An Introduction to Machine Learning and Neural Networks

—

Andrew Garrett

Agenda Today

I. Historical Context
II. What is Machine Learning?: Tasks and Goals
III. What is Machine Learning?: Implementations
    A. Simple Models
    B. Neural Networks
IV. Why Machine Learning?: Uses and Applications
V. Machine Learning as an Analog to Neurobiology
VI. Closing Thoughts

Historical Context

Machine Learning, Artificial Intelligence, and Neural Networks have been around much longer than one might expect.

Origins: Pre-1950

- The first computer was technically invented in 1834 by Charles Babbage, but was very different from any notion of computing device we know today
- Fritz Lang introduces the general public to the concept of AI in '*Metropolis*' (1927), the first film in history to depict a robot
- Alan Turing, one of the greatest mathematicians and computer scientists in modern history, develops the first mathematical model of computation known as a '**Turing Machine**,' or '**Universal Machine**' in 1936

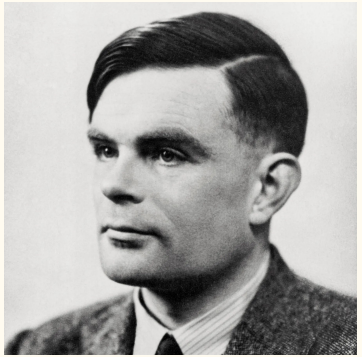

Early Progress: 1950-1990

- In 1952, IBM scientists develop the first program which ‘learns’ how to play the game Checkers
  - This would set the stage for games being a popular testing environment for ML algorithms due to their simplistic actions and complex outcomes
- Stanley Kubrick’s “*2001: A Space Odyssey*” gives a dark but not-all-that-unreasonable take on the human, AI interaction
- NETtalk, a neural network exploring human cognitive behavior, learns to pronounce 20,000 English words in the span of one week

Cementing Popular Status: 1990-2005

- In a highly publicized 1997 event, a chess-playing algorithm dubbed *Deep Blue* beats Grandmaster Garry Kasparov, making it nominally the best chess player in the world
- In 1998, AT&T Bell Labs developed the first neural network implementation which could recognize handwritten digits using letters from the US Postal Service
- Several very important advancements came algorithmically in the 1990s, notably including Long Short-Term Memory Networks and AdaBoost.

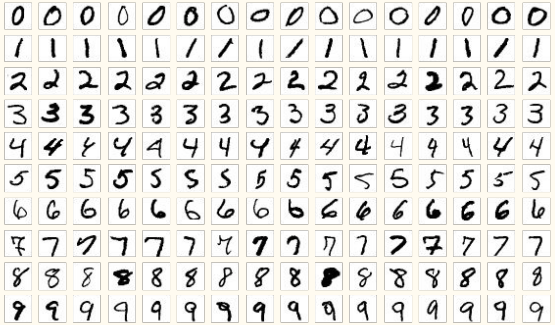

Dominating the Modern World: 2005-present

- In 2006, Neural Network research is rebranded as Deep Learning research in order to stimulate interest and funding
  - The charge is led by Geoffrey Hinton and Yann Lecun
- By 2011, Google and IBM's respective AI projects, Brain and Watson, had made huge strides in voice, text, and image recognition
- An algorithm developed by Google called AlphaGo beats world Go champion Lee Sedol 4-1 in a five game series

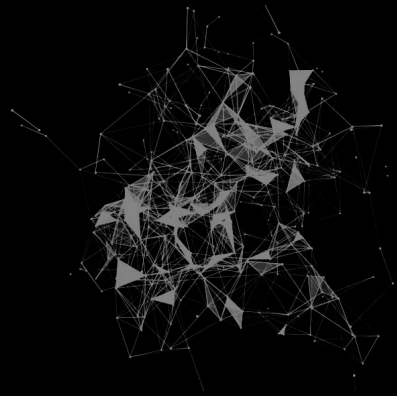

ML Tasks & Goals

Machine Learning - An algorithmic approach for computers to learn and analyze patterns in data, for robust modelling and inference.

Supervised Learning

- Supervised Learning constitutes developing some relationship between inputs and outputs when given example inputs and outputs
  - This means that a model is *trained* to develop a relationship between different features/attributes (inputs) of the data with the desired label (outputs) of the data.
- **Classification** focuses on determining some common label or identity to individual data points
- **Regression** focuses on finding a general trend between some features or attributes of data points and the desired label of the data points

Unsupervised Learning

- Unsupervised Learning uses only inputs to determine patterns and groupings in data
- This is primarily used to find clusters and other appropriate ways to group data with similar properties
- **Clustering** essentially takes in *unlabelled* features or attributes of data points and finds some general pattern, shape, or grouping of such points according to their proximity to other data points
- **Dimensionality Reduction** is a very common technique which uses patterns among different features or attributes to lessen the number of features or attributes required to accurately represent the data

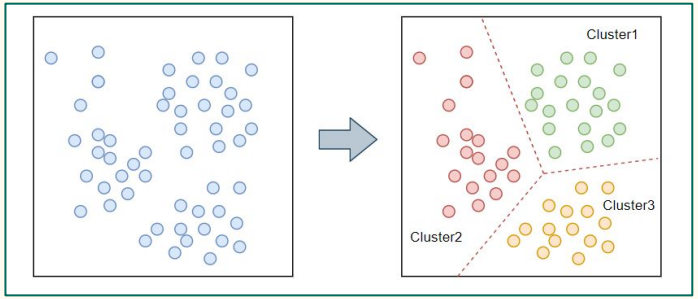

Reinforcement Learning

- Reinforcement Learning is a unique form of learning where an *agent* learns the optimal way to complete a task using *trial-and-error*.
- This means that the *agent* develops a probabilistic understanding of the consequences of an action it might make
    - Logically, this means that our *agent* will assign low probability to actions that may cause it to fail in whatever task it is looking to accomplish.
    - Furthermore, the *agent* learns to give very strong probability to actions that serve to benefit it in accomplishing its goal

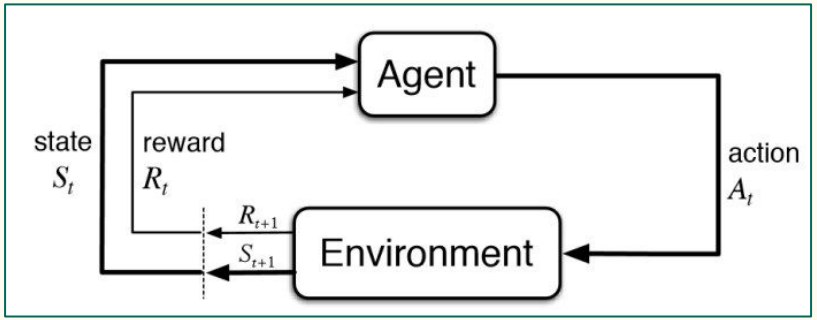

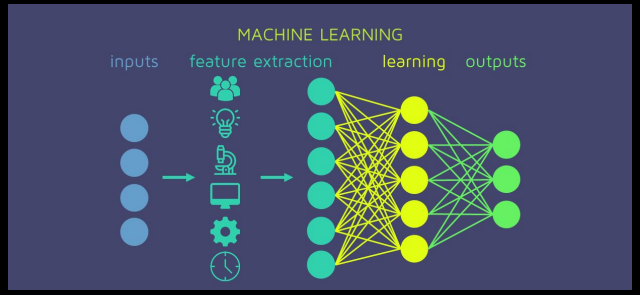

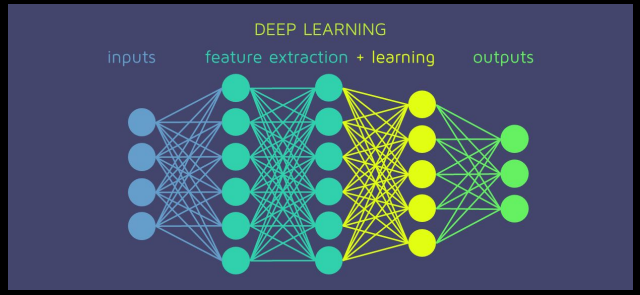

ML Implementation

Machine Learning has a huge range of approaches and implementations, with diversity in intricacy, ease of use, and adaptability.

ML Models: Classification

- K-Nearest Neighbors
- Logistic Regression
- Support Vector Machines
- Decision Trees
- **AdaBoost**
- **Random Forests**

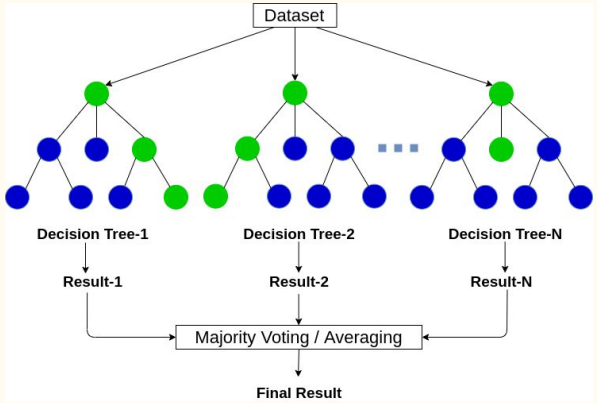

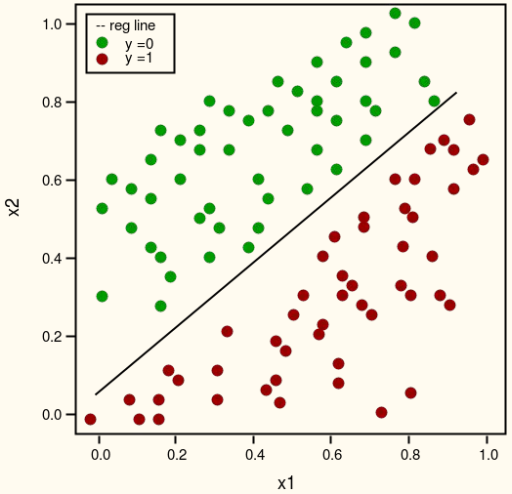

ML Models: Regression

- Linear Regression
- Lasso Regression
- Ridge Regression
- In fact, all of the aforementioned ML algorithms developed for classification have regressive counterparts
    - This reinforces the adaptivity of ML and its many uses

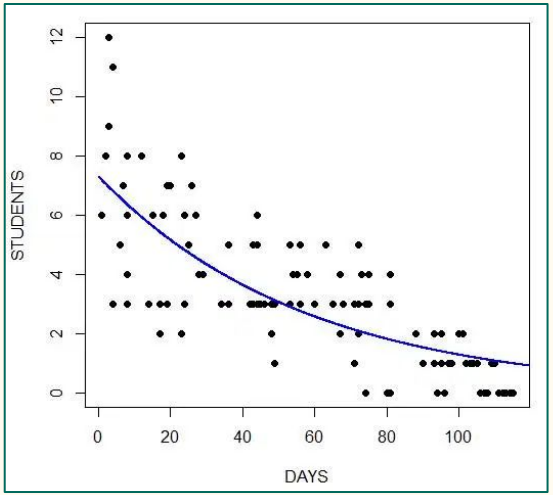

Reinforcement Learning

- Model-Free
  - Q-Learning
  - Deep Q-Learning
- Model-based
  - Sampling-based approaches

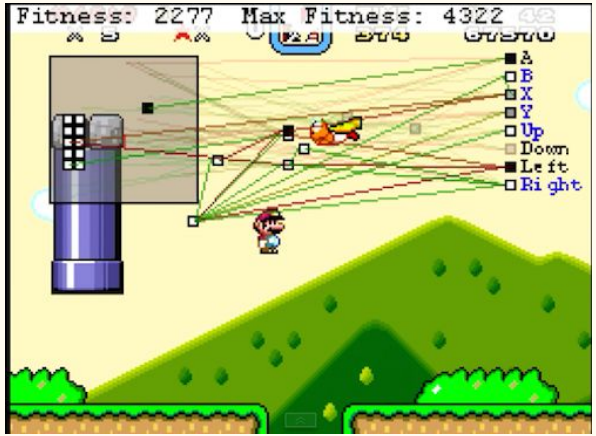

Deep Learning: Artificial Neural Networks (ANN)

- Also known as Multilayer Perceptrons (MLP), this network architecture is most commonly referred to as a '**universal function approximator**'
  - All this means is that an MLP could determine almost any relationship between variables, be it linear or very complex and non-linear
- This type of network takes in input data and feeds information forward to develop predictions for what the corresponding outputs might look like

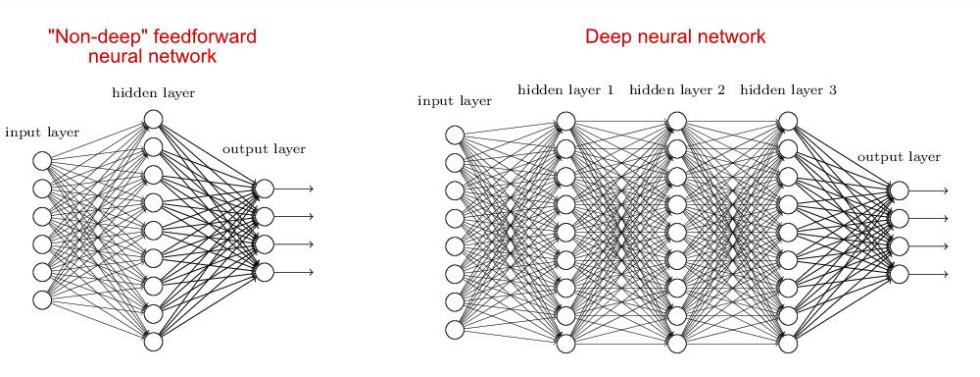

Deep Learning: Recurrent Neural Networks (RNN)

- RNNs introduce us into the extremely important field of Natural Language Processing (NLP)
  - NLP studies how a machine interprets human speech and text
- RNNs are engineered for predictive modeling, so the principal of their design is using prior data to influence the prediction of future data
- Some important variants of RNNs are LSTM and Transformers, which are both used in predictive analytics, NLP, and consumer recommendation systems.

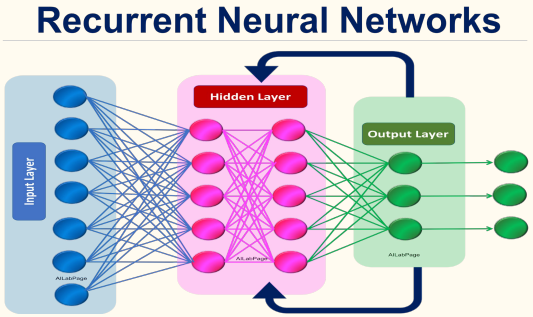

Deep Learning: Convolutional Neural Networks (CNN)

- CNNs make up the cutting edge Deep Learning architecture used in computer vision and machine perception
    - This means that CNNs are tailored to image recognition and segmentation
- CNNs use a special tool called a convolution which is a method used to break down an image into different parts or features
- CNNs have long been refined and improved, which has been very helpful in industries like self-driving cars, other autonomous vehicles, and facial recognition software

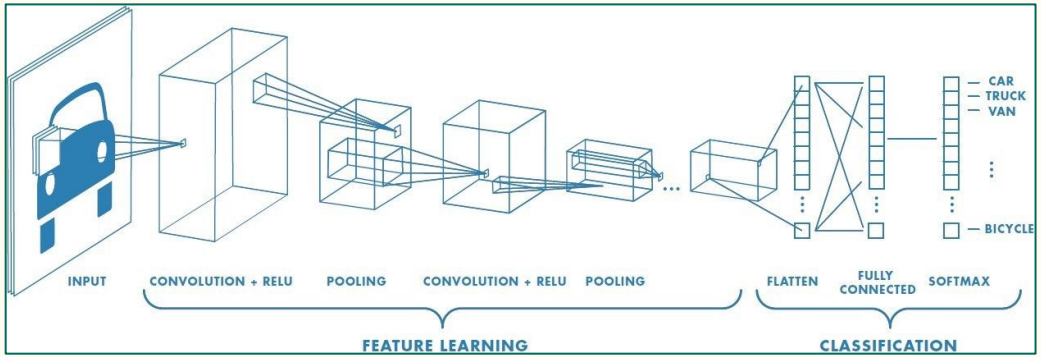

ML Uses & Applications

Machine Learning is extremely adaptable to real-world scenarios, but there are limitations and caveats.

Modern Contexts

- Much of the Deep Learning methods (Neural Networks), are used only in the quickly evolving tech industry
  - As such, Deep Learning remains a very hot topic in academic and research fields
- Almost any company in the modern world uses ML in some capacity; some important examples include:
  - Chatbots
  - Support of decision-making/policy-change
  - Customer recommendation systems
  - Market Research
  - Predictive Modeling (pricing, human behavior, etc)

Neuroscience and ML

While the guiding principles of ML lie in statistics and math, advancements are beginning to draw heavily from neurobiological inspiration.

Neural Networks: Biological inspiration

- The derivative nature of deep learning is fairly evident by the content covered thus far
  - As we have discussed previously, the terminology for biological neural networks is the same as that of the like in deep learning contexts
- The fact that our most powerful techniques in ML (and by extent DL) are networks mimicking the functions of human neurobiology gives strong weight to the convergence of computers to the human brain.

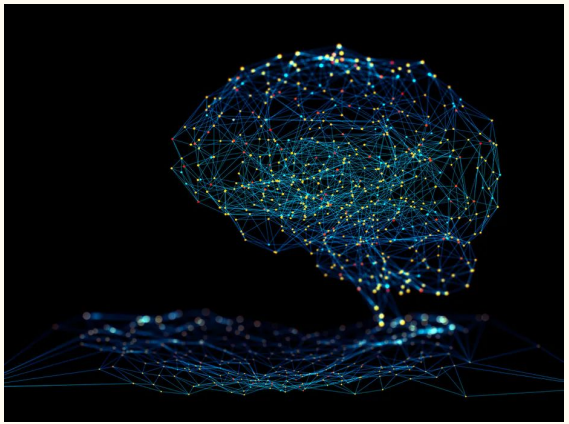

Closing Thoughts

Machine Learning is a beautiful and growing field of problem-solving. With appropriate care and patience, it's existence and evolution can bring great benefit to human society.

Questions?