architecture is composed of an input layer, hidden layers, and an output layer. The word deep means there are more than two fully connected layers.

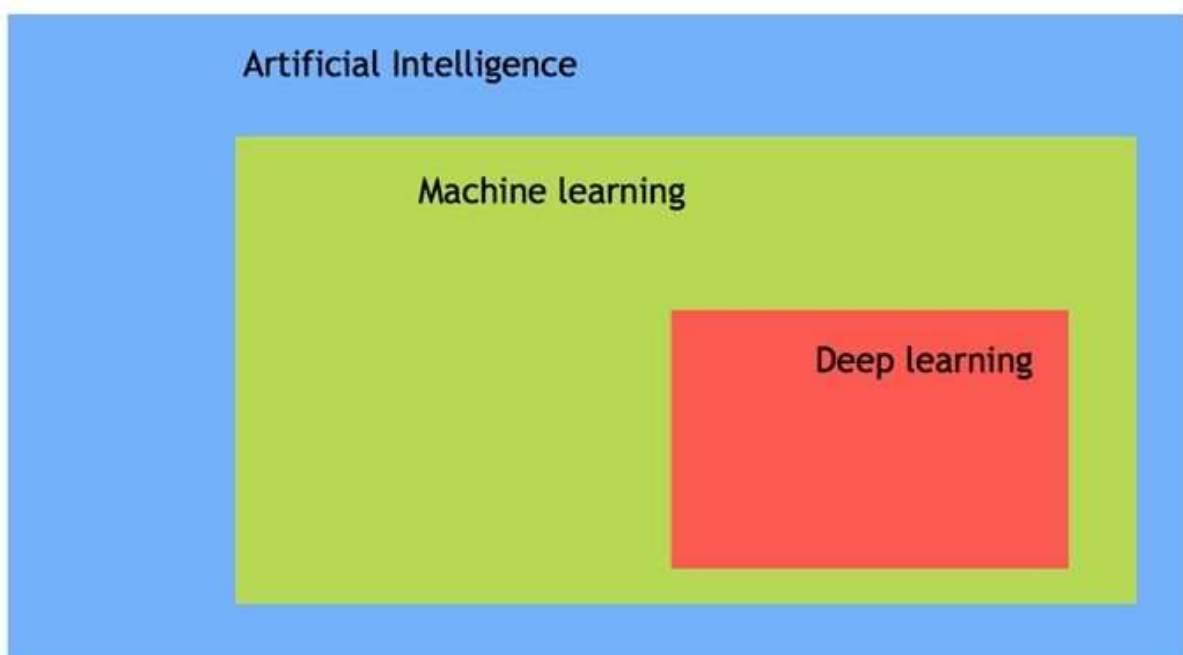
There is a vast amount of neural network, where each architecture is designed to perform a given task. For instance, CNN works very well with pictures, RNN provides impressive results with time series and text analysis.

Deep learning is now active in different fields, from finance to marketing, supply chain, and marketing. Big firms are the first one to use deep learning because they have already a large pool of data. Deep learning requires to have an extensive training dataset.

Chapter 2: Machine Learning vs Deep Learning

What is AI?

Artificial intelligence is imparting a cognitive ability to a machine. The benchmark for AI is the human intelligence regarding reasoning, speech, and vision. This benchmark is far off in the future.



AI has three different levels:

1. **Narrow AI:** A artificial intelligence is said to be narrow when the machine can perform a specific task better than a human. The current research of AI is here now
2. **General AI:** An artificial intelligence reaches the general state when it can perform any intellectual task with the same accuracy level as a human would
3. **Active AI:** An AI is active when it can beat humans in many tasks Early AI

systems used pattern matching and expert systems.

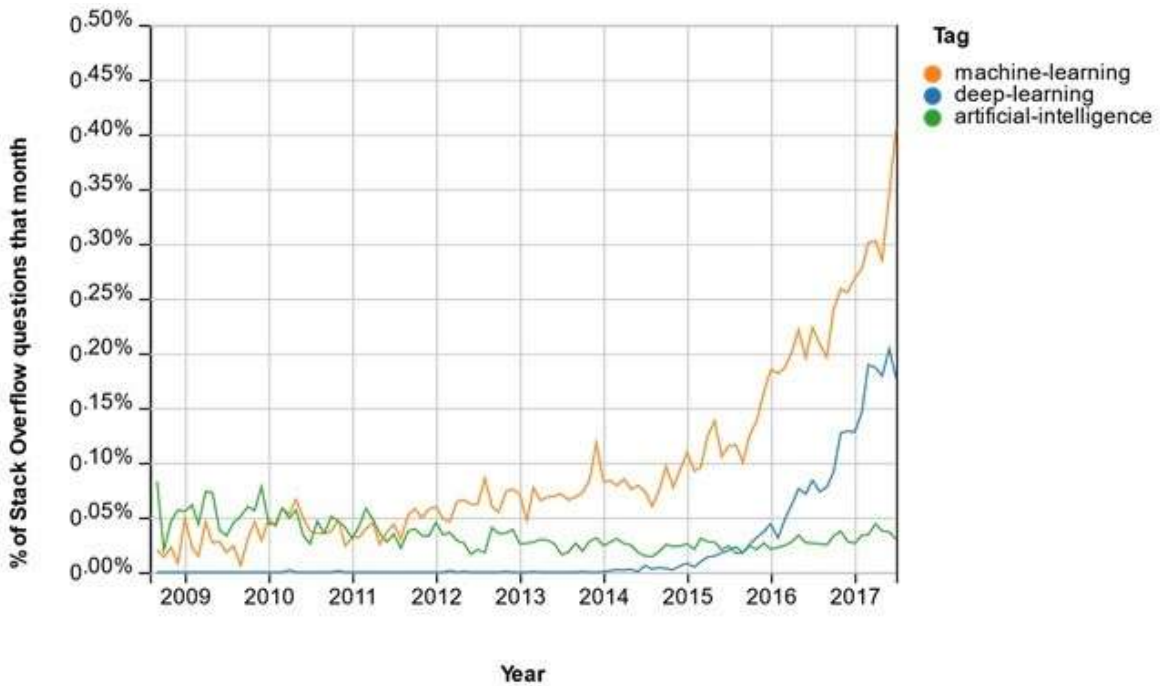
What is ML?

Machine learning is the best tool so far to analyze, understand and identify a pattern in the data. One of the main ideas behind machine learning is that the computer can be trained to automate tasks that would be exhaustive or impossible for a human being. The clear breach from the traditional analysis is that machine learning can take decisions with minimal human intervention.

Machine learning uses data to feed an algorithm that can understand the relationship between the input and the output. When the machine finished learning, it can predict the value or the class of new data point.

What is Deep Learning?

Deep learning is a computer software that mimics the network of neurons in a brain. It is a subset of machine learning and is called deep learning because it makes use of deep neural networks. The machine uses different layers to learn from the data. The depth of the model is represented by the number of layers in the model. Deep learning is the new state of the art in term of AI. In deep learning, the learning phase is done through a neural network. A neural network is an architecture where the layers are stacked on top of each other



Machine Learning Process

Imagine you are meant to build a program that recognizes objects. To train the model, you will use a **classifier**. A classifier uses the features of an object to try identifying the class it belongs to.

In the example, the classifier will be trained to detect if the image is a:

- Bicycle
- Boat Car
- Plane
-

The four objects above are the class the classifier has to recognize. To construct a classifier, you need to have some data as input and assigns a label to it. The algorithm will take these data, find a pattern and then classify it in the corresponding class.

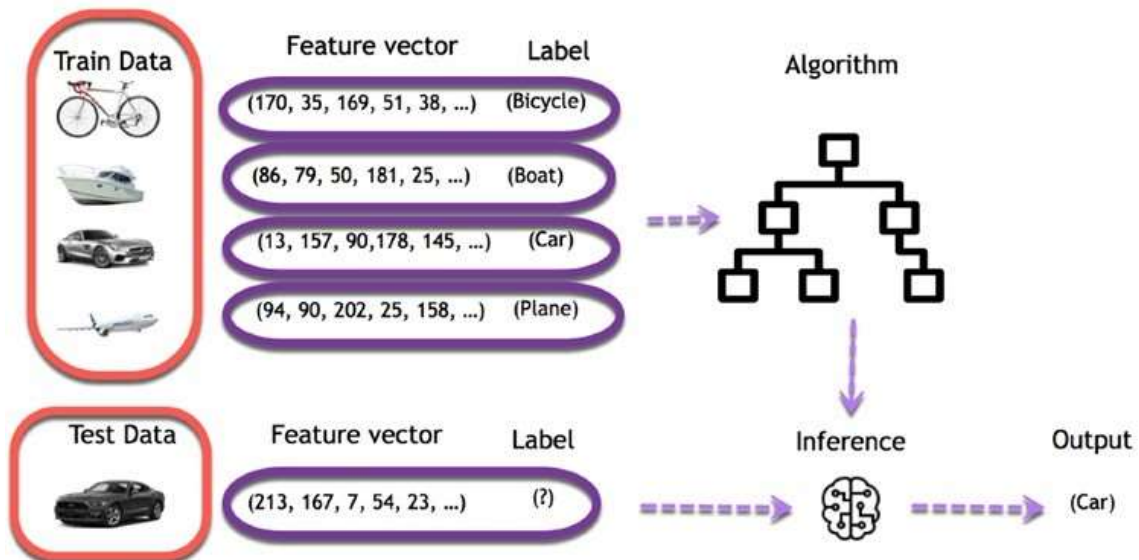
This task is called **supervised learning**. In supervised learning, the training data you feed to the algorithm includes a label.

Training an algorithm requires to follow a few standard steps:

- Collect the data Train
- the classifier Make
- predictions

The first step is necessary, choosing the right data will make the algorithm success or a failure. The data you choose to train the model is called a **feature**. In the object example, the features are the pixels of the images.

Each image is a row in the data while each pixel is a column. If your image is a 28x28 size, the dataset contains 784 columns (28x28). In the picture below, each picture has been transformed into a feature vector. The label tells the computer what object is in the image.



The objective is to use these training data to classify the type of object. The first step consists of creating the feature columns. Then, the second step involves choosing an algorithm to train the model. When the training is done, the model will predict what picture corresponds to what object.

After that, it is easy to use the model to predict new images. For each new image feeds into the model, the machine will predict the class it belongs to. For example, an entirely new image without a label is going through the model. For a human being, it is trivial to visualize the image as a car. The machine uses its previous knowledge to predict as

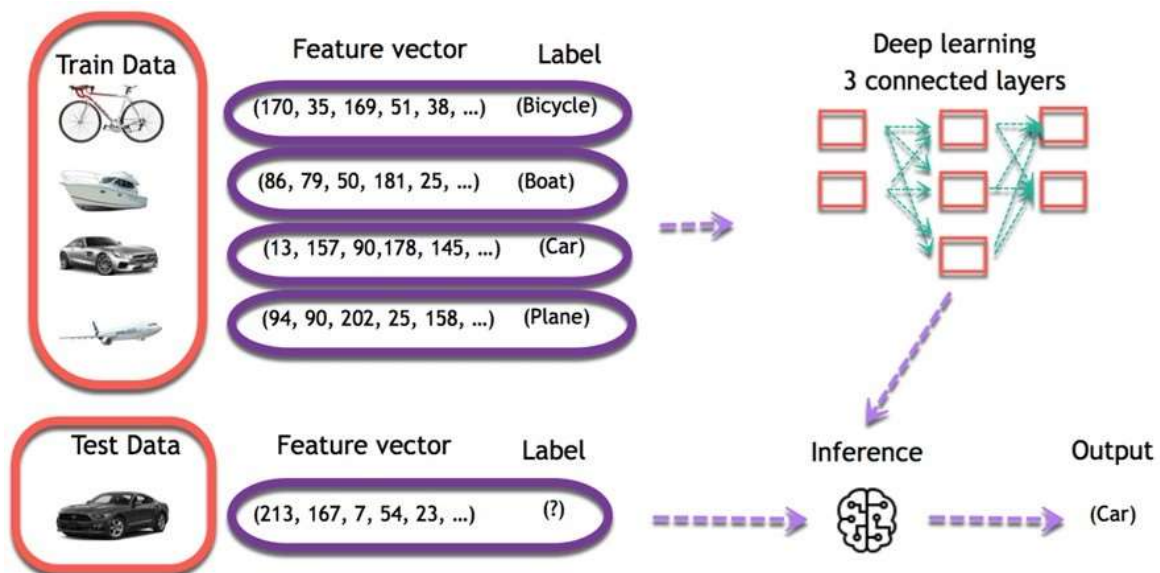
well the image is a car.

Deep Learning Process

In deep learning, the learning phase is done through a neural network. A neural network is an architecture where the layers are stacked on top of each other.

Consider the same image example above. The training set would be fed to a neural network

Each input goes into a neuron and is multiplied by a weight. The result of the multiplication flows to the next layer and become the input. This process is repeated for each layer of the network. The final layer is named the output layer; it provides an actual value for the regression task and a probability of each class for the classification task. The neural network uses a mathematical algorithm to update the weights of all the neurons. The neural network is fully trained when the value of the weights gives an output close to the reality. For instance, a well- trained neural network can recognize the object on a picture with higher accuracy than the traditional neural net.



Automate Feature Extraction using DL

A dataset can contain a dozen to hundreds of features. The system will learn from the relevance of these features. However, not all features are meaningful for the algorithm. A crucial part of machine learning is to find a relevant set of features to make the system learn something.


One way to perform this part in machine learning is to use feature extraction. Feature extraction combines existing features to create a more relevant set of features. It can be done with PCA, T-SNE or any other dimensionality reduction algorithms.

For example, in image processing, the practitioner needs to extract the feature manually in the image like the eyes, the nose, lips and so on.

Those extracted features are fed to the classification model.

Deep learning solves this issue, especially for a convolutional neural network. The first layer of a neural network will learn small details from the picture; the next layers will combine the previous knowledge to make more complex information. In the convolutional neural network, the feature extraction is done with the use of the filter. The network applies a filter to the picture to see if there is a match, i.e., the shape of the feature is identical to a part of the image. If there is a match, the network will use this filter. The process of feature extraction is therefore done automatically.

TRADITIONAL MACHINE LEARNING



DEEP LEARNING



Difference between Machine Learning and Deep Learning


	Machine Learning	Deep Learning
Data Dependencies	Excellent performances on a small/medium dataset	Excellent performance on a big dataset
Hardware dependencies	Work on a low-end machine.	Requires powerful machine, preferably with GPU: DL performs a significant amount of matrix multiplication
Feature engineering	Need to understand the features that represent the data	No need to understand the best feature that represents the data
Execution time	From few minutes to hours	Up to weeks. Neural Network needs to compute a significant number of weights
Interpretability	Some algorithms are easy to interpret (logistic, decision tree), some are almost impossible (SVM, XGBoost)	Difficult to impossible

When to use ML or DL?

In the table below, we summarize the difference between machine learning and deep learning.

	Machine learning	Deep learning
Training dataset	Small	Large
Choose features	Yes	No
Number of algorithms	Many	Few
Training time	Short	Long

With machine learning, you need fewer data to train the algorithm than deep learning. Deep learning requires an extensive and diverse set of data to identify the underlying structure. Besides, machine learning provides a faster-trained model. Most advanced deep learning architecture can take days to a week to train. The advantage of deep learning over machine learning is it is highly accurate. You do not need to understand what features are the best representation of the data; the neural network learned how to select critical features. In machine learning, you need to choose for yourself what features to include in the model.



Summary

Artificial intelligence is imparting a cognitive ability to a machine. Early AI systems used pattern matching and expert systems.

The idea behind machine learning is that the machine can learn without human intervention. The machine needs to find a way to learn how to solve a task given the data.

Deep learning is the breakthrough in the field of artificial intelligence. When there is enough data to train on, deep learning achieves impressive results, especially for image recognition and text translation. The main reason is the feature extraction is done automatically in the different layers of the network.

[Buy Now \\$9.99](#)