

CS 181: Advanced Algorithms

Course Project

1 Overview

The final project is your chance to explore an advanced algorithms topic in depth and communicate what you learned. You may work either individually or in teams of two. Your goal should be: if another CS 181 student read your write-up and saw your talk, they should come away understanding something interesting that they did not know before.

2 Learning goals

Every student should read and interpret at least one source that is more advanced than a standard undergraduate textbook. If you are presenting an algorithms problem, your goal is to extract key ideas behind the algorithm or proof and communicate them precisely and accessibly to other CS 181 students. If you're doing an implementation, you should reason about the practical effectiveness of algorithms in relation to their theoretical guarantees, and communicate this as well.

3 Deliverables

Presentation

On Tuesday, December 9th, you will lead a 15 minute presentation / discussion for the class. Try to convey the context: what is interesting about the problem, and why did you choose to study it? Why do people care about it? Then go into specifics (whatever they are) and finally conclude by again laying out context. How does this result fit in with the landscape of what is known for the problem? What might be downsides to this approach or opportunities for future study?

Report

There will be a 3-4 page report as part of your project. I am not expecting anything flowery here. It should be direct and concise, like a (well-written) research report. This is your opportunity to explain to me your topic or anything that you didn't have time to get across in the presentation. It is also an opportunity to reflect on what you learned from the class overall. Please include a section in which you discuss your overall takeaways from the class (you can have different sections for you and your teammate here, if you're working in a group). The report should be written in Latex, and everything written should be in your own words or properly cited. The page suggestion is a guideline, at the end of the day I will be looking for whatever constitutes a complete and in-depth report.

Remark. I encourage collaboration, but for fairness I will have a higher standard for team projects.

Milestones

- Tuesday, November 11th, you or your team have met with me to discuss your topic.
- Tuesday, November 18th, topic locked in.
- Tuesday, December 2nd, rough draft report due. I will give feedback.
- Tuesday, December 9th, final report due and project presentations.

4 Project Ideas

These are only suggestions. You may of course, propose your own idea.

- The Primal-dual method. Can be applied to Vertex Cover and other problems (Set Cover), see sections 11.4 and 11.5 in Kleinberg-Tardos.
- Advanced LP rounding method for Load Balancing, see section 11.7 of Kleinberg-Tardos.
- Arbitrarily good approximations, also known as Polynomial-time approximation schemes. Can be applied to Knapsack, see section 11.8 of Kleinberg-Tardos.
- The famous, super-fast randomized algorithm for global min-cut. This is Karger's algorithm, see section 13.2 in Kleinberg-Tardos.
- Solving NP-hard problems on Trees and graphs that are "close" to being trees. See section 10.4 of Kleinberg-Tardos.
- Randomized Rounding for Set Cover / Facility Location. Can these be de-randomized? (section 1.7 and 5.8 in Williamson-Shmoys)
- The prize-collecting Steiner tree problem: see section 5.7 in Williamson-Shmoys.
- Semi-definite programming for Max-Cut, section 6.2 of Williamson-Shmoys.
- Network Design: Tree and Connectivity Augmentation
- Edmonds' Non-Bipartite Perfect Matching Polytope (this is a good one).
- Algorithmic incentives and mechanism design. This is a huge topic, pick a subset of: truthful reporting in Gale-shapley, kidney-exchange, the YRMHIGYT algorithm, second price auctions and VCG.
- Learning Augmented Online Algorithms. Could be applied to caching, Ski-rental, or matchings (see me for the appropriate papers).
- Other types of clustering algorithms: Correlation Clustering, k-median, k-center with outliers

For the above topics, if there is a textbook chapter, that is a great starting point and you should absolutely focus on this first. But I am expecting a bit more discussion of surrounding ideas and state-of-the-art knowledge in your final report.