History of Machine Learning

Shrey Gupta

Applied Machine Learning (HOUSECS 59-01), Duke University

September 5, 2018

Historical Computational Advancements

- Before machine learning, there were several advances required in computing.

Historical Computational Advancements

- **1642:** Blaise Pascal invents the first mechanical calculator.
- **1679:** Gottfried Wilhelm Leibniz creates the modern binary system.
- **1834:** Charles Babbage, "father of the computer", invents the punch card.
- **1842:** Ada Lovelace becomes the first computer programmer.

Appearances in Media

- Even though machine learning and artificial intelligence weren't fields at the time, with advances in computers, glimpses of the potential appear in media.
- In 1927, Fritz Lang introduces a futuristic robot in his sci-fi film, bringing hints of AI.

Birth of Machine Learning

- The 1950s and 1960s were a crucial period for the birth of machine learning.
- The following two decades, the 1970s and 1980s, would be termed the "Artificial Intelligence winter" for a steep decline in machine learning and artificial intelligence funding and research.

Birth of Machine Learning

- **1950:** Alan Turing proposes the idea that machines can learn. He proposes the Turing test, in which a human should not be able to differentiate between a human and computer.
- **1952:** Arthur Samuel designs a program that can play checkers.
- **1957:** Frank Rosenblatt creates the "perceptron", a simple linear classifier.

Birth of Machine Learning

- **1963:** Donald Michie develops a reinforcement learning machine to play tic-tac-toe.
- **1967:** The nearest neighbor algorithm is invented, crucial to pattern recognition.

Renewal of Machine Learning

- After the Artificial Intelligence winter during the 1970s and 1980s, we see a renewal in the field during the 1990s.

Renewal of Machine Learning

- **1995:** Tin Kam Ho creates the random forest algorithm, important for interpretable machine learning.
- **1995:** Corinna Cortes and Vladimir Vapnik invent support vector machines, a crucial step forward in machine learning performance.
- **1997:** Robert Schapire and Yoav Freund invent the AdaBoost algorithm, another improvement in machine learning.

Modern Day Achievements

- **1997:** IBM's Deep Blue beats chess world champion Gary Kaspaov.
- **2009:** The Netflix Prize is won for the best recommender system in predicting user film ratings.
- **2011:** IBM's Watson is able to defeat human champions in *Jeopardy!*
- **2012:** Google Brain successfully trains a neural network to differentiate images of cats from dogs.

Modern Day Achievements

- **2014:** Facebook's DeepFace successfully uses neural networks to perform facial recognition with over 97% accuracy.
- **2014:** The "Eugene Goostman" chatbot fools a third of judges in the Turing test.
- **2016:** DeepMind develops AlphaGo and beats the top-ranked Go player. AlphaGo Zero, which is generalized to chess and other games, is developed the following year.

Other Advances

- In the meantime, there were strong advances in fields related to machine learning, including deep learning, natural language processing, and computer vision.

Deep Learning

- Deep learning involves the construction of artificial neural networks, composed of several layers of "neurons". We'll explore them later in the course.
- It has been most successfully applied to natural language processing and computer vision with recurrent neural networks (RNNs) and convolutional neural networks (CNNs), respectively.

Deep Learning

- **1951:** Marvin Minsky and Dean Edmonds design the first neural network. However, 18 years later, a book is published describing neural networks' limitations, putting them out of favor for a considerable amount of time.
- **1970:** Seppo Linnainmaa develops what becomes the "backpropogation" algorithm, a crucial algorithm for neural networks. The algorithm is fully described by David Rumelhart, Geoff Hinton, and Ronald Williams in 1986.

Natural Language Processing

- Natural language processing is a subset of artificial intelligence concerned with understanding natural language, including text and speech.
- Examples include sentiment analysis, language translation, reading comprehension, and textual question-answering.

Natural Language Processing

- **1972:** Karen Sparck Jones invents TF-IDF (term frequency-inverse document frequency), crucial to NLP. The algorithm identifies the importance of a word or phrase in a corpus.
- **1982:** John Hopfield publishes about recurrent neural networks (RNNs), crucial to NLP for "memory" in sequences.
- **1985:** Terry Sejnowski creates NetTalk, an algorithm that learns to pronounce in a similar fashion to children.
- **1997:** Sepp Hochreiter and Jurgen Schmidhuber create LSTM (long short-term memory) for RNNs, greatly improving their performance.

Computer Vision

- Computer vision is a related field that involves the understanding, processing, and reconstruction of 2- and 3-dimensional images.
- Common computer vision tasks in machine learning include classification, localization, object detection, and landmark detection.

Computer Vision

- **1998:** Yann LeCun organizes the MNIST database of handwritten digits, creating a valuable source for handwriting recognition. He develops LeNet, a convolutional neural network that can classify handwritten digits.
- **2009:** Fei-Fei Li develops the ImageNet database of images.
- **2012:** Trained on ImageNet, AlexNet makes an important step forward in image classification with great improvements in performance.
- **2015:** Joseph Redmon invents "You Only Look Once" (YOLO), performing real-time object detection with performance higher than ever before.

References

- *A Brief History of Machine Learning*
- *A History of Machine Learning*
- *Timeline of Machine Learning*