Lecture 13: Strong Composition

Lecturer: Adam Smith *Scribe: Adam Smith*

1 Strong Composition

In this lecture, we show that (ϵ, δ)-differential privacy satisfies a "strong composition" theorem, in which the ϵ parameter increases only with the square root of the number of stages of the composition.

Theorem 1 (Strong Composition) *For all $\epsilon, \delta \geq 0$ and $\delta' > 0$, the* adaptive *composition of k algorithms, each of which is (ϵ, δ)-differentially private, is $(\tilde{\epsilon}, \tilde{\delta})$-differentially private where $\tilde{\epsilon} = \epsilon\sqrt{2k\ln(1/\delta')} + k\epsilon\frac{e^\epsilon - 1}{e^\epsilon + 1}$ and $\tilde{\delta} = k\delta + \delta'$.*

If X and Y are random variables taking values in the same set (and with probabilities defined for the same collection of events), we say $X \approx_{\epsilon,\delta} Y$ if for every event E: $P_X(E) \leq e^\epsilon P_Y(E) + \delta$ and $P_Y(E) \leq e^\epsilon P_X(E) + \delta$.

We would like to characterize this relation in simpler terms. As a starting point, let's try to imagine the simplest pair of random variables that satisfies the relationship. It seems like we need one type of outcome to capture the δ additive difference in probabilities, and another type that captures the e^ϵ multiplicative change. Consider the following two special random variables, U and V, taking values in the set $\{0, 1, \text{"I am U"}, \text{"I am V"}\}$ with the probabilities

Outcome	P_U	P_V
0	$\frac{e^\epsilon(1-\delta)}{e^\epsilon+1}$	$\frac{1-\delta}{e^\epsilon+1}$
1	$\frac{1-\delta}{e^\epsilon+1}$	$\frac{e^\epsilon(1-\delta)}{e^\epsilon+1}$
"I am U"	δ	0
"I am V"	0	δ

Lemma 2 *For every pair of random variables X, Y such that $X \approx_{\epsilon,\delta} Y$, there exists a randomized map F such that $F(U) \sim X$ and $F(V) \sim Y$.*

We leave this proof as a homework problem, but provide the following pictorial hint:

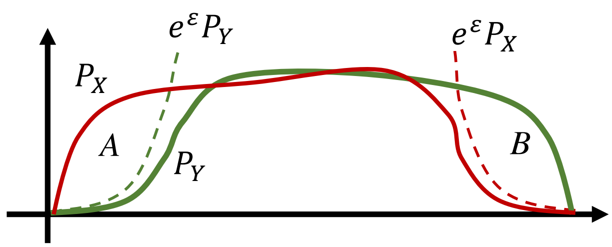

Figure 1: The "proof" of Lemma 2

The first step of the proof is to show that the areas of the regions A and B are both bounded by δ. The rest is the homework problem! It is ok to assume, for the sake of the homework problem, that X and Y take values in a discrete set.

We can now proceed to the proof of Strong Composition (Theorem 1).

Proof Fix a sequence of k mechanisms M_j, each of which takes a data set in $\mathcal{X}^n$ as well as a partial transcript $a_1, ..., a_{j-1}$ (abbreviated $\mathbf{a}_1^{j-1}$) such that, for every partial transcript, $M_j(\cdot; \mathbf{a}_1^{j-1})$ is (ϵ, δ)-differentially private. Also, fix two data sets $\mathbf{s}, \mathbf{s}'$ that differ in one entry.

For every partial transcript $\mathbf{a}_1^{j-1}$, we have $M_j(\mathbf{s}; \mathbf{a}_1^{j-1}) \approx_{\epsilon,\delta} M_j(\mathbf{s}'; \mathbf{a}_1^{j-1})$ and so there exists a randomized map $F_{\mathbf{a}_1^{j-1}}$ such that $F_{\mathbf{a}_1^{j-1}}(U)$ and $F_{\mathbf{a}_1^{j-1}}(V)$ have the same distributions as $M_j(\mathbf{s}; \mathbf{a}_1^{j-1})$ and $M_j(\mathbf{s}'; \mathbf{a}_1^{j-1})$, respectively.

This allows use to show the first important claim:

Claim 3 *There is a randomized map F^* such that the composed mechanism M satisfies:*

$$M(\mathbf{s}) \sim F^*(U_1, ..., U_k) \text{ where } U_1, ..., U_k \sim_{i.i.d.} U \text{ and} \tag{1}$$

$$M(\mathbf{s}') \sim F^*(V_1, ..., V_k) \text{ where } V_1, ..., V_k \sim_{i.i.d.} V\,. \tag{2}$$

Proof [of claim] Consider the algorithm:

Algorithm 1: $F^*(z_1, ..., z_k)$:

```
1 for j = 1 to k do
2    a_j ← F_{a_1^{j-1}}(z_j) ;
3 return (a_1, ..., a_k).
```

Since $F_{\mathbf{a}_1^{j-1}}(U_j)$ has the same distribution as $M_j(\mathbf{s}; \mathbf{a}_1^{j-1})$ for each stage j, the overall distribution of $F^*(U_1, ..., U_k)$ is the same as $M(\mathbf{s})$ (and similarly for $\mathbf{s}'$ when the inputs are i.i.d. copies of V). ■

To prove that M is $\tilde{\epsilon}, \tilde{\delta}$-differentially private, it suffices, by closure under postprocessing, to prove that $(U_1, ..., U_k) \approx_{\tilde{\epsilon},\tilde{\delta}} (V_1, ..., V_k)$.

We'll consider two "bad events": B_1 and B_2. The first, B_1, is when we see a clear signal that the input was drawn according to U:

$$B_1 = \{\mathbf{z} : \text{at least one } z_j \text{ is "I am U"}\}. \tag{3}$$

Under $\mathbf{z}$ is distributed according to either $U_1, ..., U_k$ or $V_1, ..., v_k$, the probability of B_1 is exactly $1 - (1-\delta)^k \leq k\delta$.

If $\mathbf{z} \sim U_1, ..., U_k$, then conditioned on $\bar{B}_{1,u}$ not occurring, we have $\mathbf{z} \in \{0,1\}^k$. The probability of $\mathbf{z}$ is nonzero under both U and V, and we can compute the odds ratio by taking advantage of independence:

$$\ln\left(\frac{P_U(\mathbf{z})}{P_V(\mathbf{z})}\right) = \sum_j \ln\left(\frac{P_U(z_j)}{P_V(z_j)}\right) = \sum_j \ln\left(\frac{(1-\delta)e^{\epsilon(1-z_j)}/(e^\epsilon+1)}{(1-\delta)e^{\epsilon(z_j)}/(e^\epsilon+1)}\right) = \sum_j \epsilon(-1)^{z_j}\,.$$

This log odds ratio is thus a sum of bounded, independent random variables under distribution U, with expectation

$$\mathbb{E}_{\mathbf{z}\sim(U_1,\ldots,U_k)}\left(\frac{P_U(\mathbf{z})}{P_V(\mathbf{z})}\middle|\bar{B}_1\right) = k\epsilon \cdot \mathbb{E}\left((-1)^U \middle| U \in \{0,1\}\right) = k\epsilon\frac{e^\epsilon-1}{e^\epsilon+1}\,.$$

By the Chernoff bound (Lecture 1), for any $t > 0$ we have

$$\Pr_{\mathbf{z}\sim U_1,\ldots,U_k}\left(\underbrace{\ln\left(\frac{P_U(\mathbf{z})}{P_V(\mathbf{z})}\right) > \tilde{\epsilon}}_{\text{event } B_2}\middle|\bar{B}_1\right) \leq e^{-t^2/2} \text{ where } \tilde{\epsilon} \stackrel{\text{def}}{=} k\epsilon\frac{e^\epsilon-1}{e^\epsilon+1} + t\epsilon\sqrt{k}.$$

Let B_2 be the event that $\left\{\mathbf{z} \in \{0,1\}^k : \ln\left(\frac{P_U(\mathbf{z})}{P_V(\mathbf{z})}\right) > k\epsilon\frac{e^\epsilon-1}{e^\epsilon+1} + t\epsilon\sqrt{k}\right\}$. Note that conditioned on $\bar{B}_1 \cap \bar{B}_2$, the ratio of $P_U(\mathbf{z})$ to $P_V(\mathbf{z})$ is bounded. Hence, for any event E,

$$P_U(E \cap \bar{B}_1 \cap \bar{B}_2) \leq e^{\tilde{\epsilon}} P_V(E \cap \bar{B}_1 \cap \bar{B}_2) \leq e^{\tilde{\epsilon}} P_V(E)\,.$$

This allows us to show the indistinguishability condition we want:

$$\begin{aligned} P_U(E) &\leq P_U(E \cap \bar{B}_1 \cap \bar{B}_2) + P_U(B_1) + P_U(B_2|\bar{B}_1)P_U(\bar{B}_1) \\ &\leq e^{\tilde{\epsilon}} P_V(E) + k\delta + e^{-t^2/2} \,. \end{aligned}$$

Setting $t = \sqrt{2\ln(1/\delta')}$ yields the theorem statement. ■

Exercise 1 *Use the proof strategy from the previous theorem to show that the composition of an (ϵ_1, δ_1)-DP algorithm with a (ϵ_2, δ_2)-DP algorithm is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$-DP.*

Exercise 2 *Using Lemma 2, show that if $X \approx_{\epsilon,0} Y$, then $D_{KL}(P_X \| P_Y) \leq \epsilon \frac{e^\epsilon - 1}{e^\epsilon + 1}$ (which is a tighter bound than the one we derived in earlier lectures).*

2 Notes

The first version of the strong composition theorem appeared in [**?**]. Our presentation is based on Kairouz et al. [KOV17], as well as Dwork and Roth [DR14, Sections 3.5.1–2]. The characterization of ϵ, δ indistinguishability of Lemma 2 is due to [KOV17]. Their proof is based on a much more general result of Blackwell (1953). The homework problem asks students to provide a direct proof of this special case.

References

[DR14] Cynthia Dwork and Aaron Roth. *The Algorithmic Foundations of Differential Privacy*, volume 9. Foundations and Trends® in Theoretical Computer Science, 2014.

[KOV17] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Trans. Information Theory*, 63(6):4037–4049, 2017.