Project Ideas: CSC2412 Fall 2018

Aleksandar Nikolov

If you would like to develop the following ideas into project proposals, email me as soon as possible. I will not allow multiple groups to use the same idea, so you need to claim them on a first come first served basis.

Computing k-way Marginals. Let us consider the setting in which our database x has d binary attributes, i.e. $x \in (\{0,1\}^d)^n$. One of the most important classes of counting queries on such databases are the k-way marginal queries. A k-way marginal query shows what fraction of the rows in the database have a particular setting for particular set of k attributes. Formally, a k-way marginal query $q_{S,\alpha}$ is specified by a set of k attributes $S \subseteq \{1,\dots,d\}$, $|S| = k$, and an assignment $\alpha \in \{0,1\}^S$. For any $y \in \{0,1\}^d$ we have $q_{S,\alpha}(y) = 1$ if $y_i = \alpha_i \; \forall i \in S$, and $q(y) = 0$ otherwise. As usual, this defines a counting query by $q_{S,\alpha}(x) = \frac{1}{n}\sum_{i=1}^n q_{S,\alpha}(x_i)$. So, for example, $q_{\{j,k\},(0,1)}(x)$ equals the fraction of rows x_i in x that have $x_{ij} = 0$ and $x_{ik} = 1$.

A well-known open problem in Differential Privacy is to compute k-way marginals with error as close as possible to the error achieved by the Private Multiplicative Weights (PMW) mechanism (which is optimal) and in time polynomial in n and d^k. PMW itself does not achieve this goal because it runs in time exponential in d.

This problem is solved in [DNT15] for $k = 2$. One *theoretical* goal is to improve the best known error bound achieved by an algorithm running in time polynomial in d for $k = 3$. The error bound to beat is the one in [DNT15]. An *empirical* project is to implement heuristic methods for the projection step of the projection mechanism in [DNT15] for $k = 3$ (and $k > 3$), and evaluate if the heuristics are efficient in practice.

Another possible direction is to look at 2-way marginals but when the rows of the database come from $\{0,\dots,b\}^d$ for some $b > 1$.

Tracing Attacks. The goal of a tracing attack is to identify whether a given row $y \in \mathcal{X}^n$ belongs to a private database x or not. The assumption is that y is known to the attacker, but x is not: instead the attacker has access to some aggregate statistics about x. In Homer et al. [HSR$^+$08] and Dwork et al. [DSS$^+$15] it was assumed that the aggregate statistics are 1-way marginals. Im et al. [IGNC12] extended the attacks to regression coefficients.

It would be interesting, both for theory and for practice, to extend these attacks to other classes of statistics. Some options include k-way marginals for

$k > 1$, threshold queries, range queries, or the parameters of some classifier trained on the data. Tracing attacks on such queries can be evaluated either empirically, with public or synthetic data, or theoretically, by proving bounds on the error which allows the attack to be successful, as is done in [DSS$^+$15].

Privately Generating Synthetic Data using Generative Models. It is often convenient to have a differentially private method to generate synthetic data that "looks like" the real private data and has similar statistical properties. More formally, if the private database is $x \in \mathcal{X}^n$, we want to generate a database $\widehat{x} \in \mathcal{X}^m$ which shares some statistical properties with x. Some limitations are known on generating synthetic data: we must limit the class of queries which $\widehat{x}$ needs to answer similarly to x [DMNS06]; moreover, even for 2-way marginals, generating synthetic data can be computationally infeasible in the worst case [UV10]. The Private Multiplicative Weights mechanism can be used to generate synthetic data, but its running time is at least linear in the size of the data universe $\mathcal{X}$ which is often very inefficient.

The negative results such as [UV10] are for worst-case databases which are unlikely to appear in practice. This suggests investigating methods to generate synthetic data which are useful and efficient in practice, even though they may fail in the worst case. One approach to this task is to fit the parameters of a generative model to the private database x in a differentially private way, and then generate synthetic data using the generative model. There is some very recent work in this direction using GANs [TF18]. There has also been work focused on restricted domains, for example traffic data [MMLS14]. As an empirical project, you can implement several differentially private algorithms that fit generative models to data and test them.

References

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[DNT15] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Efficient algorithms for privately releasing marginals via convex relaxations. *Discrete Comput. Geom.*, 53(3):650–673, 2015.

[DSS$^+$15] Cynthia Dwork, Adam D. Smith, Thomas Steinke, Jonathan Ullman, and Salil P. Vadhan. Robust traceability from trace amounts. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 650–669. IEEE Computer Society, 2015.

[HSR$^+$08] Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V. Pearson, Dietrich A. Stephan, Stanley F. Nelson, and David W. Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLOS Genetics*, 4(8):1–9, 08 2008.

[IGNC12] Hae Kyung Im, Eric R Gamazon, Dan L Nicolae, and Nancy J Cox. On sharing quantitative trait gwas results in an era of multiple-omics data and the limits of genomic privacy. *The American Journal of Human Genetics*, 90(4):591–598, 2012.

[MMLS14] Nick Manfredi, Darakhshan J. Mir, Shannon Lu, and Dominick Sanchez. Differentially private models of tollgate usage: The milan tollgate data set. In *BigData Conference*, pages 46–48. IEEE, 2014.

[TF18] Aleksei Triastcyn and Boi Faltings. Generating differentially private datasets using GANs, 2018.

[UV10] Jonathan Ullman and Salil P. Vadhan. Pcps and the hardness of generating synthetic data. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:17, 2010.