Representation Learning

Professor: Pramod Viswanath

TA: Ashok Vardhan

Fall 2017

What is the course about?

- Real valued, high dimensional data x
  - obtain a lower dimensional representation
  - parametric vs geometric methods
    - Reduce dimension with respect to target y

- Sometimes data is discrete/atomic
  - Example: words, computer programs, graphs
  - Question: how to represent such data?

Prerequisites

- Mathematical Maturity
  - Probability, Statistics, Optimization, Algorithms
- Programming Felicity
  - Tensorflow, Python

Schedule: 4 Modules

- Dimensionality Reduction Methods
  - geometric dimensionality reduction
- Mixture Models
  - parametric dimensionality reduction
- Neural Networks
  - the ZOO
- Representation Methods
  - word2vec

Module 1: Geometric Dimensionality Reduction

- Data is high dimensional, real valued
- Lecture 1: Principal Component Analysis (PCA)
- Lecture 2: Canonical Component Analysis (CCA)
- Lecture 3: Kernel-CCA, ACE algorithm
- Lecture 4: Nonlinear methods: Isomap, t-SNE

Module 2: Parametric Dimensionality Reduction

- Lecture 5: Nonparametric: kernel and nearest neighbor methods
- Lecture 6: family of mixture models
- Lecture 7: Gaussian mixtures and EM algorithm and method of moments
- Lecture 8: Method of Moments at scale: tensor decompositions
- Lecture 9: Mixture models for supervised learning: regression and classification mixtures
- Lecture 10: Mixture of Experts

Module 3: Neural Networks

- Lecture 11: Architectures
  - Feedforward neural networks (CNN, Resnet)
  - Recurrent neural networks (GRU, LSTM, Attention)
- Lecture 12: Algorithms
  - Backpropagation, Stochastic gradient descent

Module 4: Representation Methods

- NLP
  - Lecture 13: Language models (KN and GT smoothing)
  - Lecture 14: Word2vec (Skipgram and CBOW) and GloVe
  - Lecture 15: Doc2vec, Sentence2vec, Skip-thought
  - Lecture 16: Polysemy, Compositionality
  - Lecture 17: Deep learning pipeline for NLP

Module 4: Representation Learning

- Other representations
  - Lecture 18: Musical Instruments
  - Lecture 19: Graphs
  - Lecture 20: Computer Programs

Course Material

- Lecture notes (scribed), Slides
- Research literature
- Contemporary research papers
- Online resources including blogposts
- Tensor Flow

Grading

- Scribing
  - 3 lectures need to be scribed
  - integrate lecture references with material taught in class
- Project
  - Single or in groups of 2.
  - Report and presentation at the end of the course