

Optimal Distributed Composite Testing in High-dimensional Gaussian Models with 1-bit Communication

Szabo, Botond; Vuursteen, Lasse; Van Zanten, Harry

DOI

[10.1109/TIT.2022.3150599](https://doi.org/10.1109/TIT.2022.3150599)

Publication date

2022

Document Version

Final published version

Published in

IEEE Transactions on Information Theory

Citation (APA)

Szabo, B., Vuursteen, L., & Van Zanten, H. (2022). Optimal Distributed Composite Testing in High-dimensional Gaussian Models with 1-bit Communication. *IEEE Transactions on Information Theory*, 68(6), 4070-4084. Article 9709332. <https://doi.org/10.1109/TIT.2022.3150599>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Optimal Distributed Composite Testing in High-Dimensional Gaussian Models With 1-Bit Communication

Botond Szabó, Lasse Vuursteen^{ib}, and Harry van Zanten

Abstract—In this paper we study the problem of signal detection in Gaussian noise in a distributed setting where the local machines in the star topology can communicate a single bit of information. We derive a lower bound on the Euclidian norm that the signal needs to have in order to be detectable. Moreover, we exhibit optimal distributed testing strategies that attain the lower bound.

Index Terms—Testing, distributed algorithms, hypothesis testing, minimax lower bounds, Gaussian noise, federated learning.

I. INTRODUCTION

THE rapidly increasing amount of available data in many fields of application has triggered the development of distributed methods for data analysis. Distributed methods, besides being able to speed up computations considerably, can reduce local memory requirements and can also help in protecting privacy, by refraining from storing a whole dataset in a single central location. Moreover, distributed methods occur naturally when data is by construction observed and processed at multiple locations, for instance in astronomy, meteorology, seismology, military radar or air traffic control systems.

In the context of decentralized detection, the task of distinguishing between different signals based on information provided from a network of sensors has been investigated since the late seventies. Motivated by applications such as surveillance systems and wireless communication, statistical hypothesis testing with distributed sensors has received considerable attention over the past few decades, e.g. [1]–[4]. Most of this literature has been concerned with distinguishing between finitely many signals, by combining the decisions of sensors

or machines receiving (noisy) data from the same underlying signal [5], [6]. The statistical hypotheses considered at the time where either simple, or reduce to a simple hypothesis. Under the name of multiterminal data compression, hypothesis testing and estimation with observations in fixed finite sample spaces (alphabets) was investigated [7]–[12]. In this body of literature, each terminal observes a stream of iid observations in some alphabet, where the different terminals might receive observations from different distributions on different alphabets. Each of the terminals compresses the iid observations into a message that is sent to a central terminal, at which some inference goal is to be achieved based on the messages. For an overview, see [13].

The information theoretic aspects of distributed statistical methods have only been studied rigorously in terms of sample complexity relatively recently. The problem of distributed testing with a composite alternative hypothesis about a high-dimensional parameter with limited communication was first considered in a few recent papers. In [14], theoretical guarantees are derived for distributed uniformity testing of a discrete distribution in the case that a collection of machines each receive one observation. Here, sample complexity (in the sense of minimax convergence rate) is assessed in terms the cardinality of the sample space, the number of bits available for communication and the number of observations. In that paper, it is also shown that in the distributed setting with many local machines, testing performance can strictly improve when all machines have access to a shared source of randomness, a so-called *public coin*. Distributed uniformity testing of a discrete distribution when multiple observations per machine are available is considered in [15] and [16].

Most of the work, studying distributed inference in terms of sample complexity, up until now has focused on distributed methods for estimating a signal in the normal-means model under bandwidth, or communication restrictions (see for instance [17]–[21]). The canonical normal-means model postulates that we have an observation X satisfying

$$X = \mu + \frac{1}{\sqrt{n}}Z,$$

where $\mu \in \mathbb{R}^d$ is the unknown signal, $Z \sim N_d(0, I_d)$ is an unobserved noise vector with a d -dimensional standard Gaussian distribution, and n is the signal-to-noise ratio. Note that this is equivalent to observing n independent copies of a $N_d(\mu, I_d)$ -vector. In the distributed setting considered in the aforementioned articles, n such observations are distributed

Manuscript received December 10, 2020; revised August 26, 2021; accepted January 31, 2022. Date of publication February 9, 2022; date of current version May 20, 2022. This work was supported by the European Research Council (ERC) through the European Union's Horizon 2020 Research and Innovation Programme under Grant 101041064. (Corresponding author: Lasse Vuursteen.)

Botond Szabó is with the Department of Decision Sciences, Bocconi Institute for Data Science and Analytics (BIDSA), Bocconi University, 20100 Milano, Italy (e-mail: botond.szabo@unibocconi.it).

Lasse Vuursteen is with the Delft Institute of Applied Mathematics (DIAM), Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: l.vuursteen@tudelft.nl).

Harry van Zanten is with the Department of Mathematics, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, The Netherlands (e-mail: j.h.van.zanten@vu.nl).

Communicated by M. Lops, Associate Editor for Signal Processing and Source Coding.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2022.3150599>.

Digital Object Identifier 10.1109/TIT.2022.3150599

across m machines, which then communicate a transcript to a central machine. The central machine then forms an estimate of μ on the basis of these transcripts. The sample complexity of such a problem is expressed in n , m , the number of bits available for communication of the transcripts and the properties of μ (e.g. Euclidian norm). Related to this normal-means model, on deriving minimax lower bounds and optimal distributed estimation strategies in the context of nonparametric regression, density estimation and Gaussian signal-in-white-noise models (e.g. [20], [22]–[25]).

In this paper, we investigate the information theoretic properties of distributed methods for *testing* for the presence of a signal in the normal-means model. The theory on distributed testing in this setting is much less developed than that for estimation. Testing for the presence of a signal in the normal-means model translates to testing the null hypothesis $H_0 : \mu = 0$ that the sequence is identically equal to 0. Rejecting this hypothesis means declaring that there is a non-zero signal. A fundamental question is how large the signal should be in order to be detectable. It is well known that in this non-distributed model, the size of the signal (by which we mean the Euclidean norm) has to be of order $d^{1/4}/\sqrt{n}$ in order for the signal to be detectable (see e.g. [26]). An optimal test is for instance obtained by rejecting the hypothesis H_0 if $\|X\|^2$ is above a specific threshold, depending on d and n (see also Section III-B).

The question we address in this paper is how this changes in the distributed setting. In our analysis we consider the distributed version of the normal-means model, where the data (n independent draws from the $N_d(\mu, I_d)$ -distribution) is divided over $m \leq n$ local machines, or cores. Equivalently, we assume that at each local machine $j \in \{1, \dots, m\}$ we observe a vector $X^{(j)}$ which satisfies

$$X^{(j)} = \mu + \sqrt{\frac{m}{n}} Z^{(j)}, \quad (1)$$

where again $\mu \in \mathbb{R}^d$ and the $Z^{(j)}$ are independent $N_d(0, I_d)$ -distributed vectors. Each machine carries out a test for the hypothesis $H_0 : \mu = 0$ using its local data $X^{(j)}$, where we allow that the machines use a public coin, i.e. a common random vector U that is available to all local machines. Subsequently, the outcomes of the m local tests (which are single bits) are sent to a central machine, where they are combined into a single, overall test. We prove that in this distributed setting, the size of the signal has to be of the order $(d(m \wedge d))^{1/4}/\sqrt{n}$ in order for the signal to be detectable. Moreover, we exhibit optimal distributed tests that achieve this detection bound.

The detection bound has a remarkable “regime change” or “elbow effect” at $m = d$. As m grows from 1 to d the testing problem becomes more difficult, in the sense that the signal needs to be ever larger to be detectable using an increasing number of machines. This is intuitively understandable, since as m grows, the local signal-to-noise ratio decreases, so it is reasonable to expect that the signal needs to be larger to be able to detect it. The detection bound stops increasing if m grows above d however. In that range, the decrease of the local signal-to-noise ratios is apparently balanced by the increase of

the number of bits that are transmitted from the local machines to the central one.

The regime change is also reflected by the fact that we need different testing strategies depending on how m and d are related. If m is below some threshold, and in particular it does not increase with n , then, not surprisingly, we can simply use the classical non-distributed test mentioned above at one of the local machines. If m is larger than this threshold but $m \leq d$, then it is still possible to construct an optimal test by combining local tests that are based on the test statistics $\|X^{(j)}\|^2$, but the test needs to be constructed more carefully. In the range $m \geq d$ this strategy becomes sub-optimal and we have to adopt a different approach, using the fact that we have a public coin at our disposal. For this case we construct an optimal distributed test by appropriately combining local tests that use local test statistics of the form $U^\top X^{(j)}$, where U is a public random vector.

The approach to finding the lower bound can be summarized as follows. As a first step, we lower bound the testing risk by a type of Bayes risk, where μ is drawn from a prior such that it either belongs to the null hypothesis of the alternative, as in for example [27]. This Bayes risk can be related to the mutual information between the testing outcome and which hypothesis is selected, akin to techniques common in tackling (distributed) estimation problems through Fano-like inequalities in for example [17], [28]. In particular, the tensorization property of the mutual information is used and combined with a so called strong data processing inequality to quantify the loss incurred in the distributed setup, similar to the approaches to distributed estimation in [18], [19], and [21].

Upon completion of this work, we came across the paper [29], which considers a setting similar to ours and claims some partly overlapping results. There are also important differences between the papers, however. Most importantly perhaps, our proof strategy for the lower bound is rather different. As a result our proof is arguably easier to verify and at least provides an alternative to the approach proposed in [29]. Furthermore, the paper [29] does not allow the number of machines (our m) to vary. Essentially, only the case that $m = n$ is considered.

The remainder of the paper is organized as follows. In Section II we formally describe the model and the distributed testing problem, and introduce notations used throughout the paper. In Section III we present our main results. We state the detection lower bound in Section III-A, and we provide novel distributed tests achieving the theoretical limits in Section III-B. We provide a short simulation study demonstrating the regime change observed in the theoretical analysis in Section IV. The proofs for the distributed tests achieving the minimax rate is given in Section V, while the proofs of corresponding (technical) lemmas are deferred to Sections A-C.

II. PROBLEM SETTING AND NOTATION

We assume we have m local machines. For $j = 1, \dots, m$, we have an observation $X^{(j)}$ at machine j , which satisfies

$$X^{(j)} = \mu + \sqrt{\frac{m}{n}} Z^{(j)}.$$

Here $\mu \in \mathbb{R}^d$ is the unknown signal of interest and $Z^{(1)}, \dots, Z^{(m)}$ are independent, $N_d(0, I_d)$ -distributed vectors. We allow both the dimension $d = d_n$ and the number of machines $m = m_n$ to depend on the overall signal-to-noise ratio, or “sample size” n . In fact, the interesting cases (from an asymptotic perspective) are the ones where both d and m are tending to infinity with n . Nevertheless, we do not restrict ourselves to this asymptotic regime and cover the finite m and d cases as well.

We are interested in distributed tests for the hypotheses

$$H_0 : \mu = 0, \quad \text{against} \quad H_\rho : \|\mu\| \geq \rho, \quad (2)$$

for $\rho > 0$. We consider public coin protocols, where each machine has access to a shared random vector U , which is independent of the observations $X^{(1)}, \dots, X^{(m)}$. Each local machine j carries out a local test. Using the local data $X^{(j)}$ and the public coin U it produces a binary, $\{0, 1\}$ -valued outcome $T^{(j)}$. The outcomes $T^{(1)}, \dots, T^{(m)}$ of the local tests are transmitted to a central machine where they are aggregated into a global test, described by a $\{0, 1\}$ -valued variable T_{dist} . Schematically, a distributed test looks as follows:

$$\begin{array}{ccc} (X^{(1)}, U) & \longrightarrow & T^{(1)} \searrow \\ \vdots & \longrightarrow & \vdots \longrightarrow T_{dist} \\ (X^{(m)}, U) & \longrightarrow & T^{(m)} \nearrow \end{array} \quad (3)$$

We denote the collection of all distributed tests of this form by \mathcal{T}_{dist} .

The testing error, or risk of a distributed test T_{dist} , is defined as usual by

$$\mathcal{R}(H_\rho, T_{dist}) = P_0(T_{dist} = 1) + \sup_{\|\mu\| \geq \rho} P_\mu(T_{dist} = 0), \quad (4)$$

i.e. as the sum of the type one and type two errors of the test. (Here, and elsewhere, we denote by P_μ the underlying distribution assuming that μ is the true signal.) Uniform lower bounds for this risk express the impossibility of detecting a signal of size ρ . Indeed, fix a level $\alpha \in (0, 1)$. If $\rho > 0$ is such that $\mathcal{R}(H_\rho, T_{dist}) > \alpha$ for all $T_{dist} \in \mathcal{T}_{dist}$, then it means that there exists no consistent level- α test for testing H_0 against H_ρ . In other words, no distributed test of level α is able to detect all signals that are larger than ρ in Euclidean norm.

Our aim is to find the detection threshold, i.e. the cut-off ρ_{dist} such that no consistent level- α test exists if $\rho < \rho_{dist}$ and at least one consistent level- α test exists if $\rho \geq \rho_{dist}$. We will show that, up to constants depending on the chosen level α , the detection threshold is given by

$$\rho_{dist}^2 \asymp \min \left\{ \frac{\sqrt{dm}}{n}, \frac{d}{n} \right\}.$$

Moreover, we exhibit optimal tests for the case $\rho \geq \rho_{dist}$.

A. Notation

We write $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. For two positive sequences a_n, b_n we use the notation $a_n \lesssim b_n$ if there exists a universal positive constant C such that $a_n \leq C b_n$. We write $a_n \asymp b_n$ which holds if $a_n \lesssim b_n$ and $b_n \lesssim a_n$

are satisfied simultaneously. The Euclidean norm of a vector $v \in \mathbb{R}^d$ is denoted by $\| \cdot \|$. For absolutely continuous probability measures $P \ll Q$, we denote by $D_{KL}(P||Q) = \int \log \frac{dP}{dQ} dP$ their Kullback-Leibler divergence.

III. MAIN RESULTS

A. Lower Bound for the Detection Threshold

The following theorem establishes the detection threshold. Its proof is described in the remainder of the subsection.

Theorem 1: Fix $\alpha \in (0, 1)$ and suppose that

$$\rho^2 < c_\alpha \frac{\sqrt{d(m \wedge d)}}{n} \quad (5)$$

for $c_\alpha \leq (1 - \alpha)^2/384$. Then,

$$\inf_{T \in \mathcal{T}_{dist}} \mathcal{R}(H_\rho, T) > \alpha,$$

where infimum is over all distributed tests $T \in \mathcal{T}_{dist}$ given in (3).

The result tells us that if (5) holds, there does not exist a consistent test in \mathcal{T}_{dist} of level $\alpha \in (0, 1)$ for the hypotheses (2). In other words, no distributed test can detect all signals of size ρ . It should be noted that we did not optimize for the value of the constant c_α and the statement is likely to be still true for larger values of c_α .

The proof of the theorem relies on three key lemmas, which we state below after introducing some necessary notations. As a first step, we use the basic fact that the supremum of the probability of a type two error of a test can be lower bounded by a Bayesian type two error, i.e. for any prior distribution Π supported on H_ρ

$$\sup_{\mu \in H_\rho} P_\mu(T = 0) \geq \int_{H_\rho} P_\mu(T = 0) d\Pi(\mu).$$

To further lower bound the risk we construct an appropriate Markov chain and relate the testing problem to an information transfer problem through the chain. Consider $V \sim \text{Ber}(1/2)$, i.e. a V is 0 or 1, each with probability 1/2, independent of the public coin random vector U , such that the random vectors $X^{(j)}|V=0$, $j = 1, \dots, m$ follow (1) with $\mu = 0$ and $X^{(j)}|V=1$ follows a Gaussian mixture P_Π defined as $P_\Pi(A) = \int P_\mu(A) d\Pi(\mu)$ for all Borel sets $A \subset \mathbb{R}^d$. Let us denote by \mathbb{P} the joint probability measure describing the corresponding Markov dynamics

$$V \longrightarrow \mu \begin{array}{ccc} \nearrow & (X^{(1)}, U) & \longrightarrow T^{(1)} \searrow \\ \longrightarrow & \vdots & \longrightarrow \vdots \longrightarrow T \\ \searrow & (X^{(m)}, U) & \longrightarrow T^{(m)} \nearrow \end{array} \quad (6)$$

We then have that for any distributed test T ,

$$\begin{aligned} \mathcal{R}(H_\rho, T) &\geq \mathbb{P}(T = 1|V = 0) + \mathbb{P}(T = 0|V = 1) \\ &= 2\mathbb{P}(T \neq V). \end{aligned} \quad (7)$$

The right hand side of (7) can be further bounded from below using the *mutual information* between T and V in the chain (6), defined by

$$I_\Pi(V, T) = D_{KL}(\mathbb{P}^{V \times T} || \mathbb{P}^V \times \mathbb{P}^T),$$

where \mathbb{P}^V , \mathbb{P}^T and $\mathbb{P}^{V \times T}$ denote marginal- and joint distributions of V and T , and the subscript Π is used to indicate the dependence on the prior Π . Informally, the mutual information measures how much knowing T reduces uncertainty about V and vice versa. The lower bound based on the mutual information is given in the following lemma. The proof of the lemma is deferred to Section B-A.

Lemma 1: Let Π be a prior on H_ρ and consider the dynamics (6). We have

$$\inf_{T \in \mathcal{T}_{dist}} \mathcal{R}(H_\rho, T) \geq 1 - \sqrt{2I_\Pi(V, T)}.$$

In view of the usual data processing inequality we have $I_\Pi(V, T) \leq I_\Pi(V, (T^{(1)}, \dots, T^{(m)}))$. The following lemma asserts that, up to an additional term, this further tensorizes conditional on the public coin randomness.

Lemma 2: Consider the dynamics (6). We have

$$\begin{aligned} I_\Pi(V, (T^{(1)}, \dots, T^{(m)})) \\ \leq \sum_{j=1}^m I_\Pi(V, T^{(j)}|U) + \sum_{j=1}^m I_\Pi(\mu, T^{(j)}|U, V). \end{aligned} \quad (8)$$

The proof of this lemma is given in Section B-B. This bound, combined with Lemma 1, allows us to break down the difficulty of the 'global' testing problem in terms of the difficulty of the m 'local' testing problems, captured by the quantities $I_\Pi(V, T^{(j)}|U)$. These conditional local mutual informations quantify the capacity of the local tests to distinguish a signal drawn from the prior Π from the zero signal. The second sum in the display of the lemma captures dependency between the transcripts and the prior draw $\mu \sim \Pi$. Essentially, it captures how well the signal can be *estimated* by the local tests.

We now discuss the choice of prior distribution Π . Let us set $\epsilon = \rho/\sqrt{d}$, let R be a d -dimensional vector of independent Rademacher random variables, and define the prior Π as the distribution of ϵR . Note that Π has support contained in H_ρ . (Such choices are typically considered as least favorable priors supported on signals that are difficult to detect, see for instance Section 3.2 of [27].)

Since V , μ and $X^{(j)}$ are independent of U , conditioning on U does not disrupt the Markov chain property: we have the chain $V|U \rightarrow \mu|U \rightarrow X^{(j)}|U \rightarrow T^{(j)}|U$. Consequently, the "estimation term" $I_\Pi(\mu, T^{(j)}|U, V)$ can be handled using strong data processing techniques employed in distributed estimation, see for example Lemma 11 in [21]. For completeness, we adopted the aforementioned lemma for the above choice of prior distribution in the form of Lemma 9 in the appendix, which yields that $I_\Pi(\mu, T^{(j)}|U, V) \leq 128 \frac{n\rho^2}{dm} I_\Pi(X^{(j)}, T^{(j)}|U, V = 1)$. Using that $T^{(j)}$ is binary valued, we obtain that the second term in (8) is bounded above by $128 \frac{n\rho^2}{d}$.

The loss of information about V resulting from the compression of $X^{(j)}|U$ into $T^{(j)}|U$ in this Markov chain is quantified by Lemma 3 below. The lemma comes in the form of a strong data processing inequality for the testing problem and forms the crux of the proof of the lower bound. It captures the difficulty of the local testing problem in terms of n , m , d and ρ .

Lemma 3 (Public Coin Strong Data Processing Inequality): With Π as defined above we have

$$I_\Pi(V, T^{(j)}|U) \leq (48\beta \wedge 1) I_\Pi(X^{(j)}, T^{(j)}|U),$$

where

$$\beta = \begin{cases} \frac{n^2 \rho^4}{dm^2} & \text{if } \frac{m}{n\rho^2} < 1/2, \\ \frac{2n\rho^2}{dm} & \text{if } \frac{m}{n\rho^2} \geq 1/2. \end{cases} \quad (9)$$

We give the proof of the lemma in Section B-C. By combining the information theoretic inequalities above with the fact that $I(X^{(j)}, T^{(j)}|U) \leq H(T^{(j)}|U) \leq 1$, which is true because $T^{(j)}$ is a binary variable, we get that

$$\begin{aligned} I_\Pi(V, T) &\leq \sum_{j=1}^m I_\Pi(V, T^{(j)}|U) + \sum_{j=1}^m I_\Pi(\mu, T^{(j)}|U, V) \\ &\leq 48\beta m + 128 \frac{n\rho^2}{d}. \end{aligned}$$

Therefore, in view of Lemma 1,

$$\mathcal{R}(H_\rho, T) \geq 1 - 4\sqrt{6 \frac{n\rho^2}{d} \left(\max \left\{ \frac{n\rho^2}{m}, 2 \right\} + 8/3 \right)}.$$

For ρ satisfying (5), the right-hand side is bounded from below by α for an arbitrary distributed test $T \in \mathcal{T}_{dist}$.

B. Optimal Tests Attaining the Lower Bound

In this section, we exhibit a distributed testing procedure that is optimal in the sense that it attains the lower bound of Theorem 1. More precisely, we show that if ρ^2 is larger than a multiple of the right-hand side of (5), there exists a distributed test for H_0 against H_ρ with risk bounded by a chosen level $\alpha \in (0, 1)$. Summarising, we have the following theorem, complementing the lower bound of Theorem 1.

Theorem 2: Fix $\alpha \in (0, 1)$ and suppose that

$$\rho^2 \geq C_\alpha \frac{\sqrt{d(m \wedge d)}}{n} \quad (10)$$

for $C_\alpha > 0$ a constant depending only on α as given in (15). Then there exists a test $T \in \mathcal{T}_{dist}$ such that $\mathcal{R}(H_\rho, T) \leq \alpha$.

We prove the theorem by constructing three concrete distributed tests, for three different ranges of m . We outline the construction of these tests in this section. The detailed verification that they are consistent at the level α for testing H_0 against H_ρ in their respective ranges is deferred to Section V.

First assume that the number of machines is large enough, but does not exceed the dimension d , i.e. $M_\alpha \leq m \leq d$, for a large enough constant $M_\alpha > 0$ given in (15). In this case, Theorem 1 asserts that the detection lower bound for ρ is a multiple of $(dm)^{1/4}/n^{1/2}$. We propose the following distributed test that is able to detect signals with Euclidian norm of that order if $M_\alpha \leq m \leq d$. In this case no public coin is needed. In this setting we first compute the local test statistic $S_{m \leq d}^{(j)} = (n/m) \|X^{(j)}\|^2$ at every machine j . Under the null hypothesis, $S_{m \leq d}^{(j)}$ follows a chi-square distribution with d degrees of freedom, i.e. $S_{m \leq d}^{(j)} \sim \chi_d^2$. Then for every j

we consider the randomized test $T_{m \leq d}^{(j)}$ using Bernoulli random variables

$$T_{m \leq d}^{(j)} | S_{m \leq d}^{(j)} \sim \text{Ber} \left(F_{\chi_d^2} \left(S_{m \leq d}^{(j)} \right) \right),$$

where $F_{\chi_d^2}$ denotes the distribution function of χ_d^2 . Under the null hypothesis the $T_{m \leq d}^{(j)}$ are independent and $\text{Ber}(1/2)$. At the central machines we combine the local test in a global test $T_{m \leq d} \in \mathcal{T}_{dist}$ by setting

$$T_{m \leq d} = \mathbb{1} \left\{ \left| \sum_{j=1}^m (T_{m \leq d}^{(j)} - 1/2) \right| \geq \sqrt{m \bar{\kappa}_\alpha} \right\}, \quad (11)$$

with $\bar{\kappa}_\alpha = 3 \log(4/\alpha)$. By a standard Chernoff bound, the type one error of this test is bounded by $\alpha/2$ for m large enough. In Section V-A, we prove that the same is true for the type two error if $\|\rho\|^2 \geq C_\alpha \sqrt{dm}/n$, as posed by Theorem

Next we assume that $m \geq d \vee M_\alpha$, in which case the detection lower bound for ρ is a multiple of $\sqrt{d/n}$. When $m/d \rightarrow \infty$, tests based on the statistics $S_{m \leq d}^{(j)}$ cannot reach the detection lower bound anymore as can be observed by inspection of the variance of (18), or through Theorem 16 in [29]. We propose a novel distributed test using a public coin. Specifically, we assume all machine have access to a vector $U = (U_1, \dots, U_d)$ of independent standard normal random variables. For $j = 1, \dots, m$ we compute the local test statistics $S_{m \geq d}^{(j)} = \sqrt{\frac{n}{md}} U^T X^{(j)}$ and the corresponding local tests

$$T_{m \geq d}^{(j)} = \mathbb{1} \left\{ S_{m \geq d}^{(j)} \geq 0 \right\}.$$

Then we aggregate these local tests in the central machine to a distributed test $T_{m \geq d} \in \mathcal{T}_{dist}$ by defining

$$T_{m \geq d} = \mathbb{1} \left\{ \left| \sum_{j=1}^m (T_{m \geq d}^{(j)} - 1/2) \right| \geq \sqrt{m \bar{\kappa}_\alpha} \right\}, \quad (12)$$

with $\bar{\kappa}_\alpha = (1/3) \log(16/\alpha)$. In Section V-B we prove that this test satisfies the required error bound if $\rho^2 \geq C_\alpha d/n$.

Finally, we consider the case $m \leq M_\alpha$ for completeness. We have to treat it separately for technical reasons, although in practice we would probably simply use the first test above for all cases $m \leq d$. To achieve optimality in this case we can simply choose a single machine and conduct the hypothesis test we would use in the classical, non-distributed setting. Specifically, we can for instance use as global test $T_{m \geq 1} \in \mathcal{T}_{dist}$ the test given by

$$T_{m \geq 1} = \mathbb{1} \left\{ \frac{n}{\sqrt{dm}} \|X^{(1)}\|^2 - \sqrt{d} \geq \kappa_\alpha \right\}, \quad (13)$$

with $\kappa_\alpha = 2/\sqrt{\alpha}$. See Section V-C for details.

IV. SIMULATION EXPERIMENTS

In this section we investigate the performance of the tests $T_{m \leq d}$ and $T_{m \geq d}$ that were proposed in the preceding section, in particular with regards to the ‘‘elbow effect’’ when $m \simeq d$ that we observed in our theoretical results. We visualize the performance using two different simulation experiments.

In the first experiment we consider two fixed choices of m , n and d , one with $m < d$ and one with $m > d$. We simulate data with a signal with increasing strength ρ and then assess the performance of the tests by computing the ‘‘true positive rate’’ (TPR), i.e. the fraction of the simulations in which they correctly detect the signal.

For the second experiment we also consider two scenarios, one with $m \lesssim d$ and one with $m \gtrsim d$. But now we fix the signal strength ρ a little above the detection limit in both cases and investigate how the performance of the tests depends on the total signal-to-noise ratio, or sample size n .

A. First Experiment

In the first simulation we consider fixed values for n , m and d and we simulate data in which we let the strength $\|\mu\|$ of the unobserved signal vary between 0 and 1. We investigate how well each of the distributed tests manages to correctly reject the null hypothesis, i.e. detect the signal.

In Figure 1, we consider two different scenarios. In both scenarios we choose $n = 10^4$ and we have specified the rejection criterion for both tests such that they have a type one error probability of less than $\alpha = 0.05$. The signal μ is drawn according to

$$\mu = \frac{\rho}{\sqrt{d}} R, \quad (14)$$

where R is a vector of independent Rademacher random variables, and we let the signal strength $\rho = \|\mu\|$ vary from 0 to 1.

In the $m < d$ scenario corresponding to the plot on the left-hand side in Figure 1, we see that the test $T_{m \leq d}$ outperforms the public coin test $T_{m \geq d}$, in accordance with our theoretical results. In the $m > d$ scenario on the right-hand side we see that the test $T_{m \geq d}$ detects the presence of the signal much earlier than the test $T_{m \leq d}$.

B. Second Experiment

In the second experiment we also consider two scenarios.

The first scenario corresponds to a situation in which a fixed number of machines m receive more observations as n increases, but the dimension of the signal increases as well. Specifically, we take $d = n^{2/3}$ and $m = 500$. We set the signal strength slightly above the detection limit for this case, namely $\rho = \log(d)(dm)^{1/4}/\sqrt{n}$. Note that $\rho \rightarrow 0$ as $n \rightarrow \infty$. In view of the theory we expect that the test $T_{m \leq d}$ detects the signal consistently, whereas the public coin test $T_{m \geq d}$ should have a worse performance, since it requires that $\rho \gtrsim \sqrt{d/n}$.

In the second scenario we take $m = n/10$ and $d = 5$ and set $\rho = 2 \log(m) \sqrt{d/n}$, which is again slightly above the detecting limit. This corresponds to the situation in which for each 10 additional observations a new machine is added, whilst the dimension of the signal remains fixed. In this case we also have $\rho \rightarrow 0$ as $n \rightarrow \infty$. We expect that for large n the test $T_{m \leq d}$ will fail to consistently detect the signal, as it requires $\rho \gtrsim (md)^{1/4}/\sqrt{n}$ to do so.

The resulting plots can be found in Figure 2. The results are again in accordance with our theoretical findings. We note in particular that the public coin test $T_{m \geq d}$ is a suitable choice in




Fig. 1. The horizontal axes correspond to the Euclidian norm $\|\mu\| = \rho$ of the signal (14) used in the simulations. In the plot on the left we have $m = 50$ and $d = 500$, on the right we have $m = 5000$ and $d = 5$. For each ρ in a grid between 0 and 1, 100 datasets were simulated. The solid lines give the TPR, i.e. the fraction of the 100 runs in which the tests correctly detected the signal. The red lines correspond to $T_{m \leq d}$, the blue lines to $T_{m \geq d}$.




Fig. 2. The horizontal axes corresponds to the total sample size n . In the plot on the left we have $d = n^{2/3}$, $m = 500$, and $\rho = \log(d)(dm)^{1/4}/\sqrt{n}$. On the right, $m = n/10$, $d = 5$ and $\rho = 2 \log(m) \sqrt{d/n}$. For every n in a grid ranging from 1 to $3 \cdot 10^4$, 100 datasets were simulated. The solid lines give the TPR, i.e. the fraction of the 100 runs in which the tests correctly detected the signal. The red lines correspond to $T_{m \leq d}$, the blue lines to $T_{m \geq d}$.

the realistic scenario where the number of machines m scales with n , so batches of additional observations are distributed over additional machines, while the dimension of the signal remains fixed or at least relatively small compared to the number of observations.

V. PROOF OF THEOREM 2

Let $\alpha \in (0, 1)$ be given, recall that $\tilde{\kappa}_\alpha^2 = 3 \log(4/\alpha)$, $\tilde{\kappa}_\alpha^2 = (1/3) \log(16/\alpha)$, $\kappa_\alpha^2 = 4/\alpha$, and set

$$M_\alpha := \max \left\{ (2^5 5 \tilde{\kappa}_\alpha)^2, \bar{D}, 36 \tilde{\kappa}_\alpha^2, 4e^2 \pi \tilde{\kappa}_\alpha^2 \alpha^{1/2} \right\},$$

$$C_\alpha := \max \left\{ 2^4 M_\alpha^2 \kappa_\alpha^2, 2(1 + \sqrt{2}) \kappa_\alpha M_\alpha, 80 \tilde{\kappa}_\alpha, \right. \\ \left. 2^{12} e^2 \tilde{\kappa}_\alpha^2 \alpha^{-5/2} \right\}, \quad (15)$$

where \bar{D} is defined in Lemma 4. Then we show below that for all $d, m, n \in \mathbb{N}$ at least one of the three distributed tests $T_{m \leq d}, T_{m \geq d}, T_{m \asymp 1}$, given in (13)-(12), achieves uniform consistency at level α , i.e.

$$\mathcal{R}(H_\rho, T) \leq \alpha, \quad (16)$$

for some $T \in \{T_{m \asymp 1}, T_{m \leq d}, T_{m \geq d}\}$.

We distinguish three regimes in view of the interplay of m and d , and show in the following subsections that the corresponding tests reach uniform consistency:

- **Case 1:** $M_\alpha < m \leq d$.
- **Case 2:** $M_\alpha \vee d < m$.
- **Case 3:** $m \leq M_\alpha$.

with corresponding tests $T_{m \leq d}, T_{m \geq d}, T_{m \asymp 1}$, respectively.

A. Case 1: $M_\alpha < m \leq d$

In view of Chernoff's bound, see Lemma 6, for $m > 4\bar{\kappa}_\alpha^2$,

$$\begin{aligned} P_0(T_{m \leq d} = 1) &= P_0\left(\sum_{j=1}^m (T_{m \leq d}^{(j)} - 1/2) \geq \bar{\kappa}_\alpha/\sqrt{m}\right) \\ &\leq 2 \exp(-2\bar{\kappa}_\alpha^2/3) < \alpha/2. \end{aligned} \quad (17)$$

We now turn to bounding the type II error probability. For arbitrary $\mu = (\mu_1, \dots, \mu_d) \in H_\rho$,

$$S_{m \leq d}^{(j)} = \sum_{i=1}^d \left(\sqrt{\frac{n}{m}} \mu_i + Z_i^{(j)} \right)^2 \quad (18)$$

with $Z_i^{(j)} \stackrel{iid}{\sim} N(0, 1)$, follows a noncentral chi-square distribution with d degrees of freedom and noncentrality parameter $\delta := (n/m) \|\mu\|_2^2$, i.e. $S_{m \leq d}^{(j)} \stackrel{iid}{\sim} \chi_d^2(\delta)$.

Let us take independent random variables $V_d^\delta \sim \chi_d^2(\delta)$ and $U_d \sim \chi_d^2(0)$. Then in view of Lemma 4, for $c > 1/40$ and for all $d \geq D$,

$$E_\mu(F_{\chi_d^2}(S_{m \leq d}^{(j)})) = \Pr(V_d^\delta \geq U_d) \geq \frac{1}{2} + c \left(\frac{\delta}{\sqrt{d}} \wedge \frac{1}{2} \right),$$

where E_μ is the expectation corresponding to P_μ . This further yields

$$\begin{aligned} P_\mu \left(\left| \sum_{j=1}^m \left[T_{m \leq d}^{(j)} - \frac{1}{2} \right] \right| \leq \sqrt{m} \bar{\kappa}_\alpha \right) \\ \leq P_\mu \left(\sum_{j=1}^m \left[T_{m \leq d}^{(j)} - E_\mu(F_{\chi_d^2}(S_{m \leq d}^{(j)})) \right] \right. \\ \left. \leq \sqrt{m} \bar{\kappa}_\alpha \left(1 - \frac{c\sqrt{m}}{\bar{\kappa}_\alpha} \left(\frac{\delta}{\sqrt{d}} \wedge \frac{1}{2} \right) \right) \right). \end{aligned}$$

In view of Chernoff's bound and the inequality $E_\mu(F_{\chi_d^2}(T_{m \leq d}^{(j)})) \leq 1$, this is further bounded by $2 \exp(-\bar{\kappa}_\alpha^2/3) = \alpha/2$, given that

$$\frac{c\sqrt{m}}{\bar{\kappa}_\alpha} \left(\frac{\delta}{\sqrt{d}} \wedge \frac{1}{2} \right) \geq 2. \quad (19)$$

This last inequality follows from the assumption (10)

$$\sqrt{\frac{m}{d}} \delta = \frac{n \|\mu\|_2^2}{\sqrt{dm}} \geq C_\alpha \geq 2\bar{\kappa}_\alpha/c \quad \text{and} \quad m > 16\bar{\kappa}_\alpha^2/c^2.$$

B. Case 2: $M_\alpha \vee d < m$

For the choice and $m > 4\bar{\kappa}_\alpha^2$ the same bound as in (17) holds for $P_0(T_{m \geq d} = 1)$.

For the type II error, consider $\mu = (\mu_1, \dots, \mu_d) \in H_\rho$. Define for $\eta > 0$, $\nu > 0$ the events $D_\eta := \{\|U\|_2/\sqrt{d} \leq \eta\}$,

$$A_{\nu, \eta} := \left\{ \frac{\sqrt{n}}{\sqrt{d}} \sum_{i=1}^d \mu_i U_i > \nu \right\} \cap D_\eta$$

and

$$B_{\nu, \eta} := \left\{ \frac{\sqrt{n}}{\sqrt{d}} \sum_{i=1}^d \mu_i U_i < -\nu \right\} \cap D_\eta,$$

where $U = (U_1, \dots, U_d) \sim N(0, I_d)$ is the public coin random vector, whose probability distribution we shall denote by Q . We set

$$\eta = 4\alpha^{-1/2} \quad \text{and} \quad \nu = 4e\sqrt{2\pi}\bar{\kappa}_\alpha\sqrt{\eta}, \quad (20)$$

and note that $C_\alpha \geq 2^7\nu^2/(\pi\alpha^2)$.

Then the type II error is bounded from above as

$$\begin{aligned} E_Q P_\mu(T_{m \geq d} = 0 | U = u) \\ \leq \int_{A_{\nu, \eta}} P_\mu \left(\sum_{j=1}^m (T_{m \geq d}^{(j)} - 1/2) \leq \sqrt{m} \bar{\kappa}_\alpha \middle| U = u \right) dQ(u) \\ + \int_{B_{\nu, \eta}} P_\mu \left(\sum_{j=1}^m (T_{m \geq d}^{(j)} - 1/2) \geq -\sqrt{m} \bar{\kappa}_\alpha \middle| U = u \right) dQ(u) \\ + Q(A_{\nu, \eta}^c \cap B_{\nu, \eta}^c). \end{aligned} \quad (21)$$

We show below that each of the first two terms on the right hand side are bounded by $\alpha/8$ and the third term by $\alpha/4$, resulting our statement.

First we deal with the third term in (21). Since

$$\frac{\sqrt{n}}{\sqrt{d}} \sum_{i=1}^d \mu_i U_i \sim N(0, \frac{n}{d} \|\mu\|_2^2),$$

the set $(A_{\nu, \eta} \cup B_{\nu, \eta})^c = A_{\nu, \eta}^c \cap B_{\nu, \eta}^c$, in view of Chebyshev's inequality, assumption (10), and definitions (20), satisfies that

$$\begin{aligned} Q(A_{\nu, \eta}^c \cap B_{\nu, \eta}^c) &\leq \Pr \left(\sqrt{\frac{n}{d}} \|\mu\|_2 |Z| \leq \nu \right) + Q(D_\eta^c) \\ &\leq \frac{2\nu}{\sqrt{2\pi}C_\alpha} + \frac{2}{\eta^2} \leq \alpha/4, \end{aligned}$$

where $Z \sim N(0, 1)$.

Next we deal with the first term on the right hand side of (21). For $u \in A_{\nu, \eta}$, we have $\|u\|_2/\sqrt{d} \leq \eta$, hence for $m \geq \nu^2/(2\eta^2)$,

$$\begin{aligned} \Pr \left(-\frac{\nu}{\sqrt{m}} < \frac{1}{\sqrt{d}} \sum_{i=1}^d u_i Z_i \leq 0 \right) &\geq \Pr \left(-\frac{\nu}{\sqrt{m\eta}} < Z \leq 0 \right) \\ &\geq \frac{e^{-\frac{\nu^2}{2m\eta^2}}}{\sqrt{2\pi}} \frac{\nu}{\sqrt{m\eta}} \geq c' \frac{\nu}{\sqrt{m\eta}}, \end{aligned}$$

where $Z, Z_1, \dots, Z_d \stackrel{iid}{\sim} N(0, 1)$ and $c' = (e\sqrt{2\pi})^{-1}$. Therefore,

$$\begin{aligned} P_\mu \left(S_{m \geq d}^{(j)} \geq 0 \middle| U = u \right) \\ = \Pr \left(\frac{\sqrt{n}}{\sqrt{dm}} \sum_{i=1}^d \mu_i u_i + \frac{1}{\sqrt{d}} \sum_{i=1}^d u_i Z_i \geq 0 \right) \\ \geq \Pr \left(\frac{\nu}{\sqrt{m}} + \frac{1}{\sqrt{d}} \sum_{i=1}^d u_i Z_i \geq 0 \right) \\ = \frac{1}{2} + \Pr \left(-\frac{\nu}{\sqrt{m}} < \frac{1}{\sqrt{d}} \sum_{i=1}^d u_i Z_i \leq 0 \right) \\ \geq \frac{1}{2} + c' \frac{\nu}{\sqrt{\eta m}}. \end{aligned}$$

Then in view of Chernoff's bound

$$\begin{aligned} & \int_{A_\nu} P_\mu \left(\sum_{j=1}^m (T_{m \geq d}^{(j)} - 1/2) \leq \sqrt{m} \tilde{\kappa}_\alpha \middle| U = u \right) dQ(u) \\ & \leq \int_{A_\nu} P_\mu \left(\sum_{j=1}^m [T_{m \geq d}^{(j)} - E_\mu(T_{m \geq d}^{(j)} | U = u)] \right. \\ & \quad \left. \leq \sqrt{m} (\tilde{\kappa}_\alpha - \frac{c'\nu}{\sqrt{\eta}}) \middle| U = u \right) dQ(u). \end{aligned}$$

Since by definition (20) $\tilde{\kappa}_\alpha - c'\nu/\sqrt{\eta} = -3\tilde{\kappa}_\alpha$, in view of Chernoff's bound (Lemma 6), for $36\tilde{\kappa}_\alpha^2 < m$ the preceding display is further bounded by $2e^{-3\tilde{\kappa}_\alpha^2} = \alpha/8$.

We can deal with the second term on the right hand side of (21) similarly. First, we obtain for $u \in B_{\nu, \eta}$ that

$$P_\mu \left(S_{m \geq d}^{(j)} \geq 0 \middle| U = u \right) \leq 1/2 - c' \frac{\nu}{\sqrt{\eta m}}.$$

And then we can derive as above that

$$E_Q \mathbb{1}_{B_\nu} P_\mu \left(\sum_{j=1}^m \left[T_{m \geq d}^{(j)} - \frac{1}{2} \right] \geq -\sqrt{m} \tilde{\kappa}_\alpha \right) \leq \alpha/8,$$

concluding the proof of case 2.

C. Case 3: $m \leq M_\alpha$

Finally, for completeness, we deal with the case when the number of machines is limited and we are back basically to the single server, centralized case from a minimax rate point of view.

Under the null hypothesis, $(n/m)\|X^{(1)}\|_2^2$ follows a chi-square distribution with d degrees of freedom, so by Chebyshev's inequality

$$P_0(T_{m \leq 1} = 1) \leq 2/\kappa_\alpha^2 = \alpha/2.$$

For any $\mu \in H_\rho$, it holds that

$$\begin{aligned} & P_\mu \left(\frac{n}{\sqrt{dm}} \|X^{(1)}\|_2^2 - \sqrt{d} \leq \kappa_\alpha \right) \\ & = \Pr \left(\left| \frac{n}{\sqrt{dm}} \|\mu\|_2^2 + 2d^{-1/2} \sum_{i=1}^d \mu_i Z_i \right. \right. \\ & \quad \left. \left. + d^{-1/2} \sum_{i=1}^d (Z_i^2 - 1) \right| \leq \kappa_\alpha \right), \end{aligned}$$

where $Z_i \stackrel{iid}{\sim} N(0, 1)$, $i = 1, \dots, d$. By the reverse triangle inequality, we can bound the previous probability with

$$\begin{aligned} & \Pr \left(\left| 2d^{-1/2} \sum_{i=1}^d \mu_i Z_i \right| \geq \frac{n\|\mu\|_2^2}{2\sqrt{dm}} \right) + \\ & \Pr \left(\left| d^{-1/2} \sum_{i=1}^d (Z_i^2 - 1) \right| \geq \frac{n\|\mu\|_2^2}{2\sqrt{dm}} - \kappa_\alpha \right). \end{aligned}$$

In view of Chebyshev's inequality and $\|\mu\|_2^2 \geq \rho^2 \geq C_\alpha \sqrt{d}/n$,

$$\Pr \left(\left| 2d^{-1/2} \sum_{i=1}^d \mu_i Z_i \right| \geq \frac{n\|\mu\|_2^2}{2\sqrt{dm}} \right) \leq \frac{16m^2}{n^2\|\mu\|_2^2} \leq \frac{16m^2}{C_\alpha} \leq \frac{\alpha}{4}$$

for $C_\alpha \geq 2^6 M_\alpha^2/\alpha$. Likewise, for $C_\alpha \geq 2(1 + \sqrt{2})\kappa_\alpha M_\alpha$

$$\frac{n\|\mu\|_2^2}{2\sqrt{dm}} - \kappa_\alpha \geq \sqrt{2}\kappa_\alpha$$

and consequently, by applying again Chebyshev's inequality, the second term is bounded by

$$\Pr \left(\left| d^{-1/2} \sum_{i=1}^d (Z_i^2 - 1) \right| \geq \sqrt{2}\kappa_\alpha \right) \leq \frac{1}{\kappa_\alpha^2} \leq \frac{\alpha}{4}.$$

Hence we can conclude that for $C_\alpha \geq \max\{2^6 M_\alpha^2/\alpha, 2(1 + \sqrt{2})\kappa_\alpha M_\alpha\}$,

$$P_\mu(T_{m \leq 1} = 0) \leq \alpha/2,$$

concluding the proof of the theorem.

APPENDIX A LEMMAS FOR THEOREM 2

In this section we collect the proofs of the lemmas used to derive the minimax testing upper bound.

Lemma 4: Let U_d and $V_d^{\delta_d}$ be independent chi-square distributed random variables with d degree of freedom and non-centrality parameters zero and $\delta_d > 0$, respectively. Then for a universal $\bar{D} \in \mathbb{N}$, not depending on δ_d , we have for all $d \geq \bar{D}$ that

$$\Pr \left(V_d^{\delta_d} - U_d \geq 0 \right) \geq \frac{1}{2} + \frac{1}{40} \left(\frac{\delta_d}{\sqrt{d}} \wedge \frac{1}{2} \right). \quad (22)$$

Proof: First note that the function $\delta \mapsto \Pr(V_d^\delta - U_d \geq 0)$ is monotone increasing. Then

$$\Pr \left(V_d^{\delta_d} - U_d \geq 0 \right) \geq \Pr \left(V_d^{\delta_d \wedge \sqrt{d}/2} - U_d \geq 0 \right),$$

so without loss of generality we can assume that $\delta_d \leq \sqrt{d}/2$.

The density of $V_d^{\delta_d}$ is

$$\sum_{k=0}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} p_{d+2k},$$

where p_k denotes the χ_k^2 -density. By the independence of U_d and $V_d^{\delta_d}$,

$$\begin{aligned} & \Pr \left(V_d^{\delta_d} - U_d \leq 0 \right) \\ & = \sum_{k=0}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} \int_{\{v-u \geq 0\}} p_{d+2k}(v) p_d(u) d(v, u). \end{aligned}$$

Let $U'_d \sim \chi_d^2$ and $U''_{2k} \sim \chi_{2k}^2$ be independent from each other and U_d . For any given $k \in \mathbb{N}$, we have

$$\int_{\{v-u \geq 0\}} p_{d+2k}(v) p_d(u) d(v, u) = \Pr(U_d - U'_d \leq U''_{2k}).$$

For convenience let us introduce the notation $W_d = (U_d - U'_d)/(2\sqrt{d})$. Conditioning and using independence once more, the latter equals

$$\begin{aligned} & \int \Pr\left(W_d \leq \frac{u}{2\sqrt{d}}\right) d\mathbb{P}_{U_{2k}''}(u) \\ &= \frac{1}{2} + \int \Pr\left(0 \leq W_d \leq \frac{u}{2\sqrt{d}}\right) d\mathbb{P}_{U_{2k}''}(u). \end{aligned}$$

Since U_{2k}'' has a median larger than $2k/3$ and the map $u \mapsto \Pr\left(0 \leq W_d \leq \frac{u}{2\sqrt{d}}\right)$ is increasing, we have that the second term in the last display satisfies

$$\begin{aligned} & \int \Pr\left(0 \leq W_d \leq \frac{u}{2\sqrt{d}}\right) p_{2k}(u) du \\ & \geq \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right) \int_{[\frac{2k}{3}, \infty)} p_{2k}(u) du \\ & \geq \frac{1}{2} \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right). \end{aligned}$$

By combining the above inequalities we obtain that

$$\begin{aligned} & \Pr\left(V_d^{\delta_d} - U_d \leq 0\right) \\ & \geq \frac{1}{2} + \frac{1}{2} \sum_{k=0}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right). \quad (23) \end{aligned}$$

Assume now that $\delta_d \gtrsim 1$. Let k_d be the largest integer such that $k_d \leq 3\sqrt{d}$. We divide the sum on the right hand of the preceding display to two parts, i.e. $k < k_d$ and $k \geq k_d$. By applying Lemma 5 with $\varepsilon_d = k$, it holds that for $c_0 = e^{-9/8}/6$,

$$\begin{aligned} & \sum_{k=0}^{k_d} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right) \\ & \geq \frac{c_0}{\sqrt{d}} \sum_{k=1}^{k_d} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{(k-1)!} \\ & \geq \frac{c_0 \delta_d}{2\sqrt{d}} \sum_{k=0}^{k_d-1} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!}. \end{aligned}$$

We have $\Pr(0 \leq W_d \leq 1) \xrightarrow{d} \Pr(0 \leq Z \leq 1) > 1/3$, hence there exists a $D_1 \in \mathbb{N}$, such that for all $d \geq D_1$ we have $\Pr(0 \leq W_d \leq 1) > 1/3$. For $k > k_d$ we have $k > 3\sqrt{d}$, hence for all $d \geq D_1$,

$$\begin{aligned} & \sum_{k > k_d}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right) \\ & \geq \frac{c_0}{2} \sum_{k > k_d}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!}. \end{aligned}$$

Since $\delta_d/\sqrt{d} \leq 1/2$, we have for $d \geq D_1$,

$$\begin{aligned} & \frac{1}{2} \sum_{k=0}^{\infty} \frac{e^{-\delta_d/2} (\delta_d/2)^k}{k!} \Pr\left(0 \leq W_d \leq \frac{k}{3\sqrt{d}}\right) \\ & \geq \frac{c_0 \delta_d}{2\sqrt{d}} \left(1 - \frac{e^{-\delta_d/2} (\delta_d/2)^{k_d}}{k_d!}\right). \end{aligned}$$

The proof is finished by showing that for large enough d we have $c_0/2 - 1/40 > (\delta_d/2)^{k_d}/k_d! > 0$. Recalling that $2\sqrt{d} \leq 3\sqrt{d} - 1 \leq k_d \leq 3\sqrt{d}$ and hence $\delta_d \leq \sqrt{d}/2 \leq \sqrt{d}/4$ we get in view of Stirling's inequality, that for some universal constant $C > 0$

$$\frac{(\delta_d/2)^{k_d}}{k_d!} \leq \frac{(k_d/4)^{k_d}}{k_d!} \lesssim e^{k_d(1-\log 4)} k_d^{-1/2},$$

which is bounded from above by $c_0/2 - 1/40$ for $d \geq D_1$, for some sufficiently large $D_1 > 0$. \square

Lemma 5: Let $U_d, U'_d \stackrel{iid}{\sim} \chi_d^2$, and $0 < \varepsilon_d \leq C\sqrt{d}$. Then there exists a large enough $D_0 \in \mathbb{N}$, such that for all $d \geq D_0$

$$\Pr\left(0 \leq \frac{U_d - U'_d}{2\sqrt{d}} \leq \frac{\varepsilon_d}{\sqrt{d}}\right) \geq \frac{e^{-C^2/8}}{6} \frac{\varepsilon_d}{\sqrt{d}}.$$

Proof: The characteristic function of the random variable $W_d := (U_d - U'_d)/(2\sqrt{d})$ is

$$\begin{aligned} \phi_d(t) &= \mathbb{E}e^{itW_d} = \mathbb{E}e^{i\frac{t}{2\sqrt{d}}U_d} \mathbb{E}e^{-i\frac{t}{2\sqrt{d}}U'_d} \\ &= (1 + it/\sqrt{d})^{-d/2} (1 - it/\sqrt{d})^{-d/2} \\ &= (1 + t^2/d)^{-d/2} \xrightarrow{d \rightarrow \infty} e^{-t^2/2}. \end{aligned}$$

Using the Fourier inversion formula, the density f_{W_d} of W_d satisfies

$$f_{W_d}(v) = \frac{1}{2\pi} \int_{\mathbb{R}} e^{itv} \phi_d(t) dt = \frac{1}{2\pi} \int_{\mathbb{R}} \cos(tv) \phi_d(t) dt,$$

where the second equality follows from the symmetry of ϕ_d . Let

$$g(v) := \frac{1}{2\pi} \int_{\mathbb{R}} \cos(tv) e^{-t^2/2} dt = \frac{1}{\sqrt{2\pi}} e^{-v^2/2},$$

where the last equation follows for instance by contour integration. Then by the dominated convergence theorem

$$|f_{W_d}(v) - g(v)| \leq \frac{1}{2\pi} \int_{\mathbb{R}} |e^{-t^2/2} - \phi_d(t)| dt \xrightarrow{d \rightarrow \infty} 0.$$

By the earlier established uniform convergence we have that for every $d \geq D_0$, for some large enough D_0 ,

$$\int_0^{\frac{\varepsilon_d}{2\sqrt{d}}} f_{W_d}(v) dv \geq \frac{1}{3} e^{-\varepsilon_d^2/(8d)} \frac{\varepsilon_d}{2\sqrt{d}} = \frac{e^{-C^2/8}}{6} \frac{\varepsilon_d}{\sqrt{d}},$$

where the constant $1/3$ is arbitrary and could be taken anything smaller than $1/\sqrt{2\pi}$. \square

Lemma 6 (Chernoff's Bound): Let $B_i \stackrel{iid}{\sim} \text{Ber}(p_i)$, $i = 1, \dots, k$, and $0 < \delta < 1$. Then

$$\Pr\left(\left|\sum_{i=1}^k (B_i - \mathbb{E}B_i)\right| \geq \delta \sum_{i=1}^k p_i\right) \leq 2 e^{-(\delta^2/3) \sum_{i=1}^k p_i}. \quad (24)$$

APPENDIX B LEMNAS FOR THEOREM 1

In this section we collect the proofs of the lemmas used to derive the minimax testing lower bound.

A. Proof of Lemma 1

In view of (7) we have

$$\begin{aligned} \mathcal{R}(H_\rho, T) &\geq 1 - (\mathbb{P}(T = 0|V = 0) - \mathbb{P}(T = 0|V = 1)) \\ &\geq 1 - \left| \mathbb{P}^{T|V=0} - \mathbb{P}^{T|V=1} \right| (T = 0) \\ &\geq 1 - \|\mathbb{P}^{T|V=0} - \mathbb{P}^{T|V=1}\|_{TV}. \end{aligned}$$

By the triangle inequality,

$$\begin{aligned} &\|\mathbb{P}^{T|V=0} - \mathbb{P}^{T|V=1}\|_{TV} \\ &\leq \|\mathbb{P}^{T|V=0} - \mathbb{P}^T\|_{TV} + \|\mathbb{P}^T - \mathbb{P}^{T|V=1}\|_{TV}. \end{aligned}$$

Applying the second Pinsker bound to the two terms on the RHS and using that $2ab \leq a^2 + b^2$,

$$\begin{aligned} &\|\mathbb{P}^{T|V=0} - \mathbb{P}^{T|V=1}\|_{TV}^2 \\ &\leq D_{KL}(\mathbb{P}^{T|V=0}|\mathbb{P}^T) + D_{KL}(\mathbb{P}^{T|V=1}|\mathbb{P}^T) \\ &= 2I_\Pi(V, T), \end{aligned}$$

which completes the proof of the lemma.

B. Proof of Lemma 2

We prove a more general lemma, but before stating it we recall some information theoretic definitions and identities, see [30], [31].

For discrete random variables X and arbitrary random variable Y , define the *entropy of X* as

$$H(X) = - \sum_x \mathbb{P}(X = x) \log \mathbb{P}(X = x)$$

and the *conditional entropy of X given Y* ,

$$\begin{aligned} H(X|Y) &= \int H(X|Y = y) d\mathbb{P}^Y(y) \\ &= - \int \sum_x \mathbb{P}^{X|Y=y}(x) \log \mathbb{P}^{X|Y=y}(x) d\mathbb{P}^Y(y). \end{aligned}$$

We also recall that conditioning reduced entropy $H(X) \geq H(X|Y)$. Following from this conditioning, on an arbitrary random vector Z , also reduces conditional entropy

$$\begin{aligned} H(X|Y) &= \int H(X|Y = y) dP_Y(y) \\ &\geq \int H(X|Y = y, Z) dP_Y(y) = H(X|Y, Z). \end{aligned}$$

For random variables X, Y, Z we define the mutual information between X and Y and conditional mutual information between X and Y given Z as

$$\begin{aligned} I(X; Y) &= D_{KL}(P_{(X,Y)} \| P_X \times P_Y), \\ I(X; Y|Z = z) &= D_{KL}(P_{(X,Y)|Z=z} \| P_{X|Z=z} \times P_{Y|Z=z}), \\ I(X; Y|Z) &= \int I(X; Y|Z = z) dP_Z(z). \end{aligned}$$

Next we recall some conditions of the mutual information. First we note that $I(X, Y) = 0$ if and only if X is independent

from Y . The chain rule for the mutual information between the random vector $Y = (Y^{(1)}, \dots, Y^{(m)})$ and V is

$$I(V; Y) = \sum_{j=1}^m I(V; Y^{(j)} | Y^{(1)}, \dots, Y^{(j-1)}). \quad (25)$$

For discrete random variable X and arbitrary random variable Y

$$\begin{aligned} I(X; Y) &= E_{(X,Y)} \log \frac{dP_{(XY)}}{dP_X dP_Y} \\ &= E_{(X,Y)} \log \frac{1}{dP_X} - E_{(X,Y)} \log \frac{1}{dP_{(X|Y=y)}} \\ &= H(X) - H(X|Y). \end{aligned} \quad (26)$$

In addition, by similar arguments, for arbitrary random variable Z we have

$$I(X; Y|Z) = H(X|Z) - H(X|Y, Z). \quad (27)$$

Lemma 7: Let us assume that the discrete random variable V and the discrete random vector F are such that the pair (V, F) is independent from the random variable U and the discrete random vector $Y = (Y_1, \dots, Y_m)$ satisfies that Y_j is conditionally independent from $Y_{1:j-1} := (Y_1, \dots, Y_{j-1})$ given U and (V, F) , then

$$I(V; Y) \leq \sum_{j=1}^m I(V; Y_j|U) + \sum_{j=1}^m I(F; Y_j|U, V).$$

Proof: A non-public coin version of the lemma is given for instance in [32].

First note that in view of (26) and since conditioning reduces entropy

$$\begin{aligned} I((Y, U); V) &= H(V) - H(V|Y, U) \\ &\geq H(V) - H(V|Y) = I(Y; V). \end{aligned}$$

Furthermore, by the chain rule (25) and the independence of U and V ,

$$I((Y, U); V) = I(Y; V|U) + I(U; V) = I(Y; V|U).$$

Similarly, by the chain rule and nonnegativity of mutual information,

$$\begin{aligned} I(V; Y|U) &= I((V, F); Y|U) - I(F; Y|U, V) \\ &\leq I((V, F); Y|U). \end{aligned}$$

By the identity (27) and the chain rule (25),

$$\begin{aligned} &I((V, F); Y|U) \\ &= H(Y|U) - H(Y|V, F, U) \\ &= \sum_{j=1}^m [H(Y_j|Y_{1:j-1}, U) - H(Y_j|V, F, Y_{1:j-1}, U)]. \end{aligned}$$

Since conditioning reduces entropy we have $H(Y_j|Y_{1:j-1}, U) \leq H(Y_j|U)$. Furthermore, by the conditional independence of $Y_{1:j-1}$ and Y_j given (U, V, F) results in

$H(Y_j|V, F, Y_{1:j-1}, U) = H(Y_j|V, F, U)$. Using these two facts, we obtain that

$$\begin{aligned} I((V, F); Y|U) &\leq \sum_{j=1}^m H(Y_j|U) - H(Y_j|V, F, U) \\ &= \sum_{j=1}^m I((V, F); Y_j|U). \end{aligned}$$

Combining the above displays and again applying the chain rule we now obtain that

$$\begin{aligned} I(Y; V) &\leq \sum_{j=1}^m I((V, F); Y_j|U) \\ &= \sum_{j=1}^m [I(V; Y_j|U) + I(F; Y_j|U, V)]. \end{aligned}$$

□

C. Proof of Lemma 3

Using the sub-Gaussianity of the likelihood $P_{\Pi} = \int P_{\mu} d\Pi(\mu)$ verified in Lemma 8, we adapt the proof of Theorem 3.7 in [32] to the present continuous setting with public coin protocol.

Proof: We start by noting that if $48\beta \geq 1$, the result follows immediately from the regular data processing inequality for mutual information.

In view of the definition of the conditional mutual information and noting that $\mathbb{P}^{V|U=u}(v) = \mathbb{P}^V(v) = 1/2$ by independence of U and V , $I(V, T^{(j)}|U)$ equals

$$\frac{1}{2} \int \sum_{v \in \{0,1\}} D_{KL} \left(\mathbb{P}^{T^{(j)}|(V,U)=(v,u)}; \mathbb{P}^{T^{(j)}|U=u} \right) d\mathbb{P}^U(u). \quad (28)$$

By Lemma 13 below,

$$\begin{aligned} &D_{KL} \left(\mathbb{P}^{T^{(j)}|(V,U)=(v,u)} \| \mathbb{P}^{T^{(j)}|U=u} \right) \\ &\leq \sum_{t \in \{0,1\}} \mathbb{P}^{T^{(j)}|U=u}(t) \left(\frac{\mathbb{P}^{T^{(j)}|(V,U)=(v,u)}(t)}{\mathbb{P}^{T^{(j)}|U=u}(t)} - 1 \right)^2 \end{aligned} \quad (29)$$

\mathbb{P}^U -almost surely. Furthermore, by Bayes rule,

$$\frac{\mathbb{P}^{T^{(j)}|(U,X^{(j)})=(u,x)}(t)}{\mathbb{P}^{T^{(j)}|U=u}(t)} = \frac{d\mathbb{P}^{X^{(j)}|(U,T^{(j)})=(u,t)}}{d\mathbb{P}^{X^{(j)}|U=u}}(x) =: g_{t,u}(x), \quad (30)$$

where the equality holds in an $L_1 \left(\mathbb{P}^{X^{(j)}} \right)$ sense \mathbb{P}^U -almost surely.

For $v \in \{0, 1\}$, define

$$\mathcal{L}_v(X^{(j)}) := \frac{d\mathbb{P}^{X^{(j)}|V=v}}{d\mathbb{P}^{X^{(j)}}}(X^{(j)}).$$

Since $V \rightarrow (X^{(j)}, U) \rightarrow T^{(j)}$ forms a Markov chain, we can write

$$\begin{aligned} &\frac{\mathbb{P}^{T^{(j)}|(V,U)=(v,u)}(t)}{\mathbb{P}^{Y^{(j)}|U=u}(t)} \\ &= \int \frac{\mathbb{P}^{T^{(j)}|(U,X^{(j)})=(u,x)}(t)}{\mathbb{P}^{T^{(j)}|U=u}(t)} d\mathbb{P}^{X^{(j)}|V=v}(x) \\ &= \mathbb{E}_{X^{(j)}} \left[g_{t,u}(X^{(j)}) \mathcal{L}_v(X^{(j)}) \right]. \end{aligned}$$

Then in view of $\mathbb{E}_{X^{(j)}} \mathcal{L}_v(X^{(j)}) = 1 = \mathbb{E}_{X^{(j)}} g_{y,u}(X^{(j)})$, \mathbb{P}^U -a.s., the right hand side of (29) equals

$$\sum_{t \in \{0,1\}} \mathbb{P}^{T^{(j)}|U=u}(t) \text{Cov} \left(\mathcal{L}_v(X^{(j)}), g_{t,u}(X^{(j)}) \right)^2, \quad \mathbb{P}^U\text{-a.s.}$$

By Theorem 4.13 in [33], we have that

$$\mathbb{E}GH \leq \mathbb{E}H \log H + \log \mathbb{E}e^G$$

for any random variables G, H with $\mathbb{E}H = 1$ and $\mathbb{E}e^G < \infty$. Applying this to $G = s(\mathcal{L}_v(X^{(j)}) - 1)$, $H = g_{t,u}(X^{(j)})$ we obtain

$$\begin{aligned} &s \text{Cov} \left(\mathcal{L}_v(X^{(j)}), g_{t,u}(X^{(j)}) \right) \\ &\leq \log \mathbb{E}_{X^{(j)}} \left[e^{s(\mathcal{L}_v(X^{(j)})-1)} \right] \\ &\quad + \mathbb{E}_{X^{(j)}} \left[g_{t,u}(X^{(j)}) \log g_{t,u}(X^{(j)}) \right]. \end{aligned}$$

By display (30) and the independence of $X^{(j)}$ and U ,

$$\begin{aligned} &\mathbb{E}_{X^{(j)}} \left[g_{t,u}(X^{(j)}) \log g_{t,u}(X^{(j)}) \right] \\ &= D_{KL}(\mathbb{P}^{X^{(j)}|(U,T^{(j)})=(u,t)} \| \mathbb{P}^{X^{(j)}|U=u}). \end{aligned}$$

Furthermore, in view of Lemma 8, $\mathcal{L}_v(X^{(j)})$ is $\sqrt{24\beta}$ -sub-Gaussian, hence

$$\log \mathbb{E}_{X^{(j)}} \left[e^{s(\mathcal{L}_v(X^{(j)})-1)} \right] \leq 24\beta s^2/2.$$

Taking $s = (24\beta)^{-1} \text{Cov}(\mathcal{L}_v(X^{(j)}), g_{t,u}(X^{(j)}))$ and combining the above displays we obtain

$$\begin{aligned} &\text{Cov} \left(\mathcal{L}_v(X^{(j)}), g_{t,u}(X^{(j)}) \right)^2 \\ &\leq \frac{1}{2} \text{Cov} \left(\mathcal{L}_v(X^{(j)}), g_{t,u}(X^{(j)}) \right)^2 \\ &\quad + 24\beta D_{KL}(\mathbb{P}^{X^{(j)}|(U,T^{(j)})=(u,t)} \| \mathbb{P}^{X^{(j)}|U=u}), \end{aligned}$$

which holds \mathbb{P}^U -almost surely. We now have shown that

$$\begin{aligned} &\int D_{KL} \left(\mathbb{P}^{T^{(j)}|(V,U)=(v,u)} \| \mathbb{P}^{T^{(j)}|U=u} \right) d\mathbb{P}^U(u) \\ &\leq 48\beta \int \sum_{t \in \{0,1\}} \mathbb{P}^{T^{(j)}|U=u}(t) D(u, t) d\mathbb{P}^U(u) \\ &= 48\beta I(X^{(j)}, T^{(j)}|U), \end{aligned}$$

where

$$D(u, t) = D_{KL}(\mathbb{P}^{X^{(j)}|(U,T^{(j)})=(u,t)} \| \mathbb{P}^{X^{(j)}|U=u}).$$

In view of (28), this concludes the proof. □

D. Sub-Gaussianity Lemma

First we recall some notations from Section III-A. Let us denote by Π the distribution of the random vector ϵR , where $R = (R_1, \dots, R_d)$ has independent Rademacher marginals and $\epsilon > 0$ is small (it is taken to be $\epsilon = \rho/\sqrt{d}$). We take $V \sim \text{Ber}(1/2)$ and set $X|(V=0) \sim N(0, \sigma^2 I_d)$ and $X|(V=1) \sim P_\Pi$, where $P_\Pi = \int P_\mu d\Pi(\mu)$ and P_μ is a multivariate Gaussian distribution with mean μ and σ^2 times the identity variance. Let \mathbb{P}^X and $\mathbb{P}^{X|V}$ denote the corresponding distributions of X and $X|V$.

Definition 1: A random variable X is called β -sub-Gaussian if for all $s \in \mathbb{R}$,

$$\mathbb{E}e^{s(X-\mathbb{E}X)} \leq e^{\beta^2 s^2/2}.$$

The lemma below shows that the likelihood ratios $\frac{d\mathbb{P}^{X|V=0}}{d\mathbb{P}^X}(X)$ and $\frac{d\mathbb{P}^{X|V=1}}{d\mathbb{P}^X}(X)$ are sub-Gaussian.

Lemma 8: The likelihood ratios

$$\frac{d\mathbb{P}^{X|V=0}}{d\mathbb{P}^X}(X) \quad \text{and} \quad \frac{d\mathbb{P}^{X|V=1}}{d\mathbb{P}^X}(X)$$

are $\sqrt{24\beta}$ -sub-Gaussian with

$$\beta = \begin{cases} d\epsilon^4/\sigma^4, & \text{if } \sigma^2/\epsilon^2 < d/2, \\ 2\epsilon^2/\sigma^2, & \text{if } \sigma^2/\epsilon^2 \geq d/2. \end{cases} \quad (31)$$

Proof: Using the notation

$$\mathcal{L}_v(X) := \frac{d\mathbb{P}^{X|V=v}}{d\mathbb{P}^X}(X), \quad v \in \{0, 1\},$$

we show below that for all $t \in \mathbb{R}$,

$$\mathbb{E}_X e^{t(\mathcal{L}_v(X) - \mathbb{E}_X \mathcal{L}_v(X))} \leq e^{24\beta t^2/2}.$$

This is implied by

$$\mathbb{P}^X (|\mathcal{L}_v - \mathbb{E}_X \mathcal{L}_v| \geq s) \leq 12 \exp\left(-\frac{s^2}{2\beta}\right) \quad \text{for all } s > 0, \quad (32)$$

where the equivalence is well known, but a proof can be found in Lemma 11. Since $|\mathcal{L}_v(X) - \mathbb{E}_X \mathcal{L}_v(X)| = |\mathcal{L}_v(X) - 1| \leq 1$, it is enough to consider $0 < s < 1$.

To prove (32), let us first introduce the notation $L := \frac{dP_\Pi}{dP_0}$, and note that

$$\mathcal{L}_0 = \frac{2}{1+L} \quad \text{and} \quad \mathcal{L}_1 = \frac{2}{1+L^{-1}}.$$

Then for $x \in \{\mathcal{L}_0 - 1 \geq s\}$ we have

$$\frac{2}{1+L}(x) = \mathcal{L}_0(x) \geq s+1$$

and

$$0 \leq \frac{2L}{1+L}(x) = 1 - \frac{1-L}{1+L}(x) \leq 1-s,$$

where the last inequality follows from $\mathcal{L}_0 - 1 = \frac{1-L}{1+L}$. Consequently, $L^{-1}(x) \geq \frac{s+1}{1-s}$. Similarly, for $x \in \{\mathcal{L}_0 - 1 \leq -s\}$,

$$0 \leq \frac{2}{1+L}(x) \leq 1-s \quad \text{and} \quad \frac{2L}{1+L}(x) \geq 1+s$$

and thus $L(x) \geq \frac{s+1}{1-s}$. Combining the above bounds results in for $x \in \{|\mathcal{L}_0 - 1| \geq s\}$ that

$$|\log L(x)| \geq \log\left(\frac{1+s}{1-s}\right) \geq \frac{2s}{1+s} \geq s,$$

where the last two inequalities follow from $\log x \geq 1 - \frac{1}{x}$ and $0 < s < 1$.

Through the same computation, the above display is also true for $x \in \{|\mathcal{L}_1 - 1| \geq s\}$. Consequently, for $v = 0, 1$,

$$\begin{aligned} & \mathbb{P}^X (|\mathcal{L}_v - \mathbb{E} \mathcal{L}_v| \geq s) \\ & \leq \mathbb{P}^X (|\log L| \geq s) \\ & = \frac{1}{2} P_0 (|\log(L)| \geq s) + \frac{1}{2} P_\pi (|\log(L)| \geq s). \end{aligned}$$

Using Markov's inequality the terms on the right hand sides can be further bounded as

$$P_0 (|\nu \log(L)| \geq \nu s) \leq e^{-\nu s} (\mathbb{E}_{X|V=0} L^\nu + \mathbb{E}_{X|V=0} L^{-\nu}), \quad (33)$$

for $\nu > 0$ and

$$\begin{aligned} & P_\pi (|\lambda \log(L)| \geq \lambda s) \\ & \leq e^{-\lambda_1 s} \mathbb{E}_{X|V=1} L^{\lambda_1} + e^{-\lambda_2 s} \mathbb{E}_{X|V=1} L^{-\lambda_2} \end{aligned} \quad (34)$$

for $\lambda_1, \lambda_2 > 0$. Note that $\mathbb{E}_{X|V=1} L^\lambda = \mathbb{E}_{X|V=0} L^{\lambda+1}$, hence choosing $\lambda_1 = \nu - 1$ and $\lambda_2 = \nu + 1$ in display (34), we get by combining the above two displays that for $v \in \{0, 1\}$,

$$\begin{aligned} & \mathbb{P}^X (|\mathcal{L}_v - \mathbb{E} \mathcal{L}_v| \geq s) \\ & \leq \frac{1}{2} [1 + e^{-s} + e^s] e^{-\nu s} (\mathbb{E}_{X|V=0} L^\nu + \mathbb{E}_{X|V=0} L^{-\nu}). \end{aligned} \quad (35)$$

We proceed by bounding the expectations in the above display after which minimizing in ν gives us the result of the lemma. Recall that $X|(V=0) \sim \mathcal{N}(0, \sigma^2 I_d)$ and $X_i|(V=1) \stackrel{iid}{\sim} \frac{1}{2} \mathcal{N}(\epsilon, \sigma^2) + \frac{1}{2} \mathcal{N}(-\epsilon, \sigma^2)$, $i = 1, \dots, d$. Consequently, $L(X)$ equals

$$\begin{aligned} & \prod_{i=1}^d \left[\frac{\exp\left(-\frac{1}{2\sigma^2}(X_i - \epsilon)^2\right) + \exp\left(-\frac{1}{2\sigma^2}(X_i + \epsilon)^2\right)}{2 \exp\left(-\frac{1}{2\sigma^2}X_i^2\right)} \right] \\ & = \prod_{i=1}^d \exp\left(-\frac{1}{2}\epsilon^2/\sigma^2\right) \cosh(X_i \epsilon/\sigma^2). \end{aligned} \quad (36)$$

Then by independence of X_j , $j = 1, \dots, d$

$$\mathbb{E}_{X|V=0} L^\nu = \left(e^{-\frac{\nu}{2}\epsilon^2/\sigma^2} \mathbb{E} \cosh^\nu \left(\frac{\epsilon}{\sigma} Z \right) \right)^d,$$

where $Z \sim N(0, 1)$.

Next we distinguish two cases. Suppose first that $2/d \leq \epsilon^2/\sigma^2$. Let us take $\nu = s\sigma^4/(d\epsilon^4)$. Then $\nu\epsilon^2/\sigma^2 < 1/2$, as $0 < s < 1$, and hence in view of Lemma 10

$$\begin{aligned} & e^{-s\nu} (\mathbb{E}_0 L^\nu + \mathbb{E}_0 L^{-\nu}) \\ & \leq \exp\left(\nu^2 \frac{d\epsilon^4}{2\sigma^4} - s\nu\right) \left(1 + e^{(3/2)\nu d\epsilon^4/\sigma^4}\right) \\ & \leq \exp\left(-\frac{s^2}{2} \frac{\sigma^4}{d\epsilon^4}\right) \left(1 + e^{(3/2)s}\right). \end{aligned}$$

The remaining case is when $2/d > \epsilon^2/\sigma^2$. Choosing $\nu = s\sigma^2/(2\epsilon^2)$ results in $\nu\epsilon^2/\sigma^2 < 1/2$, hence again in view of Lemma 10

$$\begin{aligned} e^{-s\nu}(\mathbb{E}_0 L^\nu + \mathbb{E}_0 L^{-\nu}) &\leq \left(\nu^2 \frac{d\epsilon^4}{2\sigma^4} - s\nu\right) \left(1 + e^{(3/2)\nu d\epsilon^4/\sigma^4}\right) \\ &\leq \exp\left(-\frac{s^2\sigma^2}{2\epsilon^2}\right) \left(1 + e^{(3/2)s}\right). \end{aligned}$$

Hence, by plugging in the last two displays into (35), and noting that for $0 < s < 1$

$$\frac{1}{2} \left(1 + e^{(3/2)s}\right) [1 + e^{-s} + e^s] \leq 12,$$

we arrive at (32), for β given in (31), concluding the proof of the lemma. \square

APPENDIX C ADDITIONAL TECHNICAL LEMMAS

In this subsection, we collect technical lemmas and their proofs.

The following lemma is essentially Lemma 11 in [21] adopted the setting in this article.

Lemma 9 (Multivariate Gaussian Estimation SPDI): Let R a d -dimensional vector of independent Rademacher variables, V be an independent $\text{Ber}(1/2)$ random variable and let Π denote the distribution of $\mu = \epsilon VR$ where $\epsilon > 0$ is a constant. Suppose that the random vector $X = (X_1, \dots, X_d)$ satisfies $X|\mu \sim N(\mu, \sigma^2 I_d)$ and that T is a discrete random variable such that $V \rightarrow \mu \rightarrow X \rightarrow Y$ forms a Markov chain. Then,

$$I(\mu; T|V) \leq 128 \left(\frac{\epsilon}{\sigma}\right)^2 I(X; T|V = 1).$$

Proof: Write R_1, \dots, R_d for the coordinates of R and write for $k \leq d$, $R_{1:k} := (R_1, \dots, R_k)$ and $X_{1:k} = (X_1, \dots, X_k)$. Conditionally on $V = 0$, $\mu = 0$ with probability 1, so $I(\mu; T|V = 0) = 0$. Conditionally on $V = 1$, $\mu = \epsilon R$. Combining these facts with the chain rule for mutual information,

$$\begin{aligned} I(\mu; T|V) &= \frac{1}{2} I(\epsilon R; T|V = 1) \\ &= \frac{1}{2} \sum_{k=1}^d I(\epsilon R_k; T|V = 1, R_{1:k-1}) \\ &= \frac{1}{2} \sum_{k=1}^d I(\epsilon R_k; T|V = 1), \end{aligned}$$

where the last equality follows from the fact that the coordinates of R are independent, so R_{k+1} is independent of $R_{1:k}$. Furthermore, $R_k|V = 1 \rightarrow X_k|V = 1 \rightarrow T|V = 1$ forms a Markov chain with $(X_k|R_k, V = 1) \sim N(\epsilon R_k, \sigma^2)$. Consequently, by applying Lemma 14 in [21],

$$I(\epsilon R_k; T|V = 1) \leq 64 \left(\frac{2\epsilon}{\sigma}\right)^2 I(X_k; T|V = 1).$$

The proof is now finished by observing that $X_{k+1}|V = 1$ is independent of $X_{1:k}|V = 1$, so combining the above inequality

with the chain rule of mutual information, we obtain

$$\begin{aligned} I(\mu; T|V) &\leq 128 \left(\frac{\epsilon}{\sigma}\right)^2 \sum_{k=1}^d I(X_k; T|V = 1, X_{1:k-1}) \\ &= 128 \left(\frac{\epsilon}{\sigma}\right)^2 I(X; T|V = 1). \end{aligned}$$

\square

Lemma 10: Let $Z \sim N(0, 1)$ and let $\nu \in \mathbb{R}$ such that $|\nu|\epsilon^2/\sigma^2 < 1/2$. It holds that

$$\mathbb{E} \cosh^\nu \left(\frac{\epsilon}{\sigma} Z\right) \leq \exp\left(\nu \frac{\epsilon^2}{2\sigma^2} + \nu^2 \frac{3\epsilon^4}{2\sigma^4} - \mathbb{1}_{\{\nu < 0\}} \frac{3}{2} \nu \frac{\epsilon^4}{\sigma^4}\right). \quad (37)$$

Proof: First assume that $\nu \geq 0$. Using $\cosh(x) \leq e^{x^2/2}$ we find

$$\mathbb{E} \cosh^\nu \left(\frac{\epsilon}{\sigma} Z\right) \leq \mathbb{E} e^{\nu \frac{\epsilon^2}{2\sigma^2} Z^2}.$$

In view of Lemma 12,

$$\mathbb{E} e^{\lambda(Z^2-1)} \leq e^{2\lambda^2} \text{ for all } 0 \leq \lambda \leq 1/4.$$

Applying this to the second last display yields (37).

Consider now the case that $\nu < 0$. We have

$$\begin{aligned} \frac{d}{dx} \cosh^\nu \left(\frac{\epsilon}{\sigma} x\right) &= \nu \frac{\epsilon}{\sigma} \cosh^\nu \left(\frac{\epsilon}{\sigma} x\right) \tanh \left(\frac{\epsilon}{\sigma} x\right), \\ \frac{d^2}{dx^2} \cosh^\nu \left(\frac{\epsilon}{\sigma} x\right) &= \nu \frac{\epsilon^2}{\sigma^2} \cosh^\nu \left(\frac{\epsilon}{\sigma} x\right) \left[(\nu - 1) \tanh^2 \left(\frac{\epsilon}{\sigma} x\right) + 1\right] =: \tau(x) \end{aligned}$$

Since $\cosh(0) = 1$ and $\tanh(0) = 0$, a second order Taylor expansion of $x \mapsto \cosh^\nu \left(\frac{\epsilon}{\sigma} x\right)$ about 0 yields

$$\mathbb{E} \cosh^\nu \left(\frac{\epsilon}{\sigma} Z\right) = \mathbb{E} \left[1 + \frac{Z^2}{2!} \tau(r_Z Z)\right], \text{ for some } r_Z \in [0, 1].$$

Since $\tanh^2(x) \leq x^2$ and $\cosh(x) \geq 1$ for all $x \in \mathbb{R}$,

$$\begin{aligned} \mathbb{E} \frac{Z^2}{2!} \tau(r_Z Z) &\leq \nu \frac{\epsilon^2}{2\sigma^2} \left[(\nu - 1) \frac{\epsilon^2}{\sigma^2} \mathbb{E} r_Z^2 Z^4 + 1\right] \\ &\leq \nu \frac{\epsilon^2}{2\sigma^2} \left[(\nu - 1) \frac{3\epsilon^2}{\sigma^2} + 1\right]. \end{aligned}$$

Then by combining the above two displays

$$\mathbb{E} \cosh^\nu \left(\frac{\epsilon}{\sigma} Z\right) \leq \exp\left(\nu \frac{\epsilon^2}{2\sigma^2} + \nu^2 \frac{3\epsilon^4}{2\sigma^4} - \frac{3}{2} \nu \frac{\epsilon^4}{\sigma^4}\right).$$

\square

The next lemma gives a well-known sufficient (and necessary) condition for the sub-Gaussian distribution. In the literature we did not find the present, required form of the lemma, hence for completeness we also provide its proof.

Lemma 11: Let X a mean-zero random variable satisfying

$$\mathbb{P}(|X| \geq s) \leq C \exp\left(-\frac{s^2}{2\beta}\right)$$

for some $C \geq 2$, $\beta > 0$ and for all $s \in [0, \infty)$. Then,

$$\mathbb{E} e^{sX} \leq e^{2\beta C s^2/2}.$$

Proof: For $k \in \mathbb{N}$, we have

$$\mathbb{E}|X|^k = \int_0^\infty \mathbb{P}(|X|^k > t) dt \leq C \int_0^\infty \exp\left(-\frac{t^{2/k}}{2\beta}\right) dt.$$

Changing coordinates to $u = t^{2/k}/(2\beta)$ yields that the right hand side display equals

$$\frac{C}{2}(2\beta)^{k/2} k \int_0^\infty e^{-u} u^{k/2-1} du = \frac{C}{2}(2\beta)^{k/2} k \Gamma(k/2).$$

By the dominated convergence theorem, $\mathbb{E}X = 0$, and $C \geq 2$,

$$\begin{aligned} \mathbb{E}e^{sX} &= 1 + \sum_{k=2}^\infty \frac{s^k \mathbb{E}X^k}{k!} \leq 1 + \frac{C}{2} \sum_{k=2}^\infty \frac{(2\beta s^2)^{k/2} \Gamma(k/2)}{(k-1)!} \\ &\leq 1 + \sum_{k=1}^\infty \left[\frac{(C\beta s^2)^k \Gamma(k)}{(2k-1)!} + \frac{(C\beta s^2)^{k+1/2} \Gamma(k+1/2)}{(2k)!} \right]. \end{aligned}$$

Since $\Gamma(k+1/2) \leq \Gamma(k+1) = k\Gamma(k) = k!$ and $(2k)! \geq 2^k (k!)^2$, the latter is further bounded by

$$\begin{aligned} &1 + \left(1 + \sqrt{C\beta s^2}\right) \sum_{k=1}^\infty \frac{(C\beta s^2/2)^k}{k!} \\ &= e^{C\beta s^2/2} + \sqrt{C\beta s^2} (e^{\beta C s^2/2} - 1). \end{aligned}$$

Since $(e^x - 1)(e^x - \sqrt{x}) \geq 0$, we obtain that

$$\mathbb{E}e^{sX} \leq e^{\frac{2C\beta s^2}{2}}.$$

□

The following lemma is a well known result and follows from standard calculus, but we included it as we did not find a stand-alone proof.

Lemma 12: Let Z be $N(0, 1)$, $0 \leq \lambda \leq 1/4$. Then,

$$\mathbb{E}e^{\lambda(Z^2-1)} \leq e^{2\lambda^2}.$$

Proof: Using the change of variables $u = z\sqrt{1-2\lambda}$,

$$\begin{aligned} \mathbb{E}e^{\lambda(Z^2-1)} &= \frac{1}{\sqrt{2\pi}} \int e^{\lambda(z^2-1)} e^{-\frac{1}{2}z^2} dz \\ &= \frac{e^{-\lambda}}{\sqrt{2\pi(1-2\lambda)}} \int e^{-\frac{1}{2}z^2} dz = \frac{e^{-\lambda}}{\sqrt{(1-2\lambda)}}. \end{aligned}$$

The MacLaurin series of $-\frac{1}{2} \log(1-2\lambda)$ reads

$$\frac{1}{2} \sum_{k=1}^\infty \frac{(2\lambda)^k}{k},$$

which yields that the second last display equals

$$\exp\left(\frac{3}{2}\lambda^2 + \frac{1}{2} \sum_{k=3}^\infty \frac{(2\lambda)^k}{k}\right).$$

If $\lambda \leq 1/4$,

$$\sum_{k=3}^\infty \frac{(2\lambda)^k}{k} \leq \frac{(2\lambda)^3}{1-2\lambda} \leq \lambda^2,$$

from which the result follows. □

The next lemma is a standard bound for the KL-divergence, see for instance Lemma 2.7 of [34].

Lemma 13: Let P, Q probability measures on some measure space such that $Q \ll P$. Then,

$$D_{KL}(P\|Q) \leq \int \left(\frac{dP}{dQ} - 1\right)^2 dQ.$$

REFERENCES

- [1] R. R. Tenney and N. R. Sandell, "Detection with distributed sensors," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-17, no. 4, pp. 501–510, Jul. 1981.
- [2] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals Syst.*, vol. 1, no. 2, pp. 167–182, 1988. [Online]. Available: <http://link.springer.com/10.1007/BF02551407>
- [3] O. P. Kreidl, J. N. Tsitsiklis, and S. I. Zoumpoulis, "On decentralized detection with partial information sharing among sensors," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1759–1765, Apr. 2011. [Online]. Available: <http://ieeexplore.ieee.org/document/5667065/>
- [4] A. Tarighati, J. Gross, and J. Jalden, "Decentralized hypothesis testing in energy harvesting wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 65, no. 18, pp. 4862–4873, Sep. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7950975/>
- [5] P. K. Varshney, *Distributed Detection and Data Fusion*. New York, NY, USA: Springer New York, 1997, doi: [10.1007/978-1-4612-1904-0](https://doi.org/10.1007/978-1-4612-1904-0).
- [6] J.-F. Chamberland and V. V. Veeravalli, "Wireless sensors in distributed detection applications," *IEEE Signal Process. Mag.*, vol. 24, no. 3, pp. 16–25, May 2007. [Online]. Available: <http://ieeexplore.ieee.org/document/4205085/>
- [7] R. Ahlswede and I. Csiszár, "Hypothesis testing with communication constraints," *IEEE Trans. Inf. Theory*, vol. 32, no. 4, pp. 533–542, Jul. 1986. [Online]. Available: <http://ieeexplore.ieee.org/document/1057194/>
- [8] T. S. Han, "Hypothesis testing with multiterminal data compression," *IEEE Trans. Inf. Theory*, vol. 33, no. 6, pp. 759–772, Nov. 1987.
- [9] T. Han and K. Kobayashi, "Exponential-type error probabilities for multiterminal hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 35, no. 1, pp. 2–14, Jan. 1989. [Online]. Available: <http://ieeexplore.ieee.org/document/42171/>
- [10] S.-I. Amari, "Fisher information under restriction of Shannon information in multi-terminal situations," *Ann. Inst. Stat. Math.*, vol. 41, no. 4, pp. 623–648, Dec. 1989. [Online]. Available: <http://link.springer.com/10.1007/BF00057730>
- [11] H. M. H. Shalaby and A. Papamarcou, "Multiterminal detection with zero-rate data compression," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 254–267, Mar. 1992. [Online]. Available: <http://ieeexplore.ieee.org/document/119685/>
- [12] S. Watanabe, "Neyman–Pearson test for zero-rate multiterminal hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, no. 7, pp. 4923–4939, Jul. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8123940/>
- [13] T. Han and S.-I. Amari, "Statistical inference under multiterminal data compression," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2300–2324, Oct. 1998. [Online]. Available: <http://ieeexplore.ieee.org/document/720540/>
- [14] J. Acharya, C. L. Canonne, and H. Tyagi, "Inference under information constraints I: Lower bounds from chi-square contraction," *IEEE Trans. Inf. Theory*, vol. 66, no. 12, pp. 7835–7855, Dec. 2020.
- [15] I. Diakonikolas, T. Gouleakis, D. M. Kane, and S. Rao, "Communication and memory efficient testing of discrete distributions," in *Proc. 32nd Conf. Learn. Theory* (Proceedings of Machine Learning Research), vol. 99. A. Beygelzimer and D. Hsu, Eds. Phoenix, AZ, USA: PMLR, Jun. 2019, pp. 1070–1106. [Online]. Available: <http://proceedings.mlr.press/v99/diakonikolas19a.html>
- [16] O. Fischer, U. Meir, and R. Oshman, "Distributed uniformity testing," in *Proc. ACM Symp. Princ. Distrib. Comput.*, New York, NY, USA, Jul. 2018, pp. 455–464, doi: [10.1145/3212734.3212772](https://doi.org/10.1145/3212734.3212772).
- [17] Y. Zhang, J. Duchi, M. I. Jordan, and M. J. Wainwright, "Information-theoretic lower bounds for distributed statistical estimation with communication constraints," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2328–2336.
- [18] M. Braverman, A. Garg, T. Ma, H. L. Nguyen, and D. P. Woodruff, "Communication lower bounds for statistical estimation problems via a distributed data processing inequality," in *Proc. 48th Annu. ACM Symp. Theory Comput.*, Jun. 2016, pp. 1011–1020.
- [19] A. Xu and M. Raginsky, "Information-theoretic lower bounds on Bayes risk in decentralized estimation," *IEEE Trans. Inf. Theory*, vol. 63, no. 3, pp. 1580–1600, Mar. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7801953/>
- [20] L. P. Barnes, Y. Han, and A. Ozgur, "Lower bounds for learning distributions under communication constraints via Fisher information," *J. Mach. Learn. Res.*, vol. 21, no. 236, pp. 1–30, 2020.

- [21] T. Tony Cai and H. Wei, "Distributed Gaussian mean estimation under communication constraints: Optimal rates and communication-efficient algorithms," 2020, *arXiv:2001.08877*.
- [22] B. Szabó and H. van Zanten, "Adaptive distributed methods under communication constraints," *Ann. Statist.*, vol. 48, no. 4, pp. 2347–2380, Aug. 2020.
- [23] Y. Zhu and J. Lafferty, "Distributed nonparametric regression under communication constraints," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 6004–6012.
- [24] Y. Han, A. Özgür, and T. Weissman, "Geometric lower bounds for distributed parameter estimation under communication constraints," in *Proc. Conf. Learn. Theory*, 2018, pp. 3163–3188.
- [25] B. Szabo and H. van Zanten, "Distributed function estimation: Adaptation using minimal communication," 2020, *arXiv:2003.12838*.
- [26] Y. Baraud, "Non-asymptotic minimax rates of testing in signal detection," *Bernoulli*, vol. 8, no. 5, pp. 577–606, Oct. 2002. [Online]. Available: <https://projecteuclid.org/euclid.bj/1078435219>
- [27] Y. I. Ingster and I. A. Suslina, *Nonparametric Goodness-of-Fit Testing Under Gaussian Models* (Lecture Notes in Statistics), vol. 169. New York, NY, USA: Springer, 2002.
- [28] X. Chen, A. Guntuboyina, and Y. Zhang, "On Bayes risk lower bounds," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 7687–7744, 2016.
- [29] J. Acharya, C. L. Canonne, and H. Tyagi, "Distributed signal detection under communication constraints," in *Proc. 33rd Conf. Learn. Theory* (Proceedings of Machine Learning Research), vol. 125. J. Abernethy and S. Agarwal, Eds. PMLR, Jul. 2020, pp. 41–63. [Online]. Available: <http://proceedings.mlr.press/v125/>
- [30] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2012.
- [31] Y. Polyanskiy and Y. Wu, "Lecture notes on information theory," *Lect. Notes ECE563 (UIUC)*, vol. 6, nos. 2012–2016, p. 7, Feb. 2014.
- [32] M. Raginsky, "Strong data processing inequalities and ϕ -Sobolev inequalities for discrete channels," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3355–3389, Jun. 2016.
- [33] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*, 1st ed. Oxford, U.K.: Oxford Univ. Press, 2013.
- [34] A. B. Tsybakov, *Introduction to Nonparametric Estimation*. New York, NY, USA: Springer, 2009.

Botond Szabó held positions at Vrije Universiteit Amsterdam, Leiden University, and the Budapest University of Technology and Economics. He is an Associate Professor with the Department of Decision Sciences, Bocconi University, and a fellow of the Bocconi Institute for Data Science and Analytics (BIDSA). His research focuses on the theory of high-dimensional and nonparametric statistical and machine learning methods.

Lasse Vuursteen received the bachelor's degree in econometrics from the University of Groningen and the master's degree in stochastics and financial mathematics from the Vrije Universiteit Amsterdam. He is currently pursuing the graduate degree with the Delft University of Technology. His research focuses on information-theoretic properties of distributed inference and more broadly on high-dimensional and nonparametric statistics.

Harry van Zanten held positions at the Eindhoven University of Technology and the University of Amsterdam. He is a Professor of statistics with Vrije Universiteit Amsterdam. His research focuses on the design and analysis of statistical methods in complex and high-dimensional settings.