Principled Evaluation of Differentially Private Algorithms

Michael Hay*, Ashwin Machanavajjhala**, Gerome Miklau†, Yan Chen**, Dan Zhang†

*Colgate University, Department of Computer Science, mhay@colgate.edu
**Duke University, Department of Computer Science, {ashwin,yanchen}@cs.duke.edu
†UMass Amherst, College of Information and Computer Sciences, {miklau,dzhang}@cs.umass.edu

Abstract

The emergence of differential privacy as a primary standard for privacy protection has led to the development of hundreds of algorithms for various data analysis tasks. These algorithms are becoming increasingly complex, and in particular, the performance of many emerging algorithms is *data dependent*, meaning the distribution of the noise added to query answers may change depending on the input data. Theoretical analysis typically only considers the worst case, making empirical study of average case performance increasingly important.

We propose a set of evaluation principles which we argue are essential for sound evaluation. Based on these principles we propose DPBENCH, a novel evaluation framework for standardized evaluation of privacy algorithms. We then apply our benchmark to evaluate algorithms for answering 1- and 2-dimensional range queries. The result is a thorough empirical study of 15 published algorithms on a total of 27 datasets that offers new insights into algorithm behavior—in particular the influence of dataset scale and shape—and a more complete characterization of the state of the art. Our methodology is able to resolve inconsistencies in prior empirical studies and place algorithm performance in context through comparison to simple baselines. It also raises open research questions which we hope will guide future algorithm design.

The full version of this paper will appear at ACM SIGMOD 2016.