One objective that differential privacy achieves is that no individual stands out in data that is released by a private mechanism. In our discussion so far we did not make any assumption on what the data looks like and this objective was achieved purely by mechanism design. The notion of crowd-blending privacy, which we discuss next, allows for more accurate answers in cases when parts of the data itself are anonymous.

1 Crowd-blending privacy

The objective of crowd-blending privacy is to model scenarios where an individual's data attribute of interest looks identical to that of many others' in the database. For example, suppose the database contains people and their favourite colors: 40 people like blue, 46 like red, 7 like yellow, and only 2 like yellow. If I like blue, I blend in among 39 other people, so I should not suffer a particular harm if the number of blue-lovers is released. On the other hand, if I am among the rare yellow-lovers, the release of the yellow count may allow an observer to infer a fairly accurate conclusion about my participation in the database.

Crowd-blending privacy distinguishes between the case when there are many other people in the database that share my type and the case when there aren't. In the second case, my privacy should be protected in the usual way: Replacing my data with *any other* data should yield outcomes of similar probabilities. In the first case, the definition merely requires that the mechanism does not distinguish between myself and any other individual that shares my type.

We will not formalise the notion of "my type" explicitly but give an operational, indistiguishability-based definition: Two individuals are of the same type with respect to the mechanism if replacing one with the other does not affect the output by much, regardless of who else is in the database.

Definition 1. Mechanism $M\colon D^n \to R$ *ε-blends* entries $r, r' \in D$ if for every i, every $x_{-i} \in D^{[n]-\{i\}}$ and every y,

$$\Pr[M(x_{-i}, r) = y] \leq e^{\varepsilon} \Pr[M(x_{-i}, r') = y].$$

Here, $(x_{-i}, x_i = r)$ is the database that contains r as its i-th row and x_{-i} in its other rows. This definition makes sense even for $\varepsilon = 0$; we then say M blends x_i and x_i'.

The definition of crowd-blending privacy requires that an individual either blends with many other entries in the database, or if it doesn't, then it is not sensitive to the presence of that individual. For this definition, we will not fix the number of rows in the database ahead of time.

Definition 2. Mechanism $M\colon D^* \to R$ is (k, ε)-crowd blending if for every n, every $x \in D^n$, and every i,

- There exist at least k rows $j \in [n]$ such that M ε-blends x_i with x_j, *or*

- For all y, $e^{-\varepsilon}\Pr[M(x_{-i}) = y] \leq \Pr[M(x) = y] \leq e^{\varepsilon}\Pr[M(x_{-i}) = y]$.

The two possibilities are not exclusive; if we disallow the first one, we recover the usual notion of ε-differential privacy.[1]

One natural example of a (k, ε)-crowd blending mechanism that is not ε-differentially private is the following mechanisms for histograms. Recall that a histogram query q_h for a function $h\colon D \to B$ on input $x \in D^n$ is the vector of counts $k_b = |\{i\colon h(x_i) = b\}|$ as b ranges over the set of buckets B.

Mechanism $Hist_h(x)$, where $x \in D^n$ and $h\colon D \to B$:
For all $b \in B$:
Let $k_b = |\{i\colon h(x_i) = b\}|$.
If $k_b \geq k$, output k_b.
Otherwise output $\perp$.

This mechanism is clearly not differentially private as it is deterministic.

Theorem 3. *Mechanism $Hist_h$ is $(k, 0)$-crowd blending.*

Proof. Fix x and i. Let H be the set of all $j \in [n]$ such that $h(x_i) = h(x_j)$. We consider two cases.

If $|H| \geq k$, since the mechanism M blends x_i with x_j for every $j \in H$ (regardless of what the other rows are), the first condition in the definition is satisfied.

If $|H| < k$, then removing the i-th entry only decreases the number of entries in bucket b so both $Hist_h(x)$ and $Hist_h(x')$ output $\perp$ for bucket b; all the other outputs stay the same. □

2 Differential privacy from crowd-blending privacy

While crowd-blending private mechanisms are not in general differentially private, the weaker privacy notion can sometimes be used to achieve the stronger one. One such case is when the data consists of random independent samples from a large population. For example, the above histogram mechanism applied to a random sample from a sufficiently large population produces differentially private output.

Suppose M is a (k, ε)-crowd blending mechanism and consider the following mechanism M'.

Mechanism M': On input $x \in D^n$,
Choose $S \subseteq [n]$ by including each entry $i \in [n]$ independently with probability ε.
Ouput $M(x|_S)$, where $x|_S \in D^S$ is the collection of rows indexed by entries in S.

Theorem 4. *If mechanism M is (k, ε)-crowd blending private then mechanism M' is $(O(\varepsilon), e^{-\Omega(k)})$-differentially private.*

[1] To be precise, this definition talks about removing a row instead of modifying it, resulting in a factor of two loss in the privacy parameter.

To prove Theorem 4, we will show that for every x, i, and event T

$$\Pr[M'(x) \in T] \leq e^{O(\varepsilon)} \cdot \Pr[M'(x_{-i}) \in T] + e^{-\Omega(k)}$$

and

$$\Pr[M'(x_{-i}) \in T] \leq e^{O(\varepsilon)} \cdot \Pr[M'(x) \in T] + e^{-\Omega(k)}$$

where x_{-i} is the database x with row i removed. The usual requirement follows easily by combining these two inequalities. They can be summarized in the single condition

$$\big|\Pr[M'(x) \in T] - \Pr[M'(x_{-i}) \in T]\big| \leq O(\varepsilon) \cdot \min\{\Pr[M'(x) \in T], \Pr[M'(x_{-i}) \in T]\} + e^{-\Omega(k)}. \quad (1)$$

Fix x and i and let B be the set of rows $j \neq i$ such that M ε-blends x_i with x_j. When B is large, we will argue that x_i is very likely to either be absent from x_S or blend in it and conclude differential privacy by the blending condition. If not, then x_i is very unlikely to blend with k entries of x_S by the other condition in the definition of crowd-blending privacy.

Lemma 5. *For $\varepsilon \leq 1$ and every x, i, and T,*

$$\big|\Pr[M'(x) \in T] - \Pr[M'(x_{-i}) \in T]\big| \leq O(\varepsilon) \cdot \min\{\Pr[M'(x) \in T], \Pr[M'(x_{-i}) \in T]\} + e^{-\Omega(\varepsilon|B|)}.$$

Lemma 6. *If $\varepsilon|B| \leq k/2$ then for every x, i, and T,*

$$\big|\Pr[M'(x) \in T] - \Pr[M'(x_{-i}) \in T]\big| \leq O(\varepsilon) \cdot \min\{\Pr[M'(x) \in T], \Pr[M'(x_{-i}) \in T]\} + e^{-\Omega(k)}.$$

By combining these two lemmas we derive inequality (1) and prove Theorem 4.

Proof sketch of Lemma 5. Using conditional probabilities, we can write

$$\begin{aligned}\Pr[M(x_S) \in T] &= (1-\varepsilon)\Pr[M(x|_S) \in T \mid i \notin S] + \varepsilon \Pr[M(x|_S) \in T \mid i \in S] \\ &= (1-\varepsilon)\Pr[M(x_{-i}|_S) \in T] + \varepsilon \Pr[M(x|_S) \in T \mid i \in S]\end{aligned}$$

from where

$$\big|\Pr[M(x_S) \in T] - \Pr[M(x_{-i}|_S) \in T]\big| = \varepsilon\big|\Pr[M(x|_S) \in T \mid i \in S] - \Pr[M(x_{-i}|_S) \in T]\big|.$$

Let E be the event $\varepsilon|B|/2 \leq |B - S| \leq 2\varepsilon|B|$. By a multiplicative Chernoff bound (a different variant from the one in Lecture 2), $\Pr[\overline{E}] < 2^{-\Omega(\varepsilon|B|)}$. Therefore,

$$\begin{aligned}&\big|\Pr[M(x_S) \in T] - \Pr[M(x_{-i}|_S) \in T]\big| \\ &\qquad \leq \varepsilon\big|\Pr[M(x|_S) \in T \mid i \in S, E] - \Pr[M(x_{-i}|_S) \in T \mid E]\big| + 2^{-\Omega(\varepsilon|B|)}.\end{aligned}$$

To finish the proof, we will show that

$$\begin{aligned}&\Pr[M(x|_S) \in T \mid i \in S, E] = \Theta(\Pr[M(x_{-i}|_S) \in T \mid E]) \quad \text{and} \\ &\Pr[M(x|_S) \in T \mid i \in S, E] = \Theta(\Pr[M(x|_S) \in T \mid E]).\end{aligned}$$

Let $S' = S - \{i\} \cup \{i'\}$, where i' is a random element of $B - S$. By the crowd-blending privacy of M,

$$e^{-\varepsilon} \Pr[M(x|_{S'}) \in T \mid i \in S, E] \leq \Pr[M(x|_S) \in T \mid i \in S, E] \leq e^{\varepsilon} \Pr[M(x|_{S'}) \in T \mid i \in S, E]$$

so it is sufficient to show that the conditional probabilities of the events $M(x|_{S'}) \in T \mid i \in S, E$ and $M(x_{-i}|_S) \in T \mid E$ are of the same order of magnitude. To do this, we compare the probabilities of the outcomes $S = A$ and $S' = A$ under the two conditional distributions for any set $A \subseteq [n] - \{i\}$ such that $\varepsilon|B|/2 \leq |A \cap B| \leq 2\varepsilon|B|$.

Let's write A_B and A_{-B} for the sets $A \cap B$ and $A - B$. Then

$$\frac{\Pr[S = A]}{\Pr[S' = A]} = \frac{\Pr[S_B = A_B] \Pr[S_{-B} = A_{-B}]}{\Pr[S'_B = A_B] \Pr[S'_{-B} = A_{-B}]} = \frac{\Pr[S_B = A_B]}{\Pr[S'_B = A_B]}$$

because outside B, S and S' are identically distributed. Now

$$\begin{aligned}
\Pr[S'_B = A_B] &= \sum_{a \in A \cap B} \Pr[S'_B = A_B \mid i' = a] \Pr[i' = a] \\
&= \sum_{a \in A \cap B} \Pr[S_B = (A - \{a\})_B \mid i' = a] \cdot \frac{1}{|B|} \\
&= \sum_{a \in A \cap B} \Pr[S_B = (A - \{a\})_B] \cdot \frac{1}{|B|} \\
&= \sum_{a \in A \cap B} \frac{1 - \varepsilon}{\varepsilon} \cdot \Pr[S_B = A_B] \cdot \frac{1}{|B|}.
\end{aligned}$$

By the conditioning on E, $\varepsilon|B|/2 \leq |A \cap B| \leq 2\varepsilon|B|$ and so the two probabilities are within a constant factor of one another.

By a similar argument, we should be able to show that $\Pr[M(x|_S) \in T \mid i \in S, E]$ and $\Pr[M(x|_S) \in T \mid E]$ also have the same order of magnitude. We define S' as before, but now i' is chosen at random from $B - S \cup \{i\}$. The calculation should be similar. □

Proof of Lemma 6. Let E be the event $|S \cap B| < k - 1$. By a multiplicative Chernoff bound, $\Pr[\overline{E}] = e^{-\Omega(k)}$. If E holds then M ε-blends x_i with fewer than k elements of $S \cap B$, so it must be that

$$e^{-\varepsilon} \Pr[M(x_{-i}|_S) \in T \text{ and } E] \leq \Pr[M(x|_S) \in T \text{ and } E] \leq e^{\varepsilon} \Pr[M(x_{-i}|_S) \in T \text{ and } E].$$

We then have

$$\begin{aligned}
\left|\Pr[M'(x) \in T] - \Pr[M'(x_{-i}) \in T]\right| &\leq \left|\Pr[M'(x) \in T \text{ and } E] - \Pr[M'(x_{-i}) \in T \text{ and } E]\right| + \Pr[\overline{E}] \\
&\leq (e^{\varepsilon} - 1) \Pr[M'(x_{-i}) \in T \text{ and } E] + \Pr[\overline{E}] \\
&\leq (e^{\varepsilon} - 1) \Pr[M'(x_{-i}) \in T] + \Pr[\overline{E}].
\end{aligned}$$

The other inequality is proved in an analogous way. □

3 Outlier privacy

Crowd-blending privacy formalizes the requirement that individuals who blend in the crowd do not require differential privacy. Another possibility is to tailor the privacy requirement to the individual: Those database rows that do not blend well — the outliers — may require more privacy than those that do.

Definition 7. Let ε be a function of k. A mechanism M is $(\varepsilon(k), \delta)$-outlier private if for every x, i, and T

$$\Pr[M(x) \in T] \leq e^{\varepsilon(k)} \Pr[M(x_{-i}) \in T] + \delta \quad \text{and} \quad \Pr[M(x_{-i}) \in T] \leq e^{\varepsilon(k)} \Pr[M(x) \in T] + \delta.$$

where k is the number of rows j such that M blends x_i and x_j.

If $\varepsilon(k)$ is upper bounded by ε for all k then this type of mechanism is in particular ε-differentially private.

We now give an outlier private version of the histogram mechanism. The idea is to add more noise to bins that contain fewer items. Let $\varepsilon > 0$ and t, K be integer parameters and set

$$\varepsilon(k) = \begin{cases} t\varepsilon/(K-k), & \text{if } k < K - t, \\ \varepsilon, & \text{if } k \geq K - t. \end{cases}$$

Here is a graph for $t = 1/\varepsilon$ with $\varepsilon = 0.2$ and $K = 15$.

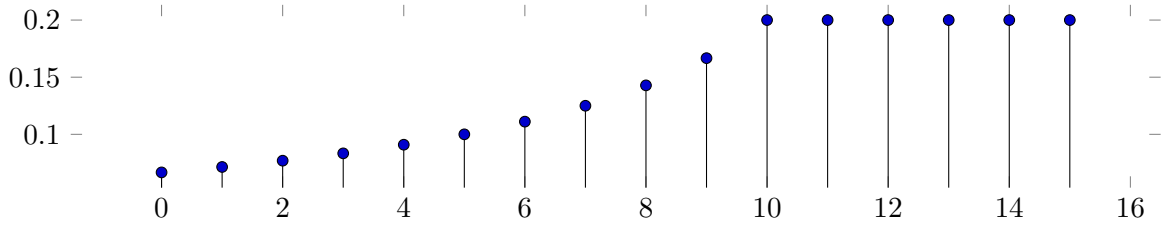

We will consider the following mechanism.

Mechanism $Out_h(x)$, where $x \in D^n$ and $h\colon D \to B$:
For all $b \in B$:
Let $k_b = |\{i\colon h(x_i) = b\}|$.
Output $k_b + N_b$, where $N_b \sim Lap(1/\varepsilon(k_b))$.

Theorem 8. *Mechanism Out_h is $(O(\varepsilon(k)), O(|B|e^{-\varepsilon t}))$-outlier private.*

Proof Sketch. Let E be the event that $N_b \leq (K - k_b)$ for all $b \in B$. By the large deviation bound for exponential random variables and the union bound, $\Pr[\overline{E}] = O(|B|e^{-\varepsilon t})$.

Now assume E holds and let $(y_b)_{b\in B}$ be an output of $Out_h(x)$. Removing the i-th row of x decreases y_b by one for $b = h(x_i)$ and does not affect the other outputs. Therefore

$$\frac{\Pr[Out(x) = y]}{\Pr[Out(x_{-i}) = y]} = \frac{\Pr_{N_b\sim\mathrm{Lap}(1/\varepsilon(k_b))}[N_b = y_b - k_b]}{\Pr_{N_b\sim\mathrm{Lap}(1/\varepsilon(k_b-1))}[N_b = y_b - (k_b-1)]} = \frac{e^{-\varepsilon(k_b)\cdot(y_b-k_b)}/Z_{k_b}}{e^{-\varepsilon(k_b-1)\cdot(y_b-k_b+1)}/Z_{k_b-1}}.$$

We will assume that $k_b \le K-t$ as the other case is straightforward. The ratio Z_{k_b-1}/Z_{k_b} is bounded by $e^{\pm O(\varepsilon(k_b-1))}$. Taking logarithms and absolute values we write

$$\begin{aligned}
\left|\ln\frac{\Pr[Out(x) = y]}{\Pr[Out(x_{-i}) = y]}\right| &\le |\varepsilon(k_b-1) - \varepsilon(k_b)|\cdot|y_b - k_b| + \varepsilon(k_b-1) + O(\varepsilon(k_b-1))\\
&\le |\varepsilon(k_b-1) - \varepsilon(k_b)|\cdot(K-k_b) + O(\varepsilon(k_b-1))\\
&= t\varepsilon\cdot\Big(\frac{1}{K-k_b} - \frac{1}{K-k_b+1}\Big)\cdot(K-k_b) + O(\varepsilon(k_b-1))\\
&= t\varepsilon\cdot\frac{1}{(K-k_b)(K-k_b+1)}\cdot(K-k_b) + O(\varepsilon(k_b-1))\\
&= O(\varepsilon(k_b-1))\\
&= O(\varepsilon(k_b)).
\end{aligned}$$

□

Regarding the utility of this mechanism, the standard deviation of the noise for a bin of size k is $(K-k)/t\varepsilon$. For $t = 1/\varepsilon$, a bin of size k will have noise on the order of $K-k$. In a typical run of the mechanism, the estimates for bins of size up to about K will be on the order of K; the amount of noise will then drop to $1/\varepsilon$ for the larger bins. In terms of information, the effect is comparable to that of suppressing the counts of the small bins.

References

Sections 1 and 2 are based on the work *Crowd-blending privacy* by Johannes Gehrke, Michael Hay, Edward Lui, and Rafael Pass.

The notion of outlier privacy in Section 3 is from the work *Outlier Privacy* by Edward Lui and Rafael Pass, but their definition is a bit different. Our outlier private histogram protocol and its analysis are also different from the ones presented in their work.