

DIFFERENTIALLY PRIVATE TESTING FOR RELEVANT DEPENDENCIES IN HIGH DIMENSIONS

BY PATRICK BASTIAN^{1,*} HOLGER DETTE^{2,†} MARTIN DUNSCH^{2,‡}

¹*Department of Mathematics, Aarhus University*

²*Fakultät für Mathematik, Ruhr-University Bochum*

We investigate the problem of detecting dependencies between the components of a high-dimensional vector. Our approach advances the existing literature in two important respects. First, we consider the problem under privacy constraints. Second, instead of testing whether the coordinates are pairwise independent, we are interested in determining whether certain pairwise associations between the components (such as all pairwise Kendall's τ coefficients) do not exceed a given threshold in absolute value. Considering hypotheses of this form is motivated by the observation that in the high-dimensional regime, it is rare and perhaps impossible to have a null hypothesis that can be modeled exactly by assuming that all pairwise associations are precisely equal to zero.

The formulation of the null hypothesis as a composite hypothesis makes the problem of constructing tests already non-standard in the non-private setting. Additionally, under privacy constraints, state of the art procedures rely on permutation approaches that are rendered invalid under a composite null. We propose a novel bootstrap based methodology that is especially powerful in sparse settings, develop theoretical guarantees under mild assumptions and show that the proposed method enjoys good finite sample properties even in the high privacy regime. Additionally, we present applications in medical data that showcase the applicability of our methodology.

1. Introduction. The ability to measure and detect statistical dependence lies at the heart of many scientific questions. From the early works of [Pearson \(1920\)](#); [Kendall \(1938\)](#) to modern machine learning methods, independence testing has been applied for this purpose across many fields such as genetics, neuroscience, medicine or economics. Classical approaches target low-dimensional settings where $p \ll n$ and generally do not perform well when p is comparable to, or even larger than n . In the era of big data, attention has shifted toward this high-dimensional regime, making classical methods inadequate. Consequently, specialized procedures for such high-dimensional settings have been developed and for a comprehensive review of the state of the art we refer the reader to the review at the end of this section.

However, to the best of our knowledge, two crucial aspects are almost entirely missing from the current literature on high-dimensional independence testing. The first one is privacy. In many fields, including medical studies, behavioral research, and other contexts where sensitive personal information is involved, statistical analysis must reconcile fundamentally competing aims: drawing meaningful conclusions from data while protecting the privacy of individuals who contribute to it. Traditional anonymization or aggregation techniques are often

*E-mail: patrick.bastian@rub.de

†E-mail: holger.dette@rub.de

‡E-mail: martin.dunsche@rub.de

§ Authors contributed equally and are listed alphabetically.

Keywords and phrases: U-statistics, Differential privacy, Relevant hypotheses, Dependence testing, high-dimensional inference.

insufficient, as subtle differences in released statistics can still reveal individual participation. Differential privacy (DP) (Dwork, 2006) addresses this challenge by providing a formal and quantifiable notion of privacy. It ensures that the inclusion or exclusion of a single observation has only a limited effect on the output, thereby bounding the information that can be inferred about any individual. This framework has become the de-facto standard in privacy-preserving data analysis, offering a principled foundation for statistical inference under privacy constraints. The second is more of statistical nature. Most existing methods aim for detecting arbitrarily small dependencies between the components of high-dimensional vectors and are based on tests for hypotheses

$$(1.1) \quad \begin{aligned} H_0 : \theta_{ij} &= 0 \quad \text{for all } 1 \leq i < j \leq d \\ H_1 : \theta_{ij} &\neq 0 \quad \text{for at least one pair } (i, j) \text{ with } 1 \leq i < j \leq d, \end{aligned}$$

where θ_{ij} is some (population) measure of dependence between i th and j th component (such as the covariance or Kendall's τ) of a d -dimensional vector. However, in the big data era, where the dimension d and the sample size is large, hypotheses of this form are usually not of interest in practice. In the high-dimensional regime it is very uncommon that all pairwise components of a vector are independent. Consequently, using a consistent test with a sufficient amount of data will virtually always result in rejection of (1.1). This point of view is in line with Tukey (1991), who succinctly stated it in the context of the comparison of multiple means: “*Statisticians classically asked the wrong question—and were willing to answer with a lie, one that was often a downright lie. They asked “Are the effects of A and B different?” and they were willing to answer “no.”*”

As pointed out by Berger and Delampady (1987) in the context of comparing two univariate means, a possible remedy for these issues is to consider the null hypothesis that all pairwise associations do not exceed a given threshold, say $\Delta > 0$. More specifically, we will consider the hypotheses

$$(1.2) \quad \begin{aligned} H_0(\Delta) : |\theta_{ij}| &\leq \Delta \quad \text{for all } 1 \leq i < j \leq d, \\ H_1(\Delta) : |\theta_{ij}| &> \Delta \quad \text{for at least one pair } (i, j) \text{ with } 1 \leq i < j \leq d. \end{aligned}$$

Thus we are interested in testing for at least one *practically significant* association, which is larger (in absolute value) than the threshold Δ , and we call the hypotheses in (1.2) *relevant* hypotheses throughout this paper. For a more thorough discussion of relevant hypotheses we refer the reader to Section 2.1. To illustrate our point of view here more concretely, we refer to an example of gene expression networks discussed in Tsaparas et al. (2006), where a 28-dimensional gene expression network is evaluated based on Pearson correlation coefficients to determine co-expressed genes. For a meaningful analysis it is necessary to prevent congelation of the network due to spurious correlations, corresponding to a threshold Δ that demarcates spurious and meaningful correlations (they choose $\Delta = 0.7$ even in this fairly low dimensional setting). We discuss this example in more detail in Section 4.

In the high-dimensional setting, either of these two aspects alone already poses considerable challenges for the development of statistical methodology (see the discussion of the related literature below). In this paper, we develop methodology that addresses both aspects. To this end, we design a framework for testing for dependencies in high-dimensional vectors under DP which accommodates a broad class of commonly used dependence measures and is therefore applicable to a wide range of settings. One particularly challenging problem of this endeavor will be the private estimation of so called “*extremal sets*” that plays a crucial role when constructing powerful tests for hypotheses of the form (1.2). As we will see in Section 3, existing approaches to related problems either perform poorly in finite samples or are not practically feasible. The main contributions of the present paper are the following:

- (1) We introduce a practical framework for testing for relevant dependencies between the components of a high-dimensional vector that achieves good finite-sample performance combined with strong privacy guarantees.
- (2) We establish rigorous statistical guarantees for our approach that enable further generalizations and extensions.
- (3) We demonstrate the versatility of the developed method on two high-dimensional data sets, illustrating that meaningful inference under privacy constraints is still possible in extremely high-dimensional settings.
- (4) The new procedure preserves privacy and additionally attains finite-sample performance that is at least comparable - if not better than that of existing non-private approaches. As the primary focus of this article is private inference we demonstrate this superiority in Section C of the online supplement.
- (5) We release an open-source implementation.¹

Related work High-dimensional (in-)dependence testing is by now a fairly mature field with a substantial amount of publications. For Gaussian data, we refer, among many others, to Jiang and Qi (2015) and Bodnar et al. (2019) who investigated the asymptotic properties of likelihood ratio tests. For more general distributions, dependence is usually quantified by different correlation measures such as Pearson’s r , Spearman’s ρ , and Kendall’s τ , and different functions are used to aggregate these estimates of the pairwise dependencies. For example, Bao et al. (2015) and Li et al. (2021) use linear spectral statistics of the matrix of the estimates of the pairwise dependencies, while, among others, Yao et al. (2018) and Leung and Drton (2018) propose tests based on the Frobenius norm. Other very popular methods of aggregating estimates of the pairwise dependencies are maximum-type tests, which have good power properties against sparse alternatives and have been investigated for various covariance/correlation statistics, see Han et al. (2017), Drton et al. (2020) and He et al. (2021) for some recent work and the references therein. A common aspect of this literature is that the authors consider the hypotheses in (1.1) of exact pairwise independence. Their methodology is based on the asymptotic distribution of a test statistic under the hypothesis of independence and can therefore not be extended to testing relevant hypotheses of the form (1.2). This problem has been addressed by Bastian et al. (2024), but, to our best knowledge, there does not exist any work on differentially private inference in this context. We substantially deviate from the methodology in the aforementioned reference, by introducing so called “*extremal sets*” and corresponding estimators. Using this information we implement a powerful parametric bootstrap test, which accommodates privacy constraints and exhibits superior performance compared to the test of Bastian et al. (2024), even in many situations where private inference is not required.

A substantial body of research has investigated differentially private inference procedures across a wide range of statistical problems (see, for example, Rogers and Kifer, 2017; Sei and Ohsuga, 2021). Most of this literature has focused on testing tasks for specific parameters, such as sample means, rather than on methods that are broadly applicable to general classes of hypothesis testing problems. Chaudhuri et al. (2024a) studied U -statistics to enhance private estimation and, in turn, inference in finite dimensional settings. In a closer vein to the present paper Liu et al. (2025) analyzed U -statistics with an emphasis on independence testing, making use of classical permutation tests. Unfortunately this approach is not applicable in the relevant hypothesis framework as the necessary permutation invariances are not valid under the null hypothesis (1.2), which we will consider.

The majority of work on DP hypothesis testing does not consider the high-dimensional regime and relies on parametric bootstrap approximations of the (quantiles of the) relevant

¹https://github.com/martindunsche/Highdimensional_U_statistics_under_privacy

statistic, or by direct analysis of the noise introduced for privacy protection (see [Dunsche et al., 2022](#), for an overview). On the other hand, [Liu et al. \(2022\)](#) proposed a general framework for DP estimation in high dimensions. While obtaining strong theoretical results, their methodology does not yield inferential guarantees, which can be used for hypotheses testing, and also may not be computationally feasible in many cases. Moreover, [Cai et al. \(2024\)](#) investigated principal component analysis (PCA) in high-dimensional spiked covariance models, while [Canonne et al. \(2020\)](#) and [Narayanan \(2022\)](#) developed a private framework for identity testing in high dimensions.

2. Background. In this section, we briefly revisit key concepts for the subsequent development of our methodology. In particular we recall the specific notions of testing relevant hypotheses, U -statistics and Differential Privacy (DP).

2.1. Relevant Hypotheses. For a d -dimensional vector $X_1 = (X_{11}, \dots, X_{1d})^\top$ let $\theta_{ij} = \theta(X_{1i}, X_{1j})$ denote a dependence measure between the i th and j th component. We propose to investigate if all associations $(\theta_{ij})_{1 \leq i < j \leq d}$ are in some sense “small” by testing the hypotheses (1.2), where $\Delta > 0$ is a given threshold, which defines when a dependence between the components i and j is considered as (scientifically) not relevant. As pointed out in the introduction, the consideration of hypotheses of this form is motivated by the observation that in many applications it is very unlikely that all pairwise associations are completely 0, in particular if the dimension d is large. We thus argue that it is more reasonable to test for at least one *practically significant* association.

An essential ingredient in this approach is the specification of the threshold Δ , and its choice depends sensitively on the particular problem under consideration. Essentially, this boils down to the important question when a correlation (or another dependence measure) is *practically significant*, which has a long history in applied statistics. For the particular case of dependence measures that we consider in the present paper several authors classify the strength of association between variables for their particular application into categories such as “small”, “medium” or “large”. The precise demarcation thresholds vary across disciplines and subject areas and we refer the interested reader to [Quintana \(2016\)](#), [Brydges \(2019\)](#) and [Lovakov and Agadullina \(2021\)](#) for a discussion of this choice for concrete applications. In this paper, we will consider two data examples, one from genomics and one from cancer research. For the genomic data and in particular gene expression networks it is customary to discard correlations below a certain threshold to facilitate a meaningful analysis. In this context [Tsaparas et al. \(2006\)](#) proposed 0.7 as a threshold for Pearson correlations and considered correlations below 0.7 as spurious. For the cancer data on the other hand, we will apply our methodology to two distinct groups within the data set - patients with cancer and those without - and then use the smallest Δ for which the null hypothesis in (1.2) is rejected as a means of structurally discriminating between the two groups. We thus obtain a natural, data dependent choice for the threshold Δ . In the following remark, we make this argument more precise and explain why this yields a valid inference procedure.

REMARK 2.1. Note that the hypotheses $H_0(\Delta)$ in (1.2) are nested and that families of test decisions $\phi(\Delta)$ for such hypotheses are often monotone in Δ . Consequently, rejecting $H_0(\Delta)$ for $\Delta = \Delta_1 > 0$ also implies rejecting $H_0(\Delta)$ for all $0 < \Delta < \Delta_1$. The sequential rejection principle then allows us to simultaneously test the hypotheses (1.2) for different choices of $\Delta > 0$ until we find the minimum value $\hat{\Delta}_\alpha$ for which $H_0(\hat{\Delta}_\alpha)$ is not rejected, that is

$$\hat{\Delta}_\alpha := \min \{ \Delta \mid \phi(\Delta) = 0 \} .$$

where we define the minimum of an empty set to be 0. Consequently, one may postpone the selection of Δ and derives a test decision in the fashion as comparing the p -value to the prescribed type I error.

2.2. U -statistics. We will phrase the testing problem (1.2) in the framework of U -statistics as many dependence measures used in practice can be expressed in this way. To be precise, let $X_1, \dots, X_n$ denote independent identically distributed d -dimensional random vectors with distribution function F . Note that formally F depends on the dimension d , which in this paper we allow to vary with n , but we will not reflect this dependence in our notation throughout this paper. For some positive integer r let

$$(2.1) \quad h = (h_1, \dots, h_p)^\top : (\mathbb{R}^d)^r \rightarrow \mathbb{R}^p$$

denote a measurable symmetric function with finite expectation

$$\theta = (\theta_1, \dots, \theta_p)^\top := \mathbb{E}_F[h(X_1, \dots, X_r)] \in \mathbb{R}^p,$$

which defines our parameter of interest. We are interested in the relevant hypotheses

$$(2.2) \quad H_0 : \|\theta\|_\infty \leq \Delta, \quad H_1 : \|\theta\|_\infty > \Delta$$

for some $\Delta > 0$, where $\|\cdot\|_\infty$ denotes the maximum-norm. In the context of testing for pairwise dependencies, the dimension p will typically be given by $p = d(d-1)/2$ as illustrated in the following example.

EXAMPLE 2.2. For the dependence measure between the i th and j th component of the vector $X_1 = (X_{11}, \dots, X_{1d})^\top$ as introduced in Section 1, we assume that

$$\theta_{ij} = \theta(X_{1i}, X_{1j}) = \mathbb{E}[\tilde{h}(X_{1i}, X_{1j}, \dots, X_{ri}, X_{rj})] \quad 1 \leq i < j \leq d.$$

Here, $\tilde{h} : (\mathbb{R}^d)^r \rightarrow \mathbb{R}$ is a kernel of order r evaluated at $(X_{1i}, X_{1j}), \dots, (X_{ri}, X_{rj})$. In this case the function $h : \mathbb{R}^{dr} \rightarrow \mathbb{R}^p$ in (2.1) is defined by

$$\begin{aligned} h(X_1, \dots, X_r) &= \text{vech}((h_{ij}(X_1, \dots, X_r))_{i,j=1,\dots,d}) \\ &= \text{vech}((\tilde{h}(X_{1i}, X_{1j}, \dots, X_{ri}, X_{rj}))_{i,j=1,\dots,d}), \end{aligned}$$

where the second equality defines the functions $h_{ij} : \mathbb{R}^{dr} \rightarrow \mathbb{R}$ in an obvious manner and $\text{vech}(\cdot)$ is the operator that stacks the columns above the diagonal of a symmetric $d \times d$ matrix as a vector with $p = d(d-1)/2$ components. Note that the index (i, j) in the definition of the function h_{ij} is only used to emphasize that each h_{ij} acts on different components of the d -dimensional vectors $X_1, \dots, X_r$. Similarly, the vector θ is defined by $\theta = \text{vech}((\theta_{ij})_{i,j=1,\dots,d})$, and the components of the vector $U = \text{vech}((U_{ij})_{i,j=1,\dots,d})$ in (2.3) are given by

$$\begin{aligned} U_{ij} &= \binom{n}{r}^{-1} \sum_{1 \leq l_1 < \dots < l_r \leq n} h_{ij}(X_{l_1}, \dots, X_{l_r}) \\ &= \binom{n}{r}^{-1} \sum_{1 \leq l_1 < \dots < l_r \leq n} \tilde{h}(X_{l_1 i}, X_{l_1 j}, \dots, X_{l_r i}, X_{l_r j}). \end{aligned}$$

Finally, we note that it is easy to see that with these notations the hypotheses (2.2) are equivalent to (1.2).

We now return to the general case discussed at the beginning of this section. In order to estimate the parameter θ we consider a U -statistic of order r given by

$$(2.3) \quad U = (U_1, \dots, U_p)^\top = \binom{n}{r}^{-1} \sum_{1 \leq l_1 < \dots < l_r \leq n} h(X_{l_1}, \dots, X_{l_r}).$$

As our primary goal is developing inferential methodology for the hypotheses (1.2), we will need estimates of the asymptotic covariance structure $\zeta_1 = (\zeta_{1,ij})_{i,j=1,\dots,p}$ of the vector U , where

$$(2.4) \quad \zeta_{1,ij} := \text{Cov}_F(h_{1,i}(X_1), h_{1,j}(X_1))$$

and

$$h_1(x) = (h_{11}(x), \dots, h_{1p}(x))^\top = \mathbb{E}_F[h(X_1, \dots, X_r) | X_1 = x]$$

is the linear part of the Hoeffding decomposition. To estimate (a multiple) of $\zeta_1 = (\zeta_{1,ij})_{i,j=1,\dots,p}$, we utilize the classical Jackknife estimator

$$(2.5) \quad \hat{\zeta}_1 := (n-1) \sum_{l=1}^n (U^{(l)} - U)(U^{(l)} - U)^\top,$$

where $U^{(l)}$ denotes the leave one out U -statistic of (2.3). In Lemma F.1 of the online supplement we show that this yields a maximum-norm consistent estimator for $r \cdot \zeta_1$ even in the ultra high-dimensional setting where $p = o(\exp(n^{1/5}))$.

2.3. Differential Privacy. We recall basic notions used throughout this paper and forward interested readers to [Dwork et al. \(2014a\)](#) and [Bun and Steinke \(2016\)](#) for an overview of DP or Zero-Concentrated Differential Privacy (zCDP), respectively.

We call two datasets X and X' neighbors (denoted $X \sim X'$ or $d_H(X, X') = 1$, where d_H is the Hamming distance) if they differ in exactly one individual. Given a possibly vector-valued query (statistic) f , its global sensitivity with respect to a norm $\|\cdot\|$ is defined as

$$\Delta f := \sup_{X \sim X'} \|f(X) - f(X')\|,$$

and quantifies the largest change in f when a single record in the data set X is modified for any X . We define for two distributions P and Q and $\alpha > 1$ their Rényi-divergence by

$$D_\alpha(P, Q) = \frac{1}{\alpha - 1} \log \left(\int p(t)^\alpha q(t)^{1-\alpha} d\mu(t) \right),$$

where p and q are densities of P and Q with respect to some dominating measure μ .

DEFINITION 2.3 (Definition 8.1 in [Bun and Steinke \(2016\)](#), Approximate zCDP). *A randomized algorithm $\mathcal{M}$ is called δ -approximate- ρ -zCDP if for all neighboring data sets X and X' , there exist events E (depending on $\mathcal{M}(X)$) and E' (depending on $\mathcal{M}(X')$) such that*

$$\mathbb{P}[E] \geq 1 - \delta \quad \text{and} \quad \mathbb{P}[E'] \geq 1 - \delta,$$

and we have for all $\alpha > 1$

$$D_\alpha(P, Q) \leq \rho\alpha \quad \text{and} \quad D_\alpha(Q, P) \leq \rho\alpha$$

where P and Q are the distributions of $\mathcal{M}(X)$ and $\mathcal{M}(X')$ conditional on E and E' , respectively.

In the case $\delta = 0$, 0-approximate- ρ -zCDP recovers the classical ρ -zCDP. By Lemma 8.2 in [Bun and Steinke \(2016\)](#), δ -approximate-zCDP satisfies the composition and post-processing property. For completeness and later use, we will also recall the definition and the privacy guarantees of the most prominent algorithm.

LEMMA 2.4. *Let T denote a $\mathbb{R}^d$ -valued statistic. The Gaussian mechanism $\mathcal{M}(X) = T(X) + \frac{\Delta_2 T}{\sqrt{2\rho}} Y$ where $Y \sim \mathcal{N}_d(0, I_{d \times d})$ and $\Delta_2 T := \sup_{X \sim X'} \|T(X) - T(X')\|_2$, preserves ρ -zCDP.*

In the context of maximum-type hypotheses such as (2.2) it will be important to extract the largest coordinates of a vector in a differentially private manner. In Algorithm 1 we therefore formulate an algorithm that generalizes the classical Report-Noisy-Max algorithm (see e.g. Dwork et al., 2014a), and the following proposition shows that this algorithm is zCDP.

PROPOSITION 2.5. *Algorithm 1 is $\epsilon^2/8$ -zCDP.*

Algorithm 1 Regularized Report-Noisy-Max (RL-GAP)

Require: vector $q = (q_1, \dots, q_p)^\top \in \mathbb{R}^p$, privacy parameter ϵ , ℓ_1 sensitivity Δ_1 , regularizer $\nu : \{1, \dots, p\} \rightarrow \mathbb{R}$
1: **return** $\arg \max_{j \in \{1, \dots, p\}} \left\{ q_j + \nu(j) + \text{Gumbel}\left(\frac{2\Delta_1}{\epsilon}\right) \right\}$

REMARK 2.6. The privacy analysis can be carried out without the regularizer ν . However, we prefer to state the slightly more general version with the regularizer ν to provide practitioners with greater flexibility. For instance, by adjusting ν , one could overweight earlier indices j and underweight later ones, which might be reasonable if the user has prior knowledge.

3. Baseline Methodology. In this section we consider the U -statistic framework introduced in Section 2.2 and lay out the challenges one encounters when trying to extend existing testing methodology for the relevant hypotheses (2.2) from the private finite dimensional to the private and high-dimensional setting. Although none of these methods discussed here can be finally used in the high-dimensional regime, some of their component techniques are useful for the development of our advanced methodology in Section 4. There, we will propose and theoretically validate a differentially private method that achieves good finite sample performance under mild assumptions - even in the high-dimensional regime. As explained in Example 2.2, the solution to the problem of testing for relevant dependencies under privacy constraints appears as a special case if the dependence measure can be expressed as a U -statistic.

3.1. Concentration-based tests and their limitations.

3.1.1. *A simple test based on a concentration inequality.* In the following paragraph we will ignore privacy aspects for simplicity of presentation, as they can easily be integrated into the discussion without changing any of the conclusions.

A first simple and very conservative approach can be based on concentration inequalities for U -statistics. More precisely, under the null hypothesis in (2.2) we have

$$(3.1) \quad \sqrt{n}(\|U\|_\infty - \Delta) \leq \sqrt{n} \max_{1 \leq i \leq p} (|U_i| - |\theta_i|) \leq \sqrt{n} \max_{1 \leq i \leq p} |U_i - \theta_i|,$$

where $\|U\|_\infty = \max_{1 \leq i \leq p} |U_i|$. For bounded kernels, the classical Hoeffding inequality (Hoeffding, 1963) then yields under the null hypothesis in (2.2) that

$$\mathbb{P}(\sqrt{n} \max_{1 \leq i \leq p} (|U_i| - \Delta) > t) \leq p \max_{1 \leq i \leq p} \mathbb{P}(\sqrt{n}(|U_i - \theta_i|) > t) \leq 2p \exp\left(\frac{-t^2}{2\|h\|_\infty r}\right),$$

where $\|h\|_\infty = \max_{i=1}^p \|h_i\|_\infty$ denotes the maximum of the sup-norms of components of the vector $h = (h_1, \dots, h_p)^\top$. Therefore, it is easy to see that the decision rule

$$\phi(x) = \begin{cases} 1, & \text{if } \max_{1 \leq i \leq p} (|U_i| - \Delta) > \sqrt{\frac{2 \log(2p/\alpha) \|h\|_\infty r}{n}}, \\ 0, & \text{otherwise} \end{cases}.$$

defines a consistent level α test. Unfortunately, it is well known that tests of this type are extremely conservative (see also our numerical results in Figure 2 in Section 5). As the performance of this approach will deteriorate even further when additional noise is introduced to ensure privacy, we will need to improve upon it.

3.1.2. Extension of existing finite dimensional methodology. As a first approach we will reflect upon the results concerning the fixed dimension setting presented in the PhD thesis of Dunsche (2025). As we will see later, the methods developed therein fail in the high-dimensional regime, but this approach nonetheless serves as a good introduction to our general methodology developed in Section 4.

Based on the observation that differentially private testing procedures inflate the variance of the test statistic in an asymptotically but not finite sample negligible manner, a parametric bootstrap procedure was constructed that takes the privatization noise into consideration. It utilizes an analog of the upper bound in (3.1) for a consistent and private estimator of $\|\theta\|_\infty$, say $\|U\|_\infty^{\text{DP}}$. The right-hand side of the corresponding inequality can be approximated by the $\|\cdot\|_\infty$ -norm of a Gaussian vector $Z \sim \mathcal{N}(0, \zeta_1)$, where $\zeta_1 = (\zeta_{1,ij})_{i,j=1,\dots,p}$ is defined in (2.4). Therefore, we define the decision rule

$$(3.2) \quad T^{\text{DP}} := \sqrt{n}(\|U\|_\infty^{\text{DP}} - \Delta) > q_{1-\alpha}^*,$$

where $q_{1-\alpha}^*$ denotes the $(1 - \alpha)$ -quantile of the distribution of $\|Z\|_\infty^{\text{DP}}$ with $Z \sim \mathcal{N}(0, \hat{\zeta}_1^{\text{DP}})$. Here, $\hat{\zeta}_1^{\text{DP}}$ is a consistent private estimator of the covariance matrix ζ_1 . More details on this parametric bootstrap test can be found in Algorithm 7 in Section B of the online supplement. For finite dimension, this yields a consistent and asymptotic level- α test under suitable regularity conditions. A precise formulation of this statement and a proof are provided in Section D of the online supplement. However, its performance deteriorates with increasing dimension, since standard private estimates of the asymptotic variance may become inconsistent in the high-dimensional regime. For instance, under the basic additive Gaussian mechanism for private covariance estimation (see Algorithm 6 in Dwork et al., 2014b), the private estimator $\hat{\zeta}_1^{\text{DP}}$ is already inconsistent in the regime $p \simeq \sqrt{n}$ with respect to the entry-wise maximum norm.

Moreover, even in the finite dimensional setting, the test (3.2) is suboptimal, including the non-private case. To highlight the difficulties one even encounters here note that the quantile for the supremum based test statistic T^{DP} in the decision rule (3.2) is calculated from a privatized maximum norm of the a p -dimensional normal distribution $\mathcal{N}(0, \hat{\zeta}_1^{\text{DP}})$ after an application of an inequality of the type (3.1). However, results on the directional differentiability of the supremum norm (see Theorem 2.1 in Cárcamo et al., 2020) show that the asymptotic distribution of the statistic T^{DP} is given by the maximum norm of a k -dimensional distribution $\mathcal{N}(0, \zeta_1(k))$, where

$$(3.3) \quad \zeta_1(k) = (\text{sign}(\theta_i \theta_j) \zeta_{1,ij})_{i,j \in \{i_1, \dots, i_k\}}$$

denotes the $k \times k$ matrix, which is obtained from the matrix $\zeta_1 = (\zeta_{1,ij})_{i,j \in \{1, \dots, p\}}$ by selecting the specific rows and columns with indices in the *extremal set*

$$(3.4) \quad \mathcal{E} := \{i_1, \dots, i_k\} := \{i = 1, \dots, p : |\theta_i| = \|\theta\|_\infty\}.$$

Thus, asymptotically we are comparing the maximum norm of a k -dimensional normal distribution with the maximum norm of a p -dimensional normal distribution, resulting in an extremely conservative test with not much power.

This phenomenon becomes particularly striking if the cardinality is substantially smaller than the dimension, that is $k \ll p$. In this case, the bootstrap test (3.2) constructs too large quantiles based on the aggregated noise of all coordinates, leading to suboptimal performance in finite samples. This situation becomes even worse when p grows with n (here neglecting the fact that private estimation already fails).

Despite these discouraging observations, this discussion also suggests a solution of this problem. Estimate the coordinates in the extremal set (4.1) and the signs of the corresponding coefficients θ_{ij} in the vector θ and apply a modified version of the test (3.2) using the quantiles of the (privatized) maximum $\max_{1 \leq i \leq k} Z_i$, where $Z \sim \mathcal{N}(0, (\hat{\zeta}_1(k))^{\text{DP}})$ follows a k -dimensional normal distribution and $\hat{\zeta}_1^{\text{DP}}(k)$ is the corresponding privatized estimator of the $k \times k$ covariance matrix (3.3). In a non-private setting this solves the problem, but unfortunately it fails when taking privacy into account for the following reasons:

1. The resulting vector of estimated relevant coordinates might still be high-dimensional in the sense that $k \gtrsim n$, which leads to inconsistent private covariance estimation.
2. Privately estimating the extremal set $\mathcal{E}$ is a highly non-trivial problem. Even if the number of elements of the set $\mathcal{E}$ were known, inferential methodology based on standard private selection methods, like top- k selection (Qiao et al., 2021) or offline sparse vector techniques (Lyu et al., 2016), perform poor in the present setting, and we refer to Section A of the online supplement for a more thorough discussion of this fact.
3. Approaches based on estimating sub-sets of the extremal set which are still sufficiently small to ensure feasibility of the bootstrap also need to be designed carefully: the quantity $\text{sign}(\theta_i \theta_j)$ in the covariance entries in (3.3) can be estimated by $\text{sign}(U_i U_j)$. However, this statistic can have a sensitivity of constant order because the sign $\text{sign}(U_i U_j)$ might flip with non-negligible probability if one of the coordinates θ_i or θ_j vanishes.

4. Extremal set estimation: balancing privacy and statistical accuracy in high dimensions. Our goal is to design a *differentially private* testing procedure for the hypotheses in (2.2) in the high-dimensional setting where p diverges with n . The discussion in the previous section suggests that this task requires an efficient procedure for estimating the extremal set $\mathcal{E}$ defined in (3.4). A canonical non-private estimator of this set is given by

$$(4.1) \quad \hat{\mathcal{E}} = \left\{ i = 1, \dots, p : |U_i| \geq \|U\|_\infty - \sqrt{\frac{\log(p) \log(n)}{n}} \right\},$$

but it is a highly non-trivial problem how to privatize this estimate. In the following discussion we will motivate an alternative approach that a) reduces the dimension, making private covariance estimation feasible and that b) retains favorable statistical properties in a wide range of real world scenarios while also maintaining appropriate privacy guarantees. We start with the most obvious method, namely the sparse vector technique (SVT) and explain why it fails in the present context. Motivated by an analysis of the failure, we then propose our approach, which includes an adaptive choice of the top- k components. Here adaptivity refers to the estimation of the cardinality $k = \#\mathcal{E}$ of the extremal set (3.4), which is a non-trivial choice to make in advance and heavily depends on the data at hand.

A natural approach to construct a private estimator of $\mathcal{E}$ is the SVT (see e.g. Lyu et al., 2016; Zhu and Wang, 2020), which is described in Section A of the online supplement for the sake of completeness. For clarity, we also provide the full algorithmic description in Section B of the online supplement. The basic idea is to privately identify all queries that surpass




Fig 1: Histogram of pair-wise (absolute) Kendall's τ coefficients between different genomes from the [1000 Genomes Project Consortium et al. \(2015\)](#) of the 21.55 Mb - 21.65 Mb window restricted to chromosome 22.

some prescribed threshold t . Setting the queries $q_i = |U_i| - \|U\|_\infty$ for all $i = 1, \dots, p$ and the threshold as $t = -\sqrt{\log(p) \log(n)/n}$, then yields a private and consistent (under mild assumptions) estimator of the extremal set $\mathcal{E}$. Unfortunately, while each individual query has small sensitivity of order $\tilde{O}(1/n)$, the finite sample performance of this approach suffers severely from the composition of the k successes (in which q_i surpasses t) and the number of queries p in the high-dimensional setting. The resulting extremal set estimator is not usable in practice as we either obtain a poor estimate or poor privacy guarantees. For an empirical illustration of these issues, we refer the reader to Section A of the online supplement.

Another common approach for estimating $\mathcal{E}$ is obtained by privately reporting the top- k values of $|U|$. When k is known in advance, Algorithm 1 in [Qiao et al. \(2021\)](#) could be applied directly to obtain a private estimate of the extremal set. However, in our setting k is unknown. As discussed at the end of the previous section, a poor choice will lead to systematic size inflation or deflation when k is chosen too small or large, respectively. Using the private top- k value algorithm is thus not feasible without a good estimate of k . Moreover, even with knowledge of k , the privacy cost of the top- k algorithm scales with $\sqrt{k}$ which is undesirable already for fairly small k .

Motivated by this discussion, we propose a two step procedure that

- (1) constructs a differentially private estimator $\hat{k}$ of k .
- (2) based on the estimator $\hat{k}$ privately estimate the coordinates of the extremal set $\{i_1, \dots, i_k\}$ in (3.4) by $\{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\}$ in a one-shot approach instead of iteratively sampling them.

(1) Adaptive top- k components: For the private estimation of k we will modify an approach for counting queries proposed by [Zhu and Wang \(2022\)](#), such that it is effective in our setting as well. Let $|U|_{(1)} \geq \dots \geq |U|_{(p)}$ denote the order statistics of the absolute values of the components of the vector $U = (U_1, \dots, U_p)^\top$ of U -statistics. We begin by choosing a suitable number of coordinates that will be included in the estimator of the set $\mathcal{E}$. The construction is motivated by the observation that there is often a small number of coordinates with a large signal that are clearly separated from the remaining bulk of coordinates. To make our point clear, instead of iteratively identifying potential coordinates of the extremal set, as it is done by SVT in Algorithm 8 of the online supplement, we aim to identify a single index that separates the extremes from the bulk of the data $(|U|_i)_{1 \leq i \leq p}$. In doing so, we reduce the complexity to a “one-dimensional” problem in a composition sense, making it very cheap from a privacy perspective. We illustrate our observation in Figure 1, which displays the pair-wise Kendall's τ 's for the genome data set discussed in Section 5.2.

For a closer alignment with the notation in [Zhu and Wang \(2022\)](#) we characterize the extremal set $\mathcal{E}$ by a p -dimensional vector of indicators $(\mathbb{1}_1, \dots, \mathbb{1}_p)^\top$, where $\mathbb{1}_j = 1$ if and only if $j \in \mathcal{E} = \{i_1, \dots, i_k\}$. To estimate the cardinality $k = \#\mathcal{E}$, let us first consider the increments of the ordered absolute values of the components of the U -statistic, that is

$$q_j := q_j(X) = |U|_{(j)} - |U|_{(j+1)} \quad j = 1, \dots, p-1.$$

Using these increments, we now estimate k by Algorithm 1, which is a differentially private maximum-selection mechanism that reports the coordinate at which the maximum increment is achieved, that is

$$\hat{k} = \text{RL-GAP}(q = (q_1, \dots, q_{p-1}), \varepsilon, 4r/n \|h\|_\infty, \nu = 0).$$

The associated privacy guarantees are stated in Proposition 2.5. In practice, this will often yield a good separation between the coordinates with a large signal and the remaining coordinates (see also Figure 1). There exist also settings, where such a clear separation is not possible, and we will see later how our method discriminates between these options in a data-adaptive way.

(2) Estimation of the extremal set: Having obtained an estimator $\hat{k}$ for k , we now move one step further and define estimates of the actual coordinates $i_1, \dots, i_k$ with privacy guarantees. For that purpose, we make use of the propose-test-release framework of [Dwork and Lei \(2009\)](#) and consider the statistic $(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top$, where

$$\hat{\mathbb{1}}_j(X) := \begin{cases} 1; & \text{if } |U_j| \text{ is among the top } \hat{k} \text{ U-statistics } |U|_{(1)}, \dots, |U|_{(\hat{k})} \\ 0; & \text{else} \end{cases},$$

and define

$$(4.2) \quad \hat{\mathcal{E}}^{\text{DP}} := \{j = 1, \dots, p \mid \hat{\mathbb{1}}_j(X) = 1\}.$$

However, as pointed out in [Zhu and Wang \(2022\)](#), the global ℓ_2 sensitivity of this query is $\sqrt{2\hat{k}}$. To solve this issue we now make use of local sensitivity and the propose-test-release framework (as introduced, for example, in [Nissim et al., 2007](#)). More specifically, we will prove in Lemma E.1 of the online supplement that the local sensitivity of $(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top$ is indeed 0 whenever there is a valid separation between the extremal coordinates and the bulk (see Assumption 4.2 below for a precise formulation). This means the following: for a fixed dataset, say X , with a sufficiently large gap (of order $O(1/n)$) between the ordered U -statistics, changing one entry actually does not change the estimated extremal set. The estimated set in equation (4.2) can then be released using the propose-test-release framework. In certain cases this set might still contain more than $O(\sqrt{n})$ points, which is problematic for the privatization of the covariance matrix. To solve this issue, we randomly select $\log(p)$ coordinates from the estimated set and proceed from there. We summarize the described procedure in Algorithm 2 below and its privacy guarantee is formalized in the following result.

THEOREM 4.1. *Algorithm 2 satisfies δ -approximate- ρ -zCDP.*

We emphasize that the propose-test-release approach outputs an estimate of the extremal set while using privacy only twice, which is in stark contrast to the SVT-approach. However, Algorithm 2 outputs $\perp$ if it does not find a set of well separated U -statistics. In this case we need a different way of proceeding, and we present two options for this purpose:

Option 1 (rely on concentration): We can proceed as in the previous section and use a privatized version of the concentration based test

$$(4.3) \quad \mathbb{1} \left\{ \|U\|_{\infty}^{\text{DP}} - \Delta > \sqrt{\log(2p/\alpha) 2 \|h\|_{\infty} r/n} \right\},$$

where $\|U\|_{\infty}^{\text{DP}} := \|U\|_{\infty} + \frac{\Delta_2 \|U\|_{\infty}}{\sqrt{2\rho}} Z$ with $Z \sim \mathcal{N}(0, 1)$, $\Delta_2 \|U\|_{\infty} = 2r \|h\|_{\infty}/n$ and ρ is the desired zCDP privacy budget. This test keeps asymptotically the nominal level and is able to detect relevant signals when the sample size is sufficiently large. However, as discussed in Section 3.1 this test is extremely conservative and - even neglecting the additional variance due to privacy - not efficient for reasonable sample sizes. We highlight this fact by an empirical study in Section 5.

Option 2 (extreme value analysis): Assuming that the sign adjusted pairwise correlations $\text{sign}(\theta_i \theta_j) \zeta_{1,ij}$ of U_i and U_j are non-negative for all statistics that satisfy

$$|\theta_i| \geq \Delta - \gamma$$

for some $\gamma > 0$, we can construct a test, that does not require knowledge about the extremal set, which is less conservative than the concentration based approach and is also asymptotically pivotal. For this purpose we use the observation that

$$(4.4) \quad \sqrt{n} \max_{1 \leq i \leq p} (|U_i| - \Delta) \leq \sqrt{n} \max_{\substack{1 \leq i \leq p \\ |\theta_i| \geq \Delta - \gamma}} \text{sign}(U_i)(U_i - \theta_i) + o_{\mathbb{P}}(1)$$

(see Bastian et al., 2024). Results on high-dimensional Gaussian approximation then yield that the right hand side of (4.4) can, in a suitable sense, be approximated by

$$\sqrt{n} \max_{\substack{1 \leq i \leq p \\ |\theta_i| \geq \Delta - \gamma}} \text{sign}(U_i) Z_i$$

for a certain Gaussian vector $Z \in \mathbb{R}^p$ that has the same covariance structure as the vector $(\text{sign}(U_i)(U_i - \theta_i))_{1 \leq i \leq p; |\theta_i| \geq \Delta - \gamma}$. Assuming that $\|h\|_{\infty} \leq L_{\infty}$ for some $L_{\infty} \in (0, +\infty)$, we further use the Bhatia-Davis inequality (Bhatia and Davis, 2000) to observe that under $H_0(\Delta)$ it holds that

$$\mathbb{V}ar(Z_i) \simeq \mathbb{V}ar(U_i) \leq L_{\infty}^2 - (\Delta - \gamma)^2,$$

for all indices $i \in \{1, \dots, p\}$ with $|\theta_i| \geq \Delta - \gamma$. Therefore, we obtain

$$\sqrt{n} \max_{\substack{1 \leq i \leq p \\ |\theta_i| \geq \Delta - \gamma}} \text{sign}(U_i) Z_i \leq \sqrt{n} \max_{1 \leq i \leq p} \tilde{Z}_i,$$

whenever the left hand side quantity is positive. Here the random variables $\tilde{Z}_i$ are defined by

$$\tilde{Z}_i = \begin{cases} \text{sign}(U_i) Z_i \frac{\sqrt{L_{\infty}^2 - (\Delta - \gamma)^2}}{\sqrt{\mathbb{V}ar(Z_i)}} & |\theta_i| \geq \Delta - \gamma \\ Y_i & |\theta_i| < \Delta - \gamma, \end{cases}.$$

where $(Y_i)_{1 \leq i \leq p}$ is a sequence of iid normal distributions with variance $1 - (\Delta - \gamma)^2$ that is independent from the data. We may then use the assumption of non-negative sign adjusted pairwise correlations, the consistency of the sign estimators and Slepian's Lemma (Slepian, 1962) to upper bound this quantity (in distribution, i.e. first order stochastic dominance) by

$$\sqrt{n} \max_{1 \leq i \leq p} Y_i$$

which converges, appropriately rescaled, to a Gumbel distribution with scale parameter $\sqrt{L_\infty^2 - (\Delta - \gamma)^2}$. We can then proceed by using the associated Gumbel quantiles for our test decision whenever Algorithm 2 outputs $\perp$.

Algorithm 2 P-REL: Adaptive private estimation of the extremal set $\mathcal{E}$

Require: Ordered $|U|_{(1)}, \dots, |U|_{(p)}$, appr. zCDP budget parameters δ, ρ and threshold $t := 4r/nL_\infty$.

Ensure: Estimated extremal set $\hat{\mathcal{E}}^{\text{DP}}$.

```

1: function P-REL( $U, \delta, \rho, t$ )
2:   Set  $q_j = |U|_{(j)} - |U|_{(j+1)}$  for  $j = 1, \dots, p-1$ .
3:   Obtain  $\hat{k}$  by invoking Algorithm 1 with  $q = (q_1, \dots, q_{p-1})$  and  $\varepsilon = 2\sqrt{\rho}$ .
4:   Set  $\sigma = t/\sqrt{\rho}$  and construct a high-probability lower bound  $\hat{q}_{\hat{k}} = q_{\hat{k}} + \mathcal{N}(0, \sigma^2) - \sigma z_{1-\delta}$  for  $q_{\hat{k}}$ .
5:   if  $\hat{q}_{\hat{k}} > t$  then
6:     if  $\hat{k} \leq \log(p)$  then
7:       return  $\hat{i}_1, \dots, \hat{i}_{\hat{k}}$ .
8:     else
9:       return Randomly draw  $\log(p)$  indices from  $\hat{i}_1, \dots, \hat{i}_{\hat{k}}$ .
10:    end if
11:  else
12:    return  $\perp$ .
13:  end if
14: end function

```

Algorithm 3 P-HD-U-TEST: Private High-Dimensional U-statistic test

Require: Data X , privacy δ, ρ, α , ℓ_2 sensitivity $\Delta_2 \hat{\zeta}_1(k)$, Bootstrap iterations B , Gumbel parameter γ .

Ensure: reject or fail to reject H_0 in (1.2) and output $\{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\}, \|U\|_\infty^{\text{DP}}$ and $\hat{\zeta}_1^{\text{DP}}$.

```

1: function P-HD-U-TEST( $X, \rho, \delta, \Delta, \alpha, \Delta_2 \hat{\zeta}_1, B, \gamma$ )
2:   Compute U-statistics  $|U| := (|U|_1, \dots, |U|_p)^\top$ .
3:   Run  $\hat{\mathcal{E}}^{\text{DP}} = \text{P-REL}(|U|, \delta, \rho/3, t = 4r/nL_\infty)$ . ▷ Algorithm 2
4:   if  $\hat{\mathcal{E}}^{\text{DP}} = \{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\}$  for some  $\hat{k} \leq \log(p)$  then
5:     Compute  $\hat{\zeta}_1(\hat{k})$  and  $\hat{S} = (\text{sign}(U_i U_j))_{i,j}$ .
6:     Obtain bootstrap quantile  $\hat{q}_{1-\alpha}^* = \text{HQU}(\hat{\zeta}_1(\hat{k}), \hat{S}, n, B, 2r/nL_\infty, \Delta_2 \hat{\zeta}_1(k), \rho/3)$ . ▷ Algorithm 4
7:     Compute private estimator  $\|U\|_\infty^{\text{DP}} = \|U\|_\infty + Z$  with  $\rho/3$ . ▷ Gaussian mechanism
8:     Derive test decision  $\text{dec} = \mathbb{1}\{\|U\|_\infty^{\text{DP}} \geq \hat{q}_{1-\alpha}^* + \Delta\}$ .
9:     if  $\text{dec} = 1$  then
10:      return Reject  $H_0$  and output  $\{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\}, \|U\|_\infty^{\text{DP}}$  and  $\hat{\zeta}_1^{\text{DP}}$ .
11:    else
12:      return Fail to reject  $H_0$  and output  $\{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\}, \|U\|_\infty^{\text{DP}}$  and  $\hat{\zeta}_1^{\text{DP}}$ .
13:    end if
14:  else
15:    return P-GUMBEL-TEST( $|U|, 2\rho/3, n, p, \gamma$ ). ▷ Algorithm 5
16:  end if
17: end function

```

We emphasize that in practice, although this method yields better results than the simple concentration bound approach, our empirical results in Section 5 demonstrate, that it is still far from being comparable to Algorithm 2 when separation is in fact possible. Thus we only recommend its stand-alone application in the case where it is clear that a separation is not possible.

The resulting procedure is summarized in Algorithm 3, which defines a test for the relevant hypotheses (1.2). In the following section we prove that this test is a valid procedure from an asymptotic point of view.

4.1. *Statistical guarantees.* We make the following assumptions.

ASSUMPTION 4.2.

(P) *There exists a set $\mathcal{B}$ with $\mathbb{P}(\mathcal{B}) = 1 - o(1)$ such that for any $l, j_1, \dots, j_l$ the equality*

$$\mathbb{P}(U_{j_1} \geq t_1, \dots, U_{j_l} \geq t_l) = \mathbb{P}(\tilde{U}_1 \geq t_1, \dots, U_{\hat{k}} \geq t_{\hat{k}} | \hat{i}_1 = j_1, \dots, \hat{i}_{\hat{k}} = j_l, \hat{k} = l)$$

holds on $\mathcal{B}$, where $\tilde{U} := (\tilde{U}_1, \dots, \tilde{U}_{\hat{k}}) = (U_{\hat{i}_1}, \dots, U_{\hat{i}_{\hat{k}}})^\top$ is the sub-vector of U defined by the output of Algorithm 2.

(V) *There exist constants $\underline{b} > 0$ and $c \in (0, \Delta)$ such that $\min_{1 \leq i \leq p, |\theta_i| > c} \zeta_{1,ii} > \underline{b}$ for all $p = p(n), n \in \mathbb{N}$, where $\zeta_{1,ii}$ is defined in (2.4). Here and in the following, a minimum over the empty set is defined as $+\infty$.*

(B) *The components of the kernel $h = (h_1, \dots, h_p)^\top$ are bounded in absolute value by a constant $L_\infty > 0$.*

(E) *There exist a constant $\gamma > 0$ such that the covariances $\zeta_{1,ij}$ defined in (2.4) satisfy*

$$\min_{1 \leq i < j \leq p, |\theta_i| \wedge |\theta_j| \geq \Delta - \gamma} \text{sign}(\theta_i \theta_j) \zeta_{1,ij} \geq 0$$

Assumptions (V) and (B) are fairly mild and standard in the context of high-dimensional dependence testing via U -statistics (see, for example, Bastian et al., 2024; Drton et al., 2020). Assumption (P) is a technical condition and needed to show that the test keeps its level at the boundary of the hypotheses in (1.2), that is $\|\theta\|_\infty = \Delta$. It ensures that whether or not and where a gap is detected has negligible impact on the distribution of the vector U . It is fulfilled in a variety of scenarios. Examples include situations where no gap can be detected or when there exists exactly one k such that

$$(4.5) \quad |\theta|_{(k)} - |\theta|_{(k+1)} \geq \max_{j \neq k} |\theta|_{(j)} - |\theta|_{(j+1)} + \sqrt{\log(n) \log(p \vee n)/n},$$

i.e. there is a largest gap that can be separated from all other gaps (here $|\theta|_{(1)} \geq \dots \geq |\theta|_{(p)}$ denote the ordered values of $|\theta|_1, \dots, |\theta|_p$). In particular, we prove in Lemma F.4 of the online supplement, that $\hat{k} = k$ with high probability in that scenario. Assumption (E) is required for using the extreme value approach instead of the concentration based approach when Algorithm 2 outputs $\perp$. It can also be weakened to allow up to $q = o(p)$ negative $\zeta_{1,ij}$.

Under the above assumptions we may now construct a test procedure based on the output of Algorithm 3. We summarize it together with its asymptotic properties in Theorem 4.3 below.

THEOREM 4.3. *Let $\log(p) = o(n^{1/5})$, assume that Assumptions (V), (E) and (B) hold. The test decision ϕ of Algorithm 3 defines a consistent and asymptotic level- α test for hypotheses (2.2). More precisely,*

- (1) *If $\|\theta\|_\infty < \Delta - \gamma$ holds for some $\gamma > 0$, we have $\lim_{n \rightarrow \infty} \mathbb{P}_\theta(\phi = 1) = 0$.*
- (2) *If $\|\theta\|_\infty = \Delta$ and Assumption (P) holds, we have $\lim_{n \rightarrow \infty} \mathbb{P}_\theta(\phi = 1) \leq \alpha$.*
- (3) *Assume that $\|\theta\|_\infty > \Delta + \gamma \sqrt{\log(p)/n}$ for sufficiently large $\gamma > 0$ and either*
 - a) *Algorithm 2 outputs $\perp$ with high probability or,*
 - b) *(4.5) is satisfied for some $k \leq \log(p)$,**holds, then $\lim_{n \rightarrow \infty} \mathbb{P}_\theta(\phi = 1) = 1$.*

Moreover, Assumption (E) may be dropped when using a concentration based instead of the extreme value based approach in the case where Algorithm 2 returns $\perp$.

Our next result provides privacy guarantees for Algorithm 3.

THEOREM 4.4. *Algorithm 3 satisfies δ -approximate- ρ -zCDP.*

REMARK 4.5.

- 1) The results can be extended in a straightforward manner to statistics T that are approximated by a U -statistic in the sense that $T = U + R_n$, where $R_n \in \mathbb{R}^p$ is a remainder that converges to 0 sufficiently fast in the maximum-norm, that is $\|R_n\|_\infty = o_{\mathbb{P}}(\sqrt{\log(p)}/n)$. Prominent examples, where this is possible, are V-statistics, certain functionals of the Kaplan-Meier estimator (see Gijbels and Veraverbeke, 1991) and Spearman's ρ .
- 2) In Algorithm 3, we allocated the privacy budget equally across all sub-procedures. Our simulations indicate that, when a largest gap of size $O(1/n)$ is present, correctly identifying it is the key step of the methodology. Therefore, one might consider allocating a larger portion of the privacy budget ρ to this step.
- 3) The methodology can be extended to U -statistics with unbounded kernels by appropriate truncation methods such as discussed in Chaudhuri et al. (2024b).

REMARK 4.6. A careful inspection of the proof of Theorem 4.3 shows that there exists parameters $\theta \in \mathbb{R}^p$ with $\|\theta\|_\infty = \Delta$ such that there is equality in part (2). For a prominent example, consider the case where the components of θ satisfy for some $\gamma > 0$

$$|\theta_i| \begin{cases} = \Delta & (1 \leq i \leq s), \\ \leq \frac{\Delta}{2+\gamma} & (s < i \leq p). \end{cases}$$

REMARK 4.7. If the null hypothesis in (2.2) is rejected, the next step in the statistical analysis is to privately identify the relevant coordinates, i.e. the set

$$\mathcal{R} := \{i \in \{1, \dots, p\} \mid |\theta_i| > \Delta\}.$$

Although this is not the main objective of this paper, we briefly discuss a first solution of this problem. For this purpose we use Algorithm 2 with the queries

$$q_i = \max\{|U|_{(i)}, \Delta\} - \max\{|U|_{(i+1)}, \Delta\} \quad (i = 1, \dots, p-1),$$

and denote the output by $(\hat{k}_{\mathcal{R}}, \hat{\mathcal{R}}^{\text{DP}})$. With this adjustment, the noisy max Algorithm 1 will output an index $\hat{k}_{\mathcal{R}}$ corresponding to a component of θ with $|\theta|_{(\hat{k}_{\mathcal{R}})} > \Delta$ with high probability, since all gaps corresponding to components with $|\theta_i| \leq \Delta$ are set to zero (asymptotically). We hence obtain that either $\hat{\mathcal{R}}^{\text{DP}} \subset \mathcal{R}$ or $\hat{\mathcal{R}}^{\text{DP}} = \mathcal{R}$. In the former case it is of course possible to recover the missing coordinates via an offline SVT (exponential mechanism) until an index exceeding Δ is found. This modification, however, introduces a trade-off. Suppose the largest gap occurs at an index $\hat{k}_{\mathcal{R}}$ separating the set $\mathcal{R}$ of relevant coordinates and the set $\mathcal{R}^c$ of non-relevant coordinates, that is $|\theta|_{(\hat{k}_{\mathcal{R}})} > \Delta$, and $|\theta|_{(\hat{k}_{\mathcal{R}})} \leq \Delta$. In this case, the gap shrinks from its original size $|U|_{(\hat{k}_{\mathcal{R}})} - |U|_{(\hat{k}_{\mathcal{R}}+1)}$ to $|U|_{(\hat{k}_{\mathcal{R}})} - \Delta$. Consequently, the probability of correctly identifying the gap decreases. To obtain the same detection accuracy one therefore has to allocate a larger portion of the privacy budget to Algorithm 2.

5. Finite sample properties. In this section we investigate the finite sample properties of the proposed procedure by means of a simulation study and illustrate its application in two data examples. Throughout this section we concentrate on the problem of detecting relevant dependencies as discussed in the introduction and in Example 2.2.

5.1. *Simulation study.* We first illustrate the performance of the set estimation in Algorithm 1 in different scenarios. Later, we will also further investigate how these properties propagate into the testing procedure.

Experimental Setup: We generate data from a d dimensional normal distribution with covariance Γ to be specified later, i.e.

$$(5.1) \quad X_1, \dots, X_n \sim \mathcal{N}_d(0, \Gamma)$$

and consider as measure of (monotone) dependence the classical Kendall's τ given by

$$\tau = (\tau_{ij})_{1 \leq i < j \leq d} = \left(\mathbb{E}[\text{sign}(X_{1i} - X_{2i})\text{sign}(X_{1j} - X_{2j})] \right)_{1 \leq i < j \leq p}.$$

An unbiased estimator of τ_{ij} is given by an U -statistic of degree $r = 2$

$$\hat{\tau}_{ij} = \hat{\tau}_{ij}(X) = \frac{2}{n(n-1)} \sum_{1 \leq k < l \leq n} \text{sign}(X_{ki} - X_{li}) \text{sign}(X_{kj} - X_{lj}),$$

with corresponding kernel

$$h_{ij}(x_1, x_2) = \tilde{h}(x_{1i}, x_{1j}, x_{2i}, x_{2j}) = \text{sign}(x_{1i} - x_{2i})\text{sign}(x_{1j} - x_{2j}).$$

The vector of U -statistics is then defined by $U = \text{vech}((\hat{\tau}_{ij})_{1 \leq i < j \leq p})$ with sensitivity

$$|\|\text{vech}(\hat{\tau}(X))\|_\infty - \|\text{vech}(\hat{\tau}(X'))\|_\infty| \leq \frac{4}{n},$$

i.e. $L_\infty = 1$. Therefore, by Lemma 2.4, we have

$$\|\hat{\tau}\|_\infty^{\text{DP}} = \|\hat{\tau}\|_\infty + Y$$

with $Y \sim \mathcal{N}(0, \frac{8}{n^2\rho})$ yields a ρ -zCDP private estimator of $\|\theta\|_\infty$.

Simulation setup: As sample size we choose $n = 250, 500$ and 1000 , while the dimension d is $\lceil \sqrt{2n} \rceil$ and n corresponding to a moderate ($p \approx n$) and high-dimensional ($p \approx n^2/2$) case, respectively. The nominal level for the test is chosen as $\alpha = 0.05$ and $B = 500$ replications are used to calculate the critical values. For the parameter ρ we choose $\rho = 0.1, 0.25$ and 1 ranging from strict to lax privacy. We repeated the experiments over 500 simulation runs. We consider Kendall's τ matrices $\mathcal{T} = (\tau_{ij})_{i,j=1,\dots,d}$, and the covariance matrix $\Gamma = (\Gamma_{ij})_{i,j=1,\dots,d}$ in (5.1) is then obtained by means of the formula

$$\Gamma_{ij} = \sin\left(\frac{\pi}{2}\tau_{ij}\right),$$

which holds for all elliptical distributions with continuous margins. We now consider two **Favorable** settings, for which it is easy to check that condition (4.5) is satisfied such that Theorem 4.3 is applicable. Here we expect a good performance of the test defined by Algorithm 3. Further simulation results can be found in Section C. There we study the robustness of our approach in two **Unfavorable** settings, for which the performance of the test cannot be predicted by theory, and demonstrate that the test has a reasonable performance in such cases as well. The two favorable designs are defined as follows.

F1) A dense signal with $\|\text{vech}(\mathcal{T})\|_\infty = 0.5$ on $p/2$ coordinates, i.e.

$$\mathcal{T} = \mathbf{I}_d + 0.5 \sum_{1 \leq i < j \leq \lfloor d/\sqrt{2} \rfloor} (e_i e_j^\top + e_j e_i^\top),$$

where $e_i \in \mathbb{R}^d$ denotes the i th unit vector and $\mathbf{I}_d$ is the $d \times d$ identity matrix.

F2) A sparse signal with $\|\text{vech}(\mathcal{T})\|_\infty = 0.5$ on only three coordinates, i.e.

$$\mathcal{T} = \mathbf{I}_d + 0.5(e_1 e_2^\top + e_2 e_1^\top) + 0.5(e_2 e_3^\top + e_3 e_2^\top) + 0.5(e_1 e_3^\top + e_3 e_1^\top)$$

We will first explain how to read the figures. Note that we display the rejection probabilities for different values of the threshold Δ , which is decreasing from the left to the right. The horizontal red dotted line corresponds to the nominal level $\alpha = 0.05$, while the vertical red dotted line demarcates our choice of $\|\text{vech}(\mathcal{T})\|_\infty = 0.5$. The values for the threshold Δ in the hypotheses (1.2) are displayed on the x -Axis (in decreasing order), such that the right part of the figures correspond to the alternative and the left to the null hypothesis. Consequently, it is desirable that the rejection curves converge to 1 in the top right quadrant and that they are close to 0 in the bottom left quadrant, crossing precisely at the intersection of the two red lines.

Comparison to Hoeffding Approach: To the best of our knowledge there exists no other differentially private method that can test hypotheses of the form (1.2) in moderate or high-dimensional settings. To demonstrate the improvement of the test defined by Algorithm 3, we thus begin with a quick comparison of our method implemented in Algorithm 3 (see Figure 2a) to the concentration-based baseline we described in Section 3.1 (see Figure 2b) that already outperforms more naive procedures like the Bonferroni correction and related procedures. Across all simulated settings our procedure attains strong empirical power close or equal to one at a signal size for which the rejection rate of concentration-based method has not even exceeded the nominal level α .




Fig 2: Empirical rejection probabilities of the test defined by Algorithm 3 and the test based on a concentration inequality for privacy parameters $\rho = 0.1$ in setting **F1)** with $n \in \{250, 500, 1000\}$ and $p = d(d-1)/2$ for $d = \lfloor \sqrt{2n} \rfloor$ (moderate dimensional regime).

Performance in the moderate dimensional regime: In Figure 3 we display rejection probabilities of the test defined by Algorithm 3 in models **F1)** and **F2)** for the moderate dimensional regime. We observe that the test keeps its size in all settings under consideration with a type I error α quickly decaying to 0 for Δ larger than 0.5. For larger sample sizes - where the gap can be detected more reliably - we see that the test approximates the nominal level more

accurate. Regarding the rejection rate under the alternative we observe that the proposed method is able to detect both dense and sparse correlations well, with a detection boundary that quickly gets sharper as the sample size increases. We further observe that the privacy constraints impact the power most notably for the lower sample size $n = 250$. Here the gap is detected less reliably in the strict privacy setting. As a consequence the test then defaults to the Gumbel approximation, leading to lower detection rates.




Fig 3: Empirical rejection probabilities of the test defined by Algorithm 3 for different privacy parameters $\rho = 0.1, 0.25, 1$ and models **F1)** (first row) and **F2)** (second row) with $n \in \{250, 500, 1000\}$, $p = d(d-1)/2 \approx n$ with $d = \lceil \sqrt{2n} \rceil$ (moderate dimensional regime).




Fig 4: Empirical rejection probabilities of the test defined by Algorithm 3 for different privacy parameters $\rho = 0.1, 0.25, 1$ and models **F1)** (first row) and **F2)** (second row) with $n \in \{250, 500, 1000\}$, $p = d(d-1)/2$ with $d = n$ (high-dimensional regime).

Performance in the high-dimensional regime: In Figure 4 we present corresponding results in the high-dimensional regime. A comparison with Figure 3 shows only a minor loss in

power, which is also predicted by our results. The main driver of this loss is that in the high-dimensional regime, the detection of the gap is more difficult due to the larger amount of coordinates present, as this yields a smaller observable gap with high probability. Whenever the gap can be reliably detected (typically for sample size $n = 500$ in low/moderate privacy regimes or for $n = 1000$ for high privacy) the behavior of the test is independent of the ambient dimension.

5.2. Real-World Examples.

Monotone dependence between genomes We analyze a data set from the [1000 Genomes Project Consortium et al. \(2015\)](#), restricted to chromosome 22. For this study we select $n = 2000$ individuals and focus on a genomic window between base pairs 21.55 Mb and 21.65 Mb, corresponding to a 100 kb region. In population genetics alleles in close physical proximity are often inherited together as they are less likely to be separated by recombination. This phenomenon is well documented ([Slatkin, 2008](#)) and known under the name linkage disequilibrium (LD). We hence expect the existence of some base pairs with high correlation and will investigate whether the proposed methodology is able to detect them or not. To be precise we are interested in detecting Kendall correlations exceeding $\Delta = 0.493$ to search for genes that may exhibit co-expression. The choice of the threshold Δ is here motivated by [Tsaparas et al. \(2006\)](#), who used a Pearson correlation threshold of 0.7 to define genes as co-expressed. This corresponds to a Kendall's τ threshold of ~ 0.493 for a bivariate elliptically distribution.

For each variant and sample we extract the GT (genotype) field, which encodes the pair of alleles as $0|0$, $0|1$, or $1|1$, where “0” denotes the reference allele and “1” the alternative. These are converted into integer allele counts

$$0|0 \mapsto 0, \quad 0|1 \mapsto 1, \quad 1|0 \mapsto 1, \quad 1|1 \mapsto 2.$$

This yields a genotype matrix $X = (X_{ij})_{i=1, \dots, n}^{j=1, \dots, d}$ with $n = 2000$ samples and $d = 750$ variants, where $X_{ij} \in \{0, 1, 2\}$ denotes the allele count for individual i at variant j .

To assess dependencies between genes, we compute pairwise Kendall's rank correlation coefficients, where we break ties (tie-adjusted versions have higher sensitivity, making private inference substantially more difficult) by adding independent normal noise with standard deviation 10^{-6} , leading to a practically negligible bias of order 10^{-12} (see [Kitagawa et al., 2018](#)). This yields a symmetric matrix of empirical Kendall's $\hat{\tau}$ coefficients

$$\hat{\tau} = (\hat{\tau}_{jk})_{1 \leq j, k \leq d} \in [-1, 1]^{d \times d},$$

and the corresponding vector $U = \text{vech}(\hat{\tau})$ of U -statistics has dimension $p = 280,875$. In this example it is possible to identify a gap between the differences $q_i = |U|_{(i)} - |U|_{(i+1)}$, see Figure 5 (a) (the threshold from Algorithm 2 is given by $t = 8/n = 0.004$ and depicted by the dotted red line). We observe that many gaps clearly exceed this bound, also already highlighting that a substantial subset of variant pairs exhibit non-zero dependence.

We now apply the test defined by Algorithm 3 with a moderate privacy level of $\rho = 1$, $\delta = 1/n$. The extremal set $\hat{\mathcal{E}}^{\text{DP}}$ identified by Algorithm 2 is a singleton containing the coordinate corresponding to $|U|_{(1)}$. Repeating Algorithm 2 for different seeds exhibited only small changes, almost always choosing one of the gaps q_i to the right of the red line. For the subsequent testing steps of our method we used the same parameters as in the simulation study ($\alpha = 0.05$, $B = 500$) and rejected the null of $\|\tau\|_\infty \leq \Delta$ for all $\Delta \leq 0.63$. We hence obtain strong evidence that within the 100 kb region of chromosome 22, there exist groups of variants with Kendall's tau larger than 0.493, indicating the existence of co-expressed genes. In particular, the high magnitude of the detected correlations is fully consistent with the presence of Linkage Disequilibrium in this genomic region. Moreover, due to the nature of our




Fig 5: Differences $q_i = |U|_{(i)} - |U|_{(i+1)}$ between successive order statistics of $U = \text{vech}(\hat{\tau})$ for the genome data of [1000 Genomes Project Consortium et al. \(2015\)](#).

approach we also immediately obtain the gene pair with the highest Kendall’s τ with no extra cost in terms of privacy. In comparison, the naive approach based on Hoeffding’s inequality only rejects $H_0(\Delta)$ for $\Delta \leq 0.47$, yielding weak evidence for the existence of co-expressed genes.

To provide a fuller picture of the capabilities of the proposed method on real data we also consider a scenario where the genomic window is larger and there is no clear gap separating the bulk and the largest correlations. To this end we chose the genomic window between base pairs 20.45 Mb and 20.65 Mb, corresponding to a 200 kb region and keep the sample size at $n = 2000$. As evidenced in Figure 5 (b), a gap is visible in the data; however, privacy constraints make reliable detection difficult due to the Gumbel noise introduced by the Report Noisy Max algorithm. Note that shifting the majority of the privacy budget to this step can mitigate this problem. However, we do not adopt this approach, as our goal is to illustrate that even without such an adjustment, Algorithm 3 still attains satisfactory performance. The resulting test rejects $H_0(\Delta)$ for all $\Delta \leq 0.53$, yielding a weaker, but qualitatively comparable result to the case where the gap is clearly visible that still outperforms the Hoeffding approach. This illustrates both the fact that identifying and using gap structures if they are present is a worthwhile endeavor, and also that the proposed method still performs acceptably even if no such gap is detected.

Mass-spectrometry re-analysis (prostate cancer vs. healthy)

We re-analyze the protein mass spectrometry data described by [Adam et al. \(2002\)](#). Each sample i provides intensities $X_{i,j}$ at many time-of-flight values t_j , with time-of-flight related to the mass-to-charge ratio (m/z) of blood-serum proteins. There are 157 healthy and 167 prostate-cancer patients. Following prior work ([Levina et al., 2008](#); [Tibshirani et al., 2005](#)), m/z sites below 2000 are discarded. Intensities are then averaged in consecutive blocks of 10, yielding $p = 23653$ correlation features per sample. We present a heatmap of the matrix of Kendall’s τ coefficients in Figure 6.

Both correlation heatmaps clearly display a banded structure but appear visually hard to distinguish. Therefore, it is essential to formulate a testing problem that can detect subtle differences consistently. Following the previous analyses, we note that the healthy patients seem to exhibit a bandedness structure that the cancer patients lack (observe the light vs dark blue color in the upper right and lower left corners). While the classical hypothesis framework is not able to demarcate these two structures, [Bastian et al. \(2024\)](#) instead consider the relevant hypotheses-pair of the form

$$H_0(\Delta) : \max_{|i-j| \geq m} |\rho_{ij}| \leq \Delta, \quad \text{vs.} \quad H_1(\Delta) : \max_{|i-j| \geq m} |\rho_{ij}| > \Delta,$$




Fig 6: *Block-compressed Kendall's τ correlation heatmap of healthy (left) group and cancer group (right).*

where ρ_{ij} is the spearman correlation and $m \in \mathbb{N}$ determines the size of the band and is chosen as $m = 125$. With this the authors were able to demarcate the two data sets in a non-private setting based on the fact that the hypothesis $H_0(0.1)$ is rejected only for the cancer patients.

We will now analyze this data set with the proposed differentially private methodology at privacy level $\rho = 0.1$ with nominal level $\alpha = 0.05$. Instead of Spearman's ρ , we will consider Kendall's τ as in our case the lower rank of the U -statistic yields a lower sensitivity bound for the test statistic, and thus requires less noise to privatize. Applying the proposed method we did not detect any gap between the signals. The increments of the ordered U -statistics are small across all differences, not exceeding 0.01 even once, whereas $t = 8/n \approx 0.048$ is the threshold in Algorithm 2.

We thus resorted to the extreme value based test in Algorithm 3, which rejects $H_0(\Delta)$ for the healthy group for any $\Delta \leq 0.11$. On the other hand, the test rejects for the cancer patients $H_0(\Delta)$ whenever $\Delta \leq 0.38$. This means that, qualitatively, we obtain the same demarcation between healthy and ill patients as Bastian et al. (2024) in the non-private setting.

6. Conclusions. In this work we consider the problem of testing for practically relevant dependencies in high-dimensional data under DP constraints. While simple concentration based approaches can work in settings with either a very strong signal or very high sample size they perform poorly in more realistic scenarios with moderate to strong signal and moderate sample size. We propose a novel private testing procedure that improves upon this baseline of performance substantially. The improvement is particularly strong for data where the largest signals are separated from the bulk by a gap. In a wide range of scenarios the method even has good finite sample performance if the size of the gap is fairly small (of order $O(1/n)$). In the case where no such gap is present in the data, the new method still improves upon approaches based on concentration bounds. As we investigate the problem in the general framework of high-dimensional U -statistics our method is also applicable for other testing problems.

We prove the validity of our approach in the asymptotic scenario, where the dimension increases exponentially with the sample size. The algorithm has excellent finite sample properties if the private and accurate estimation of an extremal set - the set of indices of the coordinates where the signal attains its maximum norm - is possible. We construct an estimator for this purpose, which is accurate in the presence of a gap in the statistics and only requires a privacy budget for the propose-test-release mechanism once. In contrast, methods which sequentially output the indices such as the Report-Noisy-Max or Sparse Vector technique, require a privacy budget that scales as $\sqrt{k}$, where k is the (unknown) cardinality of the extremal set.

Acknowledgements. MD was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972. PB and HD were partially supported by the DFG; TRR 391 *Spatio-temporal Statistics for the Transition of Energy and Transport* (520388526); Research unit 5381 *Mathematical Statistics in the Information Age* (460867398).

REFERENCES

- 1000 Genomes Project Consortium, Auton, A., Brooks, L. D., Durbin, R. M., Garrison, E. P., Kang, H. M., Korbel, J. O., Marchini, J. L., McCarthy, S., McVean, G. A., and Abecasis, G. R. (2015). A global reference for human genetic variation. *Nature*, 526(7571):68–74.
- Adam, B.-L., Qu, Y., Davis, J. W., Ward, M. D., Clements, M. A., Cazares, L. H., Semmes, O. J., Schellhammer, P. F., Yasui, Y., Feng, Z., et al. (2002). Serum protein fingerprinting coupled with a pattern-matching algorithm distinguishes prostate cancer from benign prostate hyperplasia and healthy men. *Cancer research*, 62(13):3609–3614.
- Bao, Z., Lin, L.-C., Pan, G., and Zhou, W. (2015). Spectral statistics of large dimensional spearman’s rank correlation matrix and its application. *The Annals of Statistics*, 43(6):2588–2623.
- Bastian, P., Dette, H., and Heiny, J. (2024). Testing for practically significant dependencies in high dimensions via bootstrapping maxima of U-statistics. *The Annals of Statistics*, 52(2):628 – 653.
- Berger, J. O. and Delampady, M. (1987). Testing precise hypotheses. *Statist. Sci.*, 2(3):317–335.
- Bhatia, R. and Davis, C. (2000). A better bound on the variance. *The American Mathematical Monthly*, 107(4):353–357.
- Bodnar, T., Dette, H., and Parolya, N. (2019). Testing for independence of large dimensional vectors. *The Annals of Statistics*, 47(5):2977 – 3008.
- Brydges, C. R. (2019). Effect Size Guidelines, Sample Size Calculations, and Statistical Power in Gerontology. *Innovation in Aging*, 3(4). igz036.
- Bun, M. and Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and lower bounds. *arXiv preprint arXiv:1605.02065*.
- Cai, T. T., Xia, D., and Zha, M. (2024). Optimal differentially private pca and estimation for spiked covariance matrices. *arXiv preprint arXiv:2401.03820*.
- Canonne, C. L., Kamath, G., McMillan, A., Ullman, J., and Zakynthinou, L. (2020). Private identity testing for high-dimensional distributions. *Advances in neural information processing systems*, 33:10099–10111.
- Cárcamo, J., Cuevas, A., and Rodríguez, L.-A. (2020). Directional differentiability for supremum-type functionals: Statistical applications. *Bernoulli*, 26(3):2143 – 2175.
- Chaudhuri, K., Loh, P.-L., Pandey, S., and Sarkar, P. (2024a). On differentially private u statistics. *Advances in Neural Information Processing Systems*, 37:23078–23122.
- Chaudhuri, K., Loh, P.-L., Pandey, S., and Sarkar, P. (2024b). On differentially private u statistics. In Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., and Zhang, C., editors, *Advances in Neural Information Processing Systems*, volume 37, pages 23078–23122. Curran Associates, Inc.
- Drton, M., Han, F., and Shi, H. (2020). High-dimensional consistent independence testing with maxima of rank correlations. *Annals of Statistics*, 48:3206–3227.
- Dunsche, M. (2025). *Statistical inference under differential privacy*. doctoral thesis, Ruhr-Universität Bochum, Universitätsbibliothek.
- Dunsche, M., Kutta, T., and Dette, H. (2022). Multivariate mean comparison under differential privacy. In *International Conference on Privacy in Statistical Databases*, pages 31–45. Springer.
- Dwork, C. (2006). Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer.
- Dwork, C. and Lei, J. (2009). Differential privacy and robust statistics. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC ’09, page 371–380, New York, NY, USA. Association for Computing Machinery.
- Dwork, C., Roth, A., et al. (2014a). The algorithmic foundations of differential privacy. *Foundations and trends® in theoretical computer science*, 9(3–4):211–407.
- Dwork, C., Talwar, K., Thakurta, A., and Zhang, L. (2014b). Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC ’14, page 11–20, New York, NY, USA. Association for Computing Machinery.
- Gijbels, I. and Veraverbeke, N. (1991). Almost Sure Asymptotic Representation for a Class of Functionals of the Kaplan-Meier Estimator. *The Annals of Statistics*, 19(3):1457 – 1470.
- Han, F., Chen, S., and Liu, H. (2017). Distribution-free tests of independence in high dimensions. *Biometrika*, 104(4):813–828.

- He, Y., Xu, G., Wu, C., and Pan, W. (2021). Asymptotically independent U-statistics in high-dimensional testing. *The Annals of Statistics*, 49(1):154 – 181.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Jiang, T. and Qi, Y. (2015). Likelihood ratio tests for high-dimensional normal distributions. *Scandinavian Journal of Statistics*, 42(4):988–1009.
- Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Kitagawa, T., Nybom, M., and Stuhler, J. (2018). Measurement error and rank correlations. *cemmap working paper*.
- Leung, D. and Drton, M. (2018). Testing independence in high dimensions with sums of rank correlations. *The Annals of Statistics*, 46(1):280 – 307.
- Levina, E., Rothman, A., and Zhu, J. (2008). Sparse estimation of large covariance matrices via a nested Lasso penalty. *The Annals of Applied Statistics*, 2(1):245 – 263.
- Li, Z., Wang, Q., and Li, R. (2021). Central limit theorem for linear spectral statistics of large dimensional kendall’s rank correlation matrices and its applications. *Annals of Statistics*, 49(3):1569–1593.
- Liu, X., Chen, Y., and Xu, W. (2025). Differentially private joint independence test. *arXiv preprint arXiv:2503.18721*.
- Liu, X., Kong, W., and Oh, S. (2022). Differential privacy and robust statistics in high dimensions. In *Conference on Learning Theory*, pages 1167–1246. PMLR.
- Lovakov, A. and Agadullina, E. (2021). Empirically derived guidelines for effect size interpretation in social psychology. *European Journal of Social Psychology*, 51:485–504.
- Lyu, M., Su, D., and Li, N. (2016). Understanding the sparse vector technique for differential privacy. *arXiv preprint arXiv:1603.01699*.
- Narayanan, S. (2022). Private high-dimensional hypothesis testing. In Loh, P.-L. and Raginsky, M., editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 3979–4027. PMLR.
- Nissim, K., Raskhodnikova, S., and Smith, A. (2007). Smooth sensitivity and sampling in private data analysis. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC ’07, page 75–84, New York, NY, USA. Association for Computing Machinery.
- Pearson, K. (1920). Notes on the history of correlation. *Biometrika*, 13(1):25–45.
- Qiao, G., Su, W., and Zhang, L. (2021). Oneshot differentially private top-k selection. In *International Conference on Machine Learning*, pages 8672–8681. PMLR.
- Quintana, D. (2016). Statistical considerations for reporting and planning heart rate variability case-control studies. *Psychophysiology*, 54.
- Rogers, R. and Kifer, D. (2017). A New Class of Private Chi-Square Hypothesis Tests. In Singh, A. and Zhu, J., editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 991–1000, Fort Lauderdale, FL, USA. PMLR.
- Sei, Y. and Ohsuga, A. (2021). Privacy-preserving chi-squared test of independence for small samples. *BioData Mining*, 14(1):1–25.
- Slatkin, M. (2008). Linkage disequilibrium — understanding the evolutionary past and mapping the medical future. *Nature Reviews Genetics*, 9(6):477–485.
- Slepian, D. (1962). The one-sided barrier problem for gaussian noise. *The Bell System Technical Journal*, 41(2):463–501.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(1):91–108.
- Tsaparas, P., Mariño-Ramírez, L., Bodenreider, O., Koonin, E. V., and Jordan, I. K. (2006). Global similarity and local divergence in human and mouse gene co-expression networks. *BMC Evolutionary Biology*, 6:70.
- Tukey, J. W. (1991). The philosophy of multiple comparisons. *Statistical Science*, 6(1):100–116.
- Yao, S., Zhang, X., and Shao, X. (2018). Testing mutual independence in high dimension via distance covariance. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):455–480.
- Zhu, Y. and Wang, Y.-X. (2020). Improving sparse vector technique with renyi differential privacy. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20249–20258. Curran Associates, Inc.
- Zhu, Y. and Wang, Y.-X. (2022). Adaptive private-k-selection with adaptive k and application to multi-label pate. In Camps-Valls, G., Ruiz, F. J. R., and Valera, I., editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 5622–5635. PMLR.

APPENDIX A: LIMITATIONS OF THE SPARSE VECTOR TECHNIQUE

In this section we compare the privacy guarantees for the sparse vector technique (SVT) in [Zhu and Wang \(2020\)](#) to ours in a high-dimensional regime where $p \gg n$. For that, let us first quickly recap how the SVT technique works: the basic technique lets you answer many noisy queries privately by only revealing which ones are *big enough* to cross a noisy threshold. In simple terms, it adds noise to both the threshold and the queries so you can safely say “yes” or “no” a few times without spending privacy for each comparison. However, for every positive answer the used privacy increases. First we note that, in the language of (ϵ, δ) -privacy our worst privacy bound ($\rho = 1, \delta = 1/n$) in the finite sample study relates to $\epsilon \approx 6$ and $\delta = 1/n$ for which our method performs well in all considered cases. Our method is based on (randomly) selecting at most $\log(p)$ coordinates from the extremal set $\hat{\mathcal{E}}^{\text{DP}}$ to ensure that privatization of the covariance is still feasible. Let us now consider the SVT with the bound

$$\epsilon(\delta) \leq \frac{\Delta^2}{2\sigma_1^2} + \frac{2c\Delta^2}{\sigma_2^2} + \sqrt{2\left(\frac{\Delta^2}{2\sigma_1^2} + \frac{2c\Delta^2}{\sigma_2^2}\right) \left(\log(\delta^{-1}) + \log\left(c\binom{p}{c}\right)\right)},$$

from [Zhu and Wang \(2020\)](#) with $c = \log(p)$. If the queries are shuffled in advance this is, mathematically speaking, the same as the random selection in our methodology. Here, following the notation of the paper, Δ denotes the sensitivity of the queries $q_j = |U_j| - \|U\|_\infty$, and σ_1, σ_2 are the user chosen noise variances of the queries and of the threshold, respectively (see [Zhu and Wang \(2020\)](#) for details). At a first glance the right-hand side also only depends logarithmically on the dimension p . Unfortunately the constants in this bound are too large for a fruitful application in the present context. We underpin this observation by a simple empirical illustration.

Numerical illustration. As in Section 5, we take $n = 250$, $d = n$, and the pair-wise Kendall’s τ s, i.e. $p = d(d-1)/2 = 31125$. Keep $\delta = 1/n$, $\Delta = 4r/n$, and first set $\sigma_1 = \sigma_2 = 0.1$. Note that this is quite a generous choice considering that we want to detect gaps between Kendall’s τ associations which take values in $[-1, 1]$. We then obtain the upper bound $\epsilon(\delta) \leq 24.21838$. For a more realistic scenario with good statistical results, i.e. $\sigma_1 = \sigma_2 = 0.01$, we obtain $\epsilon(\delta) \leq 449.5438$, which, for all intents and purposes, is practically equivalent to no privacy at all. Even raising the sample size, often not possible in practice without major effort, to $n = 1000$ does not really improve the results: returning to the less realistic $\sigma_1 = \sigma_2 = 0.1$, we merely obtain $\epsilon \leq 8.012233$. In contrast our algorithm performs well in these scenarios even for strong privacy regimes such as $\rho = 0.1$ (i.e. $\epsilon \approx 1$).

We conclude emphasizing one more the crucial difference between SVT and our methodology for estimating the extremal set. SVT is a composition of mechanisms, which scales (with respect to privacy) like $\sqrt{k}$ due to composition, where k is the number of positively answered queries. In contrast to that, whenever the data has some structure, such as a gap of order $O(1/n)$ in the ordered (absolute) U -statistics, we can simplify the estimation of the extremal set to a “two-dimensional” problem from a privacy perspective using Algorithm 1 and the PTR framework for the estimation of the extremal set. This is already sufficient for our purposes.

APPENDIX B: ADDITIONAL ALGORITHMS

In this section, we present the algorithms that accompany our theoretical results. That includes the resampling procedures (Algorithm 4), the Gumbel-based test (Algorithm 5) and

a private covariance estimator (Algorithm 6). We also recall the resampling procedure introduced in Dunsche (2025) in Algorithm 7. In Algorithm 8, we state the generalized SVT algorithm proposed in Zhu and Wang (2020).

Algorithm 4 High-Dimensional Quantile Monte Carlo U-Statistics (HQU)

Require: sign matrix: $\hat{S}$, covariance matrix: $\hat{\zeta}_1$, sample size n , resample parameter: B , privacy parameter: ρ , sensitivities: $\Delta_2 \|U\|_\infty, \Delta_2 \hat{\zeta}_1$

Ensure: Empirical $(1 - \alpha)$ -quantile: $q_{1-\alpha}^{*,B}$.

```

1: function HQU( $\hat{\zeta}_1, \hat{S}, n, B, \Delta_2 \|U\|_\infty, \Delta_2 \hat{\zeta}_1, \rho$ )
2:   for  $b = 1, \dots, B$  do
3:     Define  $\hat{\zeta}_1^S := \hat{S} \odot \hat{\zeta}_1$ . ▷ Hadamard product
4:     Define  $\hat{\zeta}_1^{S,DP} := \text{GAUSSCOV}(\hat{\zeta}_1^S, \rho, \Delta_2 \hat{\zeta}_1)$ 
5:     Sample  $U_b^* \sim \mathcal{N}(0, \hat{\zeta}_1^{S,DP})$ .
6:     Define  $U_b^{DP,*} := \|U_b^*\|_\infty^{DP}$ , where DP denotes the same mechanism used for  $\|U\|_\infty^{DP}$ . ▷ Lemma 2.4
7:     Define  $T_b^* := U_b^{DP,*}$ .
8:   end for
9:   Sort statistics in ascending order:  $(T_{(1)}^*, \dots, T_{(B)}^*) = \text{sort}(T_1^*, \dots, T_B^*)$ .
10:  Define  $q_{1-\alpha}^{*,B} := T_{(\lfloor (1-\alpha)B \rfloor)}^*$ .
11:  return  $q_{1-\alpha}^{*,B}$ 
12: end function

```

Algorithm 5 P-GUMBEL-TEST: Private Test for U-statistics in High Dimensions using Gumbel Approximation

Require: Data set X , privacy budgets δ, ρ , α level.

Ensure: test decision and $\|U\|_\infty^{DP}$.

```

1: function P-GUMBEL-TEST( $U, \rho, n, p, \gamma$ )
2:   Compute private estimator  $\|U\|_\infty^{DP} = \|U\|_\infty + Z$  with  $2\rho/3$ . ▷ Lemma 2.4
3:   Define  $a_p = \sqrt{2 \log(p)}$  and  $G \sim \text{Gumbel}(0, \sqrt{L_\infty - (\Delta - \gamma)^2})$ 
4:   Define  $Q := \frac{q_{1-\alpha}^G}{a_p} + a_p - \frac{\log \log(p) + \log(4\pi)}{2a_p}$ 
5:   if  $\|U\|_\infty^{DP} \geq Q/\sqrt{n} + \Delta$  then
6:     return Reject  $H_0$  and output  $\|U\|_\infty^{DP}$ .
7:   else
8:     return Fail to reject  $H_0$  and output  $\|U\|_\infty^{DP}$ .
9:   end if
10: end function

```

Algorithm 6 Private Covariance estimation with additive noise **Gausscov**

Require: $\hat{\zeta}_1 \in \mathbb{R}^{d \times d}$, privacy parameter ρ , sensitivity $\Delta_2 \hat{\zeta}_1$.

Ensure: Privatized covariance matrix $\hat{\zeta}_1^{\text{DP}}$.

- 1: **function** GAUSSCOV($\hat{\zeta}_1, \rho, \Delta_2 \hat{\zeta}_1$)
- 2: Generate $Z_{i,j} \sim \mathcal{N}(0, 1)$ for $1 \leq i \leq j \leq d$ with $Z_{j,i} = Z_{i,j}$.
- 3: Define $Z = (Z_{i,j})_{i,j=1,\dots,d}$ and compute

$$\hat{\zeta}_1^{\text{DP}} = \hat{\zeta}_1 + \frac{\Delta_2 \hat{\zeta}_1}{\sqrt{2\rho}} Z.$$

- 4: **return** $\hat{\zeta}_1^{\text{DP}}$
 - 4: **end function**
-

Algorithm 7 Multivariate Quantile Monte Carlo U-Statistics (QU) [Dunsche \(2025\)](#)

Require: private covariance matrix: $\hat{\zeta}_1^{\text{DP}}$, sample size: n , resample parameter: B , privacy parameter: ρ , sensitivity: $\Delta_2 \|U\|_\infty$.

Ensure: Empirical $(1 - \alpha)$ -quantile: $q_{1-\alpha}^{*,B}$.

- 1: **function** QU($\hat{\zeta}_1^{\text{DP}}, n, B, \Delta_2 \|U\|_\infty, \rho$)
 - 2: **for** $b = 1, \dots, B$ **do**
 - 3: Sample $U_b^* \sim \mathcal{N}_s(0, \hat{\zeta}_1^{\text{DP}}/n)$.
 - 4: Compute private estimator $U_b^{\text{DP},*} := \|U_b^*\|_\infty^{\text{DP}} = \|U_b^*\|_\infty + Z$ with $2\rho/3$. ▷ Lemma 2.4
 - 5: Define $T_b^* := \sqrt{n}(U_b^{\text{DP},*})$.
 - 6: **end for**
 - 7: Sort statistics in ascending order: $(T_{(1)}^*, \dots, T_{(B)}^*) = \text{sort}(T_1^*, \dots, T_B^*)$.
 - 8: Define $q_{1-\alpha}^{*,B} := T_{(\lfloor (1-\alpha)B \rfloor)}^*$.
 - 9: **return** $q_{1-\alpha}^{*,B}$
 - 10: **end function**
-

Algorithm 8 Generalized Sparse Vector Technique (SVT) (slightly adjusted Algorithm 2 in [Zhu and Wang \(2020\)](#))

Require: Data X , adaptive queries $q_1, q_2, \dots \in Q$ with sensitivity Δ , noise mechanisms $\mathcal{M}_\rho, \mathcal{M}_\nu$, threshold t , cut-off c , max length $k_{\max}$.

Ensure: Outputs $a_i \in \{\top, \perp\}$.

- 1: **function** SVT($X, Q, t, \Delta, c, k_{\max}$)
 - 2: Sample $\hat{t} \sim \mathcal{M}_\rho(D, t)$
 - 3: count = 0
 - 4: **for** $i = 1, 2, 3, \dots, k_{\max}$ **do**
 - 5: Sample $\hat{q}_i \sim \mathcal{M}_\nu(X, q_i)$
 - 6: **if** $\hat{q}_i \geq \hat{t}$ **then**
 - 7: Output $a_i = \top$, count = count + 1
 - 8: **if** count $\geq c$ **then**
 - 9: **abort**
 - 10: **end if**
 - 11: **else**
 - 12: Output $a_i = \perp$
 - 13: **end if**
 - 14: **end for**
 - 15: **end function**
-

APPENDIX C: ADDITIONAL SIMULATIONS

In this section we present further simulation results. First we study the robustness of the proposed method with respect to violations of the assumptions. For this purpose we consider Unfavorable scenarios, in which our theory does not predict the performance of Algorithm 3. Second, we compare our methodology to the state of the art in the non-private setting and demonstrate superior performance in some settings despite the additional cost incurred by ensuring DP.

Robustness From a theoretical perspective, the strong performance of the test defined by Algorithm 3 relies on the existence of a sufficiently large gap. While this assumption is often reasonable in practice, an obvious question is how the procedure performs in less favorable settings. For this purpose, we consider two further models. All other parameters are the same as in Section 5.1.

U1) A dense signal with $\|\text{vech}(\mathcal{T})\|_\infty = 0.5$ but no gaps, i.e. we define $m = \lfloor d/\sqrt{2} \rfloor$ and construct the matrix $\mathcal{T} \in \mathbb{R}^{d \times d}$ as follows:

$$\mathcal{T} = \begin{bmatrix} \mathbf{A} & 0 \\ 0 & \mathbf{I}_{d-m} \end{bmatrix}$$

where

$$\mathbf{A} := \mathbf{A}' - \text{diag}(\mathbf{A}') + \mathbf{I}_m$$

and $\mathbf{A}' = \mathbf{a}\mathbf{a}^\top \in \mathbb{R}^{m \times m}$. Here, $\mathbf{a} \in \mathbb{R}^m$ is an equidistant scaled vector

$$\mathbf{a} = \sqrt{\frac{0.5}{\max_{i < j} (b_i b_j)}} \mathbf{b}.$$

with $\mathbf{b} = (0.01 + (j-1)(0.99-0.01)/(m-1))_{j=1,\dots,m} \in \mathbb{R}^m$.

U2) A dense signal with $\|\text{vech}(\mathcal{T})\|_\infty = 0.5$ but two gaps, i.e. for $m = \lfloor d/\sqrt{2} \rfloor$ we take

$$\mathcal{T} = \begin{bmatrix} 0.5\mathbf{I}_m + 0.5\mathbf{J}_m & 0 \\ 0 & 0.75\mathbf{I}_{d-m} + 0.25\mathbf{J}_{d-m} \end{bmatrix}$$

where identity matrix, $\mathbf{J}_l$ is a $l \times l$ matrix with all entries equal to one.

Note that in both scenarios at least one of the assumptions is violated. We display in Figure 7 and 8 the rejection probabilities of the test in the scenarios **U1)**, **U2)** for a moderate dimensional ($p \approx n$) and high-dimension ($p \approx n^2/2$) setting, respectively.

Performance for unfavorable setups Again, the test keeps its nominal level across all cases under consideration. For the scenario **U2)**, we observe very good power which is comparable but slightly inferior to power for the scenarios **F1)** and **F2)**. As this scenario has two gaps we sometimes detect the gap between the 0 and 0.25 valued correlations, leading to more than $\log(p)$ coordinates that are deemed relevant and thus to randomly sampling from the 0.25 and 0.5 valued correlation coordinates which reduces the detection power. This effect can be, to some degree, alleviated by penalizing gaps of smaller coordinates by appropriately choosing appropriate weights $\nu(j)$ in Algorithm 1, e.g. by $c(1 - j/n)$. Some care is needed in that case, as this does not lead to uniform improvement across all possible scenarios. In scenario **U1)** there is no gap of sufficient size for reliable detection. Accordingly we do not reliably separate relevant from irrelevant coordinates, leading to use of the Gumbel approximation in most cases. Nevertheless the test works reasonably well even in this scenario.




Fig 7: Empirical rejection probabilities of the test defined by Algorithm 3 for different privacy parameters $\rho = 0.1, 0.25, 1$ and models **U1** (first row) and **U2** (second row) with $n \in \{250, 500, 1000\}$, $p = d(d-1)/2$ with $d = \lfloor \sqrt{2n} \rfloor$ (moderate dimensional regime).




Fig 8: Empirical rejection probabilities of the test defined by Algorithm 3 for different privacy parameters $\rho = 0.1, 0.25, 1$ and models **U1** (first row) and **U2** (second row) with $n \in \{250, 500, 1000\}$, $p = d(d-1)/2$ with $d = n$ (high-dimensional regime).

Performance compared to non-private state of the art In Figure 9 we display the rejection probabilities of the test defined by Algorithm 3 and the non-private test proposed in equation (2.27) of Bastian et al. (2024), where we consider the settings **F1** and **F2** from Section 5.1. In **F2** only a few coordinates carry a signal, and we observe that the new private procedure can even outperform the non-private state-of-the-art test. This superiority arises because the methodology in Bastian et al. (2024) does not attempt to estimate the extremal set, and thus suffers from the difficulty discussed in Section 3.1.2. In the dense setting **F1**, however, we begin to see the effect of privacy more clearly, particularly for smaller sample sizes. Here the

effect of not estimating the extremal set becomes less pronounced and the loss of power by the additional privacy noise becomes more clearly visible.




Fig 9: Empirical rejection probabilities of the test defined by Algorithm 3 and the test based on Bastian et al. (2024) for privacy parameters $\rho = 1$ in setting **F1** (first row) and **F2** (second row) with $n \in \{250, 500, 1000\}$ and $p = d(d-1)/2$ for $d = \lceil \sqrt{2n} \rceil$ (moderate dimensional regime).

APPENDIX D: FINITE DIMENSIONAL METHODOLOGY

In this section, we will prove that for finite dimension p , the decision rule (3.2) defines a consistent asymptotic level α test for the hypotheses (2.2). Following Dunsche (2025), we first state some assumptions.

ASSUMPTION D.1.

- (1) Let $X_1, \dots, X_n \sim F$ be iid random vectors, where F is a distribution on the d -dimensional cube $[-m, m]^d$, $m \in \mathbb{R}, d \in \mathbb{N}$.
- (2) Let $U = (U_1, \dots, U_p)^\top$, where each component U_j is a U -statistic with kernel h_j of order $r(j)$. Then, we assume that for all $j = 1, \dots, p$

$$\text{Var}(\mathbb{E}[h_j(X_1, \dots, X_r) | X_1]) = \zeta_{1,j} > b,$$

where $b > 0$ is fixed.

- (3) For any kernel $h_j : [-m, m]^d \rightarrow \mathbb{R}$, $j = 1, \dots, p$, we assume

$$\|h_j\|_\infty \leq \frac{1}{2} L_\infty,$$

where L_∞ is a constant depending on h_j, m for $j = 1, \dots, p$.

Note that the assumptions above yield for any kernel h that

$$\|h(X) - h(X')\|_2 \leq L_2, \quad \forall X, X',$$

where $\|\cdot\|_2$ denotes the euclidean norm and $L_2 = (\sum_{j=1}^p L_\infty^2)^{1/2}$. With that in hand, we can formulate a decision rule that is a consistent level α test:

THEOREM D.2 (Monte Carlo). *Assume Assumption D.1 holds and let $\Delta > 0$ be fixed. Furthermore consider a U -statistic U with $\theta = \mathbb{E}_F[U] \in \mathbb{R}^p$. We define for a standard normal random variable Y*

$$\|U\|_\infty^{DP} := \|U\|_\infty + \frac{\Delta_2 \|U\|_\infty}{\sqrt{2\rho}} Y,$$

let $q_{1-\alpha}^*$ be the theoretical $(1-\alpha)$ -quantile of T_1^* approximated by Algorithm 7 with $\hat{\zeta}_1^{DP} = \text{GaussCov}(\hat{\zeta}_1, \rho, \Delta_2 \hat{\zeta}_1)$ (Algorithm 6). Then the decision rule "reject if

$$(D.1) \quad T^{DP} := \sqrt{n}(\|U\|_\infty^{DP} - \Delta) > q_{1-\alpha}^*",$$

yields a consistent, asymptotic level- α test for hypotheses in equation (2.2). That is for all $\|\theta\|_\infty < \Delta$

$$\lim_{n \rightarrow \infty} \mathbb{P}_\theta(T^{DP} > q_{1-\alpha}^*) = 0,$$

for all $\|\theta\|_\infty = \Delta$

$$\lim_{n \rightarrow \infty} \mathbb{P}_\theta(T^{DP} > q_{1-\alpha}^*) \leq \alpha$$

(level α). For all $\|\theta\|_\infty > \Delta$

$$\lim_{n \rightarrow \infty} \mathbb{P}_\theta(T^{DP} > q_{1-\alpha}^*) = 1$$

(consistency). Furthermore, the decision rule (D.1) is 2ρ -zCDP.

PROOF OF THEOREM D.2. We first recall asymptotic results for the statistic $T_\theta^{DP} := \sqrt{n}(\|U\|_\infty^{DP} - \|\theta\|_\infty)$. First note that we have

$$T_\theta^{DP} := \sqrt{n}(\|U\|_\infty^{DP} - \|\theta\|_\infty) + o_{\mathbb{P}}(1).$$

Then using Hoeffdings CLT and the Delta method for Hadamard directional differential functionals (see Theorem 2.1 in Cárcamo et al., 2020), we get for all $\|\theta\|_\infty > 0$ that

$$(D.2) \quad T_\theta^{DP} \xrightarrow{d} L = \max_{\substack{j=1, \dots, s \\ \|\theta_j\| = \|\theta\|_\infty}} \text{sign}(\theta_j) Z_j \leq \|Z\|_\infty,$$

where $Z \sim \mathcal{N}_p(0, \zeta_1)$. For $\|\theta\|_\infty = 0$, using the multivariate CLT of Hoeffding and the continuous mapping theorem yields

$$(D.3) \quad T_0^{DP} \xrightarrow{d} \|Z\|_\infty,$$

where again $Z \sim \mathcal{N}_p(0, \zeta_1)$. Now let $q_{1-\alpha}^*$ be the theoretical quantile of T_1^* approximated by Algorithm 7. Since, we have used an upper bound $\|Z\|_\infty$ in equation (D.2) for the resampling procedure in Algorithm 7, we prove the consistency with respect to that upper bound. In particular, we know that

$$d_K(\mathcal{L}(T_1^* | X_1, \dots, X_n), \mathcal{L}(\|Z\|_\infty)) \xrightarrow{\mathbb{P}} 0,$$

where d_K denotes the Kolmogorov distance and we used that $T_1^* \sim \max_{1 \leq j \leq p} |Z'| + \sqrt{n}Y$, with $Z' \sim \mathcal{N}_p(0, \hat{\zeta}_1^{\text{DP}})$ and $Y \sim \mathcal{N}(0, (\frac{\Delta_2 \|U\|_\infty}{\sqrt{2\rho}})^2)$ conditional on $X_1, \dots, X_n$. This is indeed true due to the continuous mapping theorem and the conditional version of Slutsky's Lemma. Therefore, Lemma 23.3 in [Vaart \(1998\)](#) yields that as $n \rightarrow \infty$ we have

$$q_{1-\alpha}^* \xrightarrow{\mathbb{P}} z_{1-\alpha}$$

conditional on $X_1, \dots, X_n$, where $z_{1-\alpha}$ is the $(1-\alpha)$ -quantile of $\|Z\|_\infty$. Consequently, there also exists a subsequence for which almost sure convergence holds, i.e. $q_{1-\alpha}^* \xrightarrow{a.s.} z_{1-\alpha}$ on that subsequence. We can use the almost sure convergence of the bootstrap quantile $q_{1-\alpha}^*$ and the convergence of T_θ^{DP} combined with Slutsky's Lemma, to obtain

$$(D.4) \quad \lim_{n \rightarrow \infty} \mathbb{P}(T_\theta^{\text{DP}} > q_{1-\alpha}^*) \leq \mathbb{P}(\|Z\|_\infty > z_{1-\alpha}) = \alpha.$$

With that in hand, we can prove the statements of Theorem [D.2](#) for T^{DP} by considering four different cases:

- (1) $0 < \|\theta\|_\infty < \Delta$,
- (2) $\|\theta\|_\infty = 0$,
- (3) $\|\theta\|_\infty = \Delta$,
- (4) $\|\theta\|_\infty > \Delta$.

For that purpose let us first decompose T^{DP} :

$$(D.5) \quad T^{\text{DP}} = T_\theta^{\text{DP}} + \sqrt{n}(\|\theta\|_\infty - \Delta).$$

(1) Consider the first case $0 < \|\theta\|_\infty < \Delta$. Since $q_{1-\alpha}^*$ is stochastically bounded, we can conclude that for every $M > 0$ there exist an $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$ we have

$$\mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = \mathbb{P}(T_\theta^{\text{DP}} > q_{1-\alpha}^* - \sqrt{n}(\|\theta\|_\infty - \Delta)) \leq \mathbb{P}(T_\theta^{\text{DP}} > M),$$

where we used that $\|\theta\|_\infty - \Delta < 0$. Now taking the limit for $n \rightarrow \infty$ on both sides yields

$$\limsup_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) \leq \mathbb{P}(\|Z\|_\infty > M),$$

where $Z \sim \mathcal{N}_p(0, \zeta_1)$. Hence, taking the limit for $M \rightarrow \infty$ lets us conclude that

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = 0.$$

(2) If $\|\theta\|_\infty = 0$, we have $T_\theta = T_0$ and can follow with the same arguments that for any $M > 0$

$$\mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = \mathbb{P}(T_0^{\text{DP}} > q_{1-\alpha}^* + \sqrt{n}\Delta) \leq \mathbb{P}(T_0^{\text{DP}} > M)$$

for n sufficiently large, since $\Delta > 0$. Consequently, observing [\(D.3\)](#), the same arguments as in (1) give

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = 0.$$

Therefore, we obtain with (1) and (2) for all $\|\theta\|_\infty < \Delta$ that

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = 0.$$

(3) For $\|\theta\|_\infty = \Delta$, we have that $T^{\text{DP}} = T_\theta^{\text{DP}}$, and we obtain from [\(D.4\)](#) to obtain

$$\limsup_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = \lim_{n \rightarrow \infty} \mathbb{P}(T_\theta^{\text{DP}} > q_{1-\alpha}^*) \leq \alpha.$$

(4) For $\|\theta\|_\infty > \Delta$, we use the decomposition (D.5) and equation (D.4) to conclude that for any $M > 0$

$$\mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) = \mathbb{P}(T_\theta^{\text{DP}} > q_{1-\alpha}^* - \sqrt{n}(\|\theta\|_\infty - \Delta)) \geq \mathbb{P}(T_\theta^{\text{DP}} > -M)$$

for n sufficiently large. Taking the limit on both sides for $n \rightarrow \infty$, we have

$$\liminf_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} > q_{1-\alpha}^*) \geq \mathbb{P}(\|Z\|_\infty > -M),$$

where again we have used the upper bound in equation (D.2). Now taking the limit for $M \rightarrow \infty$ yields

$$\lim_{n \rightarrow \infty} \mathbb{P}(T^{\text{DP}} \geq q_{1-\alpha}^*) = 1,$$

where again we have used that $\|Z\|_\infty$ is tight. The privacy guarantee holds, due to the sensitivity of U derived in Lemma E.2. For the private covariance, we can use Lemma E.3. The result finally follows from the composition theorem for ρ -zCDP mechanisms. $\square$

APPENDIX E: PROOFS OF PRIVACY STATEMENTS

E.1. Auxiliary results.

LEMMA E.1. *For any fixed $l \in \{1, \dots, p-1\}$ for which $q_l(X) > 4\frac{r}{n}L_\infty$, the local ℓ_2 sensitivity of $(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top$ is 0.*

PROOF. Let $l \in \{1, \dots, p-1\}$ be arbitrary but fixed. We have to prove that $(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top$ and $(\hat{\mathbb{1}}_1(X'), \dots, \hat{\mathbb{1}}_p(X'))^\top$ remain unchanged (i.e. the set of indices of the l U-statistics with largest absolute value do not change) if we alter one entry of X . Changing any single entry of X at worst yields a change of size $2\frac{r}{n}L_\infty$ for $|U|_{(l)}$ and $|U|_{(l+1)}$. Hence, if the gap between those two is lower bounded by $q_l(X) > 4\frac{r}{n}L_\infty$, the index set does not change. Here it is important to note that the order of $|U|_{(1)}, \dots, |U|_{(l)}$ may indeed change, but the set of the corresponding indices does not. Consequently, we have

$$(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top = (\hat{\mathbb{1}}_1(X'), \dots, \hat{\mathbb{1}}_p(X'))^\top,$$

which yields the desired result

$$\max_{X', d_H(X, X')=1} \left\| (\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top - (\hat{\mathbb{1}}_1(X'), \dots, \hat{\mathbb{1}}_p(X'))^\top \right\|_2 = 0.$$

$\square$

LEMMA E.2. *Assume that Assumption 4.2 holds and let $Y \sim \mathcal{N}(0, \frac{(2L_\infty r/n)^2}{2\rho})$. Then,*

$$\|U\|_\infty^{\text{DP}} := \|U\|_\infty + Y$$

is ρ -zCDP.

PROOF. In order to bound the sensitivity $\Delta_2 \|U\|_\infty$, we first consider two neighboring data sets X, X' which differ without loss of generality in the last component. By definition, we have that

$$|\|U\|_\infty - \|U'\|_\infty| \leq \max_{1 \leq j \leq p} |U_j - U'_j| \leq \frac{2r}{n}L_\infty,$$

where we have used that, by Assumption 4.2 (B), for $j = 1, \dots, p$:

$$\begin{aligned} |U_j - U'_j| &= \binom{n}{r}^{-1} \left| \sum_{1 \leq i_1 < \dots < i_r \leq n} h_j(X_{i_1}, \dots, X_{i_r}) - \sum_{1 \leq i_1 < \dots < i_r \leq n} h_j(X'_{i_1}, \dots, X'_{i_r}) \right| \\ &= \binom{n}{r}^{-1} \left| \sum_{1 \leq i_1 < \dots < i_{r-1} < i_r = n} h_j(X_{i_1}, \dots, X_{i_{r-1}}, X_n) - h_j(X'_{i_1}, \dots, X'_n) \right| \\ &\leq \binom{n}{r}^{-1} 2L_\infty \binom{n-1}{r-1} = \frac{2r}{n} L_\infty. \end{aligned}$$

□

PROOF. Note that Algorithm 1 can be implemented using the exponential mechanism. Therefore, it suffices to analyze the sensitivity of $q_j + \nu(j)$. For two neighboring databases X, X' , we have

$$|q_j(X) + \nu(j) - (q_j(X') + \nu(j))| \leq |U_{(j)} - U'_{(j)} + U_{(j+1)} - U'_{(j+1)}| \leq 4 \frac{r}{n} L_\infty,$$

where we have used that the ℓ_1 sensitivity of an U -statistic with a kernel of order l bounded by L_∞ is bounded by $2 \frac{r}{n} L_\infty$. This yields ε -DP. For ρ -zCDP, we use the relation between ε -DP and ρ -zCDP given in Cesar and Rogers (2021). □

LEMMA E.3 (Lemma 3.4 in (Dunsche, 2025)). Assume that Assumption 4.2 holds. Then the Jackknife variance estimator defined in equation (2.5) has sensitivity

$$\Delta_2 \hat{\zeta}_1 = \frac{(n-1)r}{n(n-r)} \sum_{c=0}^r \frac{\binom{n-r+c}{r-c}}{\binom{n-1}{r}} \binom{r}{c} |cn - r^2| \sqrt{2d} L_\infty^2.$$

E.2. Proof of Theorem 4.1. For any $\alpha > 1$ we denote for two random variables Z, Y the Rényi divergence of the associated distributions by $D_\alpha(Z, Y)$. We further condition on a fixed realization of $\hat{k}$ and denote by $\mathcal{M}$ as the output of Algorithm 2. Here, we only provide the arguments for the divergence $D_\alpha(\mathcal{M}(X), \mathcal{M}(X'))$ and note that the divergence with the reverse order works analogously. Then for two neighboring databases X, X' , we have two distinct cases to investigate:

Case (1) We assume that the indicators are the same for neighboring databases X and X' , i.e.

$$(\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top = (\hat{\mathbb{1}}_1(X'), \dots, \hat{\mathbb{1}}_p(X'))^\top.$$

Recall that the possible outputs of $\mathcal{M}$ are $\perp$ or $\{\hat{i}_{(1)}, \dots, \hat{i}_{(\hat{k})}\}$ for both X and X' . Therefore, we can derive

$$D_\alpha(\mathcal{M}(X), \mathcal{M}(X')) = D_\alpha(\mathbb{1}\{\hat{q}_{\hat{k}}(X) > t\}, \mathbb{1}\{\hat{q}_{\hat{k}}(X') > t\}) \leq D_\alpha(\hat{q}_{\hat{k}}(X), \hat{q}_{\hat{k}}(X')),$$

where the first equality follows from the definition of Algorithm 2 and the second holds by the processing property of the Rényi divergence. Furthermore, by definition

$$\hat{q}_{\hat{k}}(X) \sim \mathcal{N}(q_{\hat{k}} - \sigma z_{1-\delta}, \sigma^2),$$

with $\sigma = t/\sqrt{\rho}$, $t = 4r/nL_\infty$, and from the sensitivity bound for bounded U -statistics we obtain

$$|q_{\hat{k}}(X) - q_{\hat{k}}(X')| \leq t.$$

Therefore, the Rényi divergence can be calculated explicitly and estimated as follows

$$D_\alpha(\hat{q}_{\hat{k}}(X), \hat{q}_{\hat{k}}(X')) = \frac{\alpha(q_{\hat{k}}(X) - q_{\hat{k}}(X'))^2}{2\sigma^2} = \frac{\alpha(q_{\hat{k}}(X) - q_{\hat{k}}(X'))^2}{2(t/\sqrt{\rho})^2} \leq \alpha\rho/2.$$

Consequently, we obtain δ -approximate- $\rho/2$ -zCDP for any subset E of our choice that has probability at least $1 - \delta$. We will specify a specific set in the proof of the second case.

Case (2) Now assume that

$$(E.1) \quad (\hat{\mathbb{1}}_1(X), \dots, \hat{\mathbb{1}}_p(X))^\top \neq (\hat{\mathbb{1}}_1(X'), \dots, \hat{\mathbb{1}}_p(X'))^\top,$$

then we have that $q_{\hat{k}}(X), q_{\hat{k}}(X') \leq t$, because otherwise there would be equality in (E.1). Defining the event $E = \{\hat{q}_{\hat{k}} \leq q_{\hat{k}}\}$, we have

$$\mathbb{P}(E) = \mathbb{P}(\hat{q}_{\hat{k}} \leq q_{\hat{k}}) \geq 1 - \delta,$$

where we have used that $\hat{q}_{\hat{k}} \sim \mathcal{N}(q_{\hat{k}} - \sigma z_{1-\delta}, \sigma^2)$. Therefore conditional on E , we have $\hat{q}_{\hat{k}} \leq q_{\hat{k}} \leq t$ for any neighboring X and X' , which yields

$$\mathbb{P}(\mathcal{M}(X) = \perp | E) = \mathbb{P}(\mathcal{M}(X') = \perp | E) = 1.$$

Thus, conditional on E , we have that $D_\alpha(\mathcal{M}(X) | \mathcal{M}(X')) = 0$ for all α . Therefore, we have δ -approximated- $\rho/2$ -zCDP for any $\rho > 0$.

Now, combining both Algorithm 1 $\rho/2$ and the δ -approximate- $\rho/2$ -zCDP in the two cases of this proof, we obtain by the composition δ -approximate- ρ -zCDP.

E.3. Proof of Theorem 4.4. Let X, X' be neighboring data sets. By Theorem 4.1 and its proof, selecting the set of relevant coordinates $\hat{\mathcal{E}}^{\text{DP}}$ by Algorithm 2 with privacy budget $\rho/3$ is δ -approximate- $(\rho/3)$ -zCDP, where the set E in Definition 2.3 can be chosen as $E := \{\hat{q}_{\hat{k}} \leq q_{\hat{k}}\}$ (see the proof of Theorem 4.1 for details).

Now, let S be the event that for at least one index $i \in \{1, \dots, \hat{k}\}$ one U_i and U'_i have a different sign. If i is such an index it follows that $|U|_i \leq 2r/nL_\infty$ since the sensitivity of U_i is at most $(2r/n)L_\infty$. As Algorithm 2 uses the threshold $t = (4r/n)L_\infty$, any sign change forces the bad event E^c from Theorem 4.1. In fact, when we have

$$|U|_{(\hat{k})} - |U|_{(\hat{k}+1)} \geq 4r/nL_\infty,$$

we necessarily also have $|U|_{(i)} \geq |U|_{(\hat{k})} > 2r/nL_\infty$ because $|U|_{(\hat{k}+1)} \geq 0$. Consequently, we have $S \subseteq E^c$ and

$$\mathbb{P}(S) \leq \mathbb{P}(E^c) \leq \delta.$$

On the event E (i.e., when no sign change occurs), GAUSSCOV is $(\rho/3)$ -zCDP. Hence, with probability at least $1 - \delta$, we have the following compositions:

- If the first branch of Algorithm 3 is taken, then $\|U\|_\infty$ is released with privacy cost $\rho/3$. Composing $\hat{\mathcal{E}}^{\text{DP}}$ ($\rho/3$), GAUSSCOV ($\rho/3$), and $\|U\|_\infty$ ($\rho/3$) yields total budget δ -approximate- ρ -zCDP.
- Otherwise, the alternative branch releases $\|U\|_\infty$ with cost $2\rho/3$. Composing with $\hat{\mathcal{E}}^{\text{DP}}$ ($\rho/3$) again gives total δ -approximate- ρ -zCDP.

Thus the overall mechanism is ρ -zCDP on E^c and fails only with probability at most δ , which yields the desired result.

APPENDIX F: STATISTICAL GUARANTEES

F.1. Some results on bounded U -statistics.

LEMMA F.1. *Consider a U -statistic U of fixed order r , kernel h and expected value $\theta = \mathbb{E}[U] \in \mathbb{R}^p$. If $\log(p) = o(n^{1/3})$ and $\|h\|_\infty \leq L_\infty$ we have*

$$\max_{1 \leq i \leq j \leq p} |\hat{\zeta}_{1,ij} - \zeta_{1,ij}| \lesssim L_\infty^2 \sqrt{\frac{\log(np)}{n}}$$

with probability at least $1 - o(1)$. Here ζ_1 and $\hat{\zeta}_1$ are defined in (2.4) and (2.5), respectively.

PROOF. The case where the maximum is taken only over $1 \leq i = j \leq p$ can be found in [Zhou et al. \(2019\)](#). Our case can be handled by exactly the same arguments. $\square$

LEMMA F.2. *Consider a U -statistic U of order r with kernel h and expected value $\theta = \mathbb{E}[U] \in \mathbb{R}^p$. Assume that $\|h\|_\infty \leq L_\infty$ and that $\log p = o(n^{1/3})$. Then*

(1) *we have that*

$$\|U - \theta\|_\infty \leq L_\infty \sqrt{\frac{8r \log(p \vee n)}{n}}$$

holds with probability at least $1 - o(1)$.

(2) *Suppose that $\theta = \mathbb{E}[U]$ satisfies $\min_{1 \leq i \leq p} |\theta_i| > \underline{c}$ for some $\underline{c} > 0$. Then*

$$\mathbb{P}(\text{sign}(U_i) = \text{sign}(\theta_i), i = 1, \dots, p) = 1 - o(1).$$

PROOF. Both results follow by simple applications of [Hoeffding \(1963\)](#)'s inequality and the union bound. $\square$

LEMMA F.3. *Consider a U -statistic U of order r with $\|h\|_\infty \leq L_\infty$. Assume that $\log(p) = o(n^{1/5})$ and that $\min_{1 \leq i \leq p} |\theta_i| \geq \underline{c} > 0$. Then there exists a zero-mean Gaussian vector $Z = (Z_1, \dots, Z_p)^\top$ with*

$$\text{Cov}(Z_i, Z_j) = r^2 \zeta_{1,ij} \text{sign}(\theta_i \theta_j)$$

such that, uniformly with respect to t ,

$$\mathbb{P}\left(\sqrt{n}\left(\max_{1 \leq i \leq p} |U_i| - \max_{1 \leq i \leq p} |\theta_i|\right) \geq t\right) \leq \mathbb{P}\left(\max_{1 \leq i \leq p} Z_i \geq t\right) + o(1)$$

with equality whenever $|\theta_j| = \max_{1 \leq i \leq p} |\theta_i|$ for all $j = 1, \dots, p$.

PROOF. This result is established in the course of the proof of Theorem 2.2 in [Bastian et al. \(2024\)](#). $\square$

F.2. Analysis of Algorithm 1.

LEMMA F.4. *Assume that condition (4.5) holds, and define $q_j^\theta := |\theta|_{(j)} - |\theta|_{(j+1)}$ and $k := \arg \max_{j=1, \dots, p-1} q_j^\theta$. Then the output $\hat{k}$ of Algorithm 1 fulfills*

$$\mathbb{P}(\hat{k} = k) = 1 - o(1).$$

PROOF. By Lemma F.2 there exist constants $a_n \lesssim \sqrt{\log(n \vee p)/n}$ such that

$$|U|_{(j)} \in \cup_{i=1}^p I_i, \quad j = 1, \dots, p$$

with high probability, where the (not necessarily disjoint) intervals I_i are defined by $I_i = (|\theta|_{(i)} - a_n, |\theta|_{(i)} + a_n)$. We may obtain $|U|_{(1)}$ by first picking the largest $|U_j|$ contained in the interval I_1 . By construction the remaining $|U_j|$ are contained in $\cup_{i=2}^p I_i$ with at least one element contained in I_2 . We may pick the largest (which is necessarily in I_2) to obtain $|U|_{(2)}$. Continuing like this we pick $|U|_{(i)}$ from I_i , yielding that with high probability

$$\max_{1 \leq i \leq p} ||U|_{(i)} - |\theta|_{(i)}| \leq a_n$$

In particular it follows by the triangle inequality that

$$\max_{1 \leq i \leq p-1} |q_i - q_i^\theta| = O_{\mathbb{P}}(\sqrt{\log(n \vee p)/n}).$$

By equation (4.5) we know that

$$q_k^\theta > \max_{l \neq k} q_l^\theta + \sqrt{\log(n) \log(p \vee n)/n}.$$

Consequently it holds with high probability that

$$(F.1) \quad q_k > \max_{l \neq k} q_l + 1/2 \sqrt{\log(n) \log(p \vee n)/n}.$$

Denote by G_j the Gumbel noise added to q_j in Algorithm 1. Because $\mathbb{E}[\exp(nG_j/2)] < \infty$ we have that

$$\max_{j=1, \dots, p-1} G_j = O_{\mathbb{P}}(\log(p)/n) = o_{\mathbb{P}}(\sqrt{\log(n) \log(p \vee n)/n}).$$

Consequently (up to changing the constant 1/2) (F.1) remains unaffected by the addition of Gumbel noise to the queries q_l so that with high probability q_k is the largest query, as desired. $\square$

F.3. Proof of Theorem 4.3. We distinguish two cases corresponding to the null hypothesis and alternative.

Case 1: $\|\theta\|_\infty \leq \Delta$. Due to the fact that

$$(F.2) \quad \sqrt{n}(\|\tilde{U}\|_\infty^{\text{DP}} - \|\tilde{U}\|_\infty) = O_{\mathbb{P}}(n^{-1/2}),$$

we have by Lemma C.4 in Bastian et al. (2024) for any $\delta > 0$ that

$$\begin{aligned} \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty^{\text{DP}} - \Delta) > t) &= \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) > t + \sqrt{n}(\|\tilde{U}\|_\infty - \|\tilde{U}\|_\infty^{\text{DP}})) \\ &\leq \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) > t) + O\left(n^{-1/2+\delta} \sqrt{\log(p)} + \frac{\log(np)^{5/4}}{n^{1/4}}\right) + o(1) \\ &= \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) > t) + o(1) \end{aligned}$$

We can use the same arguments to obtain the reverse inequality, which gives

$$\mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty^{\text{DP}} - \Delta) > t) = \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) > t) + o(1).$$

Note that by Lemma F.2, we have for $\|\theta\|_\infty < \Delta - \gamma$ that

$$(F.3) \quad \sqrt{n}(\|\tilde{U}\|_\infty - \Delta) = O_{\mathbb{P}}\left(\sqrt{\log(p \vee n)/n}\right) - \sqrt{n}(\Delta - \|\theta\|_\infty) \xrightarrow{\mathbb{P}} -\infty.$$

From this the first statement of the Theorem follows: the test statistic diverges to $-\infty$ and the quantiles $q_{1-\alpha}$ in Algorithm 4 and Algorithm 5 are stochastically bounded below by 0 for $\alpha < 0.5$.

Let us now consider statement (ii), that is the case $\|\theta\|_\infty \leq \Delta$. By Assumption (P) we may condition on $\hat{i}_1 = j_1, \dots, \hat{i}_{\hat{k}} = j_l, \hat{k} = l$ so that the output of Algorithm 2 is deterministic and given by (a possibly randomly selected subset of) $j_1, \dots, j_l$ or $\perp$. For the remainder of the proof we will hence simply write l instead of $\hat{k}$. In the former case we additionally condition on the random subselection, so that we may assume WLOG (the selection is independent of everything else) that $l \leq \log(p)$ and that $\hat{i}_1, \dots, \hat{i}_{\hat{k}}$ are fixed. To be precise Assumption (P) yields that with high probability

$$\begin{aligned} & \mathbb{P}(U_{j_m} \geq t_m \text{ for all } m = 1, \dots, l) \\ &= \mathbb{P}(\tilde{U}_m \geq t_m \text{ for all } m = 1, \dots, \hat{k} | \hat{i}_1 = j_1, \dots, \hat{i}_{\hat{k}} = j_l, \hat{k} = l) \end{aligned}$$

so that in all following considerations we may simply assume that $\tilde{U} = (U_{j_1}, \dots, U_{j_l})^\top$ or $\tilde{U} = U$. We will argue each case separately.

Subcase 1(a): Algorithm 2 returns $\{\hat{i}_1, \dots, \hat{i}_{\hat{k}}\} = \{j_1, \dots, j_l\}$: From now on assume that $\|\theta\|_\infty > 0$ (otherwise the previous arguments apply). We define for some $\gamma > 0$ the set

$$I_1 := \{i \in \{j_1, \dots, j_l\} \mid |\theta_i| > \Delta - \gamma\}$$

and the associated statistic

$$\hat{S} = \sqrt{n} \max_{i \in I_1} (|U_i| - \Delta).$$

whose coordinate-variances are lower bounded by Assumption 4.2(V). Lemma F.3 (with $p = |I_1|$) then yields

$$(F.4) \quad \mathbb{P}(\hat{S} > t) \leq \mathbb{P}(\max_{i \in I_1} Z_i > t) + o(1) \leq \mathbb{P}\left(\max_{i \in \{j_1, \dots, j_l\}} Z_i > t\right) + o(1).$$

where the precise definition of the Gaussian random variables $(Z_i)_{i \in I_1}$ is given in Lemma F.3. Finally we note that

$$\sqrt{n} \left(\max_{i \in \{j_1, \dots, j_l\} \setminus I_1} |U_i| - \Delta \right) \xrightarrow{\mathbb{P}} -\infty.$$

which implies

$$\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) = \max \left\{ \hat{S}, \sqrt{n} \left(\max_{i \in \{j_1, \dots, j_l\} \setminus I_1} |U_i| - \Delta \right) \right\} = \hat{S} + o_{\mathbb{P}}(1).$$

Combining this estimate with (F.2) and (F.4) gives, for any $t \in \mathbb{R}$,

$$\begin{aligned} \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty^{\text{DP}} - \Delta) > t) &= \mathbb{P}(\sqrt{n}(\|\tilde{U}\|_\infty - \Delta) > t) + o(1) \\ &= \mathbb{P}(\hat{S} > t) + o(1) \\ &\leq \mathbb{P}\left(\max_{i \in \{j_1, \dots, j_l\}} Z_i > t\right) + o(1). \end{aligned}$$

With this at our disposal we now only need to establish that, on a set with high probability, the inequality

$$\left| \mathbb{P}\left(\max_{i \in \{j_1, \dots, j_l\}} Z_i > t\right) - \mathbb{P}^*\left(\max_{i \in \{j_1, \dots, j_l\}} Z_i^{\text{DP}} > t\right) \right| \leq c_n$$

holds for some sequence $c_n \rightarrow 0$. Here $Z_i^{\text{DP}} \stackrel{iid}{\sim} \mathcal{N}_k(0, (\hat{\zeta}_1(k))^{\text{DP}})$ and $\mathbb{P}^*$ denotes the probability space obtained by conditioning on the data and the privacy noise of the covariance privatization. The desired statement is now a consequence of Theorem 2 from [Chernozhukov et al. \(2015\)](#). For the application of this result it suffices to establish that

$$(\log(l \wedge \log(p)))^2 \hat{\Delta} = o_{\mathbb{P}}(1)$$

where

$$\hat{\Delta} = \max_{1 \leq h < j \leq l \wedge \log(p)} |\hat{\zeta}_{1,hj}^{\text{DP}} - \zeta_{1,hj}|.$$

However, this statement follows from Lemma F.1 and the bound on the privatization error incurred by Algorithm 6 (similar to the proof of Lemma 3.5 in [Dunsche, 2025](#)).

Subcase 1(b): Algorithm 2 returns $\perp$: In this case we simply have $\tilde{U} = U$. By the same argument that yields (F.3) and by Lemma F.2 we obtain

$$\begin{aligned} \sqrt{n} \max_{1 \leq i \leq p} (|U_i| - \Delta) &\leq \sqrt{n} \max_{i \in \tilde{I}_1} (|U_i| - \Delta) + o_{\mathbb{P}}(n^{-1/2}) \\ &= \sqrt{n} \max_{i \in \tilde{I}_1} \text{sign}(\theta_i)(U_i - \theta_i) + o_{\mathbb{P}}(n^{-1/2}) =: T_n + o_{\mathbb{P}}(n^{-1/2}), \end{aligned}$$

where the last line defines in an obvious manner and the set $\tilde{I}_1$ is defined by

$$\tilde{I}_1 := \{i \mid 1 \leq i \leq p; |\theta_i| \geq \Delta - \gamma\}$$

Next we may apply Lemma F.3 to obtain the existence of a Gaussian vector $Z = (Z_i)_{i \in \tilde{I}_1}$ with covariances given by

$$r^2 \zeta_{1,ij} \text{sign}(\theta_i \theta_j), \quad i, j \in \tilde{I}_1$$

such that

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}(T_n \geq t) - \mathbb{P}\left(\max_{i \in \tilde{I}_1} Z_i \geq t\right) \right| = o(1).$$

By the Hoeffding decomposition and assumption (B) we have

$$\sup_{1 \leq i < j \leq p} |\text{Cov}(U_i, U_j) - r^2 \zeta_{1,ij}| = O(n^{-1}),$$

and, by Lemma 2 from [Chernozhukov et al. \(2015\)](#), there exists another Gaussian vector Z^1 of the same dimension with covariance structure given by the covariance of the vector

$$(\text{sign}(U_i)(U_i - \theta_i))_{i \in \tilde{I}_1},$$

such that

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}\left(\max_{i \in \tilde{I}_1} Z_i^1 \geq t\right) - \mathbb{P}\left(\max_{i \in \tilde{I}_1} Z_i \geq t\right) \right| = o(1).$$

Note that $|\theta_i| \geq \Delta - \gamma$ implies $\text{Var}(Z_i^1) \leq L_\infty^2 - (\Delta - \gamma)^2$ by the [Bhatia and Davis \(2000\)](#) inequality. Pad the vector Z_i^1 with some additional independent normal random variables with variances given by $L_\infty^2 - (\Delta - \gamma)^2$ and denote the resulting random vector by $\tilde{Z}$. Clearly,

$$\max_{i \in \tilde{I}_1} Z_i^1 \leq \max_{1 \leq i \leq p} \tilde{Z}_i.$$

By Slepian's Lemma and assumption (E) we then have for $t \geq 0$ that

$$\mathbb{P}\left(\sqrt{n} \max_{1 \leq i \leq p} \tilde{Z}_i \geq t\right) \leq \mathbb{P}\left(\sqrt{n} \max_{1 \leq i \leq p} Y_i \geq t\right),$$

where Y_i is a collection of iid $\mathcal{N}(0, L_\infty^2 - (\Delta - \gamma)^2)$ random variables. The right hand side maximum converges, appropriately rescaled, in distribution to a Gumbel distribution with scale parameter $\sqrt{L_\infty^2 - (\Delta - \gamma)^2}$. In particular we obtain that

$$\limsup_n \mathbb{P}(\sqrt{n} \max_{1 \leq i \leq p} (|U_i| - \Delta) > t) \leq \limsup_n \mathbb{P}(\sqrt{n} \max_{1 \leq i \leq p} Y_i \geq t)$$

for any $t \geq 0$. Letting $t = q_{1-\alpha}^G/a_p + a_p - \frac{\log \log(p) + \log(4\pi)}{2a_p}$ then yields

$$\limsup_n \mathbb{P}(\sqrt{n} \max_{1 \leq i \leq p} Y_i \geq t) = \alpha$$

as desired. Here $q_{1-\alpha}^G$ is the $(1 - \alpha)$ -quantile of a Gumbel distribution with scale parameter $\sqrt{L_\infty^2 - (\Delta - \gamma)^2}$.

Case 2: $\|\theta\|_\infty > \Delta$ If Algorithm 2 outputs $\perp$ i.e. when we use the Gumbel test, we may use exactly the same arguments as given in the proof of Theorem 2.4 in Bastian et al. (2024) for the non-private setting. By equation (F.2) we may reduce to this case and are done.

In the other case, again using (F.2) to reduce to the non-private setting, let i_0 be an index for which $\|\theta\|_\infty = |\theta_{i_0}|$ (we suppress the possible dependence on n in our notation). As (4.5) is satisfied, we have by Lemma F.4 that $\hat{k} = k \leq \log(p)$ with high probability. This implies that $i_0 = \hat{i}_1$, where we recall $\hat{i}_1$ is the index of the largest U-statistic in absolute value. From that we can obtain

$$(F.5) \quad \sqrt{n}(\|\tilde{U}\|_\infty - \Delta) \geq \sqrt{n}(|U_{\hat{i}_1}| - |\theta_{\hat{i}_1}|) + \sqrt{n}(|\theta_{\hat{i}_1}| - \Delta).$$

By Lemma F.2 we further obtain that

$$\sqrt{n} \left| |U_{\hat{i}_1}| - |\theta_{\hat{i}_1}| \right| \lesssim \sqrt{n} \|U - \theta\|_\infty \leq r L_\infty \sqrt{2 \log(p \vee n)}$$

with probability $1 - o(1)$. The second term on the right hand side of (F.5) converges to $+\infty$ at rate $c\sqrt{\log(p \vee n)}$, yielding the desired conclusion upon choosing c sufficiently large.

REFERENCES

- Bastian, P., Dette, H., and Heiny, J. (2024). Testing for practically significant dependencies in high dimensions via bootstrapping maxima of U-statistics. *The Annals of Statistics*, 52(2):628 – 653.
- Bhatia, R. and Davis, C. (2000). A better bound on the variance. *The American Mathematical Monthly*, 107(4):353–357.
- Cárcamo, J., Cuevas, A., and Rodríguez, L.-A. (2020). Directional differentiability for supremum-type functionals: Statistical applications. *Bernoulli*, 26(3):2143 – 2175.
- Cesar, M. and Rogers, R. (2021). Bounding, concentrating, and truncating: Unifying privacy loss composition for data analytics. In *Algorithmic Learning Theory*, pages 421–457. PMLR.
- Chernozhukov, V., Chetverikov, D., and Kato, K. (2015). Comparison and anti-concentration bounds for maxima of gaussian random vectors. *Probability Theory and Related Fields*, 162(1):47–70. Published online: 09 May 2014.
- Dunsche, M. (2025). *Statistical inference under differential privacy*. doctoral thesis, Ruhr-Universität Bochum, Universitätsbibliothek.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30.
- Vaart, A. W. v. d. (1998). *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Zhou, C., Han, F., Zhang, X.-S., and Liu, H. (2019). An extreme-value approach for testing the equality of large U-statistic based correlation matrices. *Bernoulli*, 25(2):1472 – 1503.
- Zhu, Y. and Wang, Y.-X. (2020). Improving sparse vector technique with renyi differential privacy. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20249–20258. Curran Associates, Inc.