CS 3750 Advanced Topics in Machine Learning (Spring 2020)

Course abstract

The goal of the field of machine learning is to build computer systems that learn from experience and that are capable to adapt to their environments. Learning techniques and methods developed by researchers in this field have been successfully applied to a variety of learning tasks in a broad range of areas, including, for example, text classification, gene discovery, financial forecasting, credit card fraud detection, collaborative filtering, design of adaptive web agents and others.

The objective of the Advances Machine Learning course is to expand on the material covered in the introductory Machine Learning course (CS2750). It focuses on special topics in ML such as exact and approximate inference in graphical models, dimensionality reduction and component analysis methods, latent variable models, models of documents and words, time series models, selected topic from deep neural networks and reinforcement learning. The course will consist of (student-lead) presentations and discussions. Students will be evaluated based on their participation in discussions, presentations and projects.

Prerequisites

CS 2750 Machine Learning, or the permission of the instructor.

Tentative topics to be covered:

- **Graphical models**
  - BBNs and MRFS
  - Exact inference
  - Approximation inference: Monte Carlo, Variational
- **Low dimensional representation of data**
  - PCA, auto-encoders
  - Probabilistic latent variable models: pPCA, CVQ, NOCA
  - Modern auto-encoders:
    - Restricted Boltzman machines
    - Variational autoencoders
- **Document analysis models**
  - SVD, LSA, non-negative matrix factorization
  - pLSA, LDA
  - Word similarity and word models: word2vec, CBOW
- **Sequence/Time series models**
  - Discrete state models: Markov models, DBNs, HMMs
  - Continuous state models: AR, Linear dynamical system (LDS)
  - Modern autoregressive models: RNNs, LSTMs and GRUs
- **Topics in deep NNs**
  - Self-attention, transformers models and all that
  - Convolutional neural networks
  - Generative adversarial networks
- **Sequential decision making**
  - MDPs: model and algorithms
  - POMDPs: model and algorithms
  - Reinforcement learning: Q learning, policy gradient, Actor critic
  - Deep reinforcement learning