

IEOR 265 – Learning and Optimization

Spring 2017

- Instructor:** Anil Aswani
4119 Etcheverry
Office hours – Tu 10-11A; Th 230-330P
aaswani@berkeley.edu
- GSI:** Pedro Hespanhol
4176-B Etcheverry
Office hours – TuTh 11-12P
pedrohespanhol@berkeley.edu
- Lectures:** TuTh 330-5P, 310 Hearst Memorial Mining Building
- Website:** <http://ieor.berkeley.edu/~ieor265/>
- Prerequisites:** Course on optimization; course on statistics or stochastic processes
- Grading:** 4 homeworks (50%); class project (50%)
- Description:** This course will cover topics related to the interplay between optimization and statistical learning. The first part of the course will cover statistical modeling procedures that can be defined as the minimizer of a suitable optimization problem. The second part of the course will discuss the formulation and numerical implementation of learning-based model predictive control (LBMPC), which is a method for robust adaptive optimization that can use machine learning to provide the adaptation. The last part of the course will deal with inverse decision-making problems, which are problems where an agent's decisions are observed and used to infer properties about the agent.
- Class Project:** The projects can be in the form of a literature review, a comprehensive application of data analysis methods, or involve the exploration of original research ideas. The project should be chosen in consultation with the course instructor, and a **project proposal (one page summary) is due on February 28, 2017**. The **project report (10-12 pages) is due on the last day of lecture May 4, 2017**. Project proposals and reports should be submitted as a PDF file emailed to the GSI and cc'ed to the Instructor. Joint projects, involving reasonably sized groups, are allowed.

Outline:

Specific topics that will be covered include:

- Regression
 - Classical M-estimators (e.g., ordinary least squares, maximum likelihood estimation)
 - High-dimensional M-estimators (e.g., lasso or L1 regression, nuclear norm minimization)
 - Collinearity (e.g., ridge regression, exterior derivative estimation)
 - Semiparametric regression of partially linear models and Nadaraya-Watson regression
- Learning-Based Model Predictive Control (LBMPC)
 - Robustness (e.g., Berge maximum theorem, Lyapunov theory, reachability analysis)
 - Consistent approximations (e.g., semicontinuity, epiconvergence)
 - Oracle design (e.g., L2 regularized Nadaraya-Watson estimator)
 - Software code generation (e.g., numerical optimization solvers, heuristic reachability analysis)
- Inverse Decision-Making Problems
 - Inverse reinforcement learning (e.g., Markov Decision Processes, apprenticeship learning)
 - Learning objective/utility functions (e.g., bilevel programming, single level reformulations using KKT conditions)
 - Learning utilities from game-theoretic equilibria described by variational inequalities (e.g., parametric utilities, nonparametric utilities)