Overview

AI researchers break down intelligence into two forms: *Fluid intelligence* is about devising new methodologies for solving problems. *Crystallized intelligence* is about applying proven methodologies and algorithms to solving problems. Machine learning (ML) is about leveraging crystallized intelligence to build models that statistically relate columns of an excel spreadsheet to other columns with a formula so that when new rows of the former columns are observed, cells in the latter columns can be predicted. There are many types of Machine learning algorithms, some are particularly effective with some types of data, such as images, or natural language. But they all adopt a statistical approach to come up with the formula, in other words, they learn from observations rather than abstract theories about matter which is the most typical epistemological approach to knowledge. Probably the most talked about Machine Learning model is the *neural* model, which mirrors in some regards how biological brains build models to make sense of the world, learn from experience, and make predictions. They usually require more observations than other models and many people feel that the resulting models are less explainable, generating fear that neural models may one day take over the world from humans. This course is an introduction to the field of Machine Learning, concentrating on feedforward, recurrent, convolutional, and Transformer-based neural models. It will focus on the machine learning methods that have proven valuable and successful in practical applications, which includes decision trees and Support Vector machines. It will also explain basic issues that confront machine learning methods, such as trust and interpretability. With the successful completion of the course, the student will be ready to apply for jobs in the life sciences, financial, advertising, and social Web industries and build Machine Learning models to accelerate their speed of business.

Pre-requisites

You should understand basic probability and statistics and college-level algebra and calculus. For example it is expected that you know about standard probability distributions (Gaussians, Poisson), and also how to calculate derivatives. Knowledge of linear algebra is also expected, and knowledge of mathematics underlying probability models will be useful. A programming background in python is required as we will build Machine Learning models with python-based frameworks. A prerequisite class is INFO 6105.

Textbook

The Science of Deep Learning, Iddo Drori, Cambridge University Press.

Course requirements and grading

The grading in the class is divided up as follows:

Assignments 30%

Mid-Term Exam 30%

Final Projects (teams) 30%

Final Exam 10%

Class Modules

Decision Trees and Random Forests

Support Vector Machines

Probabilistic Classifiers, Naive Bayes Classifiers

Feedforward and Dense Neural networks

Recurrent Neural networks

Convolutional Neural networks

Transformer Neural networks

Interpretability and Trust

Reinforcement Learning

Probabilistic Neural Networks and an introduction to Quantum Computing