An Introduction to Neural Networks

David Bisant, PhD
Central Security Service
bisant@umbc.edu

Abstract

Neural networks ares an important type of machine learning that have generated significant interest due to their success on difficult problems.. This tutorial will provide an introduction and a high-level overview of artificial neural networks. It will include the biological basis and the inspiration for these methods, and their colorful history. An explanation of deep learning and deep belief networks, hardware acceleration, and case studies will be used to illustrate how the technology works and the direction in which it is developing.

Description

Artificial neural networks are computing systems inspired by the brain's biological neural networks. There are mostly used for machine learning and recent successes on difficult problems have raised their stature. Neural networks are now the most popular form of machine learning. As a form of machine learning, they are very versatile and can be applied to a wide range of problems, such as pattern recognition, prediction, transformation, control, and filtering.

Modern applications are found in biomedicine, robotics, text mining, diagnostic problems, telecommunications, power systems, signal processing; medical and health applications, text, video, and multi-media mining, E-commerce and web data, financial data analysis, cyber security, remote sensing, earth sciences, bioinformatics, and astronomy.

Deep belief or deep learning neural networks are larger and more complex forms of neural networks which have generated a lot of interest over the last 10 years due to their success on speech recognition, image recognition, and natural language processing.

Neural networks tend to be computationally intensive, especially when learning. An entire industry has arisen to provide the computational resources needed for future applications. These include systems based on GPUs and other GPU- like accelerators. Additionally, a race is on for neural network chips which will be needed for the self driving cars and autonomous systems that will make our future world.

Outline

Why Neural Networks?
Types of Machine Learning
Biological Neurons, Networks, and Processing
Artificial Neurons, Networks, and Processing
Historical Developments

Machine Learning Problem Areas
Supervised, Reinforcement, Unsupervised Learning
Training Neural Networks
The Machine Learning Process
Data Preparation and Preprocessing
Model Induction and Generalization
Model Interpretation and Validation

Ensembles
Deep, Convolutional, and Recurrent Networks
Versus Other Methods
Neural Network Algorithms
Hardware Acceleration
Future Developments
Resources

References

Ham, Fredric M., and Ivica Kostanic. Principles of neurocomputing for science and engineering. McGraw-Hill, 2000.

Xu, Rui, and Don Wunsch. Clustering. Vol. 10. John Wiley & Sons, 2008.

Haykin, Simon S., et al. Neural networks and learning machines. Vol. 3. Upper Saddle River: Pearson education, 2009.

Haykin, Simon. Neural networks: a comprehensive found. Prentice Hall, 1994.