Detection of Sparse Mixtures with Differential Privacy

Ruizhi Zhang

***Abstract*—Detection of sparse signals arises in many modern applications such as signal processing, bioinformatics, finance, and disease surveillance. However, in many of these applications, the data may contain sensitive personal information, which is desirable to be protected during the data analysis. In this article, we consider the problem of (ϵ, δ)-differentially private detection of a general sparse mixture with a focus on how privacy affects the detection power. By investigating the nonasymptotic upper bound for the summation of error probabilities, we find any (ϵ, δ)-differentially private test cannot detect the sparse signal if the privacy constraint is too strong or if the model parameters are in the undetectable region in [2]. Moreover, we study the private clamped log-likelihood ratio test proposed in [3] and show it achieves vanishing error probabilities in some conditions on the model parameters and privacy parameters. Then, for the case when the null distribution is a standard normal distribution, we propose an adaptive (ϵ, δ)-differentially private test, which achieves vanishing error probabilities in the same detectable region in [2] when the privacy parameters satisfy certain sufficient conditions. Several numerical experiments are conducted to verify our theoretical results and illustrate the performance of our proposed test.**

I. Introduction

Detection of sparse mixtures, with the goal to determine the existence of signals in a small fraction of a noisy dataset, has many important applications such as signal processing [4], finance [5], industrial quality control [6], and disease surveillance [7]. For example, for image-based quality inspection, the photoelasticity test is a nondestructive evaluation method used for stress and strain analysis of translucent parts, or material [8]. The output is often presented by a colormap, which is used to detect sparse anomalies with high-tensile stress. In many cases, such image colormap data may have complicated spatial correlation structures, and some decorrelation techniques methods, for example, [9]–[11], can be used. Then, the model residual or residual image, which is often assumed to be i.i.d for the normal pixels, is used to detect sparse anomalies. However, in many of these applications of detection of sparse mixtures, the data may contain sensitive personal information such as financial or medical records. Thus, procedures with good detection ability for the sparse signal while preserving individual information are highly desirable.

This work was supported in part by the U.S. National Science Foundation under grant ECCS-2236565, through the University of Georgia.

Ruizhi Zhang is with the Department of Statistics, University of Georgia, Athens, Georgia, 30602 USA, Email: ruizhi.zhang@uga.edu.

An earlier version of this paper was presented in part at the 2023 IEEE International Symposium on Information Theory (ISIT) [1]. (Corresponding author: Ruizhi Zhang.)

The study of the detection of sparse Gaussian mixtures where the fraction of signal is close to zero was first investigated by Ingster [12], where the detectable region is discovered. It has been shown that for any test, the sum of the probabilities of type I and type II errors are always bounded away from 0 outside the detectable region, which implies no tests can detect the sparse signal asymptotically. Moreover, although the likelihood ratio test can be applied to distinguish the two hypotheses efficiently in the detectable region, it may not be applied directly since it involves parameters in the alternative distribution, which may be unknown in practice. Later on, many efficient tests, which do not involve the information of parameters in the alternative hypothesis, were proposed and proved to be optimal in the detectable region, such as the higher criticism test [13], [14], goodness-of-fit test based on Φ-divergences [15], [16] and the Berk-Jones test [17]. Moreover, the detection boundary for the heterogeneous and heteroscedastic Gaussian mixture model was studied in [18], and a double-sided version of the higher criticism test was proved to be optimal in the detectable region. [2] developed a general detectable region and has shown the higher criticism test is adaptive optimal in such a detectable region. Later on, [19] conducted a rate analysis on the error probabilities for the log-likelihood ratio test to detect the general sparse mixtures. However, none of these works consider the privacy-preserving guarantees.

Since Dwork *et al.*'s pioneer work on differential privacy [20], differential privacy has garnered much attention. A wide variety of differentially private procedures with theoretical efficiency guarantees have been developed for many statistical problems such as point estimation [21], [22] and hypothesis testing [3], [23]–[25]. Informally, the field of differential privacy provides a systematic tool to construct private algorithms or procedures by adding designed random noise such that the output has a similar distribution with or without data for each individual participant, which helps to protect the information of individuals in the dataset. In particular, in the area of hypothesis testing, [23] proposed differentially private tests for categorical data. [24] focused on the differentially private test of the Gaussian distribution and proposed a test based on the likelihood ratio test. Under the differential privacy context, [22] proposed a general framework to construct robust hypothesis tests based on M-estimation. For the general simple hypothesis testing setting, [3] derived an upper bound for the sum of type I and type II error probabilities of simple hypotheses for general ϵ-differentially private tests and then proposed an optimal test by the clamped likelihood ratio statistics. [25] proposed differentially private algorithms for

controlling the false discovery rate in multiple hypothesis testings. The problem of private identity testing for high-dimensional distributions was studied in [26]. [1] studied the problem of detection of the sparse Gaussian mixture under the ϵ-differentially private constraint with a focus on how the privacy parameter ϵ affects the detection ability for some differentially private tests.

In this paper, we focus on the problem of differentially private detection of a general sparse mixture and make three major contributions. First, we study the detection ability of a general sparse mixture under the framework of (ϵ, δ)-differential privacy. Motivated by the detection boundary of the sparse Gaussian mixtures proposed by [12], [13] and the general sparse mixture derived by [2] under the non-privacy-preserving scenario, we further take the privacy parameters ϵ, δ into account and derive a nonasymptotic bound on the summation of type I and type II of any (ϵ, δ)-DP test. Based on such bound, we can see that if $\epsilon = o(1/n), \delta = o(1/n)$, or if the model parameters are in the undetectable region derived by [2] and $\delta/\epsilon = O(1)$, any (ϵ, δ)-DP test cannot have vanishing type I and type II probabilities asymptotically and thus cannot detect the presence of the sparse signal. This result is consistent with the fact that strong privacy may damage the accuracy of the procedure. Second, we study the detection power of the noisy clamped log-likelihood ratio test proposed by [3], which is optimal in the sense of requiring the smallest sample size to achieve designed detection power under the ϵ-differentially privacy constraint for the simple hypothesis testing problem where the null and the alternative distributions do not involve the sample size n. In this paper, we further consider the clamped log-likelihood ratio test under the (ϵ, δ)-DP constraint and show that the test has good detection efficiency in the sense of achieving vanishing probability of errors when the model parameters are inside of a detectable region and the privacy parameters satisfy $\frac{\sqrt{\log(1.25/\delta)\log(n)}}{\epsilon} = O(1)$. Third, for the Gaussian null hypothesis, we propose an adaptive differentially private test, which achieves vanishing probability of errors when model parameters are inside the detectable region of [2] and $\sqrt{\frac{\log(\log n)}{\log(1.25/\delta)}}\epsilon \to \infty$. Our proposed test is based on one variant of higher criticism statistic, which can separate the null hypothesis and the alternative hypothesis asymptotically without the privacy constraint [2], [13]. We show that this variant of higher criticism statistics has a relatively small sensitivity. Thus, we can construct an efficient (ϵ, δ)-differentially private test by adding a small noise. We should emphasize that our proposed test does not involve parameters in the alternative distribution, which are usually unknown to statisticians in practice.

The organization of this article is as follows. In Section II, we introduce the preliminaries and background on the detection of general sparse mixtures and (ϵ, δ)-differentially private tests. In Section III, we study the fundamental bounds of general (ϵ, δ)-DP tests for detecting the sparse mixture and investigate the detection property of the clamped likelihood ratio test by adding an independent Gaussian noise. We then describe an adaptive (ϵ, δ)-DP test in Section IV and show simulation results in Section V. The conclusion and further discussions are provided in Section VI. The proofs of the main theorems are postponed in the Appendix.

II. Preliminaries

In this section, we provide the necessary preliminaries and background on the detection of sparse mixtures and differentially private tools with an emphasis on the hypothesis testing problems.

A. Detection of Sparse Mixtures

Suppose we have n independent and identically distributed (i.i.d.) random samples $X_1, X_2, \cdots, X_n$. We consider the following hypothesis testing problem with the null hypothesis as

$$\mathcal{H}_0 : X_i \overset{i.i.d.}{\sim} p, \tag{1}$$

and the alternative hypothesis as

$$\mathcal{H}_1^{(n)} : X_i \overset{i.i.d.}{\sim} (1-\lambda)p + \lambda g_n. \tag{2}$$

Here, p denotes the p.d.f of the null distribution, g_n denotes the p.d.f of the non-null effects (signal) distribution, which depends on the sample size n. Roughly speaking, under the null hypothesis $\mathcal{H}_0$, all data are i.i.d from the distribution p. Under the alternative hypothesis $\mathcal{H}_1^{(n)}$, a fraction λ of the data come from another distribution g_n and the remaining $1-\lambda$ fraction of the data are still from the null distribution p.

Clearly, if λ, p, and g are fixed and known, the optimal procedure is simply the likelihood ratio test. In this paper, we follow the literature on detecting the sparse mixture [2], [12], [13] and assume $\lambda = n^{-\beta}$ for some exponent $\frac{1}{2} < \beta < 1$, so that the fraction of nonzero means is small but not vanishingly small. In the special case of detection of sparse Gaussian mixtures when the null distribution $p \sim N(0,1)$ and $g_n \sim N(\sqrt{2r\log(n)}, 1)$, for a positive constant $r > 0$, [12], [13] found the sharp detection boundary in the (r, β) plane, which is given by

$$\rho^*(\beta) = \begin{cases} \beta - \frac{1}{2}, & \text{if } 1/2 < \beta \le 3/4, \\ (1-\sqrt{1-\beta})^2, & \text{if } 3/4 < \beta < 1. \end{cases} \tag{3}$$

That is if $r < \rho^*(\beta)$, $\mathcal{H}_0$ and $\mathcal{H}_1^{(n)}$ are asymptotically unseparatable in the sense that the sum of type-I and II error probabilities for any tests goes to 1. Otherwise, if $r > \rho^*(\beta)$, then $\mathcal{H}_0$ and $\mathcal{H}_1^{(n)}$ are asymptotically detectable. The strict epigraph $\{(r, \beta) : r > \rho^*(\beta)\}$ is known as the detectable region. Moreover, inside the detectable region, several adaptive optimal tests, which have vanishing Type I and Type II error probabilities and do not rely on the information of r and β, are proposed in the literature. For example, [13] showed that Tukey's higher criticism test can separate $\mathcal{H}_0$ and $\mathcal{H}_1^{(n)}$ asymptotically when $r > \rho^*(\beta)$. The test is based on Tukey's higher criticism statistic, which is defined by

$$\mathrm{HC}_n^* = \sup_{t \in (0,1)} \frac{[\sum_{k=1}^n I(\pi_k \le t) - nt]}{\sqrt{nt(1-t)}}, \tag{4}$$

where $\pi_k = \mathbf{P}(N(0,1) \ge X_k)$, $I(.)$ is the indicator function. Tukey's higher criticism test rejects the null hypothesis if

HC_n^* in (4) is greater than some critical values. There are several variants of HC statistic, see [13], [14]. In this paper, we choose the variant as in (4) and construct a differentially private test based on the discrete version of (4) since it can satisfy the privacy guarantee while keeping good detection efficiency. Besides the higher criticism test, the Berk-Jones test can also be used to detect the sparse Gaussian mixtures [27]. Let $u_i = \mathbf{P}(N(0,1) \le X_i)$ and $u_{(i)}$ denote its sorted order statistics in increasing order. Then, let the corresponding one-sided p-values be denoted as $p_i = \mathbf{P}(\mathrm{Beta}(i, n-i+1) < u_{(i)})$, where $\mathrm{Beta}(i,j)$ denotes the Beta distribution with parameters i, j. The exact Berk-Jones statistics M_n is defined by

$$\begin{aligned} M_n^+ &= \min_{1\le i\le n} p_i, \quad M_n^- = \min_{1\le i\le n}(1-p_i), \\ M_n &= \min\{M_n^+, M_n^-\}, \end{aligned} \tag{5}$$

and the BJ test rejects the null hypothesis if the test statistic M_n in (5) is smaller than some critical values. The optimality of the BJ test on detecting the sparse Gaussian mixtures has been studied in [17]. However, in the simulation study of this paper, we will show to meet the privacy constraint, the BJ-based private test has low detection efficiency since it may add too much noise compared with the signal or the BJ test statistics .

In general, [2] established an explicit expression for the detection boundary of the hypothesis testing problem in (1) and (2) under mild regularity conditions. Let $H^2(p,q)$ denote the Hellinger distance of two distributions p, q, which is defined by

$$H^2(p,q) = \int \left(\sqrt{p(x)} - \sqrt{q(x)}\right)^2 dx. \tag{6}$$

Note $H^2(p,q)$ takes values in the interval $[0,2]$. Denote the Hellinger distance between the null distribution p and the alternative distribution $(1-n^{-\beta})p + n^{-\beta}g_n$ in (1) and (2) by

$$H_n^2(\beta) = H^2(p, (1-n^{-\beta})p + n^{-\beta}g_n)). \tag{7}$$

It has been shown in [2] that $H_n^2(\beta)$ is decreasing with β. Moreover, define $\bar{\beta}^* = \inf\{\beta \ge 0 : nH_n^2(\beta) \to 0\}, \underline{\beta}^* = \sup\{\beta \ge 0 : nH_n^2(\beta) \to \infty\}$. [2] showed that $0 \le \underline{\beta}^* \le \bar{\beta}^* \le 1$ and when $\beta > \bar{\beta}^*$, all tests cannot separate the null and the alternative hypotheses since the summation of the type I and type II errors converge to 1 for any tests. Moreover, when $\beta < \underline{\beta}^*$, there exists a sequence of tests with vanishing type I and type II error probabilities. Furthermore, under mild regularity conditions on the distribution of $\log(g_n/p)$, [2] found $\bar{\beta}^* = \underline{\beta}^*$, which yield the explicit expressions for the detection boundary. As an example of Gaussian mixtures when $p \sim N(0,1)$ and $g_n \sim N(\sqrt{2r\log(n)}, 1)$, the detection boundary obtained by the condition $\beta = \bar{\beta}^* = \underline{\beta}^*$ is the same boundary as in (3).

[19] derived another set of sufficient conditions to find the detection boundary of the hypothesis testing problems in (1), (2) and also studied the optimal rate of decay of type I and type II error probabilities for the log-likelihood test in the detectable region.

B. Differentially Private Hypothesis Test

Considering a random algorithm that maps from a database space $\mathbb{R}^n$ with n entries, where each entry belongs to the set $\mathbb{R}$, to some measurable output space, we say the algorithm is differentially private if the outputs have similar distributions for neighboring datasets that we want to make it hard to distinguish. Here, two databases D, D' are neighboring if they differ in at most one entry. We now introduce the formal definition of differential privacy [20].

Definition 1. *A randomized algorithm $T : \mathbb{R}^n \to \mathbb{R}$ is (ϵ, δ)-differentially private (DP) if for every pair of neighboring databases $D, D' \in \mathbb{R}^n$, and for every subset of possible events $\mathcal{S} \subseteq \mathbb{R}$, $\mathbf{P}(T(D) \in \mathcal{S}) \le e^\epsilon \mathbf{P}(T(D') \in \mathcal{S}) + \delta$.*

In particular, for the special case of hypothesis testings with binary output $\{0,1\}$, we say a test with the test function $T : \mathbb{R}^n \to \{0,1\}$ is (ϵ,δ)-DP if $\mathbf{P}(T(D)=0) \le e^\epsilon \mathbf{P}(T(D')=0) + \delta$, and $\mathbf{P}(T(D)=1) \le e^\epsilon \mathbf{P}(T(D')=1) + \delta$ for every pair of neighboring databases $D, D' \in \mathbb{R}^n$. Here, $T(D)=1$ implies the test will reject $\mathcal{H}_0$ and $T(D)=0$ implies the test will accept $\mathcal{H}_0$ based on observed dataset D. Moreover, we call the (ϵ,δ)-DP test as ϵ-DP if $\delta = 0$.

For a regular test constructed by some test statistics L, e.g.,

$$T(D) = \begin{cases} 1, & \text{if } L(D) \ge c, \\ 0, & \text{otherwise,} \end{cases}$$

where c is the critical value controlling the type I error of the test, one common technique to achieve (ϵ,δ) differential privacy is by adding a Gaussian noise [28]. Specifically, we define the sensitivity of a real-valued function L as $\Delta(L) = \max_{D,D' \text{are neighbors}} |L(D) - L(D')|$. Then, for $0 < \epsilon < 1, 0 < \delta < 1$, we can get an (ϵ,δ)-DP test by adding an independent Gaussian random noise $N(0, 2(\Delta(L)/\epsilon)^2 \log(1.25/\delta))$ to the realization of the statistic $L(D)$, i.e., the resulting (ϵ,δ)-DP test will reject the null hypothesis if $L(D) + N(0, 2(\Delta(L)/\epsilon)^2\log(1.25/\delta)) \ge c$.

III. Fundamental Bounds for (ϵ,δ)-Differentially Private Tests

In this section, we focus on the fundamental bound for (ϵ,δ)-DP tests for the hypotheses (1) and (2). In particular, we will investigate the bound of the summation of type I and type II error probabilities for any private tests, which yields sufficient conditions on the regime of parameters so that any (ϵ,δ)-DP test cannot separate $\mathcal{H}_0$ and $\mathcal{H}_1^{(n)}$ asymptotically. Moreover, we will present the detection property of the noisy clamped log-likelihood ratio test proposed by Canonne et al. [3] and study the regime where it has vanishing type I and type II error probabilities.

Theorem 1. *For the hypothesis testing problem in (1) and (2), any (ϵ,δ)-DP test T satisfies*

$$\begin{aligned} &|\mathbf{P}(T(X)=1|\mathcal{H}_0) + \mathbf{P}(T(X)=0|\mathcal{H}_1^{(n)}) - 1| \\ &\le \frac{e^{\epsilon n}-1}{e^{\epsilon n}+1}\left(1 + 2\frac{\delta}{e^\epsilon - 1}\right)\sqrt{\min\left(2, nH_n^2(\beta)\right)}, \end{aligned} \tag{8}$$

where $H_n^2(\beta) = H^2(p, (1-n^{-\beta})p + n^{-\beta}g_n))$ is the Hellinger distance between the null distribution p in (1) and the alternative distribution $(1-n^{-\beta})p + n^{-\beta}g_n$ in (2).

The proof of Theorem 1 is postponed to the Appendix. By Theorem 1, if $n\epsilon \to 0, n\delta \to 0$, then $\frac{e^{\epsilon n}-1}{e^{\epsilon n}+1}\left(1 + 2\frac{\delta}{e^{\epsilon}-1}\right) \to 0$, equivalently, the summation of type I and type II for the test converges to 1, which implies that any (ϵ, δ)-differentially private tests cannot have vanishing error probabilities asymptotically and thus cannot separate the null hypothesis $\mathcal{H}_0$ in (1) and the alternative hypothesis $\mathcal{H}_1^{(n)}$ in (2) asymptotically for any value of β. Moreover, define $\bar{\beta}^* = \inf\{\beta \geq 0 : nH_n^2(\beta) \to 0\}, \underline{\beta}^* = \sup\{\beta \geq 0 : nH_n^2(\beta) \to \infty\}$. If $\beta > \bar{\beta}^*$ and $\delta/\epsilon = O(1)$, we can also get the summation of type I and type II for the test converges to 1. In summary, we have any (ϵ, δ)-differentially private tests cannot separate the null hypothesis $\mathcal{H}_0$ in (1) and the alternative hypothesis $\mathcal{H}_1^{(n)}$ in (2) asymptotically if $n\epsilon \to 0, n\delta \to 0$, or $\delta/\epsilon = O(1), \beta > \bar{\beta}^*$.

Next, we study the detection property of the clamped log-likelihood ratio test proposed by [3], which is optimal for fixed ϵ and simple hypotheses in the sense that it requires the smallest number of samples to achieve certain error probabilities. Here, we modify the clamped log-likelihood ratio test to make it into an (ϵ, δ)-DP test by adding an independent Gaussian noise. Specifically, given n i.i.d. samples $\mathbf{X} = (X_1, X_2, \cdots, X_n)$, the clamped log-likelihood ratio statistic is defined by

$$L_c(\mathbf{X}) = \sum_{i=1}^{n} \left[\log(1 - \lambda + \lambda\frac{g(X_i)}{p(X_i)})\right]_{-c}^{c}, \tag{9}$$

where $[\cdot]_a^b$ denotes the projection onto the interval $[a, b]$ (that is, $[z]_a^b = \max(a, \min(z, b))$). Then, we can modify the noisy clamped log-likelihood ratio test proposed by [3] to satisfy the (ϵ, δ)-differential private constraint by adding an independent Gaussian noise and obtain the following test function

$$\psi_c(\mathbf{X}) = \begin{cases} 1, & \text{if } L_c(\mathbf{X}) + Y \geq 0, \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

where $Y \sim N(0, 2(2c/\epsilon)^2 \log(1.25/\delta))$ is independent to $\mathbf{X}$. In this section, we will investigate its performance under our asymptotic settings, where the parameters $\lambda = n^{-\beta}, \epsilon, \delta$ are allowed to change in n. Note the test $\psi_{(\epsilon)}(\mathbf{X})$ in (10) could also be written as

$$\psi_c(\mathbf{X}) = \begin{cases} 1, & \text{with probability } h(\mathbf{X}) \\ 0, & \text{with probability } 1 - h(\mathbf{X}), \end{cases} \tag{11}$$

where $h(\mathbf{X}) = Q(-\frac{\epsilon}{2c\sqrt{2\log(1.25/\delta)}}L_c(\mathbf{X}))$, $Q(x)$ is the tail probability of the standard normal distribution $Q(x) = \mathbf{P}(N(0,1) \geq x)$. Since $Q(x) \leq e^{-x^2/2}$ for $x \geq 0$, we have if $L_c(\mathbf{X}) < 0$,

$$\begin{aligned} & Q\left(-\frac{\epsilon}{2c\sqrt{2\log(1.25/\delta)}}L_c(\mathbf{X})\right) \\ \leq & \exp\left(-\frac{\epsilon^2}{16c^2\log(1.25/\delta)}L_c(\mathbf{X})^2\right) \\ \leq & \exp\left(\frac{L_c(\mathbf{X})}{2} + \frac{c^2\log(1.25/\delta)}{\epsilon^2}\right). \end{aligned}$$

If $L_c(\mathbf{X}) > 0$,

$$\begin{aligned} & \exp\left(\tfrac{L_c(\mathbf{X})}{2} + \tfrac{c^2\log(1.25/\delta)}{\epsilon^2}\right) \geq 1 \\ & \geq Q\left(-\tfrac{\epsilon}{2c\sqrt{2\log(1.25/\delta)}}L_c(\mathbf{X})\right). \end{aligned}$$

Thus we have the following bounds for the error probabilities of the test $\psi_{(\epsilon)}(\mathbf{X})$:

$$\begin{aligned} & \mathbf{P}(\psi_{(\epsilon)}(\mathbf{X}) = 1|\mathcal{H}_0) = \mathbf{E}_0(h(\mathbf{X})) \\ \leq & \mathbf{E}_0 \exp\left(\frac{L_c(\mathbf{X})}{2} + \frac{c^2\log(1.25/\delta)}{\epsilon^2}\right), \\ & \mathbf{P}(\psi_{(\epsilon)}(\mathbf{X}) = 0|\mathcal{H}_1^{(n)}) = \mathbf{E}_1(1 - h(\mathbf{X})) \\ \leq & \mathbf{E}_1 \exp\left(-\frac{L_c(\mathbf{X})}{2} + \frac{c^2\log(1.25/\delta)}{\epsilon^2}\right). \end{aligned}$$

Note if we choose $c \geq -\log(1-\lambda)$, then $\log(1 - \lambda + \lambda\frac{g(x)}{p(x)}) \geq \log(1-\lambda) \geq -c$ for any $x \in \mathbb{R}$, which yields the following theorem about the detectability of noisy clamped log-likelihood ratio test $\psi_{(c)}$.

Theorem 2. *For any $c \geq -\log(1-\lambda)$, the error probabilities of the noisy clamped log-likelihood ratio test $\psi_{(c)}$ defined in (10) satisfy*

$$\begin{aligned} & \mathbf{P}(\psi_{(c)}(\mathbf{X}) = 1|\mathcal{H}_0) \\ \leq & \left(1 - \frac{H_n^2(\beta)}{2}\right)^n \exp(\frac{c^2\log(1.25/\delta)}{\epsilon^2}), \\ & \mathbf{P}(\psi_{(c)}(\mathbf{X}) = 0|\mathcal{H}_1^{(n)}) \\ \leq & \left(1 - \frac{H_n^2(\beta)}{2} + \lambda e^{-c/2}\right)^n \exp(\frac{c^2\log(1.25/\delta)}{\epsilon^2}), \end{aligned}$$

where $H_n^2(\beta) = H^2(p, (1-n^{-\beta})p + n^{-\beta}g))$.

The proof of Theorem 2 is postponed to the Appendix. By Theorem 2, we can see if $nH_n^2(\beta) \to \infty$, or equivalently $\beta < \bar{\beta}^*$, and $\frac{\sqrt{\log(1.25/\delta)}}{n^{\beta}\epsilon} = O(1)$, the type I error probability $\mathbf{P}(\psi_{(c)}(\mathbf{X}) = 1|\mathcal{H}_0) \to 0$. Moreover, if we further let $e^{-c/2} = \lambda$, or equivalently $c = 2\beta(\log n)$, we have $\left(1 - \frac{H_n^2(\beta)}{2} + \lambda^2\right)^n \leq \frac{1}{1+n(H_n^2(\beta)/2-\lambda^2)}$, which converges to 0 as $n \to \infty$. Thus, the type II error probability $\mathbf{P}(\psi_{(c)}(\mathbf{X}) = 0|\mathcal{H}_1^{(n)}) \to 0$ if $\beta < \bar{\beta}^*$, and $\frac{\sqrt{\log(1.25/\delta)\log(n)}}{\epsilon} = O(1)$.

We should emphasize that by adding an independent Laplace noise $\text{Lap}(2c/\epsilon)$ to the clamped log-likelihood ratio statistic in (9), we can obtain an $(\epsilon, 0)$-DP test or ϵ-DP test easily. The optimality in terms of the required sample size to achieve the designed detection power of this noisy clamped likelihood ratio test has been studied in [3] for the simple hypothesis testing problem when both the null distribution and the alternative distribution are fully specified and are not depend on the sample size n. Moreover, the detection performance of this ϵ-DP test for detecting Gaussian sparse mixtures was studied in [1]. It is not surprising that the Laplace mechanism has a better detection power than the Gaussian mechanism we introduced in this paper because the variance of the added Laplace noise is $2(2c/\epsilon)^2$, which is smaller than the variance of the added Gaussian noise $2(2c/\epsilon)^2\log(1.25/\delta)$. However, it is still interesting to study

how the privacy parameters ϵ and δ affect the detection power of the Gaussian mechanism. In particular, to find the precise detection boundary on the space of $(\beta, \epsilon, \delta)$.

Since the clamped log-likelihood ratio test involves parameters in the alternative hypothesis, which may be unknown in practice, in the next section, we consider a special case when the null distribution p follows a standard normal distribution. We construct an adaptive (ϵ, δ)-DP test, which does not rely on the information of the parameters in the alternative hypothesis and also provides a sufficient condition for ϵ, δ about when it can asymptotically separate $\mathcal{H}_0$ and $\mathcal{H}_1^{(n)}$.

IV. Differentially private higher criticism test

In this section, we construct an adaptive (ϵ, δ)-differentially private test for the detection of the general sparse mixtures when the null distribution is a standard normal distribution. Our test is based on the discrete version of the higher criticism test in (4), which achieves vanishing probability of errors inside the detectable region in [2].

Given n i.i.d. samples in the database $\mathbf{x} = (x_1, \cdots, x_n) \in \mathbb{R}^n$, let $\pi_i = \mathbf{P}(N(0,1) \geq x_i)$ and

$$\mathrm{HC}_n = \max_{1 \leq i \leq n-1} \frac{[\sum_{k=1}^n I(\pi_k \leq i/n) - i]}{\sqrt{i(1-i/n)}}. \tag{12}$$

Note the higher criticism statistic HC_n in (12) can be thought of as a deterministic function of the input database $\mathbf{x} = (x_1, \cdots, x_n)$. By Lemma 1 below, we have the sensitivity of HC_n in (12) is upper bounded by $\sqrt{n/(n-1)}$. Therefore, we can construct an (ϵ, δ)-DP test by adding an independent Gaussian noise $Y \sim N(0, \frac{2n}{(n-1)\epsilon^2} \log(1.25/\delta))$ to the statistic HC_n. Specifically, our test function is defined by

$$T(\mathbf{X}) = \begin{cases} 1, & \text{if } \mathrm{HC}_n + Y \geq c, \\ 0, & \text{otherwise,} \end{cases} \tag{13}$$

where c is the critical value controlling the type I error probability.

We should emphasize that the higher criticism statistic HC_n^* in (4) is also often written by

$$\mathrm{HC}_n^* = \max_{1 \leq i \leq n} \frac{[i - n\pi_{(i)}]}{\sqrt{n\pi_{(i)}(1-\pi_{(i)})}}, \tag{14}$$

where $\pi_{(1)} \leq \pi_{(2)} \leq \cdots \leq \pi_{(n)}$ are the order statistics of π_i. However, it is not easy to construct an efficient ϵ-DP test by adding a small Laplace noise to (14) since the sensitivity of HC_n^* is large (as the order of n.).

Lemma 1. *HC_n defined in (13) is a deterministic function of the database $\mathbf{x} = (x_1, x_2, \cdots, x_n)$ with the global sensitivity upper bounded by $\sqrt{n/(n-1)}$, i.e.,*

$$\Delta(HC_n) \leq \sqrt{n/(n-1)}.$$

Proof of Lemma 1. Let $f_i(\mathbf{x}) = \frac{[\sum_{k=1}^n I(\pi_k \leq i/n) - i]}{\sqrt{i(1-i/n)}}$. We will use f_i to represent $f_i(\mathbf{x})$ for simplification. Clearly, we have $\mathrm{HC}_n = \max_{1 \leq i \leq n-1} f_i$. Let $\mathbf{x}'$ be a neighboring dataset of $\mathbf{x}$. Denote f_i' as $f_i(\mathbf{x}')$ and $\mathrm{HC}_n' = \max_{1 \leq i \leq n-1} f_i'$. Let f_{i^*} be the maximum of $f_1, f_2, \cdots, f_{n-1}$, and f_{j^*}' be the maximum of $f_1', f_2', \cdots, f_{n-1}'$, i.e., $\mathrm{HC}_n = f_{i^*}$ and $\mathrm{HC}_n' = f_{j^*}'$. Therefore,

$$\begin{aligned} \mathrm{HC}_n - \mathrm{HC}_n' &= f_{i^*} - f_{j^*}' \leq f_{i^*} - f_{i^*}' + (f_{i^*}' - f_{j^*}') \\ &\leq f_{i^*} - f_{i^*}' \leq \Delta(f_{i^*}), \qquad (15) \\ \mathrm{HC}_n - \mathrm{HC}_n' &= f_{i^*} - f_{j^*}' \geq (f_{i^*} - f_{j^*}) + f_{j^*} - f_{j^*}' \\ &\geq f_{j^*} - f_{j^*}' \geq -\Delta(f_{j^*}). \qquad (16) \end{aligned}$$

Note $\Delta(f_i) = 1/\sqrt{i(1-i/n)} \leq \sqrt{n/(n-1)}$ for all $1 \leq i \leq n-1$. Therefore, we have $\Delta(\mathrm{HC}_n) \leq 1/\sqrt{i(1-i/n)} \leq \sqrt{n/(n-1)}$. □

Then, following [2], we assume the null distribution $p \sim N(0,1)$. The distribution of the non-null effect is absolutely continuous with respect to the null distribution and has the pdf g_n. Denote the log-likelihood ratio by $\ell = \log(g_n/p)$. Suppose that

$$\lim_{n \to \infty} \frac{\ell(u\sqrt{2\log n})}{\log n} = \alpha(u) \tag{17}$$

holds uniformly in $u \in \mathbb{R}$ for some measurable function $\alpha : \mathbb{R} \to \mathbb{R}$. Then by Theorem 1 of [2], $\underline{\beta}^* = \bar{\beta}^* = \beta^*$, which characterizes the detection boundary under the non-privacy-perserving scenario. In particular,

$$\beta^* = 1/2 + 0 \vee \operatorname*{ess\,sup}_{u \in \mathbb{R}} \{\alpha(u) - u^2 + \frac{u^2 \wedge 1}{2}\}. \tag{18}$$

Moreover, the following theorem shows the detection power of our proposed (ϵ, δ)-DP test $T(\mathbf{X})$.

Theorem 3. *For any $1/2 < \beta < \beta^*$, if $\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)} \to +\infty$, by choosing the critical value $c = \sqrt{3\log\log n}$, both type I and II error probabilities of our proposed test $T(\mathbf{X})$ in (13) converge to 0.*

From Theorem 3, we can see if ϵ, δ are fixed constants, then as $n \to \infty$, our proposed test will achieve vanishing probability errors for all values of $1/2 < \beta < \beta^*$ inside the detectable region in [2]. If we allow ϵ, δ to change with the sample size n, then our test still has good detection efficiency as long as $\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)} \to +\infty$.

V. Simulation

In this section, we conduct two numerical experiments to validate our theoretical results and illustrate the detection performance of our proposed test.

In the first experiment, we show the detection power of our private adaptive test T at different privacy levels ϵ when the $p \sim N(0,1), g_n \sim N(\mu, 1)$. We take $n = 10^4$ and two different choices of the sparsity parameter $\beta = 0.6, 0.75$. The corresponding $\lambda = n^{-\beta}$ and detection boundary $\mu^* = \sqrt{2\rho^*(\beta)\log(n)}$ are $(\lambda, \mu^*) = (0.004, 1.357), (0.001, 2.146)$. Then, we choose two true signal strengths $\mu = 2, 3$ to pass the detection boundary under the non-privacy-preserving context. In summary, we simulate data in two settings: $(\beta, \mu) = (0.6, 2), (0.75, 3)$. We consider the performance of our proposed test $T(\mathbf{X})$ in (13) with fixed $\alpha = 0.1$ and four differential privacy levels $\epsilon = 10^{-4}, 0.5, 1, 5$. The Receiver

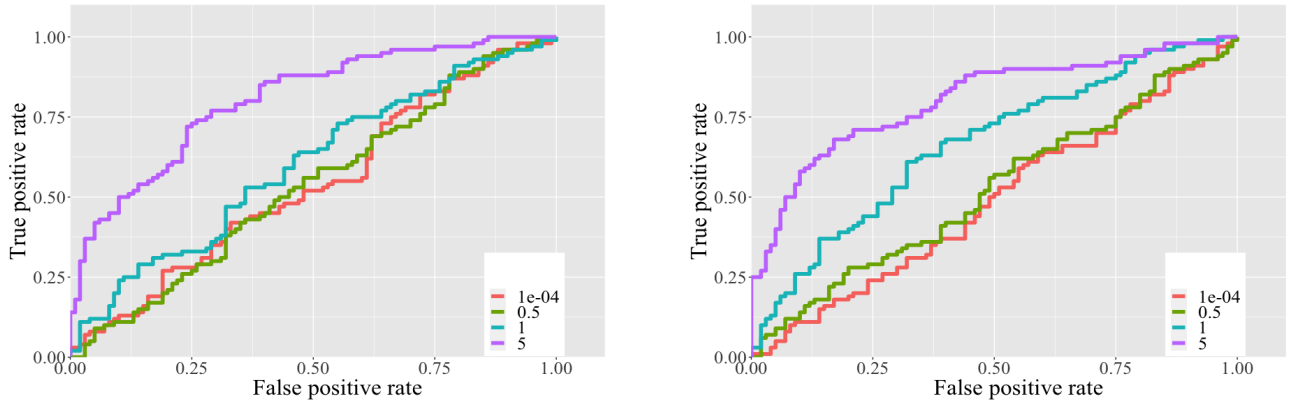

Fig. 1: ROC curves for our test T with $n = 10^4$ samples, $\alpha = 0.1$, $\epsilon = 10^{-4}, 0.5, 1, 2$, and with different β and μ. Left: $\beta = 0.6, \mu = 2$. Right: $\beta = 0.75, \mu = 3$.

Operating Characteristic (ROC) curves for our proposed test with different choices of (β, μ) are shown in Figure 1.

From these ROC curves, we can see when ϵ is very small, our test behaves like a random guess and cannot separate $\mathcal{H}_0$ from $\mathcal{H}_1^{(n)}$. This observation matches our theoretical results in Theorem 1 that any (ϵ, δ)-DP test cannot separate the null and the alternative hypothesis if the privacy parameter ϵ is too small even if the model parameter β is in the detectable region in [2]. As ϵ increases, our test shows a better detection performance. The result is not surprising since we only need to add a smaller noise to satisfy the privacy constraint when ϵ is larger. Finally, although in Theorem 3, we have shown asymptotically as $\sqrt{\log(\log n)}\epsilon \to \infty$, our test achieves the optimal detection performance, for finite sample size n, there is still room to improve the detection power. We then fix $\epsilon = 1$, let $\delta = 0.01, 0.1, 1$ and repeat the experiment. The resulting ROC plots are shown in Figure 2. From these ROC curves, we can observe a similar phenomenon: If the privacy parameter α is very small, our test cannot separate the null and alternative hypotheses. As δ increases, it will be easier to detect the sparse mixture.

We further report the AUC (Area Under the Curve) values for these methods in Table I and Table II. Based on these AUC values, we can get consistent results: if the privacy parameters ϵ or α are very small, our test cannot detect the sparse mixture even if the model parameter (β, μ) is inside of the detectable region under the non-privacy setting [2], [13]. As the privacy parameters become larger, it will be easier to detect the sparse mixture.

	ϵ value			
	10^{-4}	0.5	1	5
$\beta = 0.6, \mu = 2$	0.54	0.54	0.6	0.81
$\beta = 0.75, \mu = 3$	0.49	0.52	0.67	0.8

TABLE I: AUC values for our adaptive private test T in (13) when $\alpha = 0.1$.

In the second experiment, we compare the performance of the following tests:

1) T: our proposed adaptive $(\epsilon.\delta)$-DP test T in (13).

	α value		
	0.01	0.1	1
$\beta = 0.6, \mu = 2$	0.61	0.66	0.81
$\beta = 0.75, \mu = 3$	0.65	0.59	0.78

TABLE II: AUC values for our adaptive private test T in (13) when $\epsilon = 0.1$.

2) ncLLR: The noisy clamped log-likelihood ratio test $\psi_{(c)}(\mathbf{X})$ in (10) with $c = 2\beta \log(n)$.
3) HC: the HC test using the statistic HC_n^* in (14) without privacy constraint.
4) nBJ: the noisy BJ test by adding a $N(0, 2/\epsilon^2 \log(1.25/\delta))$ random noise, i.e., reject the null hypothesis if $M_n + N(0, 2/\epsilon^2 \log(1.25/\delta))$ exceeds some critical value.

We use a similar setup as the previous experiment, where $n = 10^4$, $(\beta, \mu, \epsilon, \delta) = (0.6, 2, 5, 0.1), (0.75, 3, 1, 1)$. But now we set the non-null distribution $g_n \sim N(\mu, 2)$. We then simulate the ROC curves by these four methods. The results are presented in Figure 3. From these figures, we can see that for the same privacy levels (ϵ, δ), the noisy BJ test always has the worst detection performance, implying that too much noise is added to keep the data private. Although the sensitivities of the BJ statistic and the HC statistic are close to 1, for the non-privacy-preserving case, as shown in [17], by choosing the critical value as the order of $\frac{1}{\log n}$, the type I and type II error probabilities of the BJ test converges to 0. Moreover, as shown in [2], by choosing the critical value as the order of $\sqrt{\log(\log n)}$, the type I and type II error probabilities of the HC test converges to 0. Therefore, although the sensitivities for both statistics are close to 1, the relative noise added to the BJ statistic is much larger than the noise added to the HC statistic, which causes a significantly different performance in the simulation study. It is not surprising that the non-privacy-preserving HC test has the best performance. Surprisingly, the adaptive (ϵ, δ)-DP test has a better detection performance than the noisy clamped log-likelihood ratio test, which implies more noise may be added to the clamped log-likelihood ratio test.

We further report the AUC (Area Under the Curve) values for these methods in Table III. Based on these AUC values, we can get consistent results: the non-privacy-preserving HC

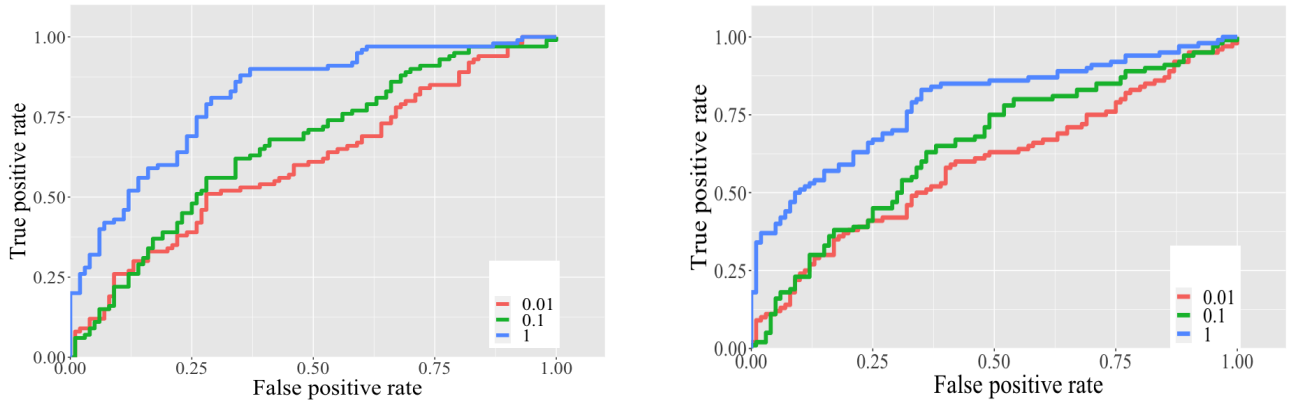

Fig. 2: ROC curves for our test T with $n = 10^4$ samples, $\epsilon = 1$, $\delta = 0.01, 0.1, 1$, and with different β and μ. Left: $\beta = 0.6, \mu = 2$. Right: $\beta = 0.75, \mu = 3$.

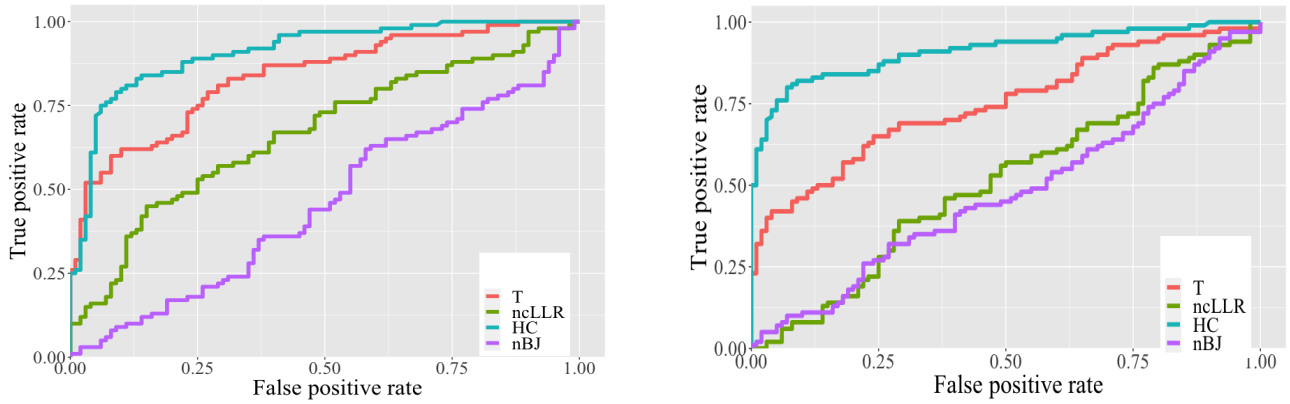

Fig. 3: ROC curves for four different tests with $n = 10^4$ samples, and with different (β, μ) and (ϵ, δ). Left: $\beta = 0.6, \mu = 2, \epsilon = 5, \delta = 0.1$. Right: $\beta = 0.75, \mu = 3, \epsilon = 1, \delta = 1$

test has the best performance. The adaptive (ϵ, δ)-DP test has a better performance than the noisy clamped log-likelihood ratio test. The noisy BJ test has the worst performance.

	T	ncLLR	HC	nBJ
$\beta = 0.6, \mu = 2, \epsilon = 5, \delta = 0.1$	0.84	0.61	0.91	0.46
$\beta = 0.75, \mu = 3, \epsilon = 1, \delta = 1$	0.75	0.51	0.91	0.48

TABLE III: AUC values for four different tests

VI. Conclusion

In conclusion, we investigate the problem of differentially private detection of sparse mixtures and provide a non-asymptotic upper bound of the sum of type I and type II error probabilities for any (ϵ, δ)-DP tests. Our result reveals the fact that no (ϵ, δ)-DP test can detect the sparse mixture efficiently if the privacy constraints are too strong such that $n\epsilon \to 0, n\delta \to 0$ or if the model parameters are in the undetectable region in [2] such that $\delta/\epsilon = O(1), \beta > \bar{\beta}^*$. We also study the performance of the noisy clamped likelihood ratio test in the context of detecting sparse mixtures and find when $\beta < \bar{\beta}^*$, and $\frac{\sqrt{\log(1.25/\delta)\log(n)}}{\epsilon} = O(1)$, the test has vanishing type I and type II error probabilities. Additionally, for the detection of general sparse mixtures when the null distribution is a standard normal distribution, we propose a new $(\epsilon, \delta)-$ differentially private test by using a specific variant of the higher criticism statistic. Then, we find when the model parameter $\beta < \bar{\beta}^*$ is in the detectable region in [2] and the privacy parameters satisfy $\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)} \to +\infty$, the test has vanishing type I and type II error probabilities. We also conducted several simulations and the numerical results are consistent with our theoretical results.

An interesting future direction would be to explore the precise detectable boundary for the model and privacy parameters: That means whether we can find a region of these parameters such that any (ϵ, δ)-DP tests cannot separate the null hypothesis and the alternative hypothesis outside of the region; while there exist some tests having vanishing error probabilities inside of the region.

Acknowledgments

The authors greatly thank the editor, the associate editor, and anonymous referees for many constructive comments and suggestions, which greatly improved the quality of the paper.

Appendix

Proof of Theorem 1. Suppose we have n independent and identically distributed (i.i.d.) random samples

$X_1, X_2, \cdots, X_n$. We consider the following hypothesis testing problem with the null hypothesis as

$$\mathcal{H}_0 : X_i \overset{i.i.d.}{\sim} p, \tag{19}$$

and the alternative hypothesis as

$$\mathcal{H}_1^{(n)} : X_i \overset{i.i.d.}{\sim} (1-\lambda)p + \lambda g, \tag{20}$$

where $\lambda = n^{-\beta}$. Let $q = (1-\lambda)p + \lambda g$.

Then we have for any (ϵ, δ)-DP test T,

$$\begin{aligned}
&\mathbf{P}(T(X) = 1|\mathcal{H}_1^{(n)}) - \mathbf{P}(T(X) = 1|\mathcal{H}_0) \qquad (21)\\
&= \int_A \mathbf{P}_T(T(\mathbf{x}) = 1)\left(q(x_1)...q(x_n) - p(x_1)...p(x_n)\right) d\mathbf{x}\\
&- \int_{A^c} \mathbf{P}_T(T(\mathbf{x}) = 1)\left(p(x_1)...p(x_n) - q(x_1)...q(x_n)\right) d\mathbf{x},
\end{aligned}$$

where $A := \{(x_1, x_2, \cdots .x_n) : q(x_1)...q(x_n) > p(x_1)...p(x_n)\}, \mathbf{P}_T(.)$ denotes that the probability is over the randomness of the randomized algorithm T. Suppose $\sup_{\mathbf{x}} \mathbf{P}_T(T(\mathbf{x}) = 1) = a \in [0, 1]$ and $\mathbf{P}_T(T(\mathbf{x}^*) = 1) = a$. Then for any dataset $\mathbf{x}$, we can find a sequence of $n+1$ datasets $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \cdots, \mathbf{x}^{(n)}$ such that $\mathbf{x}^{(0)} = \mathbf{x}^*, \mathbf{x}^{(n)} = \mathbf{x}$, and each pair of datasets $(\mathbf{x}^{(i)}, \mathbf{x}^{(i+1)})$ are adjacent dataset, i.e., have at most one different element. Then by the property of (ϵ, δ)-DP, we have for any $0 \le i \le n-1$,

$$\mathbf{P}_T(T(\mathbf{x}^{(i)}) = 1) \le e^{\epsilon}\mathbf{P}_T(T(\mathbf{x}^{(i+1)}) = 1) + \delta,$$

or equivalently,

$$(\mathbf{P}_T(T(\mathbf{x}^{(i)}) = 1) + \frac{\delta}{e^\epsilon - 1})e^{-\epsilon} \le \mathbf{P}_T(T(\mathbf{x}^{(i+1)}) = 1) + \frac{\delta}{e^\epsilon - 1},$$

which yields

$$(a + \frac{\delta}{e^\epsilon - 1})e^{-\epsilon n} - \frac{\delta}{e^\epsilon - 1} \le \mathbf{P}_T(T(\mathbf{x}) = 1) \le a.$$

Moreover, we have $\inf_{\mathbf{x}} \mathbf{P}_T(T(\mathbf{x}) = 0) = 1 - \mathbf{P}_T(T(\mathbf{x}^*) = 1) = 1 - a \in [0, 1]$. Similarly, we can get for any dataset $\mathbf{x}$,

$$1 - a \le \mathbf{P}_T(T(\mathbf{x}) = 0) \le e^{\epsilon n}(1 - a + \frac{\delta}{e^\epsilon - 1}) - \frac{\delta}{e^\epsilon - 1}.$$

Therefore,

$$\begin{aligned}
a \ge& \mathbf{P}_T(T(\mathbf{x}) = 1)\\
\ge& \max\big((a + \frac{\delta}{e^\epsilon - 1})e^{-\epsilon n} - \frac{\delta}{e^\epsilon - 1},\\
& 1 - e^{\epsilon n}(1 - a + \frac{\delta}{e^\epsilon - 1}) + \frac{\delta}{e^\epsilon - 1}\big) \qquad (22)
\end{aligned}$$

Note the two terms on the right-hand side of (22) are linear functions of α, which implies the maximum value depends on the value of α so that the two terms are equal. Specifically, we have

$$\begin{aligned}
&\mathbf{P}_T(T(\mathbf{x}) = 1) \ge 1 - e^{\epsilon n}(1 - a + \frac{\delta}{e^\epsilon - 1}) + \frac{\delta}{e^\epsilon - 1},\\
&\text{if } a \ge \frac{e^{\epsilon n}}{1 + e^{\epsilon n}} + \left(\frac{\delta}{e^\epsilon - 1}\right)\frac{e^{\epsilon n} - 1}{1 + e^{\epsilon n}}.\\
&\mathbf{P}_T(T(\mathbf{x}) = 1) \ge (a + \frac{\delta}{e^\epsilon - 1})e^{-\epsilon n} - \frac{\delta}{e^\epsilon - 1},\\
&\text{if } a \le \frac{e^{\epsilon n}}{1 + e^{\epsilon n}} + \left(\frac{\delta}{e^\epsilon - 1}\right)\frac{e^{\epsilon n} - 1}{1 + e^{\epsilon n}}.
\end{aligned}$$

Thus, by (21),

$$\begin{aligned}
&\mathbf{P}(T(X) = 1|\mathcal{H}_1^{(n)}) - \mathbf{P}(T(X) = 1|\mathcal{H}_0)\\
&\le \left(a - 1 + e^{\epsilon n}(1 - a + \frac{\delta}{e^\epsilon - 1}) - \frac{\delta}{e^\epsilon - 1}\right) TV(p^n, q^n),\\
&\text{if } a \ge \frac{e^{\epsilon n}}{1 + e^{\epsilon n}} + \left(\frac{\delta}{e^\epsilon - 1}\right)\frac{e^{\epsilon n} - 1}{1 + e^{\epsilon n}}.\\
&\mathbf{P}(T(X) = 1|\mathcal{H}_1^{(n)}) - \mathbf{P}(T(X) = 1|\mathcal{H}_0)\\
&\le \left(a - (a + \frac{\delta}{e^\epsilon - 1})e^{-\epsilon n} + \frac{\delta}{e^\epsilon - 1}\right) \text{TV}(p^n, q^n),\\
&\text{if } a \le \frac{e^{\epsilon n}}{1 + e^{\epsilon n}} + \left(\frac{\delta}{e^\epsilon - 1}\right)\frac{e^{\epsilon n} - 1}{1 + e^{\epsilon n}},
\end{aligned}$$

which implies

$$\begin{aligned}
&\mathbf{P}(T(X) = 1|\mathcal{H}_1^{(n)}) - \mathbf{P}(T(X) = 1|\mathcal{H}_0)\\
\le& \frac{e^{\epsilon n} - 1}{e^{\epsilon n} + 1}\left(1 + 2\frac{\delta}{e^\epsilon - 1}\right) \text{TV}(p^n, q^n).
\end{aligned}$$

Here $\text{TV}(p, q)$ denote the total variation distance of distributions p, q, which is defined by $\text{TV}(p, q) := \frac{1}{2}\int |p(x) - q(x)|dx$. Let $H^2(p, q)$ denote the Hellinger distance of two distributions p, q, which is defined by

$$H^2(p, q) = \int \left(\sqrt{p(x)} - \sqrt{q(x)}\right)^2 dx. \tag{23}$$

Note $H^2(p, q)$ satisfies the following relationship [29]:

$$\frac{1}{2}H^2(p, q) \le \text{TV}(p, q) \le H(p, q)\sqrt{1 - \frac{H^2(p, q)}{4}} \le 1. \tag{24}$$

By the fact that the Hellinger distance tensorizes under the product measures in [2], we can get

$$H^2(p^n, q^n) = 2 - 2\left(1 - \frac{H^2(p, q)}{2}\right)^n. \tag{25}$$

Denote

$$H_n^2(\beta) = H^2(p, (1 - n^{-\beta})p + n^{-\beta}g_n)). \tag{26}$$

By (24) and (25), we have

$$\text{TV}(p^n, q^n) \le H(p^n, q^n) \le \sqrt{\min\left(2, nH_n^2(\beta)\right)}, \tag{27}$$

which completes the proof. □

Proof of Theorem 2. Let $c = \frac{\lambda}{1-\lambda}$, we have the log-likelihood ratio $\log(1 - \lambda + \lambda\frac{g(x)}{p(x)}) \ge -c$. Thus we have $\mathbf{E}_0 \exp\left(\frac{L_c(\mathbf{X})}{2} + \frac{c^2 \log(1.25/\delta)}{\epsilon^2}\right) \le \mathbf{E}_0 \exp\left(\frac{L_\infty(\mathbf{X})}{2} + \frac{c^2 \log(1.25/\delta)}{\epsilon^2}\right)$. Note $L_\infty(\mathbf{X})$ is the log-likelihood ratio statistic, we have $\mathbf{E}_0 \exp\left(\frac{L_\infty(\mathbf{X})}{2}\right) = (1 - H_n^2(\beta)/2)^n$. Thus, we have the type I error probability satisfies

$$\mathbf{P}(\psi_{(c)}(\mathbf{X}) = 1|\mathcal{H}_0) \le \left(1 - H_n^2(\beta)/2\right)^n \exp(\frac{c^2 \log(1.25/\delta)}{\epsilon^2}).$$

For the type II error, let $\ell_c(X_i) = \left[\log(1 - \lambda + \lambda\frac{g(X_i)}{p(X_i)})\right]_{-c}^{c}$. By the change of measure, we have

$$\mathbf{E}_1(e^{-\frac{1}{2}L_c(\mathbf{X})}) = \mathbf{E}_0(e^{L_\infty(\mathbf{X}) - \frac{1}{2}L_c(\mathbf{X})}) = \Pi_{i=1}^n \mathbf{E}_0(e^{\ell_\infty(X_i) - \frac{1}{2}\ell_c(X_i)})$$

Note

$$
\begin{aligned}
&\mathbf{E}_0(e^{\ell_\infty(X_i)-\frac{1}{2}\ell_c(X_i)})\\
=&\int_{\{x_i:\ell_\infty(x_i)\le c\}}(e^{\ell_\infty(x_i)-\frac{1}{2}\ell_c(x_i)})p(x_i)dx_i\\
&+\int_{\{x_i:\ell_\infty(x_i)\ge c\}}(e^{\ell_\infty(x_i)-\frac{1}{2}\ell_c(x_i)})p(x_i)dx_i\\
=&\int_{\{x_i:\ell_\infty(x_i)\le c\}}(e^{\frac{1}{2}\ell_\infty(x_i)})p(x_i)dx_i\\
&+\int_{\{x_i:\ell_\infty(x_i)\ge c\}}(e^{\ell_\infty(X_i)-\frac{1}{2}c})p(x_i)dx_i\\
=&\int_{\{x_i:\ell_\infty(x_i)\le c\}}(e^{\frac{1}{2}\ell_\infty(x_i)})p(x_i)dx_i+e^{-c/2}\mathbf{P}_1(\ell_\infty(X_i)\ge c)\\
=&\mathbf{E}_0(e^{\frac{1}{2}\ell_\infty(X_i)})+e^{-c/2}\mathbf{P}_1(\ell_\infty(X_i)\ge c)\\
&-\int_{\{x_i:\ell_\infty(x_i)\ge c\}}(e^{\frac{1}{2}\ell_\infty(x_i)})p(x_i)dx_i\\
\le&\mathbf{E}_0(e^{\frac{1}{2}\ell_\infty(X_i)})+e^{-c/2}\mathbf{P}_1(\ell_\infty(X_i)\ge c)\\
&-e^{c/2}\mathbf{P}_0(\ell_\infty(X_i)\ge c)\\
\le&\mathbf{E}_0(e^{\frac{1}{2}\ell_\infty(X_i)})+e^{-c/2}\left((1-\lambda)\mathbf{P}_0(\ell_\infty(X_i)\ge c)+\lambda\right)\\
&-e^{c/2}\mathbf{P}_0(\ell_\infty(X_i)\ge c)\\
=&\mathbf{E}_0(e^{\frac{1}{2}\ell_\infty(X_i)})+\left(e^{-c/2}(1-\lambda)-e^{c/2}\right)\mathbf{P}_0(\ell_\infty(X_i)\ge c)\\
&+\lambda e^{-c/2}\\
\le&1-\frac{H_n^2(\beta)}{2}+\lambda e^{-c/2}.
\end{aligned}
$$

Therefore, the type II error probability satisfies

$$
\begin{aligned}
&\mathbf{P}(\psi_{(c)}(\mathbf{X})=0|\mathcal{H}_1^{(n)})\\
\le&\left(1-\frac{H_n^2(\beta)}{2}+\lambda e^{-c/2}\right)^n\exp(\frac{c^2\log(1.25/\delta)}{\epsilon^2}).
\end{aligned}
$$

□

Proof of Theorem 3. We first bound the difference between HC_n and HC_n^*. Define $f(t)=\frac{[\sum_{k=1}^n I(\pi_k\le t)-nt]}{\sqrt{nt(1-t)}}$. We have

$$\mathrm{HC}_n=\max_{1\le i\le n-1}(f(i/n))\le\sup_{t\in(0,1)}f(t)=\sup_{1\le i\le n}f(\pi_{(i)})=\mathrm{HC}_n^*.$$

Denote $i^*=\arg\max_{1\le i\le n}f(\pi_{(i)})$ and assume $k/n\le\pi_{(i^*)}<(k+1)/n$ for some $k\in\{1,\cdots,n-1\}$. Then, we have

$$
\begin{aligned}
0 &\le \mathrm{HC}_n^*-\mathrm{HC}_n\le f(\pi_{(i^*)})-f(\frac{k+1}{n})\\
&\le \frac{i^*-n\pi_{i^*}}{\sqrt{n\pi_{i^*}(1-\pi_{i^*})}}-\frac{i^*-(k+1)}{\sqrt{(k+1)(1-\frac{k+1}{n})}}\\
&\le \sqrt{n/(n-1)}.
\end{aligned}
$$

Let Y denote the Gaussian random variable $N(0,\frac{2n}{(n-1)\epsilon^2}\log(1.25/\delta)$ that is independent to the database $\mathbf{X}$. By the asymptotic property of the higher criticism statistic HC_n^* [2], under the null hypothesis, the Type I error

$$
\begin{aligned}
&\mathbf{P}(\mathrm{HC}_n+Y\ge\sqrt{3\log\log n}|\mathcal{H}_0)\\
\le&\ \mathbf{P}\left(\mathrm{HC}_n^*\ge\sqrt{2.5\log\log n}-\sqrt{n/(n-1)}|\mathcal{H}_0\right)\\
+&\ \mathbf{P}\left(Y\ge(\sqrt{3}-\sqrt{2.5})\sqrt{\log\log n}\right)\\
\le&\ \exp\left(-\frac{(\sqrt{3}-\sqrt{2.5})^2(n-1)}{4n}\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)}\right),
\end{aligned}
$$

which converges to 0 as $\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)}\to+\infty$. Moreover, as shown in [2], if $1/2<\beta<\beta^*$, $\mathbf{P}(\mathrm{HC}_n^*\le 2\sqrt{\log\log n}|\mathcal{H}_1^{(n)})\to 0$. Therefore, the Type II error

$$
\begin{aligned}
&\mathbf{P}(\mathrm{HC}_n+Y\le\sqrt{3\log\log n}|\mathcal{H}_1^{(n)})\\
\le&\ \mathbf{P}\left(\mathrm{HC}_n^*\le 2\sqrt{\log\log n}+\sqrt{n/(n-1)}|\mathcal{H}_1^{(n)}\right)\\
+&\ \mathbf{P}\left(Y\le-(2-\sqrt{3})\sqrt{\log\log n}\right)\\
\le&\ \exp\left(-\frac{(2-\sqrt{3})^2(n-1)}{4n}\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)}\right),
\end{aligned}
$$

which converges to 0 as long as $\frac{\log(\log n)\epsilon^2}{\log(1.25/\delta)}\to+\infty$. □

REFERENCES

[1] R. Zhang and S. Cao, "Differentially private detection of sparse gaussian mixtures," in *2023 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2023, pp. 2637–2642.

[2] T. T. Cai and Y. Wu, "Optimal detection of sparse mixtures against a given null distribution," *IEEE Transactions on Information Theory*, vol. 60, no. 4, pp. 2217–2232, 2014.

[3] C. L. Canonne, G. Kamath, A. McMillan, A. Smith, and J. Ullman, "The structure of optimal private tests for simple hypotheses," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019, pp. 310–321.

[4] Y. Ingster and I. A. Suslina, *Nonparametric goodness-of-fit testing under Gaussian models*. Springer Science & Business Media, 2012, vol. 169.

[5] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial fraud: a review of anomaly detection techniques and recent advances," *Expert systems With applications*, vol. 193, p. 116429, 2022.

[6] H. Yan, K. Paynabar, and J. Shi, "Anomaly detection in images with smooth background via smooth-sparse decomposition," *Technometrics*, vol. 59, no. 1, pp. 102–114, 2017.

[7] M. Kulldorff, R. Heffernan, J. Hartman, R. Assunçao, and F. Mostashari, "A space–time permutation scan statistic for disease outbreak detection," *PLoS medicine*, vol. 2, no. 3, 2005.

[8] R. Prasath, K. Skenes, and S. Danyluk, "Comparison of phase shifting techniques for measuring in-plane residual stress in thin, flat silicon wafers," *Journal of electronic materials*, vol. 42, pp. 2478–2485, 2013.

[9] Y. Xie, J. Huang, and R. Willett, "Change-point detection for high-dimensional time series with missing data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 1, pp. 12–27, 2012.

[10] H. Yan, K. Paynabar, and J. Shi, "Real-time monitoring of high-dimensional functional data streams via spatio-temporal smooth sparse decomposition," *Technometrics*, vol. 60, no. 2, pp. 181–197, 2018.

[11] P. Qiu, W. Li, and J. Li, "A new process control chart for monitoring short-range serially correlated data," *Technometrics*, vol. 62, no. 1, pp. 71–83, 2020.

[12] Y. I. Ingster, "Minimax detection of a signal for ℓ_n^p-balls," *Mathematical Methods of Statistics*, vol. 7, no. 4, pp. 401–428, 1998.

[13] D. Donoho and J. Jin, "Higher criticism for detecting sparse heterogeneous mixtures," *The Annals of Statistics*, vol. 32, no. 3, pp. 962–994, 2004.

[14] ——, "Higher criticism thresholding: Optimal feature selection when useful features are rare and weak," *Proceedings of the National Academy of Sciences*, vol. 105, no. 39, pp. 14 790–14 795, 2008.

[15] L. Jager and J. A. Wellner, "Goodness-of-fit tests via phi-divergences," *The Annals of Statistics*, vol. 35, no. 5, pp. 2018–2053, 2007.

[16] M. Ditzhaus, "Signal detection via phi-divergences for general mixtures," *Bernoulli*, vol. 25, no. 4A, pp. 3041–3068, 2019.

[17] A. Moscovich, B. Nadler, and C. Spiegelman, "On the exact berk-jones statistics and their *p*-value calculation," *Electronic Journal of Statistics*, vol. 10, no. 2, pp. 2329–2354, 2016.

[18] T. Tony Cai, X. Jessie Jeng, and J. Jin, "Optimal detection of heterogeneous and heteroscedastic mixtures," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 73, no. 5, pp. 629–662, 2011.

[19] J. G. Ligo, G. V. Moustakides, and V. V. Veeravalli, "Rate analysis for detection of sparse mixtures," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4244–4248.

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

[21] T. T. Cai, Y. Wang, and L. Zhang, "The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy," *The Annals of Statistics*, vol. 49, no. 5, pp. 2825–2850, 2021.

[22] M. Avella-Medina, "Privacy-preserving parametric inference: a case for robust statistics," *Journal of the American Statistical Association*, pp. 1–15, 2020.

[23] Y. Wang, J. Lee, and D. Kifer, "Revisiting differentially private hypothesis tests for categorical data," *arXiv preprint arXiv:1511.03376*, 2015.

[24] K. H. Degue and J. Le Ny, "On differentially private gaussian hypothesis testing," in *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2018, pp. 842–847.

[25] C. Dwork, W. Su, and L. Zhang, "Differentially private false discovery rate control," *Journal of Privacy and Confidentiality*, vol. 11, no. 2, 2021.

[26] C. L. Canonne, G. Kamath, A. McMillan, J. Ullman, and L. Zakynthinou, "Private identity testing for high-dimensional distributions," *Advances in Neural Information Processing Systems*, vol. 33, pp. 10 099–10 111, 2020.

[27] R. H. Berk and D. H. Jones, "Goodness-of-fit test statistics that dominate the kolmogorov statistics," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 47, no. 1, pp. 47–59, 1979.

[28] 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.

[29] L. Le Cam, *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.