**Abstract**

We consider the problem of answering queries about a sensitive dataset subject to differential privacy. The queries may be chosen adversarially from a larger set $Q$ of allowable queries in one of three ways, which we list in order from easiest to hardest to answer:

- *Offline:* The queries are chosen all at once and the differentially private mechanism answers the queries in a single batch.
- *Online:* The queries are chosen all at once, but the mechanism only receives the queries in a streaming fashion and must answer each query before seeing the next query.
- *Adaptive:* The queries are chosen one at a time and the mechanism must answer each query before the next query is chosen. In particular, each query may depend on the answers given to previous queries.

Many differentially private mechanisms are just as efficient in the adaptive model as they are in the offline model. Meanwhile, most lower bounds for differential privacy hold in the offline setting. This suggests that the three models may be equivalent. We prove that these models are all, in fact, distinct. Specifically, we show that there is a family of statistical queries such that exponentially more queries from this family can be answered in the offline model than in the online model. We also exhibit a family of search queries such that exponentially more queries from this family can be answered in the online model than in the adaptive model. We also investigate whether such separations might hold for simple queries like threshold queries over the real line.