

Welcome to
Sample Efficient Optimization
80618A

Outline

- Introductions
- Zone Cours overview: course contents, grading
- General “Sample Efficient Optimization” introduction
- Gaussian Processes

A little bit about myself

- HEC, Department of Decision Sciences
- MIT faculty : 2010-2020
- PhD Mathematics, EPFL, Switzerland
- MSc. Statistics, University College London
- Bachelor Applied Mathematics & Computer Science, ENSIMAG, France
- Operations Research: mathematical modeling and optimization
- Urban mobility and transportation problems: simulation-based optimization, probabilistic modeling

Collaborators & Sponsors



- Holder of the Scale AI Chair in Scale AI Research Chair in Artificial Intelligence for Urban Mobility and Logistics
- Staff Research Scientist at Google Research
- Past affiliations include: Visiting Faculty at Google Maps; Consultant at Alphabet company Sidewalk Labs

Introductions

- Your name
- Your institution, academic program and year
- Your academic and professional background
- Your academic and professional interests
- Why you are taking this class ?
- Do you work with any problems where sample efficiency is important ?

Topics

- Gaussian processes
- Bayesian Optimization
- Stochastic Simulation Principles
- Variance Reduction
- Dimensionality Reduction
- Surrogate Modeling / Metamodeling
- Continuous Simulation-based Optimization (SO)
- Discrete SO

Perspectives from the field

- Harvey Cheng
 - Founding engineer at Distributional: builds software stack for testing modern AI applications.
 - Formerly: lead AI research scientist at SigOpt, a black-box optimization software startup for enterprises. Productionized Bayesian Optimization and other solvers for sequential decision-making.
- Talk:
 - Exploring the spectrum of black-box optimization
 - In this talk, we explore the spectrum of employing black-box Bayesian optimization in real-world applications, from tight-knit collaboration with materials scientists on using sample-efficient optimization techniques in simulation and experimentation, to deploying completely black-box software as a service to enterprises. We discuss the different challenges that we've encountered in each scenario, and share some lessons that we've learned along the way.



Perspectives from the field

- Pablo Rodríguez Fernández
 - Research Scientist, MIT Plasma Science and Fusion Center
 - Leads the Magnetic Fusion Experiments (MFE) Integrated Modeling (IM) group at MIT
- Talk:
 - On the application of surrogate-based optimization techniques to solve complex physics problems in nuclear fusion research
 - Fusion energy has long promised a sustainable, clean future. Despite the challenges that still exist, we are closer than ever to realizing fusion as an energy source, with billions of dollars being poured into an emerging fusion energy industry aiming to bring fusion to realization in the 2030s. To achieve this, a deep understanding of physics is required to optimize reactor designs, which can only be accomplished through the modeling of complex physical processes. This talk will present examples of the application of surrogate-based optimization techniques to solve fusion-relevant modeling problems in a timely manner. The use of Gaussian Process Regression within a Bayesian optimization framework is demonstrated to accelerate the design process of fusion reactors through the development of computationally efficient modeling workflows.



Perspectives from the field

- Zi Wang
 - Senior Research Scientist, Google Deep Mind
- Talk:
 - Transfer Learning for Bayesian Optimization on Homogeneous and Heterogeneous Search Spaces
 - Bayesian Optimization (BO) is a powerful approach for optimizing expensive black-box functions. However, its success depends heavily on expert-specified Gaussian Process (GP) priors, which can be difficult to construct in practice. This talk explores the use of pre-training to specify these functional priors in two settings: 1) homogeneous: the "training" functions share the same domain as the "test" function (black-box function to be optimized), and 2) heterogeneous: the domains of all training and test functions can be different. In the homogeneous setting, we introduce HyperBO, which pre-trains a GP and uses it for optimizing the test function. This approach leads to bounded posterior predictions and near-zero regrets without requiring the knowledge of the ground truth GP prior. For the heterogeneous case, we introduce a method for model pre-training on heterogeneous domains (MPHD). MPHD enjoys the benefits of HyperBO and further enhances it with domain adaptability. Using a neural net pre-trained in existing domains, MPHD generates customized hierarchical GP specifications based on domain descriptions. We discuss both the theoretical and practical advantages, and demonstrate their superior performance on a wide range of real-world tuning tasks.



Perspectives from the field

- Eytan Bakshy
 - Director and Research Scientist, Meta
 - Lead of the Adaptive Experimentation Team, Meta
- Talk:
 - TBD: Sample Efficient Optimization and Learning at Meta



Perspectives from the field

- Required preparatory readings prior to guest lectures
- Sign up to lead the discussion: email me w/ your name + paper title

Grading

- Test: 30%
- Written Project Report: 40%
- Project Presentation + Slides: 30%

Important dates

- Project proposal submission: Oct. 10 (Session #6)
- In-class test: Nov. 7, 8:30-10am (Session #10)
- Written project report: Nov. 21 (Session #12)
- Project slides: Nov. 28 (Session #13)
- Project presentation: Nov. 28 (Session #13)

Important dates

- In-class test:
 - Covers all class material seen up until the session before the test.
 - Open book.
 - Calculations: only w/ simple calculator; no computers for calculations.
 - Access to computers: to access your class materials offline, no internet.
- Written project report + project slides:
 - Upload before start of the corresponding class (8:30am)
- In-class test + project report + project slides:
 - Upload digital version in pdf format.
 - Clearly explain your reasoning, detail your answers.
- All submissions: Use of generative AI is prohibited.
- Individual work. No team work.

Class logistics

- My office hours: Fridays 3-4pm
- Class logistics inputs: wooclap.com/PSEOZO
- Would you like the synchronous session to be made available in real-time over Zoom (connect through ZoneCours)
 - Excludes guest lectures
 - No listeners, no audits.

After class

- Send me an email to lead a paper discussion.
 - See “State of the Art - Research & Practice” session in ZoneCours for available papers.
- Start thinking about your class project.
 - Can be based on your own research/projects.
 - If you want me to give you an Transportation / Urban Mobility project: email me.
- Install Sumo Traffic Simulation software.
 - See “Learning Material” session in ZoneCours for guidelines.