

MCS 549 – Mathematical Foundations of Data Science
Fall 2023
Class Presentations

Lev Reyzin

Part of your grade for this class is to present a paper on a mathematical data science topic. I will put example papers on topics that we have learned so far. You are free to choose from among those or to suggest a different paper with sufficient theoretical content to present, subject to my approval. The list of example papers and available times can be found on:

http://homepages.math.uic.edu/~lreyzin/f23_mcs549/papers.html.

Example venues where suitable papers can usually be found include SODA, STOC, FOCS, ICALP, RANDOM, COLT, and ALT.

Your first task is to email me (at lreyzin@uic.edu) your choice of paper and the 4 ordered time-slots (both assigned first-come, first-serve) by the end of the day on the 13th of November. The available days are November 27th, November 29th, and December 1st. Please check the website for up-to-date available times. If you do not email me by then, a time/topic will be arbitrarily assigned to you.

Presentations will be done on the white-board or Zoom and (unlike class) will begin strictly on time. You will have 20 minutes to present the paper, and 5 additional minutes will be used for questions by the instructor (definitely) and your fellow students (if they wish)¹. You must have a deeper understanding of the paper than you are able to present, and this will be elucidated in the question period. Many of the possible papers are complicated, so you will not be able to present them in full, or even in much detail. Some papers are so difficult that you may not be able to understand them fully, and that is okay too – the more difficult the paper, the more leeway you will have.

Nonetheless, your goal will be to 1) clearly explain the problem being solved, 2) explain the state of the art before the paper, 3) explain what advance the paper made, 4) in as much (or little) detail as time allows, give an overview of how the result was obtained, and 5) be prepared to answer questions from your audience. You will not be responsible for knowing any advances that were made after the appearance of the paper you are presenting

The idea of this exercise is for you to get a chance to engage with modern data science research, go through the process of reading and understanding a research paper, and give a “conference talk,” as is customary in mathematical computer science (though admittedly on someone else’s work).

¹Do not plant questions (if you are the speaker) nor try to trap the speaker (if you are in the audience).