Oxford University, Hilary term 2020, Syllabus for: Advanced Econometrics

Foundations of machine learning

instructor	Maximilian Kasy
email	teachingmaxkasy@gmail.com
class time	Th 16:00-17:30 and Fr 11:30-13:00
location	Seminar room C, Manor Road building
webpage	https://maxkasy.github.io/home/ML_Oxford_2020/

Overview and Objectives

This class will cover some of the theoretical foundations of machine learning. The class starts with a review of statistical decision theory, which provides the conceptual framework for the rest of the course. We then consider regularization and data-driven choice of tuning parameters. We will discuss the canonical normal means model. In this model, we will motivate shrinkage estimators in different ways, and will prove the famous result that shrinkage estimators can uniformly dominate conventional estimators. As an example of a supervised learning method that builds on these ideas, we will discuss (deep) neural nets, including some numerical methods used for training them, such as stochastic gradient descent.

After that, we will discuss methods for active learning in the context of multi-armed bandit settings. We will review some theoretical results providing performance guarantees (regret bounds) for algorithms used for learning in bandit settings. We will then turn to a generalization of bandit problems, Markov decision problems, and will discuss reinforcement learning approaches for solving these.

Practice problems I have posted example exam question on my course homepage. Additionally, the slides contain a lot of practice problems, which you will have to solve in class. The idea is to have you complete most of the proofs, after I pointed you in the right direction. After a few minutes, we will discuss the solutions to these problems. These problems provide good guidance for what you might expect from the exam.

If you need any special accommodations for physical or medical reasons, please see me after class or send me an email.

Outline of the course

Decision theory

- Basic definitions
- Optimality criteria
- Relationships between optimality criteria
- Analogies to microeconomics
- Two justifications of the Bayesian approach

Shrinkage in the normal means model

- Setup: the normal means model $\boldsymbol{X} \sim N(\boldsymbol{\theta}, I_k)$ and the canonical estimation problem with loss $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2$.
- The James-Stein (JS) shrinkage estimator.
- Three ways to arrive at the JS estimator (almost):
  1. Reverse regression of θ_i on X_i.
  2. Empirical Bayes: random effects model for θ_i.
  3. Shrinkage factor minimizing Stein's Unbiased Risk Estimate.
- Proof that JS uniformly dominates $\boldsymbol{X}$ as estimator of $\boldsymbol{\theta}$.
- The normal means model as asymptotic approximation.

Deep neural nets

- What are neural nets?
- Network design:
  Activation functions, network architecture, output layers.
- Calculating gradients for optimization:
  Backpropagation, stochastic gradient descent.
- Regularization using early stopping.

Bandit problems

- Setup: The multi-armed bandit problem.
  Adaptive experiment with exploration / exploitation trade-off.
- Two popular approximate algorithms:
  1. Thompson sampling
  2. Upper Confidence Bound algorithm
- Characterizing regret.
- Characterizing an exact solution: Gittins Index.
- Extension to settings with covariates (contextual bandits).

Reinforcement learning

- Markov decision problems.
- Expected updates – dynamic programming.
- Sample updates:
  - On policy: Sarsa.
  - Off policy: Q-learning.
- Approximation:
  - On policy: Semi-gradient Sarsa.
  - Off policy: Semi-gradient Q-learning.
  - Deep reinforcement learning.

References

Review of decision theory

Robert, C. (2007). *The Bayesian choice: from decision-theoretic foundations to computational implementation.* Springer Verlag, chapter 2.

Shrinkage in the normal means model

Wasserman, L. (2006). *All of nonparametric statistics.* Springer Science & Business Media, chapter 7.

Stigler, S. M. (1990). The 1988 Neyman memorial lecture: a Galtonian perspective on shrinkage estimators. *Statistical Science*, pages 147–155.

Morris, C. N. (1983). Parametric empirical Bayes inference: Theory and applications. *Journal of the American Statistical Association*, 78(381):pp. 47–55.

Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 9(6):1135–1151.

van der Vaart, A. W. (2000). *Asymptotic statistics.* Cambridge University Press, chapter 7.

Hansen, B. E. (2016). Efficient shrinkage in parametric models. *Journal of Econometrics*, 190(1):115–132.

Abadie, A. and Kasy, M. (2019). Choosing among regularized estimators in empirical economics - the risk of machine learning. *Review of Economics and Statistics*, forthcoming.

Deep neural nets

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning.* MIT Press, chapters 6-8.

Bottou, L., Curtis, F. E., and Nocedal, J. (2018). Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311

Bandit problems

Bubeck, S. and Cesa-Bianchi, N. (2012). Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122.

Russo, D. J., Roy, B. V., Kazerouni, A., Osband, I., and Wen, Z. (2018). A Tutorial on Thompson Sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96.

Weber, R. et al. (1992). On the Gittins index for multiarmed bandits. *The Annals of Applied Probability*, 2(4):1024–1033.

Kasy, M. and Sautmann, A. (2019). Adaptive treatment assignment in experiments for policy choice. *Working Paper*.

Reinforcement learning

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., and Pineau, J. (2018). An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354.