

Algorithmic Aspects of Machine Learning: Final Project

Instructor: Ankur Moitra

Due: May 16th

The last assignment for the semester is to write a 4 – 6 page final paper. You can work in groups or you can work alone, and there are two options:

(1) You can write a literature review on some topic related to the material that we covered in class. In this class, we only focused on problems where there are *provable guarantees* so you should certainly choose a topic where there are provable guarantees. The goal of the writeup is to survey what is known, some main ideas of the proof in the papers you choose to read, and identify important open questions. You should think of this project as if you were going to give 1 – 2 lectures in the course, and how you would explain some topic beyond what we covered in class, to other students. We covered a lot of material in class, but there is still much more out there and I hope to make your literature reviews available on the course website so that they can be a useful reference for others. Also, you should choose a somewhat focused topic because if you choose a topic that is too broad, it will be difficult to get into the details.

(2) You can do original research. This is the more challenging type of project, but it is also the most open ended. You should choose some open question either explicitly mentioned in class, or connected to any of the topics we covered. You should try to give new *provable guarantees* for some problem, either by giving a new algorithm or giving an improved analysis of an existing algorithm. As with all research, it is never clear whether you will reach the goal that you set out. And so you should make sure that there is some partial progress that you can make along the way, that you can write up in your final paper. Even if you end up not proving what you set out to prove, you can still write up some preliminary ideas along the way or at the very least write up a literature review focused on the open question you worked on. If you choose this type of project, it is important that you come talk to me about it so that I can talk to you about your ideas or point you to other papers that might be relevant.

Please email me your final paper by Tuesday, May 16th. This is the last day of classes. Many of you have already come to talk to me about topics for your final project, and if you have ideas of your own of what you would like to work on, that is great! This is especially good because the purpose of the class was to expose you to some of the provable guarantees that are known for various problems in machine learning, and if you can find topics in your own research area where there

are interesting avenues to explore, then you are more likely to use the material from this course in your own research even after the semester. That was one of my main goals in teaching this. If you do not have an idea, I am here to help. Here are some suggested topics, for writing a literature review on. If you are interested in one of these topics, email me to claim it and set up a meeting so that we can talk about it. If you reserve a topic by yourself, but are interested in working with another student, I will put a ‘?’ next to your name. And if we run out of topics, I will come up with more suggestions.

- **Spectral Clustering**

“Improved Spectral Norm Bounds for Clustering”, Awasthi, Sheffet.

- **Stability**

“Are Stable Instances Easy?”, Bilu, Linial

“Bilu-Linial Instances of MAX CUT and Minimum Multiway-cut”, Makarychev, Makarychev, Vijayaraghavan.

- **Mixture Models**

“On Learning Mixtures of Heavy-Tailed Distributions”, Dasgupta, Hopcroft, Kleinberg, Sandler

- **Smoothed Analysis**

“Learning Mixtures of Gaussians in High Dimensions”, Ge, Huang, Kakade.

- **Mixtures of Linear Regressions**

“Small Covers for Near-Zero Sets of Polynomials and Learning Latent Variable Models”, Diakonikolas, Kane.

- **Exact Recovery in SBM**

“Achieving Exact Cluster Recovery Threshold via Semidefinite Programming”, Hajek, Wu, Xu.

- **Lower Bounds in RL**

“Computational-Statistical Gaps in Reinforcement Learning”, Kane, Liu, Lovett, Mahajan.

- **Alternating Minimization**

“Understanding Alternating Minimization for Matrix Completion”, Hardt.

- **Expectation Maximization**

“Statistical Guarantees for the EM Algorithm: From Population to Sample-based Analysis”, Balakrishnan, Wainwright, Yu.

- **Samples/Memory Tradeoffs**

“Efficient Convex Optimization Requires Superlinear Memory”, Marsden, Sharan, Sidford, Valiant.

- **Multimodality**
“Beyond Log-concavity: Provable Guarantees for Sampling Multi-modal Distributions using Simulated Tempering Langevin Monte Carlo”, Ge, Lee, Risteski.
- **Score Matching**
“Statistical Efficiency of Score Matching: The View from Isoperimetry”, Heckett, Koehler, Risteski.
- **Trace Reconstruction**
“New upper bounds for trace reconstruction”, Chase.
- **Privacy**
“Robustness Implies Privacy in Statistical Estimation”, Hopkins, Kamath, Majid, Narayanan.