

Delta Calculation for Thresholding

This document contains technical statements (complete with proofs) of when thresholding can be used in our differential privacy library.

Thresholding For Histograms of Positive Values

We consider the case histograms are constructed from input data of the form (ID, partition, value). An example form of this setting is where each ID corresponds to a user and we wish to count the number of unique users in each partition. In the example case, value is always 1. However, in this section, we will assume something a bit more general: user contributions to each partition are positive and bounded. In this setting, we wish to keep the presence or absence of any particular ID in the database private.

More formally, we assume:

- Each ID is associated with at most B_0 partitions.
- The contribution of each ID to a single partition is non-negative and bounded by B_∞ .

We build a differentially private mechanism in two parts: $M := M_2 \circ M_1$. The mechanism M_1 adds noise *independently* to each coordinate in the support of the histogram with noise distribution given by a continuous cumulative density function Φ . The mechanism M_2 thresholds histogram entries with a fixed $\tau > 0$: for any histogram partition with value less than τ , the partition is removed entirely.

Fix two histograms H and H' that differ by one ID. We divide the histogram coordinates into two sets:

- U is the set of partitions that appear in one but not both of H and H' .
- \bar{U} is the set of partitions that appear in both histograms.

For any set of partitions X , we denote by H_X the restriction of the histogram H to the partitions in X . We assume throughout that $M_1(H)_U$ is (ϵ, δ_1) -indistinguishable from $M_1(H')_U$. That is, we assume for every (measurable) S a subset of the range of $M_1|_{\bar{U}}$,

$$P[M_1(H)_{\bar{U}} \in S] \leq e^\epsilon P[M_1(H')_{\bar{U}} \in S] + \delta_1 \quad \text{and} \quad P[M_1(H')_{\bar{U}} \in S] \leq e^\epsilon P[M_1(H)_{\bar{U}} \in S] + \delta_1 .$$

We first demonstrate that for sufficiently large τ , $M_2 \circ M_1$ leaks partitions of U with small δ_2 probability. We then show that this implies the overall mechanism M is $(\epsilon, \delta_1 + \delta_2)$ differentially private on the full histogram.

Throughout, we assume that $\epsilon > 0$, $\delta_1 \in (0, 1]$, and $\delta_2 \in (0, 1]$.

Lemma 1 If $\tau \geq B_\infty + \Phi^{-1}[(1 - \delta_2)^{1/B_0}]$, then $P[M(H)_U \neq 0 \text{ or } M(H')_U \neq 0] \leq \delta_2$.

Proof. Without loss of generality, we assume that H' contains one fewer ID than H . Thus, U consists of partitions in H not present in H' , and $M(H')_U = 0$ deterministically. It suffices to show $P[M(H)_U \neq 0] \leq \delta_2$.

$$\begin{aligned}
P[M(H)_U = 0] &= \prod_{i \in U} P[M(H)_i = 0] \\
&= \prod_{i \in U} (1 - P[M(H)_i \neq 0])
\end{aligned}$$

Abusing the notation, we denote by $M_1(x)$ the application of the noise distribution to a single scalar value x . Thus $M_1(H)_i$ and $M_1(H_i)$ have the same distribution.

Note that $M(H)_i \neq 0$ if and only if $M_1(H)_i \geq \tau$. Because $P[M_1(x) > \tau]$ is an increasing function with respect to x and $H_i \leq B_\infty$ for all $i \in U$, we see

$$\begin{aligned}
P[M(H)_U = 0] &\geq \prod_{i \in U} (1 - P[M_1(B_\infty) > \tau]) \\
&= \Phi(\tau - B_\infty)^{|U|}.
\end{aligned}$$

Since Φ is nondecreasing, we substitute $\tau \geq B_\infty + \Phi^{-1}[(1 - \delta_2)^{1/B_0}]$ to obtain

$$P[M(H)_U = 0] \geq (1 - \delta_2)^{|U|/B_0} \geq 1 - \delta_2.$$

Taking the complement yields $P[M(H)_U \neq 0] \leq \delta_2$. \square

Theorem 1. If $\tau \geq B_\infty + \Phi^{-1}[(1 - \delta_2)^{1/B_0}]$, then M is $(\varepsilon, \delta_1 + \delta_2)$ differential private.

Proof. Let S be a measurable subset of the range of M . We let E_0 denote the event that both $M(H)_U = 0$ and $M(H')_U = 0$. Partitioning the probability space, we see:

$$\begin{aligned}
P[M(H) \in S] &= P[M(H) \in S | E_0] P[E_0] + P[M(H) \in S | \bar{E}_0] P[\bar{E}_0] \\
&\leq P[M(H) \in S | E_0] P[E_0] + P[\bar{E}_0]
\end{aligned}$$

By Lemma 1, $P[\bar{E}_0] \leq \delta_2$. As such,

$$P[M(H) \in S] \leq P[M(H) \in S | E_0] P[E_0] + \delta_2 \tag{1}$$

Let $S_0 := \{h_{\bar{U}} \in S | h_U = 0\}$. Since M_1 (and hence M) is (ε, δ_1) -indistinguishable on \bar{U} , we see:

$$\begin{aligned}
P[M(H) \in S | E_0] &= P[M(H)_{\bar{U}} \in S_0 | E_0] \\
&\leq e^\varepsilon P[M(H')_{\bar{U}} \in S_0 | E_0] + \delta_1 \\
&= e^\varepsilon P[M(H') \in S | E_0] + \delta_1
\end{aligned}$$

Continuing from (1), we obtain:

$$\begin{aligned}
P[M(H) \in S] &\leq (e^\varepsilon P[M(H') \in S | E_0] + \delta_1) P[E_0] + \delta_2 \\
&\leq e^\varepsilon P[M(H') \in S \text{ and } E_0] + \delta_1 + \delta_2 \\
&\leq e^\varepsilon P[M(H') \in S] + \delta_1 + \delta_2 \quad \square
\end{aligned}$$