HOMEWORK 625
PAPER PRESENTATION *

10-425/10-625 INTRODUCTION TO CONVEX OPTIMIZATION
http://425.mlcourse.org

OUT: 11/18/23
DUE: 12/04/23
TAs: Akash, Asad, Roochi, & Tiancheng

1 Paper Presentation

For this assignment you will read a recent (or classical) optimization paper and video record a (short, 5-10 minute) presentation about it.

We will follow the role-playing methodology from Collin Raffel and Alec Jacobson.

https://colinraffel.com/blog/role-playing-seminar.html

You may select one of the following roles and present your paper in that format. (The instructions below were copied and adapted from 10-719's use of the role-playing paper presentation approach.)

1. **Scientific Peer Reviewer:** Complete a full review of the paper as if it were submitted to a conference. Follow the guidelines for NeurIPS reviewers to produce your review. In particular, please answer Questions 1 to 10 under "Review Form", including assigning an overall score.
2. **Archaeologist:** Determine where this paper sits in the context of previous and subsequent work. Find and briefly report on both: (1) a prior paper that substantially influenced the current paper, and (2) a more recent paper that cites this current paper.
3. **Academic Researcher:** You're an academic researcher working on a new project in this area. Propose an imaginary follow-up project that builds on the current paper. Pretend that this new project has been successful, and write up a brief introduction for a paper about your project using the five-point structure provided here (under "The Introduction"). You do not need to actually *write* the introduction, but instead should present in the style of an introduction to this new project.
4. **Industry Practitioner:** You work at a company or organization developing an application or product of your choice (one that has not already been suggested in a prior session). Describe the application/product in detail, and bring a convincing pitch for why you should be paid to implement the method in the paper for this particular application.
5. **Private Investigator:** You are a detective who needs to run a background check on one of the paper's authors. Where have they worked? What did they study? What previous projects might have led to working on this one? What motivated them to work on this project?

*Compiled on Saturday 18th November, 2023 at 21:53

6. **Social Impact Assessor:** Identify how this paper self-assesses its positive or negative impact on the world. Have any additional positive social impacts been left out? What are possible negative social impacts that were overlooked or omitted? Please read this short paper to see examples.

A few notes:

- For each of the above roles, your presetation *must* start with a brief summary of the paper.
- You are welcome to record with your webcam turned off (Khan Academy style).
- After submission, we will share all the HW625 video presentations with the rest of class via Piazza, so others can learn from you as well.
- You will make slides and submit them alongside your video recording.

2 Select a paper

To select a paper, you should pick of the ones below. If you would like to present an optimization paper not on this list, send a Private Note to the "Instructors" with the paper and 1-2 sentences about why you are interested in it.

- Liu & Nocedal (1981). On the limited memory BFGS method for large scale optimization. [paper]
- Zinkevich (2003). Online Convex Programming and Generalized Infinitesimal Gradient Ascent [paper]
- Chambolle & Pock (2010). A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging [paper]
- Boyd et al. (2010). Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers [paper] *(Sections 1 - 3 only)*
- Niu et al. (2011). HOGWILD!: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent [paper]
- Snoek, Larochelle, & Adams (2012). Practical Bayesian Optimization of Machine Learning Algorithms. [paper]
- Johnson & Zhang (2013). Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. [paper]
- Allen-Zhu & Orecchia (2016). Linear Coupling: An Ultimate Unification of Gradient and Mirror Descent [paper]
- Allen-Zhu, Li, Song (2019). A Convergence Theory for Deep Learning via Over-Parameterization [paper]
- Loshchilov & Frank (2019). Decoupled Weight Decay Regularization. [paper]
- Li, Wei, & Ma (2020). Towards Explaining the Regularization Effect of Initial Large Learning Rate in Training Neural Networks [paper]

3 Video Recording and Submission

Video Recording For the presentation, you will record yourself and then send us the link to the recording. Here are the steps you should take:

1. Open Zoom and start your "Personal Meeting". Be sure to turn on your microphone.

2. Click ”Record” and then ”Record to the Cloud”.
3. Introduce yourself by name. Then go ahead and present for 5-10 minutes.
4. Extremely important: Now click ”End the Meeting”. Your recording will be uploaded to the cloud and processed.
5. Go to https://cmu.zoom.us/recording and wait (10 minutes) for your recording to finish processing.
6. Click the ”Share...” button on the left of the recording, then click ”Copy Sharing Information” and paste the clipboard text into Gradescope.

Submit to Gradescope You should submit your presentation video and slides to Gradescope.

7. Open the HW625 assignment in Gradescope.
8. Paste the video recording information that you copy/pasted above into the free-text field on the assignment.
9. Then upload your slides as a PDF on the file-upload question.
10. In Gradescope, save and submit this assignment.