

Lecture 6

Lecturer: Aaron Roth

Scribe: Aaron Roth

The Net Mechanism: A Partial Converse

Finishing up from last time.

Last time we showed:

Theorem 1 (Net Mechanism Theorem) For any class of counting queries C the Net Mechanism is $(2\alpha, \beta)$ -useful for any α such that:

$$\alpha \geq \frac{2}{\epsilon n} \log \frac{N_\alpha(C)}{\beta}$$

and:

Theorem 2 For any finite class of counting queries C :

$$|N_\alpha(C)| \leq |X|^{\frac{\log |C|}{\alpha^2}}$$

We can now state our utility theorem for the net mechanism.

Theorem 3 $\text{NetMechanism}(D, C, \epsilon, \frac{\alpha}{2})$ is (α, β) -accurate for any α such that:

$$\alpha \geq \left(\frac{16 \log |X| \log |C| + 2 \log \left(\frac{1}{\beta} \right)}{\epsilon n} \right)^{1/3} \quad (1)$$

Equivalently, the required size of the database D is:

$$n \geq \frac{16 \log |X| \log |C| + 4 \log \left(\frac{1}{\beta} \right)}{\epsilon \alpha^3} \quad (2)$$

Proof By Theorem 1, the mechanism is (α, β) -useful for any α such that:

$$\frac{\alpha}{2} \geq \frac{2}{\epsilon n} \log \frac{N_\alpha(C)}{\beta}$$

By Theorem 2, it is sufficient to take:

$$\frac{\alpha}{2} \geq \frac{8}{\epsilon n} \left(\frac{\log |X| \log |C|}{\alpha^2} + \log \frac{1}{\beta} \right)$$

Solving for α completes the proof. ■

The accuracy bound we have just proven depends only on the *cardinality* of the class of functions C . But this does not seem like the right measure of complexity. What if the class just contains a single query Q over and over again? And for super-exponentially sized concept classes, the bound we have proven gives no guarantee. It turns out that we can instead use another measure of function complexity.

Definition 4 (Shattering) A class of predicates P shatters a collection of points $S \subseteq X$ if for every $T \subseteq S$, there exists an $\varphi \in P$ such that $\{x \in S : \varphi(x) = 1\} = T$. That is, P shatters S if for every one of the $2^{|S|}$ subsets T of S , there is some predicate in P that labels exactly those elements as positive, and does not label any of the elements in $S \setminus T$ as positive.

We can now define our complexity measure for counting queries.

Definition 5 (VC-Dimension) *A collection of predicates P has VC-dimension d if there exists some set $S \subseteq X$ of cardinality $|S| = d$ such that P shatters S , and P does not shatter any set of cardinality $d + 1$. We can denote this quantity by $VC-DIM(P)$. We abuse notation and also write $VC-DIM(C)$ where C is a class of counting queries, to denote the VC-dimension of the corresponding collection of predicates.*

It turns out that we can essentially replace the $\log |C|$ term in the above bounds with $VCDIM(C)$. Note that for finite concept classes, we always have $VCDIM(C) \leq \log |C|$, and so this is always a stronger bound.

We can replace the sampling lemma from last time with the following lemma:

Lemma 6 *For any $D \in \mathbf{N}^{|\mathcal{X}|}$ and for any collection of counting queries C , there exists a database D' of size*

$$|D'| = O(VCDIM(C) \log(1/\alpha)/\alpha^2)$$

such that:

$$\max_{Q \in C} |Q(D) - Q(D')| \leq \alpha$$

This lemma straightforwardly gives an analogue of Theorem 2:

Theorem 7 *For any class of counting queries C :*

$$|N_\alpha(C)| \leq |X|^{O(VCDIM(C) \log(1/\alpha)/\alpha^2)}$$

Finally, we can instantiate Theorem 1 to give our main utility theorem for the Net Mechanism.

Theorem 8 *For any class of counting queries C the Net Mechanism is (α, β) -useful for any α such that:*

$$\alpha \geq O\left(\frac{1}{\epsilon \alpha^2 n} (VCDIM(C) \log |X| \log(1/\alpha) + \log 1/\beta)\right)$$

Solving for α , the Net Mechanism is (α, δ) -useful for:

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

Theorem 8 shows that a database of size $\tilde{O}\left(\frac{\log X VCDIM(C)}{\alpha^3 \epsilon}\right)$ is sufficient in order to output a set of points that is α -useful for a concept class C , while simultaneously preserving ϵ -differential privacy. If we were to view our database as having been drawn from some distribution \mathcal{D} , this is only an extra $\tilde{O}\left(\frac{\log X}{\alpha \epsilon}\right)$ factor larger than what would be required to achieve α -usefulness with respect to \mathcal{D} , even without any privacy guarantee!

Lets take a step back and restate our net mechanism theorem with different notation.

First let us define the query metric with respect to C :

Definition 9 (Query Metric) *With respect to a class of functions C , the query metric $d_C : \mathbf{N}^{|\mathcal{X}|} \times \mathbf{N}^{|\mathcal{X}|} \rightarrow \mathbf{R}$ is defined as:*

$$d_C(D, D') = \max_{Q \in C} |Q(D) - Q(D')|$$

Note that with respect to any set of datastructures R on which we can evaluate queries C (i.e. such that there exists a function $EVAL : R \times C \rightarrow \mathbf{R}$ and a convention that $Q(r) = Eval(r, Q)$) we can abuse notation and apply the query metric d_C to pairs $D \in \mathbf{N}^{|\mathcal{X}|}$ and $r \in R$.

Definition 10 The α -ball around a point x with respect to a class of queries C is:

$$B_\alpha(x) = \{D \in \mathbf{N}^{|\mathcal{X}|} : d_C(x, D) \leq \alpha\}$$

We recall our definition of useful mechanisms.

Definition 11 A mechanism $M : \mathbf{N}^{|\mathcal{X}|} \rightarrow R$ is (α, β) -useful if for every $D \in \mathbf{N}^{|\mathcal{X}|}$, we have that $\Pr[M(D) \in B_\alpha(D)] \geq 1 - \beta$.

Recall that with respect to a class of counting queries C , we defined an α -net:

Definition 12 (α -net) An α -net of databases with respect to a class of queries C is a set $N \subset \mathbf{N}^{|\mathcal{X}|}$ such that for all $D \in \mathbf{N}^{|\mathcal{X}|}$, there exists an element of the α -net $D' \in N$ such that $D \in B_\alpha(D')$. We write $N_\alpha(C)$ to denote an α -net of minimum cardinality among the set of all α -nets for C .

We proved:

Theorem 13 (Net Mechanism Theorem) For any class of counting queries C the Net Mechanism is $(2\alpha, \beta)$ -useful for any α such that:

$$\alpha \geq \frac{2}{\epsilon n} \log \frac{N_\alpha(C)}{\beta}$$

Rephrased:

Theorem 14 (Net Mechanism Theorem Rephrased) For any class of counting queries C , if α is such that:

$$|N_\alpha(C)| \leq \beta \exp\left(\frac{\epsilon \alpha n}{2}\right)$$

then there exists an ϵ -differentially private release mechanism that is $(2\alpha, \beta)$ -useful with respect to C .

That is, the existence of small nets certifies the existence of accurate private algorithms. Today we will prove a partial converse of this theorem: the existence of accurate, private algorithms for a class of queries C certifies that there exist small nets for C .

Theorem 15 If there exists an ϵ -differentially private mechanism that is (α, β) -useful with respect to a class of queries C , then:

$$|N_{2\alpha}| \leq \frac{\exp(\epsilon n)}{1 - \beta}$$

In other words, the accuracy with which answers to queries C can be released while preserving privacy is closely related to the compressibility of those answers.

To prove the converse theorem, we will consider the following thought experiment. Suppose we have some subset of databases $S \subset \mathbf{N}^{|\mathcal{X}|}$ of n elements, and we draw a database $D \in S$ uniformly at random. Now, suppose we compute $M(D) = y$, where $M : \mathbf{N}^{|\mathcal{X}|} \rightarrow R$ is some ϵ -differentially private mechanism. What is the posterior probability of having selected $D \in S$ conditioned on seeing the output y of the mechanism?

Lemma 16 Let $S \subset \mathbf{N}^{|\mathcal{X}|}$ be some subset of databases such that for all $D \in S$, $|D| \leq n$. If $D \in S$ is drawn uniformly at random, then for all $D' \in S$ and for all $y \in R$:

$$\frac{\Pr[D = D' | M(D) = y]}{\Pr[D = D']} \leq \exp(\epsilon n)$$

Proof

$$\begin{aligned}
\frac{\Pr[D = D' | M(D) = y]}{\Pr[D = D']} &= \frac{\Pr[M(D) = y | D = D']}{\Pr[M(D) = y]} \\
&\leq \frac{\Pr[M(D) = y | D = D']}{\Pr[M(D) = y | D = D^*]} \quad \text{For some } D^* \in S \\
&\leq \exp(\epsilon \|D - D^*\|_1) \\
&\leq \exp(\epsilon n)
\end{aligned}$$

The equality follows from Bayes rule. The first inequality follows from averaging, the second follows from the definition of differential privacy, and the last inequality follows from the fact that all databases in S are at distance at most n . ■

Corollary 17 *If $D \in S$ is drawn uniformly at random, then for any subset of databases $T \subseteq S$, and for all $y \in R$:*

$$\Pr[D \in T | M(D) = y] \leq \exp(\epsilon n) \cdot \frac{|T|}{|S|}$$

Proof

$$\begin{aligned}
\Pr[D \in T | M(D) = y] &= \sum_{D' \in T} \Pr[D = D' | M(D) = y] \\
&\leq \exp(\epsilon n) \sum_{D' \in T} \Pr[D = D'] \\
&= \exp(\epsilon n) \Pr[D \in T] \\
&= \exp(\epsilon n) \cdot \frac{|T|}{|S|}
\end{aligned}$$

where the inequality follows from Lemma 16. ■

We will now argue that due to the accuracy of the mechanism, the set $B_\alpha(y) \cap S$ is in expectation large with respect to the choice of y .

Lemma 18 *Let $S \subset \mathbf{N}^{|X|}$ be some subset of databases such that for all $D \in S$, $|D| \leq n$. If $D \in S$ is drawn uniformly at random and y is drawn according to $M(D)$, then:*

$$\mathbb{E}_y[\Pr_D[D \in B_\alpha(y) \cap S | M(D) = y]] \geq 1 - \beta$$

Proof The accuracy guarantee of the mechanism promises that for all $D \in S$, $\Pr_y[D \in B_\alpha(y)] \geq 1 - \beta$. Of course we always have $D \in S$. In particular:

$$\begin{aligned}
1 - \beta &\leq \Pr_{D,y}[D \in B_\alpha(y) \cap S] \\
&= \sum_{y' \in T} \Pr_D[D \in B_\alpha(y') \cap S | M(D) = y'] \Pr_y[y = y'] \\
&= \mathbb{E}_y[\Pr_D[D \in B_\alpha(y) \cap S | M(D) = y]]
\end{aligned}$$

■

On the other hand, we know:

$$\mathbb{E}_y[\Pr_D[D \in B_\alpha(y) \cap S | M(D) = y]] \leq \exp(\epsilon n) \mathbb{E}_y \left[\frac{|B_\alpha(y) \cap S|}{|S|} \right]$$

So putting these two facts together, we know:

$$\mathbb{E}_y[|B_\alpha(y) \cap S|] \geq (1 - \beta) \frac{|S|}{\exp(\epsilon n)}$$

Of course, whenever $D \in B_\alpha(y)$, for any $D' \in B_\alpha(y)$ we also have: $D' \in B_{2\alpha}(D)$ (by the triangle inequality). Therefore we have proven:

Theorem 19 *If there exists an ϵ -differentially private mechanism that is (α, β) -useful with respect to a class of counting queries C , then the following holds. For any subset of databases $S \subset \mathbf{N}^{|\mathcal{X}|}$, when D is selected uniformly at random from S :*

$$\mathbb{E}_D[|B_{2\alpha}(D) \cap S|] \geq (1 - \beta) \frac{|S|}{\exp(\epsilon n)}$$

We can now finish up our theorem. We need one more definition.

Definition 20 (α -packing) *A set $P \subset \mathbf{N}^{|\mathcal{X}|}$ is an α -packing if for all $D, D' \in P$: $D \notin B_\alpha(D')$. We write $P_\alpha(C)$ to denote the α -packing of maximum cardinality.*

To complete the proof, we can relate the size of packings and nets:

Lemma 21

$$N_{\alpha/2}(C) \geq P_\alpha(C) \geq N_\alpha(C)$$

Proof For the first inequality: by definition, for each $D \in P_\alpha(C)$, there is some $x \in N_{\alpha/2}(C)$ such that $d_C(D, x) \leq \alpha/2$. For each $D, D' \in P_\alpha(C)$, these x, x' must be distinct, because $d(x, D') \geq d(D, D') - d(x, D) > \alpha - \alpha/2 > \alpha/2$.

For the second inequality: Let $N = P_\alpha(C)$. Suppose N is not an α -net: i.e. there is some $D \in \mathbf{N}^{|\mathcal{X}|}$ such that for all $D' \in N$: $D \notin B_\alpha(D')$. Then $P_\alpha(C) \cup D$ is still an α -packing, contradicting the fact that $P_\alpha(C)$ is a maximum cardinality packing. Therefore $|N_\alpha(C)| \leq |N| = |P_\alpha(C)|$ ■

Finally, to complete the proof:

Proof [of Theorem 15] Let $S = P_{2\alpha}(C)$. By definition, for all $D \in S$: $|B_{2\alpha}(D) \cap S| = 1$. We also know:

$$1 = \mathbb{E}_D[|B_{2\alpha}(D) \cap S|] \geq (1 - \beta) \frac{|P_{2\alpha}(C)|}{\exp(\epsilon n)} \geq (1 - \beta) \frac{|N_{2\alpha}(C)|}{\exp(\epsilon n)}$$

So:

$$N_{2\alpha}(C) \leq \frac{\exp(\epsilon n)}{1 - \beta}$$

■

Bibliographic Information