Lecture 13. Review and Advanced Topics

Advanced Optimization (Fall 2025)

Peng Zhao

zhaop@lamda.nju.edu.cn

Nanjing University

Part 1. Offline Optimization

- Lipschitz Optimization (Lec 3)
  - Gradient Descent: first/second gradient descent lemma
  - Polyak step size

- Smooth Optimization (Lec 4)
  - Gradient Descent (one-step improvement lemma)
  - Polyak momentum
  - Nesterov's acceleration method
  - Composite Optimization

How: use gradient information to establish convergence

Part 2. Online Optimization

- From a fixed objective $F(\mathbf{x})$ to an interactive objective $f_t(\mathbf{x})$

- Problem-independent Regret (Lec 5 & 6 & 7)
  - Lec 4: Online Gradient Descent
  - Lec 5&6: Online Mirror Descent
    - for PEA & exp-concave; and for general geometry
  - Lec 4: (Weighted) Online-to-Batch Conversion
    - $1/\sqrt{T}$ for Lipschitz and convex; $1/T$ for Lipschitz and strongly convex

How: handle interaction and exploit geometry to optimize regret

Part 2. Online Optimization

- From a fixed objective $F(\mathbf{x})$ to an interactive objective $f_t(\mathbf{x})$

- Problem-dependent Regret (Lec 8 & 9)
    - Adaptive Online Learning: small-loss PEA, self-confident tuning
    - Lec 9: Optimistic OMD
        - general optimistic bound
        - small-loss, gradient-variance, gradient-variation
    - Lec 9: (Stabilized) Online-to-Batch Conversion
        - GV regret+O2B → $1/T$; and GV regret+stabilized O2B → $1/T^2$ for smooth functions

How: leverage predictive info to achieve faster rate guarantees

Part 3. Bandit Optimization

- From gradient feedback to partial-information feedback

- Adversarial Bandits (Lec 10)
  - Multi-armed Bandits: importance-weighting (IW) loss estimator
  - Bandit Convex Optimization: gradient estimator based on perturbation
  - More advanced: anisotropy exploration based on self-concordant barrier
  - Exploration-exploitation in adversarial bandits

How: use single function value to estimate the gradient

Part 3. Bandit Optimization

- From gradient feedback to partial-information feedback

- Stochastic Bandits (Lec 11 & Lec 12)
  - Multi-armed Bandits
    - ETC, ϵ-greedy, Upper Confidence Bound, TS
  - Linear Bandits
    - least-square estimator, UCB based self-normalized concentration
  - Exploration-exploitation in stochastic bandits

How: adaptive exploration based on the statistical estimation

Many yet to cover

- Variance Reduction (HW1 mentions a bit)
- Non-convex Optimization (HW1 mentions a bit)
- Adaptive Optimization like AdaGrad/Adam (HW2 mentions a bit)
- Online MDPs/Control/RL (HW2 mentions a bit)
- Online Games/Minmax Optimization
- Contextual Bandits…

Four Talks

- Bandit Optimization
  - Yu-Jie Zhang, Sheng-An Xu, Peng Zhao, and Masashi Sugiyama. Generalized Linear Bandits: Almost Optimal Regret with One-Pass Update. ***NeurIPS 2025***.
  - Jing Wang, Yu-Jie Zhang, Peng Zhao, and Zhi-Hua Zhou. Heavy-Tailed Linear Bandits: Huber Regression with One-Pass Update. ***ICML 2025***.
- Online Optimization
  - Yu-Hu Yan, Peng Zhao, and Zhi-Hua Zhou. A Simple and Optimal Approach for Universal Online Learning with Gradient Variations. ***NeurIPS 2024***.
  - Yuheng Zhao, Yu-Hu Yan, Kfir Yehuda Levy, Peng Zhao. Gradient-Variation Online Adaptivity for Accelerated Optimization with Hölder Smoothness. ***NeurIPS 2025***.

腾讯匿名问卷反馈

高级优化2025 课程反馈

诚邀您填写本问卷，扫码即可！