Robust Convergence of Machine Learning Algorithms

Ryan Rice [1]

DIMACS REU

Mentor: Dr. Pranjal Awasthi

[1] Work supported by NSF Grant CCF-1559855

Mean Estimation

We draw iid (independent, identically distributed) samples $x_1, \ldots, x_n \sim D$ where D is an unknown distribution, how can we find out the mean μ of D?

Maximum likelihood estimation gives sample mean $\hat{\mu} = \sum_{i \in [n]} x_i$ which concentrates well

Robust learning

We have a good understanding of how to deal with these iid samples, but what about if there is some noise in our samples?

Define adversarial noise as able to arbitrarily change up to an ϵ-fraction of data, then:

- Maximum likelihood estimation very sensitive

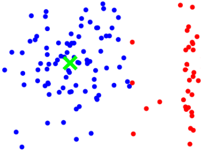

Demonstrating adversarial noise, graphic from Steinhardt et al [SCV17]

- Past robust statistical methods are hard to compute or have errors that are dimension dependent

Robust Mean estimation

Recent results give polynomial time ways to robustly estimate mean with error independent of dimension (with high probability) while having few assumptions on the underlying distribution.

However these are not necessarily fast algorithms for many use cases

Open areas

Will robust mean estimation compose for larger problems (k-means) to create other robust algorithms?

Can we speed up robust mean estimation (for practical use)?

Questions?

Reference

Jacob Steinhardt, Moses Charikar, and Gregory Valiant, *Resilience: A criterion for learning in the presence of arbitrary outliers*, CoRR **abs/1703.04940** (2017).