

What's Next?

David Rosenberg

New York University

July 26, 2017

Machine Learning

- Our course was heavy on model-free methods (i.e. no assumptions about $\mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$.)
- Worth looking more into probabilistic approaches
 - Joan Bruna's DS-GA 1005: graphical models, probabilistic learning and inference
 - Petrov's CSCI-GA 3033-001: Statistical NLP
 - books by David Barber and Kevin Murphy
- Learn neural networks!
 - Cho's DS-GA 3001: Neural networks for NLP
 - Bowman's LING-GA 3340-002: Seminar in Semantics: Artificial NNs
 - LeCun's Neural networks / vision course

Machine Learning

- 1 Look at other course notes at this level.
 - Every course covers a different subset of topics.
 - Different perspectives. (e.g. Bayesian / Probabilistic)
- 2 Read on some “second semester” topics
 - LDA / Topic Models (DS-GA 1005)
 - Sequence models: Hidden Markov Models / MEMMs / CRFs (DS-GA 1005)
 - Bayesian methods
 - Collaborative Filtering / Recommendations
 - Ranking
 - Bandit problems (Thompson sampling / UCB methods)
 - Gaussian processes

Other Stuff To Learn

- Statistics
- Data Structures & Algorithms (Theoretical)
- Some production programming language (e.g. Java, C++)