

Data-Driven Stoch Opt

Data-Driven Stochastic Optimization

(PhD School on Intersections of Algorithms and Machine Learning Theory)

Course Resources

- Links to Slides: [Lec 1](#), [Lec 2](#), [Lec 3](#), [Lec 4](#), [Lec 5](#), [Lec 6](#)
- Video Recordings: [Lec 1](#), [Lec 2](#), [Lec 3](#), [Lec 4](#), [Lec 5](#), [Lec 6](#)
- Link to [Exercises](#)
- Prereqs for Summer School: Many of the topics and relevant background is available at this [course website](#). The video recordings for this GT course is [available here](#). Some specific topics of interest are:
 - Basic concentration bounds: [Lecture 2](#)
 - Markov Decision Process: [Lecture 14](#)
 - Stochastic Multi-armed Bandits: [Thomas' notes](#),
 - Pandora's Box: [Lecture 15](#)
 - Stochastic Probing: [Lecture 17](#)

Lec 1: What is Stochastic Optimization?

- Change inputs from being given to random variables from given distributions.
- We will often assume n independent random vars and maximize expected value.
- Examples: Revenue maximization, Optimal Stopping, Pandora's box
- Goal given distributions: Find a (near-)optimal algorithm:
 - Think of a decision tree/MDP.
 - Alternatively, think of it as a class of policies. Find a near-optimal soln.
 - Computation is the only bottleneck.
- Techniques:
 - Closed form solutions
 - Dynamic programming
 - Linear programming
 - Index-based policies
- But what if the input distributions are unknown and we only have historical data?
 - We want generic techniques that apply to many pbs (lots of pb specific work)
 - Next 3 lectures on general techniques for full and limited feedback.
 - Last 2 lectures on the important Online Resource Allocation problem.

Lec 2: What is Data-Driven Stoch Optimization?

- What if the input distributions are unknown and we only have historical data? We want to still obtain the best possible approximation.
- What does “historical data” mean?
 1. Offline Learning: Given a small number of full samples (full-feedback)
 2. Online Learning: We learn as we go with full/partial feedback.
- Simple bounds 1: Number of algorithms is small
 - Class of algos is actually small or we can consider a smaller interesting subclass
 - Apply Chernoff bounds.
 - Nice but often not a polytime algo. Did not exploit product dist.
- Simple bounds 2: Distributions have a small support
 - Basics of TV distance.
 - Learn each D_i in a small TV distance.
 - This implies we can learn D in a small TV distance.
 - Optimize for the learnt distribution.
- Challenges in extending to continuous distribution:
 - Can't learn continuous distrib in TV distance
 - However, we can learn an “envelope”

Lec 3: Learning Strongly Monotone Stochastic Problems

- Previous algo either not polytime or assumed discrete distrib.
- Even though “fine discret” often works, it's suboptimal to use above general bounds.
- Hope: Can we just run an algo for the product empirical distrib?
- Guo et al show: yes, for any strongly monotone problem
- Let's again aim for $T = n^2/\epsilon^2$ since we only use TV distance. Algo is to again take the empirical distribution E and feed it to blackbox.
- **Proof:** First observe that for all t , $|F_E(t) - F_D(t)| \leq 1/\sqrt{T}$. Now we need to show $D^+ \geq E \geq D^-$ where the TV distance between $\Delta(D^-, D) \leq 1/\sqrt{T}$. This can be achieved by choosing the cdf of D^- to be $1/\sqrt{T} + F_D(t)$. The TV distance is small by considering the coupling $X^- = X \cdot 1[X < \text{Top}(1-\epsilon)]$.
- Can we improve the bounds? Yes, we can always improve n^2 to n by working with Hellinger distance (see [paper](#) for details). In general this is tight, but for some problems we can remove n , like for optimal stopping.

Lec 4: Online Learning with Partial Feedback

- What is online learning for stochastic optimization?
- Bandit feedback: Simple bounds via Guo et al.
 - Use some rounds for exploring and some for exploiting
 - In general we cannot obtain \sqrt{T} with bandit feedback, but for some problems we can: e.g., optimal stopping and pandora's box in [this paper](#).
- What is semi-bandit feedback?
- General results for discrete distributions with semi-bandit feedback.

Lec 5: Online Resource Allocation

- What is Online Resource Allocation
- Why do we need some underlying distributions?
- How to solve this stochastic problem when distributions are known?
- What if the inputs are from an unknown identical distribution?
- An approach via Multiplicative Weights algorithm.

Lec 6: Online Resource Allocation for Non-Identical Distributions

- Recall Online Resource Allocation
- What if the distributions are not identical?
- Single-sample algorithms
- Exponential pricing algorithm (more details in [this paper](#))
- Discuss exercises from the summer school.