Differentially Private Set Union: Supplementary Materials

Sivakanth Gopi [1] **Pankaj Gulhane** [1] **Janardhan Kulkarni** [1] **Judy Hanwen Shen** [1 2] **Milad Shokouhi** [1]
Sergey Yekhanin [1]

A. Proofs of Policy Algorithms (Theorems 3.1 and 4.1)

Let $\mathcal{D}$ denote the collection of all databases. We say that $D, D' \in \mathcal{D}$ are neighboring databases, denoted by $D \sim D'$, if they differ in exactly one user.

Definition A.1. For $p \geq 0$, the ℓ_p-sensitivity of $f : \mathcal{D} \to \mathbb{R}^k$ is defined as $\sup_{D \sim D'} \|f(D) - f(D')\|_{\ell_p}$ where the supremum is over all neighboring databases D, D'.

The noise that we add is sampled either from Laplace or Gaussian (Normal) distribution. The probability density functions of these distributions are given by:

$$\mathsf{Lap}(\mu, \lambda) = \mathsf{Laplace}(\mu, \lambda) = \frac{1}{2\lambda} \exp\left(-\frac{|x-\mu|}{\lambda}\right)$$

$$\mathcal{N}(\mu, \sigma^2) = \mathsf{Normal}(\mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

We will need the following standard DP mechanisms.

Proposition A.1 (The Laplace Mechanism (Dwork et al., 2006))**.** Suppose $f : \mathcal{D} \to \mathbb{R}^k$ is a function with ℓ_1 sensitivity Δ_1. For any $\varepsilon \geq 0$, the Laplace mechanism $M(x) = f(x) + (Y_1, Y_2, \ldots, Y_k)$ is $(\varepsilon, 0)$-DP when $Y_1, Y_2, \ldots, Y_k$ are i.i.d. random variables drawn from $\mathsf{Lap}(0, \Delta_1/\varepsilon)$.

Proposition A.2 (Gaussian Mechanism (Balle & Wang, 2018))**.** Suppose $f : \mathcal{D} \to \mathbb{R}^d$ is a function with ℓ_2-sensitivity Δ_2. For any $\varepsilon \geq 0$ and $\delta \in [0, 1]$, the Gaussian mechanism $M(x) = f(x) + Z$ with $Z \sim \mathcal{N}(0, \sigma^2 I)$ is (ε, δ)-DP if and only if

$$\Phi\left(\frac{\Delta_2}{2\sigma} - \frac{\varepsilon\sigma}{\Delta_2}\right) - e^{\varepsilon}\Phi\left(-\frac{\Delta_2}{2\sigma} - \frac{\varepsilon\sigma}{\Delta_2}\right) \leq \delta$$

where Φ is the cdf of standard normal distribution.

*Equal contribution [1]Microsoft [2]Work done as part of the Microsoft AI Residency Program. Correspondence to: Sivakanth Gopi <sigopi@microsoft.com>, Janardhan Kulkarni <jakul@microsoft.com>, Judy Hanwen Shen <hashe@microsoft.com>.

Proceedings of the 37^{th} International Conference on Machine Learning, Vienna, Austria, PMLR 119, 2020.

Definition A.2. We say that two distributions P, Q on a domain X are (ε, δ)-close to each other, denoted by $P \approx_{\varepsilon,\delta} Q$, if for every $S \subset X$, we have

1. $\Pr_{x\sim P}[x \in S] \leq e^{\varepsilon} \Pr_{x\sim Q}[x \in S] + \delta$ and
2. $\Pr_{x\sim Q}[x \in S] \leq e^{\varepsilon} \Pr_{x\sim P}[x \in S] + \delta$.

We say that two random variables X, Y are (ε, δ)-close to each other, denoted by $X \approx_{\varepsilon,\delta} Y$, if their distributions are (ε, δ)-close to each other.

We will need the following lemmas which is useful to prove (ε, δ)-DP.

Lemma A.1. Let P, Q be probability distributions over a domain X. If there exists an event E s.t. $P(E) = 1 - \delta'$ and $P|_E \approx_{\varepsilon,\delta} Q$, then $P \approx_{\varepsilon,\delta+\delta'} Q$.

Proof. Fix some subset $S \subseteq X$.

$$\begin{aligned}
\Pr_{x\sim P}[x \in S] &= P(\bar{E}) \Pr_{x\sim P}[x \in S|\bar{E}] + P[E] \Pr_{x\sim P}[x \in S|E] \\
&\leq P(\bar{E}) + \Pr_{x\sim P}[x \in S|E] \\
&= \delta' + \Pr_{x\sim P|_E}[x \in S] \\
&\leq \delta' + e^{\varepsilon} \Pr_{x\sim Q}[x \in S] + \delta
\end{aligned}$$

We now prove the other direction.

$$\begin{aligned}
\Pr_{x\sim Q}[x \in S] &\leq e^{\varepsilon} \Pr_{x\sim P|_E}[x \in S] + \delta \\
&\leq e^{\varepsilon} \frac{\Pr_{x\sim P}[x \in S]}{P(E)} + \delta \\
&= e^{\varepsilon} \frac{\Pr_{x\sim P}[x \in S]}{1-\delta'} + \delta \\
&= e^{\varepsilon} \Pr_{x\sim P}[x \in S] + \delta' \left(\frac{e^{\varepsilon} \Pr_{x\sim P}[x \in S]}{1-\delta'}\right) + \delta
\end{aligned}$$

Now if $e^{\varepsilon} \Pr_{x\sim P}[x \in S] \leq 1 - \delta'$, then we have $Pr_{x\sim Q}[x \in S] \leq e^{\varepsilon} \Pr_{x\sim P}[x \in S] + \delta' + \delta$. Otherwise, trivially

$$\Pr_{x\sim Q}[x \in S] \leq 1 \leq e^{\varepsilon} \Pr_{x\sim P}[x \in S] + \delta' + \delta.$$

□

We will also need the fact that if $X \approx_{\varepsilon,\delta} Y$, then after postprocessing they also remain (ε, δ)-close.

Lemma A.2. If two random variables X, Y are (ε, δ)-close and M is any randomized algorithm, then $M(X) \approx_{\varepsilon,\delta} M(Y)$.

Proof. Let $M(z) = F(z, R)$ for some function F where R is the random bits used by M. For any subset S of the possible outputs of M,

$$\begin{aligned}
\Pr[M(X) \in S] &= \Pr_{X,R}[F(X,R) \in S] \\
&= \sum_r \Pr[R = r] \Pr_X[F(X,r) \in S] \\
&\le \sum_r \Pr[R = r] \left(e^\varepsilon \Pr_Y[F(Y,r) \in S] + \delta \right) \\
&= e^\varepsilon \sum_r \Pr[R = r] \Pr_Y[F(Y,r) \in S] + \delta \\
&= e^\varepsilon \Pr_{X,R}[F(Y,R) \in S] + \delta.
\end{aligned}$$

The other direction holds by symmetry. □

Proof of Theorem 3.1. Suppose D_1 and D_2 are neighboring databases where D_1 has one extra user compared to D_2. Let P and Q denote the distribution of output of the algorithm when the database is D_1 and D_2 respectively. We want to show that $P \approx_{\varepsilon,\delta} Q$. Let E be the event that $A \subset \mathrm{supp}(H_2)$.

Claim A.1. $P|_E \approx_{\varepsilon,0} Q$

Proof. Let H_1 and H_2 be the histograms generated by the algorithm from databases D_1 and D_2 respectively. And $\hat{H}_1$ and $\hat{H}_2$ be the histograms obtained by adding $\mathsf{Lap}(0, 1/\varepsilon)$ noise to each entry of H_1 and H_2 respectively. For any possible output A of Algorithm 4, we have

$$Q(A) = \Pr[A = \{u \in \mathrm{supp}(H_2) : \hat{H}_2[u] > \rho_{\mathsf{Lap}}\}] \text{ and}$$
$$P|_E(A) = \Pr[A = \{u \in \mathrm{supp}(H_2) : \hat{H}_1[u] > \rho_{\mathsf{Lap}}\}].$$

So $A \sim P|_E$ is obtained by postprocessing $\hat{H}_1|_E$ and $A \sim Q$ is obtained by postprocessing $\hat{H}_2$. Since postprocessing only makes two distributions closer (Lemma A.2), it is enough to show that the distributions of the $\hat{H}_1|_{\mathrm{supp}(H_2)}$ and $\hat{H}_2$ are $(\varepsilon, 0)$-close to each other. Because the histogram building algorithm (Algorithm 2) has ℓ_1-sensitivity of at most 1 by hypothesis, $\left\|H_1|_{\mathrm{supp}(H_2)} - H_2\right\|_{\ell_1} \le 1$. Therefore $P|_E \approx_{\varepsilon,0} Q$ by the properties of Laplace mechanism (Proposition A.1). □

By Lemma A.1, it is enough to show that $P(E) \ge 1 - \delta$. Let $T = \mathrm{supp}(H_1) \setminus \mathrm{supp}(H_2)$. Note that $|T| \le \Delta_0$ and $H_1[u] \le \frac{1}{|T|}$ for $u \in T$.

$$\begin{aligned}
P(\bar{E}) &= \Pr[\exists u \in T \mid \hat{H}_1[u] > \rho_{\mathsf{Lap}}] \\
&= 1 - \Pr[\forall u \in T \;\; \hat{H}_1[u] \le \rho_{\mathsf{Lap}}] \\
&= 1 - \prod_{u \in T} \Pr[H_1[u] + X_u \le \rho_{\mathsf{Lap}}] \\
&\qquad\qquad (X_u \sim \mathsf{Lap}(1/\varepsilon)) \\
&\le 1 - \prod_{u \in T} \Pr\left[X_u \le \rho_{\mathsf{Lap}} - \frac{1}{|T|} \right] \\
&\qquad\qquad (H_1[u] \le \tfrac{1}{|T|} \text{ for } u \in T) \\
&= 1 - \left(1 - \frac{1}{2} \exp\left(-\varepsilon \rho_{\mathsf{Lap}} + \varepsilon \frac{1}{|T|} \right) \right)^{|T|} \qquad (1)
\end{aligned}$$

Thus for

$$\rho_{\mathsf{Lap}} \ge \max_{1 \le t \le \Delta_0} \frac{1}{t} + \frac{1}{\varepsilon} \log\left(\frac{1}{2\left(1 - (1-\delta)^{1/t}\right)} \right),$$

we have $P(\bar{E}) \le \delta$. Therefore the DP Set Union algorithm (Algorithm 1) is (ε, δ)-DP. □

Proof of Theorem 4.1. Suppose D_1 and D_2 are neighboring databases where D_1 has one extra user compared to D_2. Let P and Q denote the distribution of output of the algorithm when the database is D_1 and D_2 respectively. We want to show that $P \approx_{\varepsilon,\delta} Q$. Let E be the event that $A \subset \mathrm{supp}(H_2)$.

Claim A.2. $P|_E \approx_{\varepsilon,\delta/2} Q$

Proof. Let H_1 and H_2 be the histograms generated by the algorithm from databases D_1 and D_2 respectively. And $\hat{H}_1$ and $\hat{H}_2$ be the histograms obtained by adding $\mathcal{N}(0, \sigma^2)$ noise to each entry of H_1 and H_2 respectively. By the postprocessing lemma (Lemma A.2), it is enough to show that the distributions of the $\hat{H}_1|_{\mathrm{supp}(H_2)}$ and $\hat{H}_2$ are $(\varepsilon, \delta/2)$-close to each other. Because the histogram building algorithm (Algorithm 2) has ℓ_2-sensitivity of at most 1 by hypothesis, $\left\|H_1|_{\mathrm{supp}(H_2)} - H_2\right\|_{\ell_2} \le 1$. Therefore by properties of Gaussian mechanism (Proposition A.2), it is enough to choose σ as in the statement of the theorem. □

By Lemma A.1, it is enough to show that $P(E) \ge 1 - \delta/2$. Let $T = \mathrm{supp}(H_1) \setminus \mathrm{supp}(H_2)$. Note that $|T| \le \Delta_0$ and $H_1[u] \le \frac{1}{\sqrt{|T|}}$ for $u \in T$.

$$
\begin{aligned}
P(\bar{E}) &= \Pr[\exists u \in T \mid \hat{H}_1[u] > \rho_{\mathsf{Gauss}}] \\
&= 1 - \Pr[\forall u \in T \;\; \hat{H}_1[u] \le \rho_{\mathsf{Gauss}}] \\
&= 1 - \prod_{u \in T} \Pr[\hat{H}_1[u] \le \rho_{\mathsf{Gauss}}] \\
&= 1 - \prod_{u \in T} \Pr[H_1[u] + X_u \le \rho_{\mathsf{Gauss}}] \\
&\qquad\qquad (X_u \sim \mathcal{N}(0, \sigma^2)) \\
&\le 1 - \prod_{u \in T} \Pr\left[X_u \le \rho_{\mathsf{Gauss}} - \frac{1}{\sqrt{|T|}}\right] \\
&\qquad\qquad \left(H_1[u] \le \tfrac{1}{\sqrt{|T|}} \text{ for } u \in T\right) \\
&= 1 - \prod_{u \in T} \Phi\left(\frac{\rho_{\mathsf{Gauss}}}{\sigma} - \frac{1}{\sqrt{|T|}}\right)^{|T|} \qquad (2)
\end{aligned}
$$

Thus for

$$
\rho_{\mathsf{Gauss}} \ge \max_{1 \le t \le \Delta_0} \left(\frac{1}{\sqrt{t}} + \sigma \Phi^{-1}\left(\left(1 - \frac{\delta}{2}\right)^{1/t}\right)\right),
$$

we have $P(\bar{E}) \le \delta/2$. Therefore the DP Set Union algorithm (Algorithm 1) is (ε, δ)-DP. □

B. Bounded Sensitivity implies DP (Proof of Theorem 1.2)

We will now prove a formal version of Theorem 1.2, i.e., if the histogram output by Algorithm 2 has bounded ℓ_p-sensitivity (for $p \in \{1, 2\}$), then by adding appropriate noise and setting an appropriate threshold, Algorithm 1 for DP set union can be made differentially private. The lower bounds on the threshold (ρ) that we obtain in this generality are only slightly worse compared to the corresponding bounds in Theorems 3.1 and 4.1.

Theorem B.1. Suppose the histogram output by Algorithm 2 has ℓ_1-sensitivity 1. Then Algorithm 1 is (ε, δ)-DP when the Noise distribution is $\mathsf{Lap}(0, \lambda)$ where $\lambda = 1/\varepsilon$ and the threshold

$$
\rho \ge \max_{1 \le t \le \Delta_0} 1 + \frac{1}{\varepsilon} \log\left(\frac{1}{2\left(1 - (1-\delta)^{1/t}\right)}\right).
$$

Proof. Proof of Theorem B.1 is extremely similar to the proof of Theorem 3.1. The only place where it differes is in Equation (1) where we bound $H_1[u] \le 1$ instead of $H_1[u] \le 1/|T|$. □

Theorem B.2. Suppose the histogram output by Algorithm 2 has ℓ_2-sensitivity 1. Then Algorithm 1 is (ε, δ)-DP when the Noise distribution is $\mathcal{N}(0, \sigma^2)$ where σ and the threshold ρ are chosen s.t.

$$
\Phi\left(\frac{1}{2\sigma} - \varepsilon\sigma\right) - e^{\varepsilon}\Phi\left(-\frac{1}{2\sigma} - \varepsilon\sigma\right) \le \frac{\delta}{2} \text{ and}
$$

$$
\rho \ge \max_{1 \le t \le \Delta_0} \left(1 + \sigma\Phi^{-1}\left(\left(1 - \frac{\delta}{2}\right)^{1/t}\right)\right).
$$

Proof. Proof of Theorem B.2 is extremely similar to the proof of Theorem 4.1. The only place where it differes is in Equation (2) where we bound $H_1[u] \le 1$ instead of $H_1[u] \le 1/\sqrt{|T|}$. □

C. Proof of Lemma 4.1

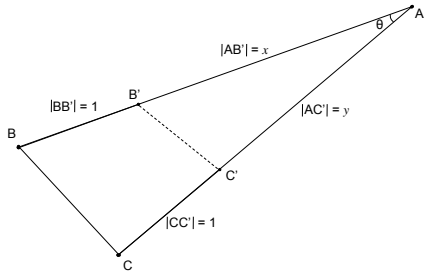

Figure 1. Geometric explanation of Lemma 4.1 when $|AB|, |AC| > 1$.

Proof of Lemma 4.1. Let us first assume that both $|AB|, |AC| > 1$. Let θ be the angle at A and let $|AB'| = x, |AC'| = y$ as shown in Figure 1. Then by the cosine formula,

$$
\begin{aligned}
|BC|^2 &= |AB|^2 + |AC|^2 - 2|AB||AC|\cos\theta \\
&= (x+1)^2 + (y+1)^2 - 2(x+1)(y+1)\cos\theta \\
&= x^2 + y^2 + 2xy\cos\theta + 2(x+y+1)(1-\cos\theta) \\
&\ge x^2 + y^2 + 2xy\cos\theta \qquad (\cos\theta \le 1) \\
&= |B'C'|^2.
\end{aligned}
$$

If $|AB|, |AC| \le 1$, then $B' = C' = A$ and then the

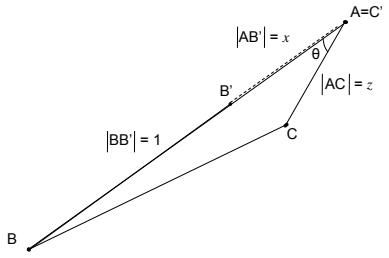

Figure 2. Geometric explanation of Lemma 4.1 when $|AB| > 1, |AC| \le 1$.

claim is trivially true. Suppose $|AB| > 1, |AC| \leq 1$. Now $C' = A$. Let $|AB'| = x, |AC| = z \leq 1$ and θ be the angle at A as shown in Figure 2. Then by the cosine formula,

$$\begin{aligned}|BC|^2 &= |AB|^2 + |AC|^2 - 2|AB||AC|\cos\theta \\ &= (x+1)^2 + z^2 - 2(x+1)z\cos\theta \\ &= x^2 + 2x(1 - z\cos\theta) + (z - \cos\theta)^2 + (1 - \cos^2\theta) \\ &\geq x^2 = |AB'|^2 = |B'C'|^2. \\ &\qquad\qquad (0 \leq z \leq 1, |\cos\theta| \leq 1)\end{aligned}$$

By symmetry, the claim is also true when $|AC| > 1, |AB| \leq 1$. □

D. Privacy analysis of Weighted Laplace and Guassian Algorithms

D.1. Weighted Laplace

Algorithm 1 LAPLACE weighted update

Input: H: Current histogram
W: A subset of U of size at most Δ_0
Output: H: Updated histogram
for u in W **do**
  H[u] ← H[u] + $\frac{1}{|W|}$
end for

Theorem D.1. The WEIGHTED LAPLACE algorithm (Algorithm 1) is (ε, δ)-DP when

$$\rho_{\mathsf{Lap}} \geq \max_{1 \leq t \leq \Delta_0} \frac{1}{t} + \frac{1}{\varepsilon}\log\left(\frac{1}{2\left(1 - (1-\delta)^{1/t}\right)}\right).$$

Proof. Proof is exactly the same as that of Theorem 3.1. □

D.2. Weighted Gaussian

Algorithm 2 GAUSSIAN weighted update

Input: H: Current histogram
W: A subset of U of size at most Δ_0
Output: H: Updated histogram
for u in W **do**
  H[u] ← H[u] + $\sqrt{\frac{1}{|W|}}$
end for

Theorem D.2. The WEIGHTED GAUSSIAN algorithm (Algorithm 2) is (ε, δ)-DP if $\sigma, \rho_{\mathsf{Gauss}}$ are chosen s.t.

$$\Phi\left(\frac{1}{2\sigma} - \varepsilon\sigma\right) - e^{\varepsilon}\Phi\left(-\frac{1}{2\sigma} - \varepsilon\sigma\right) \leq \frac{\delta}{2} \text{ and}$$

$$\rho_{\mathsf{Gauss}} \geq \max_{1 \leq t \leq \Delta_0}\left(\frac{1}{\sqrt{t}} + \sigma\Phi^{-1}\left(\left(1 - \frac{\delta}{2}\right)^{1/t}\right)\right).$$

Proof. Proof is exactly the same as that of Theorem 4.1. □

E. Greedy Policy

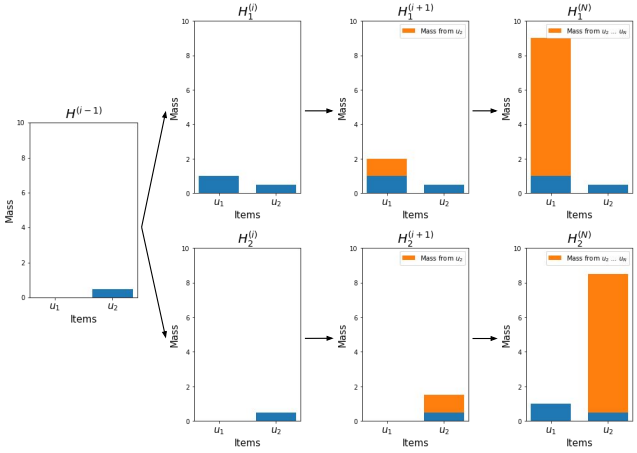

Figure 3. Visualization of greedy update example where the final l_1 sensitivity is larger than 1.

In this section, we give a simple counter example to illustrate how the sensitivity of a greedy policy algorithm can be unbounded.

Algorithm 3 GREEDY POLICY update

Input: H: Current histogram
W: A subset of U of size at most Δ_0
Γ: cutoff parameter
Output: H: Updated histogram
// Build cost dictionary G
$G = \{\}$ // Empty dictionary
for $u \in W$ **do**
  if $H[u] < \Gamma$ **then**
    // Gap to cutoff for items below cutoff Γ
    $G[u] \leftarrow \Gamma - H[u]$
  end if
end for
budget ← 1 // Each user gets a total budget of 1
// Sort in increasing order of the gap $\Gamma - H[u]$
$G \leftarrow \text{sort}(G)$
// Let $u_1, u_2, \ldots, u_{|G|}$ be the sorted order
for $j = 1$ to $|G|$ **do**
  if $G[u_j] \leq$ budget **then**
    $H[u_j] \leftarrow H[u_j] + G[u_j]$
    budget ← budget - $G[u_j]$
  else
    $H[u_j] \leftarrow H[u_j]$+ budget
    break
  end if
end for

Suppose there are N user let u_1 and u_2 be two items in the universe. We will denote the weight of item u after user

i's contribution as $H^{(i)}[u]$. Suppose user i has only item u_1 while users $i+1, i+2, \ldots, N$ have both items. Let H_1 be the histogram generated with all N users while H_2 be the histogram generated without user i. Let $\Delta_0 = 2$ and $H^{(i-1)}[u_1] < H^{(i-1)}[u_2] < 1 + H^{(i-1)}[u_1]$. According to the greedy update described in Algorithm 3, in H_1, user i will add weight 1 to u_1 and users $i+1, i+2, \ldots, N$ will also to u_1 since $H^{(i)}[u_1] > H^{(i)}[u_2]$. In H_2, users $i+1, i+2, \ldots, N$ will add to u_2 since $H^{(i-1)}[u_1] < H^{(i-1)}[u_2]$. This process is described in figure 3. Therefore the ℓ_1-sensitivity of the histogram built using Greedy Policy update (Algorithm 3) can be $\Omega(\Gamma, N)$.

F. Dataset Details

Using a log-log scale, the frequency of users for each unigram vs. the rank of the unigram is linear (Figure 4). In other words, the lowest ranked (most common) unigrams are used by almost all users while the highest ranked (least common) unigrams are used by very few users.

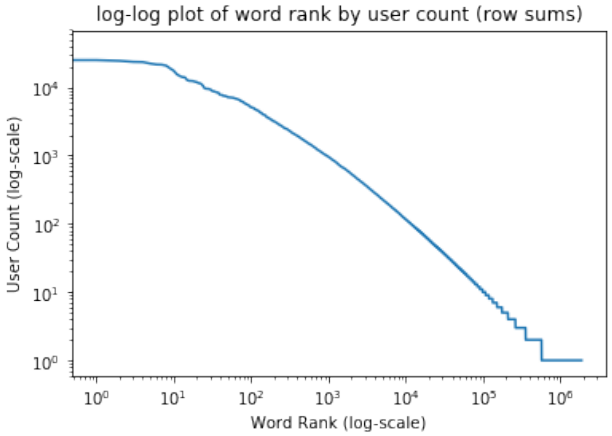

Figure 4. Frequency (i.e. number of users who use the unigram) vs. rank of the unigram (based on frequency) on a log-log scale. This linear relationship shows that the frequency of unigrams among users also follows Zipf's law (power law), i.e., $\text{count} \propto 1/\text{rank}^\alpha$ for some constant $\alpha > 0$. The α in this case is ≈ 1.

The distribution of how many unigrams each user uses also follows a long tail distribution. While the top 10 users contribute between 850 and 2000 unique unigrams, most users (93.1%) contribute less than 100 unique unigrams. Table 1 summarizes the percentage of users with a unique vocabulary smaller than each threshold T provided.

Table 1. Percentage of users with unique unigram count of less than or equal to T. The vast majority of user have less than 100 unique unigrams.

THRESHOLD (T)	USERS WITH $\lvert W_i \rvert \leq$ T
1	2.78%
10	29.82%
50	79.16%
100	93.13%
300	99.59%

G. Additional Experiments

G.1. Multiple passes through each user

In the experiments described thus far, each user contributes items once within the budget constraints. We also investigate whether the output of set union increases in size when each user contributes the same budget over multiple passes (e.g. user 1 contributes half of their budget each time over 2 passes), we compare POLICY LAPLACE and POLICY GAUSSIAN outputs. Table 2 summarizes the results showing that there is not strong evidence suggesting that running multiple passes through the users improves the size of the output set.

G.2. Selecting α: parameter to set threshold Γ

Figure 5 shows the number of unigrams released by POLICY LAPLACE and POLICY GAUSSIAN for various values of α. We observe that the number of unigrams released increases sharply until $\alpha = 4$, then remains nearly constant and then slowly decreases. This choice of α only affects the policy algorithms since the weighted and count algorithms do not use a threshold.

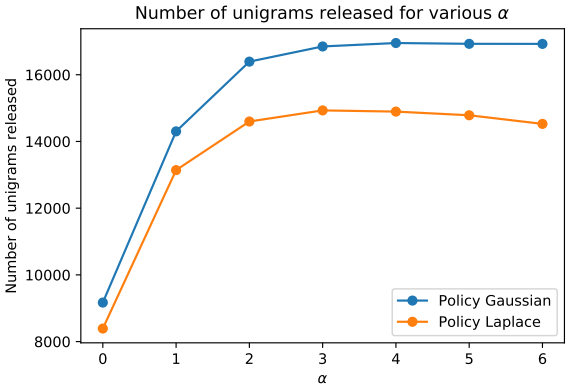

Figure 5. Number of unigrams released for various values of α. The number of unigrams released increases sharply until about $\alpha = 2$, then remains nearly constant and then decreases. Here we fixed $\Delta_0 = 100$ and $\varepsilon = 3$.

Table 2. Count of unigrams released POLICY LAPLACE and POLICY GAUSSIAN algorithms for single and double passes over users. Results are averaged and rounded across 5 shuffles of user order. The privacy parameters are $\varepsilon = 3$ and $\delta = \exp(-10)$. $\alpha = 2$ is chosen for the threshold parameter. Significant p-values for a two-sided independent t-test are bolded.

	POLICY LAPLACE			POLICY GAUSSIAN		
Δ_0	1 PASS	2 PASSES	P-VAL	1 PASS	2 PASSES	P-VAL
1	4236 ± 14	4257 ± 17	0.083	3135 ± 25	3131 ± 20	0.829
10	$\mathbf{12452 \pm 31}$	12389 ± 17	**0.008**	10784 ± 22	10817 ± 54	0.293
50	15056 ± 35	15080 ± 21	0.262	15763 ± 33	15809 ± 45	0.139
100	14562 ± 50	14567 ± 24	0.846	14562 ± 50	14568 ± 24	0.846
200	14005 ± 33	13979 ± 31	0.271	14005 ± 33	13979 ± 31	0.271
300	13702 ± 37	13678 ± 47	0.448	13702 ± 37	13678 ± 47	0.447

THE EFFECT OF ε

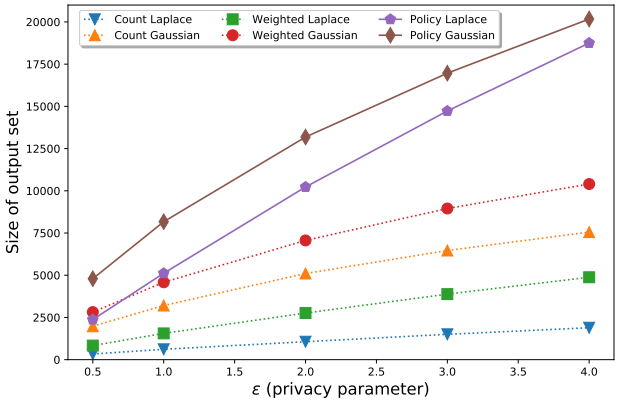

Figure 6. Number of unigrams released for various values of ε. Here we fixed $\Delta_0 = 100$ and $\alpha = 5$.

We use $\varepsilon = 3$ for the experiments in table 1. At this value of ε our policy algorithms perform much better than previous count and weighted algorithms. To check whether this result holds with smaller ε, we also run these algorithms on various values of ε. Figure 6 shows that for $\varepsilon \geq 1$ our policy algorithms always perform better.

References

Balle, B. and Wang, Y.-X. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*, pp. 403–412, 2018.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pp. 265–284. Springer, 2006.