Learning to Reconstruct: Statistical Learning Theory and Encrypted Database Attacks

Paul Grubbs, Cornell Tech

We show that the problem of reconstructing encrypted databases from access pattern leakage is closely related to statistical learning theory. This new viewpoint enables us to develop broader attacks that are supported by streamlined performance analyses.

As an introduction to this viewpoint, we first present a general reduction from reconstruction with known queries to PAC learning. Then, we directly address the problem of ϵ-approximate database reconstruction (ϵ-ADR) from range query leakage, giving attacks whose query cost scales only with the relative error ϵ, and is independent of the size of the database, or the number NN of possible values of data items. This already goes significantly beyond the state of the art for such attacks, as represented by Kellaris et al. (ACM CCS 2016) and Lacharité et al. (IEEE S&P 2018).

We also study the new problem of ϵ-approximate order reconstruction (ϵ-AOR), where the adversary is tasked with reconstructing the order of records, except for records whose values are approximately equal. We show that as few as O(ϵ^−1logϵ^−1) uniformly random range queries suffice. Our analysis relies on an application of learning theory to PQ-trees, special data structures tuned to compactly record certain ordering constraints.

We then show that when an auxiliary distribution is available, ϵ-AOR can be enhanced to achieve ϵ-ADR; using real data, we show that devastatingly small numbers of queries are needed to attain very accurate database reconstruction.

Finally, we generalize from ranges to consider what learning theory tells us about the impact of access pattern leakage for other classes of queries, focusing on prefix and suffix queries. We illustrate this with both concrete attacks for prefix queries and with a general lower bound for all query classes.