CIS 800/002 The Algorithmic Foundations of Data Privacy October 4, 2011

Lecture 7

Lecturer: Aaron Roth *Scribe: Aaron Roth*

Iterative Database Construction

We've seen "The Net Mechanism", which shows that by using correlated noise, we can answer queries with much lower noise rates than we can by using the Laplace mechanism alone. The net mechanism was mostly an information theoretic upper bound, however. In this lecture, we'll start introducing the machinery to rederive the same bounds in a more algorithmic way.

This will have several benefits:

1. It will lead to more efficient algorithms, and the possibility of deriving very efficient algorithms.
2. By deriving the bounds using multi-stage algorithms, we can apply composition theorems and get *improved* accuracy for (ϵ, δ)-differential privacy.
3. We will be able to adapt these algorithms to work in the interactive setting.

The algorithms we present will be based on an abstraction called a "database update algorithm".

Roughly, such an algorithm works by maintaining a sequence of data structures $D^1, D^2, \ldots$ that give increasingly good approximations to the input database D (in a sense that depends on the DUA). Moreover, these mechanisms produce the next data structure in the sequence by considering only one query Q that *distinguishes* the real database in the sense that $Q(D^t)$ differs significantly from $Q(D)$.

Syntactically, we will consider functions of the form $U : \mathbf{N}^{|\mathcal{X}|} \times \mathcal{C} \times \mathbf{R} \to \mathcal{D}$, where $\mathcal{D}$ represents a class of datastructures on which queries in $\mathcal{C}$ can be evaluated. The inputs to U are a data structure in $\mathcal{D}$, which represents the current data structure D^t; a query Q, which represents the distinguishing query, and may be restricted to a certain set $\mathcal{C}$; and also a real number. which estimates $Q(D)$. Formally, we define a *database update sequence*, to capture the sequence of inputs to U used to generate the database sequence $D^1, D^2, \ldots$.

Definition 1 (Database Update Sequence) *Let $D \in \mathbf{N}^{|\mathcal{X}|}$ be any database and let $\{(D^t, Q_t, v_t)\}_{t=1,\ldots,L} \in (\mathcal{D} \times \mathcal{C} \times \mathbf{R})^L$ be a sequence of tuples. We say the sequence is an $(U, D, \mathcal{C}, \alpha, L)$-*database update sequence *if it satisfies the following properties:*

1. *$D^1 = U(\emptyset, \cdot, \cdot)$,*
2. *for every $t = 1, 2, \ldots, L$, $|Q_t(D) - Q_t(D^t)| \geq \alpha$,*
3. *for every $t = 1, 2, \ldots, L$, $|Q_t(D) - v_t| < \alpha$,*
4. *and for every $t = 1, 2, \ldots, L-1$, $D^{t+1} = U(D^t, Q_t, v_t)$.*

We note that for all of the DUAS we consider, the approximate answer v_t is used only to determine the *sign* of $Q_t(D) - Q_t(D^t)$, which is the motivation for requiring that v_t have error smaller than α. The main measure of efficiency we're interested in from a DUA is the maximum number of updates we need to perform before the database D^t approximates D well with respect to the queries in $\mathcal{C}$. To this end we define a Database Update Algorithm as follows:

Definition 2 (Database Update Algorithm (DUA)) *Let $U : \mathcal{D} \times \mathcal{C} \times \mathbf{R} \to \mathcal{D}$ be an update rule and let $B : \mathbf{R} \to \mathbf{R}$ be a function. We say U is a $B(\alpha)$-*DUA for query class $\mathcal{C}$ *if for every database $D \in \mathbf{N}^{|\mathcal{X}|}$, every $(U, D, \mathcal{C}, \alpha, L)$-database update sequence satisfies $L \leq B(\alpha)$.*

Note that the definition of an $B(\alpha)$-DUA implies that if U is a $B(\alpha)$-DUA, then given any maximal $(U, D, \mathcal{C}, \alpha, L)$-database update sequence, the final database D^L must satisfy $\max_{Q \in \mathcal{C}} \left|Q(D) - Q(D^L)\right| \leq \alpha$ or else there would exist another query satisfying property 2 of Definition 1, and thus there would exist a $(U, D, \mathcal{C}, \alpha, L+1)$-database update sequence, contradicting maximality.

A Database Update Algorithm will be useful together with a corresponding distinguisher, for finding queries Q to use in the next step of the database update sequence.

Definition 3 **($(F(\epsilon), \gamma)$-Private Distinguisher)** *Let $\mathcal{C}$ be a set of queries, let $\gamma \geq 0$ and let $F(\epsilon) : \mathbf{R} \to \mathbf{R}$ be a function. An algorithm $\text{Distinguish}_\epsilon : \mathbf{N}^{|\mathcal{X}|} \times \mathcal{D} \to \mathcal{C}$ is an $(F(\epsilon), \gamma)$-Private Distinguisher for $\mathcal{C}$ if for every setting of the privacy parameter ϵ, it is ϵ-differentially private with respect to D and if for every $D \in \mathbf{N}^{|\mathcal{X}|}$, $D' \in \mathcal{D}$ it outputs a $Q^* \in \mathcal{C}$ such that $|Q^*(D) - Q^*(D')| \geq \max_{Q \in \mathcal{C}} |Q(D) - Q(D')| - F(\epsilon)$ with probability at least $1 - \gamma$.*

For those paying attention, a distinguisher is just an agnostic learning algorithm for $\mathcal{C}$.

Algorithm 1 The Iterative Construction (IC) Mechanism. It takes as input a parameter ϵ_0, an $(F(\epsilon_0), \gamma)$-Private Distinguisher Distinguish for $\mathcal{C}$, together with an $B(\alpha)$-iterative database construction algorithm U for $\mathcal{C}$.

```
IC($D, \alpha, \epsilon_0, \text{Distinguish}, U$):
  Let $D^0 = U(\emptyset, \cdot, \cdot)$.
  for $t = 1$ to $B(\alpha/2)$ do
    Let $Q^{(t)} = \text{Distinguish}(D, D^{t-1})$
    Let $\hat{v}^{(t)} = Q^{(t)}(D) + \text{Lap}\left(\frac{1}{\epsilon_0 n}\right)$.
    if $|\hat{v}^{(t)} - Q^{(t)}(D^{t-1})| < 3\alpha/4$ then
      Output $D' = D^{t-1}$.
    else
      Let $D^t = U(D^{t-1}, Q^{(t)}, \hat{v}^{(t)})$.
    end if
  end for
  Output $D' = D^{B(\alpha/2)}$.
```

The analysis of this algorithm just involves checking the technical details of a simple intuition. Privacy will follow because the algorithm is just the composition of $2B(\alpha)$ steps, each of which is ϵ_0-differentially private. Accuracy follows because we are always outputting the last database in a maximal database update sequence. If the algorithm has not yet formed a maximal DUS, then the distinguishing algorithm will find a distinguishing query to add another step to the sequence.

Theorem 4 *The IC algorithm is ϵ-differentially private for $\epsilon_0 \leq \epsilon/2B(\alpha/2)$. The IC algorithm is (ϵ, δ)-differentially private for $\epsilon_0 \leq \frac{\epsilon}{4\sqrt{B(\alpha/2)\log(1/\delta)}}$.*

Proof The algorithm runs at most $2B(\alpha/2)$ compositions of ϵ_0-differentially private algorithms. Recall that ϵ_0 differentially private algorithms are $2k\epsilon_0$ differentially private under $2k$-fold composition, and are (ϵ', δ) private for $\epsilon' = \sqrt{4k\ln(1/\delta')}\epsilon_0 + 2k\epsilon_0(e^{\epsilon_0} - 1)$. Plugging in the stated values for ϵ_0 proves the claim. ■

Now we analyze the utility of the IC mechanism:

Theorem 5 *Given an $(F(\epsilon), \gamma)$-private distinguisher and a $B(\alpha)$-DUA, the Iterative Construction mechanism is (α, β) accurate for:*

$$\alpha \geq \max\left[\frac{4\log(2B(\alpha/2)/\beta)}{\epsilon_0 n}, 2F(\epsilon_0)\right]$$

so long as $\gamma \leq \beta/(2B(\alpha/2))$.

Proof The analysis is straightforward.

Recall that if $Y_i \sim \text{Lap}(1/(\epsilon n))$, we have: $\Pr[|Y_i| \geq t/(\epsilon n)] = \exp(-t)$. By a union bound, if $Y_1, \ldots, Y_k \sim \text{Lap}(1/(\epsilon n))$, $\Pr[\max_i |Y_i| \geq t/(\epsilon n)] \leq k \exp(-t)$. Therefore, because we make at most $B(\alpha/2)$ draws from $\text{Lap}(1/(\epsilon_0 n))$, except with probability at most $\beta/2$, for all t:

$$|\hat{v}^{(t)} - Q^{(t)}(D)| \leq \frac{1}{\epsilon_0 n} \log \frac{2B(\alpha/2)}{\beta} \leq \frac{\alpha}{4}$$

Note that by assumption, $\gamma \leq \beta/(2B(\alpha/2))$, so we also have that except with probability $\beta/2$:

$$|Q^{(t)}(D) - Q^{(t)}(D^{t-1})| \geq \max_{Q \in \mathcal{C}} |Q(D) - Q(D^{t-1})| - F(\epsilon_0) \geq \max_{Q \in \mathcal{C}} |Q(D) - Q(D^{t-1})| - \frac{\alpha}{2}$$

For the rest of the argument, we will condition on both of these events occurring, which is the case except with probability β.

There are two cases. Either a database $D' = D^{B(\alpha/2)}$ is output, or database $D' = D^{t-1}$ for $t \leq B(\alpha/2)$ is output. First, suppose $D' = D^{B(\alpha/2)}$. Since for all t it must have been such that $|\hat{v}^{(t)} - Q^{(t)}(D^{t-1})| \geq 3\alpha/4$ and by our conditioning, $|\hat{v}^{(t)} - Q^{(t)}(D)| \leq \frac{\alpha}{4}$, we know for all t: $|Q^{(t)}(D) - Q^{(t)}(D^{t-1})| \geq \alpha/2$. Therefore, the sequence $(D^t, Q^{(t)}, \hat{v}^{(t)})$, formed a maximal $(U, D, \mathcal{C}, \alpha/2, B(\alpha/2))$-Database Update Sequence. Therefore, we have that $\max_{Q \in \mathcal{C}} |Q(D) - Q(D')| \leq \alpha/2$ as desired.

Next, suppose $D' = D^{t-1}$ for $t < B(\alpha/2)$. Then it must have been the case that for t, $|\hat{v}^{(t)} - Q^{(t)}(D^{t-1})| < 3\alpha/4$. By our conditioning, in this case it must be that $|Q^{(t)}(D) - Q^{(t)}(D^{t-1})| < \alpha/2$, and that therefore by the properties of an $(F(\epsilon_0), \gamma)$-distinguisher:

$$\max_{Q \in \mathcal{C}} |Q(D) - Q(D')| < \alpha/2 + F(\epsilon_0) \leq \alpha$$

as desired. ■

Note that we can use the exponential mechanism as a distinguisher: take the domain to be $\mathcal{C}$, and let the quality score be: $q(D, Q) = |Q(D) - Q(D^t)|$, which has sensitivity $1/n$. Applying the exponential mechanism utility theorem, we get:

Theorem 6 *The exponential mechanism is an* $(F(\epsilon), \gamma)$ *distinguisher for:*

$$F(\epsilon) = \frac{2}{n\epsilon}\left(\log \frac{\mathcal{C}}{\gamma}\right)$$

Therefore, using the exponential mechanism as a distinguisher, Theorem 5 gives:

Theorem 7 *Given a* $B(\alpha)$*-DUA, the Iterative Construction mechanism is* (α, β) *accurate for:*

$$\alpha \geq \max\left[\frac{4\log(2B(\alpha/2)/\beta)}{\epsilon_0 n}, \frac{4}{n\epsilon_0}\left(\log \frac{\mathcal{C}}{\gamma}\right)\right]$$

so long as $\gamma \leq \beta/(2B(\alpha/2))$.

Plugging in our values of ϵ_0:

Theorem 8 *Given a* $B(\alpha)$*-DUA, the Iterative Construction mechanism is* (α, β) *accurate and* ϵ*-differentially private for:*

$$\alpha \geq \frac{8B(\alpha/2)}{n\epsilon}\left(\log \frac{\mathcal{C}}{\gamma}\right)$$

and (ϵ, δ)*-differentially private for:*

$$\alpha \geq \frac{16\sqrt{B(\alpha/2)\log(1/\delta)}}{n\epsilon}\left(\log \frac{\mathcal{C}}{\gamma}\right)$$

so long as $\gamma \leq \beta/(2B(\alpha/2))$.

We will see next time that database update algorithms actually exist, and recover the net-mechanism bounds, and even improve them for the case of (ϵ, δ)-differential privacy.

Bibliographic Information The abstraction of a "Database Update Algorithm" was given by Gupta, Roth, and Ullman in "Iterative Database Constructions and Private Data Release", 2011. Algorithms based on this framework were given initially in the interactive setting by Roth and Roughgarden in "Interactive Privacy Via the Median Mechanism", 2010, and then Hardt and Rothblum, "A Multiplicative Weights Mechanism for Privacy Preserving Data Analysis", 2010. The non-interactive mechanism paired with an agnostic learner was given by Gupta, Hardt, Roth, and Ullman in "Privately Releasing Conjunctions and the Statistical Query Barrier", 2011.