Locally Private Mean Estimation: Z-test and Tight Confidence Intervals

Marco Gaboardi, Ryan Rogers, Or Sheffet

Mean estimation (ME)

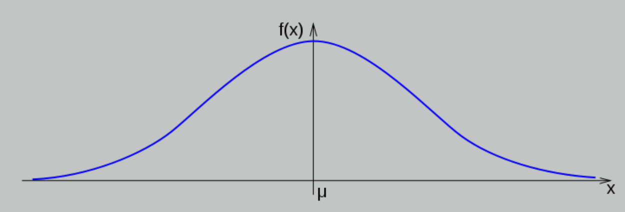

- Setting: We have n samples drawn from a Gaussian
$$X_1, \ldots, X_n \sim_{\text{i.i.d}} \mathcal{N}(\mu, \sigma^2)$$
such that $\mu \in [-R, R]$ for some known bound R, and σ is either provided as an input (known variance case) or left unspecified (unknown variance case).
- **Goal**: Determine an estimate of μ useful for Z-test and for releasing confidence intervals:
$$\underset{\mathbf{X} \overset{i.i.d.}{\sim} \mathcal{N}(\mu,\sigma^2), \mathcal{M}(\mathbf{X})}{\mathbb{P}} [\mu \in \mathcal{M}(\mathbf{X})] \geq 1 - \beta$$

The Need for Privacy

- Data may contain sensitive information.
- Releasing the result may leak information

Modified Goal: Determine an estimate of μ which preserves the privacy of those in the study and that is useful for Z-test and for releasing confidence interval.

Local Differential Privacy [4] (LDP)

- Central Model: Data is submitted in the clear to a trusted curator and the output of a statistic on the data is privatized.
- Local Model: No trusted curator - data is privatized and then collected.
- An algorithm $M : \mathcal{X} \to \mathcal{O}$ is ϵ-differentially private if for all inputs, x, x' and outcome sets $S \subseteq \mathcal{O}$:
$$\mathbb{P}[M(x) \in S] \leq e^{\epsilon} \mathbb{P}[M(x') \in S].$$
- Local model of differential privacy is used in practice.

LDP randomizers properties

We use the following mechanisms

- **Gaussian Noise** [2]: Suppose each datum is sampled from an interval I of length ℓ. Then we add independent noise
$$\mathcal{N}(0, 2\ell^2 \ln(2/\delta)/\epsilon^2)$$
to each datum guaranteeing (ϵ, δ)-differential privacy.
- **Randomized Response** [5]: Suppose each datum is a bit $\{0, 1\}$ and on each datum we operate independently, applying $\mathrm{RR}_\epsilon : \{0,1\} \to \{0,1\}$ where
$$\mathrm{RR}_\epsilon(b) = \begin{cases} b & \text{w.p. } \frac{e^\epsilon}{1+e^\epsilon} \\ 1-b & \text{else} \end{cases}$$
- **Bit Flipping algorithm** [1]: Suppose each datum x_i is a d-dimensional vector indicating its type using a standard basis vector. The Bit Flipping mechanism now runs d independent randomized response mechanism for each coordinate separately with parameter $\epsilon/2$:
$$\mathrm{BF}(x_i^1, \ldots, x_i^d) = (\mathrm{RR}_{\epsilon/2}(x_i^1), \ldots, \mathrm{RR}_{\epsilon/2}(x_i^d))$$

LDP ME - Known Variance

Our approach is inspired by the work of Karwa and Vadhan [3]. We adapt it to the local model.

KnownVar $(\mathbf{X}; \sigma, \beta, \epsilon, n, R)$ (sketch)

1. Find a bin of length σ most likely to hold μ
2. Construct an interval of length $4\sigma + 2\sigma\sqrt{2\log(8n/\beta)}$ centered at this bin
3. Project all remaining points onto this interval and add ind. Gaussian noise.

KnownVar properties

- **Privacy**: **KnownVar** is (ϵ, δ)-LDP.
- **Confidence Interval**: If $n \geq 1600 \left(\frac{e^{\epsilon/2}+1}{e^{\epsilon/2}-1}\right)^2 \log\left(\frac{8d}{\beta}\right)$, then **KnownVar** returns an interval I such that: $\underset{\mathbf{X}, \mathbf{KnownVar}}{\mathbb{P}}[\,mu \in I] \geq 1 - \beta$. whose size is:
$$|I| = O\left(\sigma \cdot \frac{\sqrt{\log(n/\beta) \cdot \log(1/\beta) \cdot \log(1/\delta)}}{\epsilon\sqrt{n}}\right)$$

Locally Private Z-test

- For any interval on the reals I we can associate a likelihood of $p_I \overset{\text{def}}{=} \underset{X \sim \mathcal{P}}{\mathbb{P}}[X \in I]$, and we know that w.p. $p_I \pm \beta$ it indeed holds that $\mu \in I$.
- This mimics the power of a Z-test — in particular we can now compare two intervals as to which one is more likely to hold μ, compare populations, etc.
- Below are results showing the empirical p-values and power averaged over 100 trials for various privacy parameters.

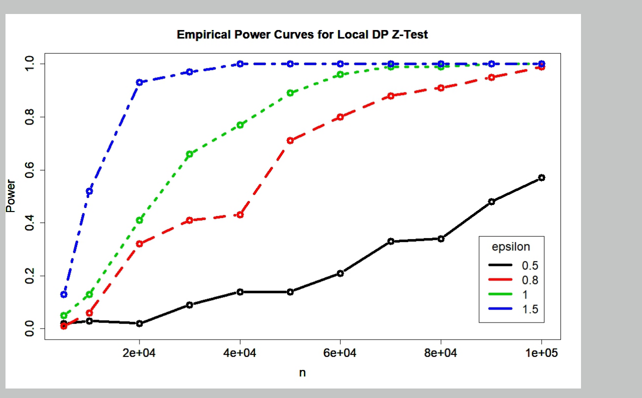

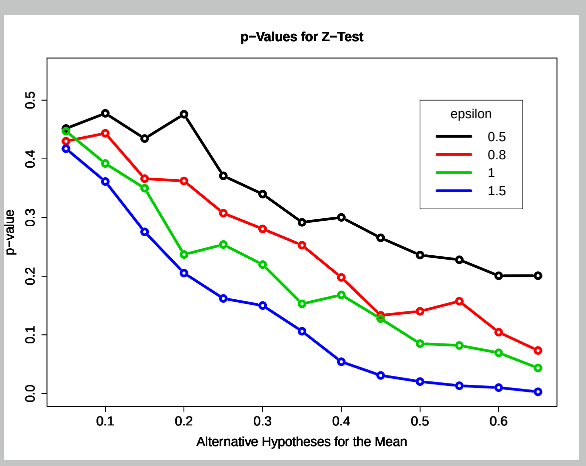

Lower Bounds

- **Main Lemma:** Let $\mathcal{M}$ be a one-shot (each individual is presented with a single query) local ϵ-differentially private mechanism. Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions, with $\Delta \overset{\text{def}}{=} d_{\mathrm{TV}}(\mathcal{P}, \mathcal{Q})$. Fix any $0 < \delta < e^{-1}$ and set $\epsilon^* = 8\epsilon\Delta\sqrt{n}\left(\sqrt{\frac{1}{2}\ln(2/\delta)} + 16\epsilon\Delta\sqrt{n}\right)$. Then, for any set S of outputs,
$$\underset{\mathbf{X} \overset{\text{i.i.d}}{\sim} \mathcal{P};\, \mathcal{M}}{\Pr}[\mathcal{M}(\mathbf{X}) \in S] \leq e^{\epsilon^*} \underset{\mathbf{X} \overset{\text{i.i.d}}{\sim} \mathcal{Q};\, \mathcal{M}}{\Pr}[\mathcal{M}(\mathbf{X}) \in S] + \delta$$
- **Lower bound:** Any one-shot local differentially private algorithm must return an interval of length
$$\Omega\left(\frac{\sigma\sqrt{\log(1/\beta)}}{\epsilon\sqrt{n}}\right)$$
- **Lower bound:** Let $\mathcal{M}$ be a ϵ-LDP mechanism which is $(\alpha_{\text{dist}}, \alpha_{\text{quant}}, \beta)$-useful for the p-quantile problem over $\mathcal{P}$, given that the true p-quantile lies in the interval $[-R, R]$. Then, for any $\beta < \frac{1}{6}$ it must hold that $n \geq \Omega(\frac{1}{\alpha_{\text{quant}}^2 \epsilon^2} \cdot \ln(\frac{R}{\alpha_{\text{dist}}\beta}))$.

LDP ME - Unknown Variance

Our approach mimics the same approach from Algorithm **KnownVar** but without the knowledge of the variance.

- Goal: Find a suitably large yet sufficiently tight interval $[s_1, s_2]$.
- Problem: This cannot be done using the off-the-shelf Bit Flipping mechanism as that required we know the granularity of each bin in advance.
- Solution: We abandon the idea of finding a histogram on the data. Instead, we propose finding a good approximation for σ using a **quantile estimation** based on a binary search, using the following algorithm.

Algorithm BinQuant

Require: Data $\{x_1, \cdots, x_N\}$, target quantile p^*; ϵ, $[Q_{\min}, Q_{\max}]$, λ, T.

Initialize $j = 0$, $n = N/T$, $s_1 = Q_{\min}$, $s_2 = Q_{\max}$.

for $j = 1, \cdots, T$ **do**

  Select users $\mathcal{U}^{(j)} = \{j \cdot n + 1, j \cdot n + 2, \cdots, (j+1) \cdot n\}$

  Set $t^{(j)} \leftarrow \frac{s_1+s_2}{2}$

  Denote $\phi^{(j)}(x) = \mathbb{1}\{x < t^{(j)}\}$.

  Run randomized response on $\mathcal{U}^{(j)}$ and obtain
$$Z^{(j)} = \tfrac{1}{n}\hat{\theta}_{\mathrm{RR}}(n, \phi^{(j)}).$$

  if $(Z^{(j)} > p^* + \frac{\lambda}{2})$ **then**

    $s_2 \leftarrow t^{(j)}$

  else if $(Z^{(j)} < p^* - \frac{\lambda}{2})$ **then**

    $s_1 \leftarrow t^{(j)}$

  else

    break

Ensure: $t^{(j)}$

Our algorithm **UnkVar** uses the quantile estimation *twice*: once for $p^* = \frac{1}{2}$ where $t^* = \mu$, and once for the value of $p^* = \Phi(1) \approx 0.8413$ for which the corresponding threshold is $t^* = \mu + \sigma$. Using these two values we obtain estimations for μ, σ and we apply a similar approach to Algorithm **KnownVar**.

UnkVar properties

- **Privacy**: **UnkVar** is (ϵ, δ)-LDP.
- **Confidence Interval**: Let $\mathbf{X} \sim \mathcal{N}(\mu, \sigma^2)$ i.i.d. Fix parameters ϵ, $\beta \in (0, 1/2)$. Given that $\mu \in [-R, R]$ and that $\sigma_{\min} \leq \sigma \leq \sigma_{\max} \leq 2R$, if
$$n \geq 1500 \log_2\left(\tfrac{16R}{\sigma_{\min}}\right) \cdot \left(\frac{e^\epsilon+1}{e^\epsilon-1}\right)^2 \cdot \ln\left(\frac{16\log_2(16R/\sigma_{\min})}{\beta}\right)$$
then the interval $\hat{I}$ returned by Algorithm **UnkVar** satisfies that $\underset{\mathbf{X},\, \mathbf{UnkVar}}{\mathbb{P}}\left[\hat{I} \ni \mu\right] \geq 1 - \beta$, and moreover
$$\hat{I} = O\left(\sigma \cdot \frac{\sqrt{\log(n/\beta)\log(1/\beta)\log(1/\delta)}}{\epsilon\sqrt{n}}\right)$$
- **Very large variance case**: If $\sigma > R$ we give a different algorithm, based on matching quantiles. We estimate $p_- = \Pr[X < -R]$ and $p_+ = \Pr[X < R]$, then plot the Gaussian based on the quantiles of $\mathcal{N}(0, 1)$ obtaining p_- and p_+.

References

[1] Bassily and Smith. Local, private, efficient protocols for succinct histograms. In *STOC'15*.

[2] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT06*, 2006.

[3] V. Karwa and S. P. Vadhan. Finite sample differentially private confidence intervals. In *ITCS18*.

[4] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. D. Smith. What can we learn privately? In *FOCS08*.

[5] S. L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60:63–69, 1965.