Differentially Private Mean Estimation in the Local Model

Humza Haider
ID: 1535075

April 22, 2018

1 Introduction

The use of personalized information has been the driver of many fields including various medicinal fields, psychology, and business analytics. As is often the case, this information may be compiled into aggregate statistics such as averages/means, variance estimates, counts, and proportions which may be released to the general public to convey some finding. While not immediately obvious, these aggregated statistics can often leak privacy at the individual level. For example, numerous occurrences of individual identification by the use of published, aggregated statistics has occurred in the field of genome-wide association studies (GWAS) [1, 2, 3].

Historically, many different schemas have evolved for protecting privacy, including removing identifiers (*e.g.* name, zip code, date of birth), limiting the type and sheer quantity of data released (as discussed in the GWAS field by Craig et al. [3]), and perturbing the data by adding some level of noise. More recently, the notion of *differential privacy* has been introduced by Dwork et al.[4] and has since gained traction and popularity among academic and industry communities alike.

We focus on a specific aspect of differential privacy known as *local* differential privacy which is focused on privatizing data at an individual level. For example, the key difference between typical differential privacy and local differential privacy is that in the local model there is no need for a trusted curator to release private, aggregated statistics, as the data itself is inherently private. This notion of local privacy continues to grow in importance as data leaks are becoming ever more frequent in the growing age of technology.

As there is no trusted curator and all data is inherently private, every individual's data must have a robust level of noise added to achieve differential privacy. In doing so, statistical estimation becomes increasingly prone to error as the noise increases. In this paper we consider the estimation of a locally, differentially private mean. As is typical with mean estimators we derive confidence intervals as a measure of utility and comparison among estimators. Additionally, we give a brief discussion of runtime among mechanisms for estimating a locally private mean.

In the following sections we present the formal definition of differential privacy in addition to examining the tools required for our estimators of the mean. Specifically the tools used are the Gaussian mechanism given by Dwork et al. [5], and the technique of randomized response, first introduced by Warner in 1965 [6].

2 The Notion of Differential Privacy

To understand the robustness of differential privacy we present the mathematical definition here. However, to define differential privacy we first need the preliminary definition of two *neighboring* datasets, D and D'.

Definition 1: We say $D \sim D'$ are two neighboring datasets if they differ by exactly one row/individual. Specifically, if we let $D = (d_1, d_2, \ldots, d_i \ldots d_n)$ then D' could be defined as $D' = (d_1, d_2, \ldots, d'_i \ldots d_n)$ where $d_i \neq d'_i$. Alternatively, we could also add or subtract one individual from D and still satisfy the definition of a neighboring dataset.

Using this we now define differential privacy as given by Dwork et al. [4].

Definition 2: Letting $\mathcal{U}^n$ represent the universe of datasets and O as potential outcomes we say a randomized algorithm $\mathcal{A} : \mathcal{U}^n \to O$ is (ϵ, δ)-differentially private for $\epsilon, \delta > 0$ if for any neighboring datasets $D \sim D'$ and any subset of outcomes, $S \subset O$ we have

$$\Pr[\mathcal{A}(D) \in S] \leq e^{\epsilon} \Pr[\mathcal{A}(D') \in S] + \delta.$$

The robustness of this definition is straightforward; we bound the probability that the output will change for any two neighboring datasets by a multiplicative term involving ϵ and the additive term δ. Here, δ should be thought of as the probability that a catastrophic event (in terms of privacy) occurs, *e.g.* releasing the true data, and ϵ as the privacy level each individual receives. Note that an algorithm with $\delta = 0$ is simply referred to as ϵ-differentially private and is often thought of as "true" privacy.

As we have previously mentioned, we focus on local differential privacy, of which a robust definition can be found in the work by Kasiviswanathan et al. [7]. For this work, local differential privacy can simply be thought of as a (ϵ, δ)-differentially private mechanism which runs on a dataset of size $n = 1$, *i.e.* privacy occurs at the level of the individual as opposed to the level of the aggregated output of the mechanism. Further, by using local privacy, individuals retain plausible deniability that the result the mechanism returned was not the value the individual inputted. In turn this should assure users that there is little risk when running their data through a locally, differentially private mechanism.

2.1 The Gaussian Mechanism

The first tool we use to attain differential privacy is known as the Gaussian mechanism, and can be thought of as adding noise derived from a normal distribution with mean 0 and some variance σ^2. To derive the value of σ^2 we first introduce *global sensitivity* which can be interpreted as the maximal impact that any one individual can have on the function being applied to the data. For example, in the case on a counting query, the function is simply the count of individuals meeting some criteria. Further, any one individual can only change the answer to the counting query by 1, thus the global sensitivity would equal 1. Mathematically, for a given function f and for any neighboring datasets, $D \sim D'$ we have

$$GS(f) = \max_{D \sim D'} ||f(D) - f(D')||^2.$$

Given a function, $f : \mathcal{U}^n \to \mathbb{R}^d$, the global sensitivity of a function, $GS(f)$, the dataset D, and privacy parameters (ϵ, δ) the Gaussian mechanism is completed in two steps:

1. Let $X_i \overset{i.i.d}{\sim} \mathcal{N}(0, \sigma^2)$ where $\sigma^2 = \frac{2\ln(\frac{2}{\delta})}{\epsilon^2} GS(f)^2$.
2. Return $f(D) + (X_1, X_2, \ldots, X_d)$.

A proof regarding the differential privacy and utility of this mechanism is available from Dwork et al. [5]. Note that in the local model we have datasets of size 1 so the global sensitivity varies across all potential individuals as opposed to datasets. In Section 3.1 we apply the Gaussian mechanism to attain a locally private mean estimate.

2.2 Randomized Response

The next tool we introduce for differential privacy, known as randomized response, was first given in 1965 by Warner [6], long before differential privacy was introduced in 2006 by Dwork et al [4]. The initial intention of randomized response was to avoid an evasive answer bias (resulting from respondents not giving true answers to very personal questions) but has since been adapted to provide differential privacy. The implementation of randomized response is straightforward; suppose our universe consists of 0's and 1's, *i.e.* $\mathcal{U} = \{0, 1\}$ and $D \in \mathcal{U}^n$. Then for every individual i in D we sample $b_i \in \{0, 1\}$ independently, such that for some probability $p < \frac{1}{2}$,

$$\Pr(b_i = x) = \begin{cases} \frac{1}{2} + p & x = D_i \\ \frac{1}{2} - p & x = 1 - D_i \end{cases}$$

A proof of the differential privacy of a simplified version of randomized response is made available by Dwork et al. [5] which can be extended to the general case to show that randomized response is $\left(\frac{4p}{1-2p}\right)$-differentially private, where $\delta = 0$. Given the setup of randomized response a natural value to estimate is the number of 1's contained in D. Denoting $RR(D)$ as the output of the randomized response mechanism, Warner [6] showed that the maximum likelihood estimator for the proportion of 1's (θ_1), is given by

$$\hat{\theta_1} = \frac{1}{2p}\left(\frac{\text{Number of 1's in } RR(D)}{n} - \left(\frac{1}{2} - p\right)\right). \tag{1}$$

Additionally, Warner went on to show that $\hat{\theta_1} \sim \mathcal{N}\left(\theta_1, \frac{1}{4p^2n^2}\left(\frac{n}{4} - np^2\right)\right)$.

As stated, randomized response handles only two types, 0 and 1. However, randomized response can be extended to T types through the work of Bassily and Smith [8]. We start by considering individuals to be vectors of length T, *e.g.* an individual of type t is given by $(0, 0, \ldots 0, 1, 0, \ldots, 0)$, where the 1 is in the t^{th} coordinate. Following the vector notation, each individual runs randomized response independently on every value of the vector with parameter p such that $\frac{\frac{1}{2}+p}{\frac{1}{2}-p} \leq e^{\frac{\epsilon}{2}}$.

The proof that this mechanism is ϵ-differentially private comes about by noting that only two points in a vector differ among neighboring datasets. First, fix $D \sim D'$, differing on a single individual, i, who switches from type t to t'. Since all other vectors of D and D' are equal, the responses from randomized response are distributed the same across these individuals. Further, by our definition of neighbors we have that the differing vector of D and D' differ only in the t^{th} and $t'^{\,\text{th}}$ position. Denoting $\mathcal{M}(i)$ as the mechanism running on the i^{th} individual, we have for any vector $b \in \{0, 1\}^T$,

$$\begin{aligned} \frac{\Pr(\mathcal{M}(i) = b \,|\, \text{user } i \text{ is type } t)}{\Pr(\mathcal{M}(i) = b \,|\, \text{user } i \text{ is type } t')} &= \frac{\Pr[RR(1) = b_t] \cdot \Pr[RR(0) = b_{t'}]}{\Pr[RR(0) = b_t] \cdot \Pr[RR(1) = b_{t'}]}, \\ &= \frac{\frac{1}{2}+p}{\frac{1}{2}-p} \cdot \frac{\frac{1}{2}+p}{\frac{1}{2}-p}, \\ &\leq e^{\frac{\epsilon}{2}} \cdot e^{\frac{\epsilon}{2}}, \\ &= e^{\epsilon}. \end{aligned}$$

Next, make note that the estimator given in equation (1) can still be used for estimating the number of individuals of any type t by using all individual's t^{th} coordinate. Further, this estimator retains the same variance and unbiasedness as given in Section 2.2.

In the following section we utilize the Gaussian mechanism and the Bassily-Smith modification of randomized response to build a locally private mean estimator. We first use the Gaussian mechanism to create a basic model and follow with the more complicated method by utilizing randomized response.

3 Mean Estimation

Mean estimators are one of the most commonly sought after statistics for data in general. While extremely common, they are also very non-private. Consider the simple example of a knowing the mean of incomes with and without a single individual, μ_A and μ_{-A}, respectively, and the number of individuals used to calculate μ_A, denoted n. Given these two values we have that $A = \mu_A\, n - \mu_{-A}\,(n-1)$, showing that the mean itself is not differentially private. Here we begin by adding noise to A by using the Gaussian mechanism and then take a more sophisticated approach by applying randomized response to all the possibilities that A could take on.

3.1 A Baseline Model Using Gaussian Noise

While using Gaussian noise to attain local differentially privacy is a fairly straightforward approach, it requires some assumptions to become a usable model. Suppose that we want to estimate the mean of some variable given to us by a total of n individuals. We choose to assume that (1) this collected variable follows a normal distribution, *i.e.* $X_i \overset{i.i.d}{\sim} \mathcal{N}(\mu, \sigma^2)$, and (2) X_i has known σ^2. Note that the function we are applying to X_i is the identity function, *i.e.* $f(x) = x$. Further, the global sensitivity will of the identity function will correspond to the support for the distribution of X_i. Specifically, the normal distribution has support $(-\infty, \infty)$ thus causing an infinite global sensitivity. To rectify this, we make the assumption that $\mu \in [-R, R]$ for some $R \in \mathbb{R}^+$. While not immediately solving the issue of the global sensitivity, recall that for the normal distribution, we have that with probability $\geq (1-\delta)$,

$$X_i \in \left[-\mu - \sigma\sqrt{2\ln\left(\frac{2n}{\delta}\right)},\ \mu + \sigma\sqrt{2\ln\left(\frac{2n}{\delta}\right)} \right].$$

By plugging in our bound on μ we have, instead, with probability $\geq (1-\delta)$,

$$X_i \in \left[-R - \sigma\sqrt{2\ln\left(\frac{2n}{\delta}\right)},\ R + \sigma\sqrt{2\ln\left(\frac{2n}{\delta}\right)} \right].$$

Note that this bound is only possible as we have assumed a known σ^2. As δ represents the probability of catastrophe, our mechanism will stop and return nothing if a value of X_i is sampled outside of our bound. As we have bounded values for X_i with probability $\geq (1-\delta)$, we now have a finite global sensitivity for the identity function,

$$GS(f) = 2\left(R + \sigma\sqrt{2\ln\left(\frac{2n}{\delta}\right)}\right)$$

Thus for every individuals response we add Gaussian noise, Y_i, such that $Y_i \overset{i.i.d}{\sim} \mathcal{N}\left(0, \frac{2\ln\left(\frac{2}{\delta}\right)}{\epsilon^2} GS(f)^2\right) = \mathcal{N}\left(0, \tilde{\sigma}^2\right)$. Then since X_i follows a normal distribution we have our private response, $Z_i = (X_i + Y_i) \overset{i.i.d}{\sim} \mathcal{N}\left(\mu, \sigma^2 + \tilde{\sigma}^2\right)$. Here, we can simply take the sample average, $\bar{Z}$, to attain an unbiased estimator of μ. More so, by an application of the central limit theorem we have that a $(1-\alpha)\%$ confidence interval for μ is given by

$$\mu \in \left[\bar{Z} - z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sigma^2 + \tilde{\sigma}^2}{n}},\ \ \bar{Z} + z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sigma^2 + \tilde{\sigma}^2}{n}}\right],$$

where z^* is the associated z-score for a given value of α.

Make note that this confidence interval expands rapidly as the bound on μ increases. For a small n this confidence interval is worthless as it gives us approximately the original bound of $[-R, R]$, since $\tilde{\sigma^2}$ is $O(R^2)$. Further, we have made three very strong assumptions, (1) the underlying sampling distribution is normal, (2) we have a known variance, and (3) we have that $\mu \in [-R, R]$ for a known R.

While this mechanism is not diverse in application, when we assume a constant runtime for sampling Y_i, the mechanism can run in $O(n)$ time as we simply add noise to n observations and take the average. Thus when the assumptions are satisfied and we have a relatively tight bound on μ, this mechanism will be suffice for a estimate of μ.

An alternative to bounding μ is to bound X_i, resulting in a truncated Gaussian distribution. The resulting distribution of adding a truncated Gaussian and a Gaussian has been studied by Kim [9], however, recovering μ through an unbiased estimator is not straightforward. For this reason, we still consider an approach by bounding X_i, but instead use randomized response to attain differential privacy.

3.2 Using Randomized Response for Mean Estimation

Suppose we collect n samples of a random variable, X_i, following an unknown distribution with mean μ such that $X_i \in [L, U]$ for $L, U \in \mathbb{R}, L \leq U$. This can be bound can be made for any random variable via truncation, *e.g.* if we are sampling individual's income we simply ask respondents for their income unless they are above or below the threshold and instead they return U and L respectively. Before we can apply randomize response we first discretize the space into evenly sized intervals resulting in T types (Figure 1). Note that if we divide the space using intervals of size c, we have that $T = \frac{U-L}{c}$. Further, by the condition of evenly sized intervals, c must be a divisor of $U - L$. Returning to the income example, we alter the question such that now individuals choose the interval in which their income lies, *e.g.* \$0 - \$20,000. Following this, each respondent will run the Bassily-Smith modification of randomized response across their vector of types, where each type is one of the T intervals. Differential privacy of this mechanism follows immediately from the privacy of the Bassily-Smith modification of randomized response.

By discretizing the space in this way we have transformed our distribution into a multinomial distribution with T types and probability vector $\Theta = (\theta_1, \theta_2, \ldots, \theta_T)$, with θ_i equaling the total probability mass for the i^{th} interval of the original distribution. For example, from Figure 1, θ_2 would represent the probability mass within the second interval.

The multinomial distribution is typically used for categorical data and estimating proportions of each type, however, this is not the intention of our discretization. We instead wish to estimate

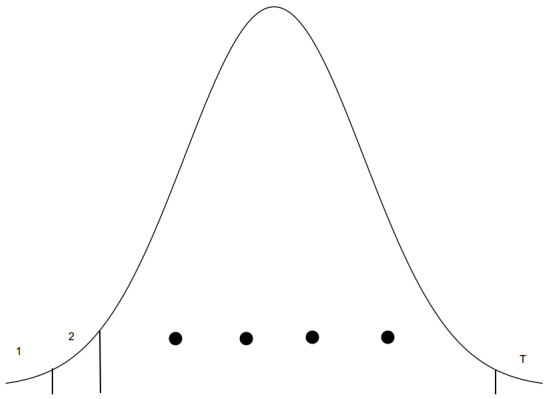

Figure 1: Depiction of discretization of a normal distribution into T types for implementation of randomized response.

the average type, ν, giving us an approximation for μ. Denoting the midpoint of the i^{th} interval as m_i, we define the average type as

$$\nu = \sum_{i=1}^{T} m_i \, \theta_i.$$

First make note that ν is (typically) not an unbiased estimator of μ, and is dependent on T. This is best seen through an example; suppose we have a normal distribution with mean 50 and standard deviation 2 and discretize it using an upper bound of 75 and a lower bound of -75 with an interval size of 75. This leaves us with two intervals, $(-75, 0]$ and $(0, 75]$[1], with respective midpoints -37.5 and 37.5. Note that due to the density function of the normal distribution, nearly all density will be in the second interval, leaving us with $\nu = -37.5 \cdot 0 + 37.5 \cdot 1 = 37.5$, resulting in an error of 12.5. Further, this estimate will not get better as n increases, up until we also increase T. Once T reaches a sufficient quantity, increasing n will better the estimate. Further, notice that this bias cannot be more than half the interval size, that is $\frac{c}{2}$.

While we do not give a formal proof due to complexities arising from no assumptions on the underlying distribution, notice that ν is a consistent estimator as both n and T increase as it will converge to $\int_L^U f(x)\, x \, dx$, the integral of the density function multiplied by the type, which is the definition of the mean, μ, after truncation of the original distribution.

Now that we have a definition of ν, we build our private estimator, $\hat{\nu}$, using the tools developed in Section 2.2. Recall our estimator from Equation 1, $\hat{\theta}_i$, for θ_i. One may recall that typical estimators for θ_i in the multinomial distribution are *dependent* with one another since they must sum to 1. However, this is not the case for $\hat{\theta}_i$ since randomized response is run for every coordinate of the response vector independently and there is no condition that the $\hat{\theta}_i$'s must sum to 1. Therefore, since each $\hat{\theta}_i$ is run on separate coordinates and there is no condition on the sum we have independence between the $\hat{\theta}_i$'s. Defining $\hat{\nu} = \sum_{i=1}^{T} m_i \, \hat{\theta}_i$, we have the following equalities:

[1]Note that the inclusion/exclusion of endpoints is arbitrary as we only use the midpoints for calculations.

$$
\begin{aligned}
\mathbb{E}[\hat{\nu}] &= \mathbb{E}\left[\sum_{i=1}^{T} m_i\,\hat{\theta}_i\right], \\
&= \sum_{i=1}^{T} m_i\,\mathbb{E}\left[\hat{\theta}_i\right], \\
&= \sum_{i=1}^{T} m_i\,\theta_i, && \text{(Due to unbiasedness of } \hat{\theta}_i\text{)} \\
&= \nu, \\
\mathrm{Var}[\hat{\nu}] &= \mathrm{Var}\left[\sum_{i=1}^{T} m_i\,\hat{\theta}_i\right], \\
&= \sum_{i=1}^{T} m_i^2\,\mathrm{Var}\left[\hat{\theta}_i\right], && \text{(Independence of the } \hat{\theta}_i\text{'s)} \\
&= \sum_{i=1}^{T} m_i^2\,\frac{1}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right), && \text{(From Section 2.2)} \\
&= \frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right).
\end{aligned}
$$

Additionally, due to the fact that $\hat{\theta}_i \sim \mathcal{N}\left(\theta_i, \frac{1}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)\right)$, we have that $\hat{\nu} \sim \mathcal{N}\left(\nu, \frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)\right)$. Using this a $(1-\alpha)\%$ confidence interval is immediately given by

$$
\nu \in \left[\hat{\nu} - z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)},\ \ \hat{\nu} + z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)}\right].
$$

Further, by again noting that ν differs from μ by at most half the interval size, $\frac{c}{2}$, we have a $(1-\alpha)\%$ confidence interval[2] for μ in the *worst case* to be

$$
\mu \in \left[\hat{\nu} - \frac{c}{2} - z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)},\ \ \hat{\nu} + \frac{c}{2} + z^*_{\frac{\alpha}{2}}\sqrt{\frac{\sum_{i=1}^{T} m_i^2}{4\,p^2\,n^2}\left(\frac{n}{4} - np^2\right)}\right].
$$

Given the confidence interval of μ it should be straightforward that there is a significant bias-variance tradeoff by the manipulation of the size of the intervals, c and the total number of intervals, T. By having too large a c or equivalently, too few intervals, ν becomes a more biased estimator of μ, but has reduced variance due to fewer midpoints being utilized. Finding an optimal value of c is not immediately obvious and is an open problem for future research of this mechanism.

As compared to the confidence interval derived using the Gaussian mechanism we have instead bounded the value of X_i as opposed to μ which should, in practice, be more obvious. Further,

[2]While this confidence interval is not in terms of the privacy parameter ϵ, recall that p is chosen for a given ϵ.

the margin of error is no longer in terms of the bound on the mean, but instead by a sum of the midpoints of the intervals. Due to this change, immediate comparison between the two mechanisms is non-obvious, however, the mechanism using randomized response has made fewer and less strong assumptions than did the mechanism utilizing the Gaussian mechanism.

While the mechanism using randomized response is generally more applicable, notice the increased runtime of the mechanism. Randomized response must be run across every coordinate of every individual, leading to a total runtime of $O(n\,T)$. If the distribution of X_i is large, a large number of both intervals and total individuals will be required to have a meaningful margin of error, which will increase the runtime such that this mechanism may no longer be suitable. Careful consideration regarding the data size and distribution should be made before using this mechanism.

A natural follow up to our locally private mean estimation is to ask whether or not we can acquire a locally private variance estimation using the same mechanism. Similar to how we defined ν, an intuitive estimate of the underlying variance, σ^2, would be $\lambda^2 = \frac{1}{n-1}\sum_{i=1}^{T} n\,\theta_i(\,m_i - \nu)^2$. As before, this a biased estimator but consistent as n and T approach infinity. Unfortunately, estimating this quantity does not yield a useful result; recall the definition of $\hat{\theta}_i$ and recognize that this estimate can take on negative values. In the event that $\hat{\theta}_i$ takes a negative value (which increases in probability as p shrinks), it is possible to get a negative estimate of λ^2, which is useless as a variance estimate. For this reason, we leave locally private variance estimation as an open problem requiring further research.

4 Conclusion

We have given two different locally, differentially private mean estimators using the Gaussian mechanism and the Bassily-Smith modification of randomized response. The approach using the Gaussian mechanism is simple, but requires strong assumptions including sampling from an underlying normal distribution and having a known variance of that normal distribution. Confidence intervals for the mean using this mechanism expand rapidly as the bound on the mean increases. However, having prior knowledge about the bound in which the mean lies may result in a relatively small value of, giving comparatively smaller confidence intervals than the mechanism derived using randomized response.

As opposed to using the Gaussian mechanism, the approach using randomized response requires fewer assumptions, but in turn requires a longer runtime. Further, finding optimal meta parameters (the interval size and number of intervals) may become difficult, as there is a bias-variance tradeoff occurring as the number of intervals/coordinates increases or as the interval size decreases. Future work can be done on this mechanism to theoretically or empirically show how one should go about choosing the number of intervals. Further, we briefly discussed that variance estimation failed with using the intuitive estimator from the randomized response mechanism. Altering the current estimator or deriving a new mechanism for estimating the population variance is an important next step for local differential privacy. Finding a robust estimation of population variance in the local model would allow for a large number of private hypothesis tests by using local private mean and variance estimations together.

Future work on local differentially private mean estimation may want to consider alternative mechanisms for comparison of the margin of error. Here we have only considered two mechanisms so there is much room for growth and development of tighter estimations, allowing for a great deal of utility even in the local, differentially private model.

References

[1] Y. Wang, X. Wu, and X. Shi, "Using aggregate human genome data for individual identification," in *Bioinformatics and Biomedicine (BIBM), 2013 IEEE International Conference on*, pp. 410–415, IEEE, 2013.

[2] N. Homer, S. Szelinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig, "Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays," *PLoS genetics*, vol. 4, no. 8, p. e1000167, 2008.

[3] D. W. Craig, R. M. Goor, Z. Wang, J. Paschall, J. Ostell, M. Feolo, S. T. Sherry, and T. A. Manolio, "Assessing and managing risk when sharing aggregate genetic variant data," *Nature Reviews Genetics*, vol. 12, no. 10, p. 730, 2011.

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

[5] C. Dwork, A. Roth, *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[6] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965.

[7] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, "What can we learn privately?," *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.

[8] R. Bassily and A. Smith, "Local, private, efficient protocols for succinct histograms," in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 127–135, ACM, 2015.

[9] H.-J. Kim, "On the distribution and its properties of the sum of a normal and a doubly truncated normal," *Communications for Statistical Applications and Methods*, vol. 13, no. 2, pp. 255–266, 2006.