Machine Learning (Reinforcement learning, Neural networks and Deep Learning)

Prof K R Chowdhary

MBM University

September 13, 2024

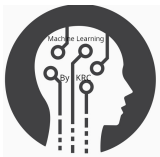

Active vs. Passive learning

Learning paradigms can vary by the role played by the learner. An *active learner* interacts with the environment at training time, say, by posing queries or performing experiments, while a passive learner only observes the information provided by the environment (or the teacher) without altering it or directing it to other resource. In an active setting, e.g., one could think of asking users to label specific e-mails chosen by the learner, or even composed by the learner, to enhance its understanding.

The learner of a spam filter is usually passive – waiting for users to mark the e-mails coming to them. This learning is used by the system to mark emails as spam for all the future emails.

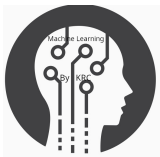

Data for Supervised and Unsupervised Learning

Both supervised and unsupervised learning tasks requires representation of input data, in the form of a table. Each data point that we want to reason about (each email, each customer, each transaction) is a row, and each property that describes that data point (say, the customer age or location of a transaction) is a column.

Image can be described by, e.g., image of a tumor, by the gray-scale values of each pixel, or maybe by using the size, shape, and color of the tumor.

Each *entity* or *row*, is known as a sample (or data point) in machine learning, while the columns – the properties that describe these entities –are called *features*.

Building a good representation of data is called *feature extraction* or feature engineering. No machine learning algorithm can make a prediction on data for which it has no information.

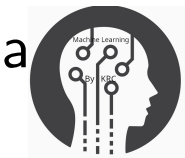

Reinforcement learning

A simple example of reinforcement learning (RL) is the way an animal is trained to take actions based on rewards. If a dog gets treat every time it obeys the master's command, it will carry on actions each time it is commanded.

Thus RL is a *behavioral learning* where algorithm receives feedback by the analysis of data so that the machine is guided towards the best outcome.

RL differs from supervised learning, as in RL, the system is not trained with the sample data set, rather, the system learns through trial and error, and a sequence of successful decisions will result in process being "reinforced" to best solve the problem at hand.

RL is used in robotics. A training robot that navigates through floor filled with obstacles changes its approach to navigating through obstacles based on the outcome of its actions. If it collides with an object, it backtracks and re-calibrates path, thus trained by trial and error.

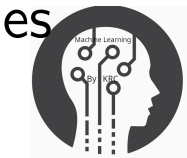

Learning through Neural Networks

An average five-year-old child can easily recognize difference between his teacher madam's face and that of his mother. But, a computer with conventional programming, cannot do it. However, the Neural networks can do it very easily.

A neural network consists of three or more layers: an input layer, one or more hidden layers, and an output layer. Data is input through the input layer, then it is modified in the hidden layer(s) and are output to output layer based on the weights applied to these nodes at one or more hidden layers.

Using an iterative approach, a neural network continuously adjusts and makes inferences until some stopping point is reached.

The Neural Networks have applications in image recognition, natural language and speech processing, and computer vision.

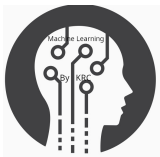

Deep Learning Networks

Deep learning networks are complex neural networks, designed to emulate human brain. Computers using these can be trained to deal with *abstractions* and problems that are *poorly defined*.

A typical neural network may consist of thousands or even millions of simple processing nodes that are densely interconnected.

Deep learning uses multiple neural network layers to learn from data, in an iterative manner. They are useful in learning patterns from unstructured data.

Deep learning uses hierarchical neural networks to learn from a combination of *unsupervised* and *supervised* algorithms. Typically, deep learning learns from unlabeled and unstructured data. The more complex the problem, the more hidden layers are there.