

Gerome Miklau

The Matrix Mechanism: Optimizing Linear Counting Queries Under Differential Privacy

Differential privacy is a rigorous privacy standard that protects against powerful adversaries, offers precise accuracy guarantees, and has been successfully applied to a range of data analysis tasks. The original algorithm for achieving differential privacy, commonly called the Laplace mechanism, returns the true answer after the addition of random noise drawn from a Laplace distribution. If an analyst requires only the answer to a single query about the database, then (a discrete version of) the Laplace mechanism is known to be optimal. But the Laplace mechanism can be highly suboptimal when a set of correlated queries are submitted, and despite much recent work, optimal strategies for answering a collection of correlated queries are not known in general.

In this talk I will describe the "matrix mechanism", a new algorithm for answering complex workloads of predicate counting queries under differential privacy. Given a workload, the mechanism first requests answers to a different set of queries, called a query strategy, which are answered using the standard Laplace mechanism. Noisy answers to the workload queries are then inferred from the noisy answers to the strategy queries. When the strategy queries are chosen appropriately, this two stage process increases accuracy (with no cost in privacy) by answering the workload queries using a more complex, correlated noise distribution.

After describing the basic operation of the matrix mechanism (originally presented at PODS 2010), I will describe recent results on improved inference using non-negativity constraints, a lower bound on the error of optimal strategies, and an efficient approximation algorithm for strategy selection.