CIS 800/002 The Algorithmic Foundations of Data Privacy October 6, 2011

Lecture 8

Lecturer: Aaron Roth *Scribe: Aaron Roth*

Database Update Algorithms: The Median Mechanism

We'll recall some definitions from last time:

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* $t = 1, 2, \ldots, L-1$, $D^{t+1} = U(D^t, Q_t, v_t)$.

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 $L \leq B(\alpha)$.*

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 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 $1 - \gamma$.*

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 $D' = D^{B(\alpha/2)}$.

Using these objects we proved the following theorem:

Theorem 4 *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 $\gamma \leq \beta/(2B(\alpha/2))$.

This theorem will become more interesting if we can show that $B(\alpha)$-DUAS and $F(\epsilon)$-distinguishers actually exist for reasonable values of $B(\alpha)$, $F(\epsilon)$. Actually – we already know of a good distinguishing algorithm – the exponential mechanism!

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 5 *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 4 gives:

Theorem 6 *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 $\gamma \leq \beta/(2B(\alpha/2))$.

Recall that we also showed that for $(\epsilon, 0)$-differential privacy, we showed we can take:

$$\epsilon_0 = \epsilon/2B(\alpha/2)$$

and for (ϵ, δ)-differential privacy, we showed we can take:

$$\epsilon_0 = \frac{\epsilon}{4\sqrt{B(\alpha/2)\log(1/\delta)}}$$

Plugging in our values of ϵ_0:

Theorem 7 *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 $\gamma \leq \beta/(2B(\alpha/2))$.

Now, we'll give a conceptually simple DUA based again on the existence of small nets, and then we will give a more sophisticated DUA for linear queries that has better running time.

First, we'll give the "Median Mechanism". It won't operate on sequences of data sets, but instead on sequences of "median data structures":

Definition 8 (Median Data Structure) *A median data structure $\mathbf{D}$ is a collection of databases: $\mathbf{D} \subset \mathbf{N}^{|\mathcal{X}|}$. Any query Q can be evaluated on a median datastructure as follows: $Q(\mathbf{D}) = Median(\{Q(D) : D \in \mathbf{D}\})$.*

Algorithm 2 The Median Mechanism (MM) Algorithm. It inputs and outputs a median data structure. It is instantiated with an α-net $\mathcal{N}_\alpha(\mathcal{C})$ for a query class $\mathcal{C}$, and its initial state is $\mathbf{D} = \mathcal{N}_\alpha(\mathcal{C})$

$MM_{\alpha,\mathcal{C}}(\mathbf{D}^t, Q_t, v_t)$:
  if $\mathbf{D}^t = \emptyset$ **then**
    Output $\mathbf{D}^1 \leftarrow N_\alpha(\mathcal{C})$.
  end if
  if $v_t < Q_t(\mathbf{D}^t)$ **then**
    Output $\mathbf{D}^{t+1} \leftarrow \mathbf{D}^t \setminus \{D \in \mathbf{D} : Q_t(D) \geq Q_t(\mathbf{D^t})\}$.
  else
    Output $\mathbf{D}^{t+1} \leftarrow \mathbf{D}^t \setminus \{D \in \mathbf{D} : Q_t(D) \leq Q_t(\mathbf{D^t})\}$.
  end if

In words, a median data structure is just a set of databases. To evaluate a query on it, we just evaluate the query on every database in the set, and then return the median value.

The median mechanism is then very simple:

The intuition for the median mechanism is as follows. It maintains a set of databases that are consistent with the answers to the distinguishing queries it has seen so far. Whenever it receives a query and answer that differ substantially from the real database, it updates itself to remove all of the databases that are inconsistent with the new information. Because it always chooses its answer as the median database among the set of consistent databases it is maintaining, every update step removes at least half of the consistent databases! Moreover, because the set of databases that it chooses initially is an α-net with respect to $\mathcal{C}$, there is always some database that is never removed, because it remains consistent on all queries. This limits how many update rounds the mechanism can perform. How does the median mechanism do?

Theorem 9 *For any class of queries $\mathcal{C}$, The Median Mechanism is a $B(\alpha)$-database update algorithm for $B(\alpha) = \log |N_\alpha(\mathcal{C})|$.*

Proof We must show that any sequence $\{(D^t, Q_t, v_t)\}_{t=1,\ldots,L}$ with the property that $|Q^t(\mathbf{D}^t) - Q^t(D)| > \alpha$ and $|v_t - Q^T(D)| < \alpha$ cannot have $L > \log |\mathcal{N}_\alpha(\mathcal{C})|$. First observe that because $\mathbf{D}^0 = \mathcal{N}_\alpha(\mathcal{C})$ is an α-net for $\mathcal{C}$, by definition there is at least one D' such that $D' \in \mathbf{D}^t$ for all t (i.e. since it has error less than α on all queries, it is never removed by an update step). Thus, we can always answer queries with $\mathbf{D}^t$, and for all t, $|\mathbf{D}^t| \geq 1$. Next observe that for each t, $|\mathbf{D}^t| \leq |\mathbf{D}^{t-1}|/2$. This is because each update step removes at least half of the elements: all of the elements at least as large as, or at most as large as the median element in $\mathbf{D}^t$ with respect to query Q_t. Therefore, after L update steps, $|\mathbf{D}^L| \leq 1/2^L \cdot |\mathcal{N}_\alpha(\mathcal{C})|$. Setting $L > \log |\mathcal{N}_\alpha(\mathcal{C})|$ gives $|\mathbf{D}^L| < 1$, a contradiction. ■

Recall that for linear queries, we have already upper bounded $|N_\alpha(\mathcal{C})|$ so we have:

Theorem 10 *For any class of linear queries $\mathcal{C}$, the Median Mechanism is a $B(\alpha)$-DUA for $B(\alpha) = \frac{\log |\mathcal{X}| \cdot \log |\mathcal{C}|}{\alpha^2}$.*

Plugging this in to the IC mechanism, we get:

Theorem 11 *Instantiated with the median mechanim, the Iterative Construction mechanism is (α, β) accurate and ϵ-differentially private for:*

$$\alpha \geq \frac{32 \log |\mathcal{X}| \log |\mathcal{C}|}{n\epsilon\alpha^2} \left(\log \frac{\mathcal{C}}{\gamma}\right)$$

$$\alpha = \tilde{O}\left(\frac{\log |\mathcal{X}| \log^2 |\mathcal{C}| \log(1/\beta)}{\epsilon n}\right)^{1/3}$$

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

$$\alpha \geq \frac{32\sqrt{\log|\mathcal{X}|\log|\mathcal{C}|\log(1/\delta)}}{n\epsilon\alpha}\left(\log\frac{\mathcal{C}}{\gamma}\right)$$

$$\alpha \geq \tilde{O}\left(\frac{(\log|\mathcal{X}|\log^3|\mathcal{C}|\log(1/\delta)\log^2(1/\beta))^{1/4}}{\sqrt{\epsilon n}}\right)$$

Note that we recover our bounds from the net mechanism, and improve on those bounds for (ϵ, δ)-differential privacy. This is the advantage of a mechanism which works as a composition of many smaller steps: we can apply our composition theorems.

Next time we will give a more sophisticated database update algorithm for linear queries. It will work by maintaining a distribution $\hat{D}^t$ over the data universe X.

Bibliographic Information The median mechanism first appeared in the context of an interactive query release mechanism, in "Interactive Privacy via the Median Mechanism", by Roth and Roughgarden.