

On combinatorial testing problems

Louigi Addario-Berry
McGill University

Nicolas Broutin
INRIA

Luc Devroye
McGill University

Gábor Lugosi
ICREA and Pompeu Fabra University *

January 16, 2010

Abstract

We study a class of hypothesis testing problems in which, upon observing the realization of an n -dimensional Gaussian vector, one has to decide whether the vector was drawn from a standard normal distribution or, alternatively, whether there is a subset of the components belonging to a certain given class of sets whose elements have been “contaminated,” that is, have a mean different from zero. We establish some general conditions under which testing is possible and others under which testing is hopeless with a small risk. The combinatorial and geometric structure of the class of sets is shown to play a crucial role. The bounds are illustrated on various examples.

1 Introduction

In this paper we study the following hypothesis testing problem introduced by Arias-Castro, Candès, Helgason and Zeitouni [3]. One observes an n -dimensional vector $\mathbf{X} = (X_1, \dots, X_n)$. The null hypothesis H_0 is that the components of \mathbf{X} are independent and identically distributed (i.i.d.) standard normal random variables. We denote the probability measure and expectation under H_0 by \mathbb{P}_0 and \mathbb{E}_0 , respectively.

To describe the alternative hypothesis H_1 , consider a class $\mathcal{C} = \{S_1, \dots, S_N\}$ of N sets of indices such that $S_k \subset \{1, \dots, n\}$ for all $k = 1, \dots, N$. Under H_1 , there exists an $S \in \mathcal{C}$ such that

$$X_i \text{ has distribution } \begin{cases} \mathcal{N}(0, 1) & \text{if } i \notin S \\ \mathcal{N}(\mu, 1) & \text{if } i \in S \end{cases}$$

where $\mu > 0$ is a positive parameter. The components of \mathbf{X} are independent under H_1 as well. The probability measure of \mathbf{X} defined this way by an $S \in \mathcal{C}$ is denoted by \mathbb{P}_S . Similarly, we write \mathbb{E}_S for the expectation with respect to \mathbb{P}_S . Throughout we will assume that every $S \in \mathcal{C}$ has the same cardinality $|S| = K$.

A test is a binary-valued function $f : \mathbb{R}^n \rightarrow \{0, 1\}$. If $f(\mathbf{X}) = 0$ then we say that the test accepts the null hypothesis, otherwise H_0 is rejected. One would like to design tests such that H_0 is accepted with a large probability when \mathbf{X} is distributed according to \mathbb{P}_0 and it is rejected when the distribution of \mathbf{X} is \mathbb{P}_S for some $S \in \mathcal{C}$. Following Arias-Castro, Candès, Helgason, and Zeitouni [3], we consider the risk of a test f measured by

$$R(f) = \mathbb{P}_0\{f(\mathbf{X}) = 1\} + \frac{1}{N} \sum_{S \in \mathcal{C}} \mathbb{P}_S\{f(\mathbf{X}) = 0\}. \quad (1)$$

*The third author acknowledges support by the Spanish Ministry of Science and Technology grant MTM2006-05650 and by the PASCAL Network of Excellence under EC grant no. 506778.

This measure of risk corresponds to the view that, under the alternative hypothesis, a set $S \subset \mathcal{C}$ is selected uniformly at random and the components of \mathbf{X} belonging to S have mean μ . In the sequel, we refer to the first and second terms on the right-hand side of (1) as the type I and type II errors, respectively.

We are interested in determining, or at least estimating the value of μ under which the risk can be made small. Our aim is to understand the order of magnitude, when n is large, as a function of n , K , and the structure of \mathcal{C} , of the value of the smallest μ for which risk can be made small. The value of μ for which the risk of the best possible test equals $1/2$ is called *critical*.

Typically, the n components of \mathbf{X} represent weights over the n edges of a given graph G and each $S \in \mathcal{C}$ is a subgraph of G . When $X_i \sim \mathcal{N}(\mu, 1)$ then the edge i is “contaminated” and we wish to test whether there is a subgraph in \mathcal{C} that is entirely contaminated.

In [3] two examples were studied in detail. In one case \mathcal{C} contains all paths between two given vertices in a two-dimensional grid and in the other \mathcal{C} is the set of paths from root to a leaf in a complete binary tree. In both cases the order of magnitude of the critical value of μ was determined. Arias-Castro, Candès, and Durand [4] investigate another class of examples in which elements of \mathcal{C} correspond to clusters in a regular grid. Both [3] and [4] describe numerous practical applications of problems of this type.

Some other interesting examples are when \mathcal{C} is

- * the set of all subsets $S \subset \{1, \dots, n\}$ of size K ;
- * the set of all cliques of a given size in a complete graph;
- * the set of all bicliques (i.e., complete bipartite subgraphs) of a given size in a complete bipartite graph;
- * the set of all spanning trees of a complete graph;
- * the set of all perfect matchings in a complete bipartite graph;
- * the set of all sub-cubes of a given size of a binary hypercube.

The first of these examples, which lacks any combinatorial structure, has been studied in the rich literature on multiple testing, see, for example, Ingster [24], Baraud [7], Donoho and Jin [14] and the references therein.

As pointed out in [3], regardless of what \mathcal{C} is, one may determine explicitly the test f^* minimizing the risk. It follows from basic results of binary classification that for a given vector $\mathbf{x} = (x_1, \dots, x_n)$, $f^*(\mathbf{x}) = 1$ if and only if the ratio of the likelihoods of \mathbf{x} under $(1/N) \sum_{S \in \mathcal{C}} \mathbb{P}_S$ and \mathbb{P}_0 exceeds 1. Writing

$$\phi_0(\mathbf{x}) = (2\pi)^{-n/2} e^{-\sum_{i=1}^n x_i^2/2} \quad \text{and} \quad \phi_S(\mathbf{x}) = (2\pi)^{-n/2} e^{-\sum_{i \in S} (x_i - \mu)^2/2 - \sum_{i \notin S} x_i^2/2}$$

for the probability densities of \mathbb{P}_0 and \mathbb{P}_S , respectively, the likelihood ratio at \mathbf{x} is

$$L(\mathbf{x}) = \frac{\frac{1}{N} \sum_{S \in \mathcal{C}} \phi_S(\mathbf{x})}{\phi_0(\mathbf{x})} = \frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu x_S - K\mu^2/2},$$

where $x_S = \sum_{i \in S} x_i$. Thus, the optimal test is given by

$$f^*(\mathbf{x}) = \mathbb{1}_{\{L(\mathbf{x}) > 1\}} = \begin{cases} 0 & \text{if } \frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu x_S - K\mu^2/2} \leq 1 \\ 1 & \text{otherwise.} \end{cases}$$

The risk of f^* (often called the Bayes risk) may then be written as

$$R^* = R_{\mathcal{C}}^*(\mu) = R(f^*) = 1 - \frac{1}{2} \mathbb{E}_0 |L(\mathbf{X}) - 1| = 1 - \frac{1}{2} \int \left| \phi_0(\mathbf{x}) - \frac{1}{N} \sum_{S \in \mathcal{C}} \phi_S(\mathbf{x}) \right| d\mathbf{x}.$$

We are interested in the behavior of R^* as a function of \mathcal{C} and μ . Clearly, R^* is a monotone decreasing function of μ . (This fact is intuitively clear and can be proved easily by differentiating R^* with respect to

μ .) For μ sufficiently large, R^* is close to zero while for very small values of μ , R^* is near its maximum value 1, indicating that testing is virtually impossible. Our aim is to understand for what values of μ the transition occurs. This depends on the combinatorial and geometric structure of the class \mathcal{C} . We describe various general conditions in both directions and illustrate them on examples.

Remark. (AN ALTERNATIVE RISK MEASURE.) Arias-Castro, Candès, Helgason, and Zeitouni [3] also consider the risk measure

$$\bar{R}(f) = \mathbb{P}_0\{f(\mathbf{X}) = 1\} + \max_{S \in \mathcal{C}} \mathbb{P}_S\{f(\mathbf{X}) = 0\} .$$

Clearly, $\bar{R}(f) \geq R(f)$ and when there is sufficient symmetry in f and \mathcal{C} , we have equality. However, there are significant differences between the two measures of risk. The alternative measure \bar{R} obviously satisfies the following monotonicity property: for a class \mathcal{C} and parameter $\mu > 0$, let $\bar{R}_{\mathcal{C}}^*(\mu)$ denote the smallest achievable risk. If $\mathcal{A} \subset \mathcal{C}$ are two classes then for any μ , $\bar{R}_{\mathcal{A}}^*(\mu) \leq \bar{R}_{\mathcal{C}}^*(\mu)$. In contrast to this, the ‘‘Bayesian’’ risk measure $R(f)$ does not satisfy such a monotonicity property as is shown in Section 5. In this paper we focus on the risk measure $R(f)$.

Remark. Throughout the paper we assume, for simplicity, that each set $S \in \mathcal{C}$ has the same cardinality K . We do this partly in order to avoid technicalities that are not difficult but make the arguments less transparent. At the same time, in many natural examples this condition is satisfied. If \mathcal{C} may contain sets of different size such that all sets have approximately the same number of elements then all arguments go through without essential changes. However, if \mathcal{C} contains sets of very different size then the picture may change because large sets become much easier to detect and small sets can basically be ignored. Another approach to handle sets of different size, adopted by Arias-Castro, Candès, and Durand [4], is to change the model of the alternative hypothesis such that the level μ of contamination is appropriately scaled depending on the size of the set S .

Plan of the paper. The paper is organized as follows. In Section 2 we briefly discuss two suboptimal but simple and general testing rules (the *maximum test* and the *averaging test*) that imply sufficient conditions for testability that turn out to be useful in many examples.

In Section 3 a few general sufficient conditions are derived for the impossibility of testing under symmetry assumptions for the class.

In Section 4 we work out several concrete examples, including the class of all K -sets, the class of all cliques of a certain size in a complete graph, the class of all perfect matchings in the complete bipartite graph, and the class of all spanning trees in a complete graph.

In Section 5 we show that, perhaps surprisingly, the optimal risk is not monotone in the sense that larger classes may be significantly easier to test than small ones, though monotonicity holds under certain symmetry conditions.

In the last two sections of the paper we use techniques developed in the theory of Gaussian processes to establish upper and lower bounds related to geometrical properties of the class \mathcal{C} . In Section 6 general lower bounds are derived in terms of random subclasses and metric entropies of the class \mathcal{C} . Finally, in Section 7 we take a closer look at the type I error of the optimal test and prove an upper bound that, in certain situations, is significantly tighter than the natural bound obtained for a general-purpose maximum test.

2 Simple tests and upper bounds

As mentioned in the introduction, the test f^* minimizing the risk is explicitly determined. However, the performance of this test is not always easy to analyze. Moreover, efficient computation of the optimal test is often a non-trivial problem though efficient algorithms are available in many interesting cases. (We discuss computational issues for the examples of Section 4.) Because of these reasons it is often useful to consider simpler, though suboptimal, tests. In this section we briefly discuss two simplistic tests, a test based on averaging and a test based on maxima. These are often easier to analyze and help understand the behavior of the optimal test as well. In many cases one of these tests turn out to have a near-optimal performance.

A simple test based on averaging

Perhaps the simplest possible test is based on the fact that the sum of the components of \mathbf{X} is zero-mean normal under \mathbb{P}_0 and has mean μK under the alternative hypothesis. Thus, it is natural to consider the *averaging test*

$$f(\mathbf{x}) = \mathbb{1}_{\{\sum_{i=1}^n X_i > \mu K/2\}} .$$

Proposition 1 *Let $\delta > 0$. The risk of the averaging test f satisfies $R(f) \leq \delta$ whenever*

$$\mu \geq \sqrt{\frac{8n}{K^2} \log \frac{2}{\delta}} .$$

Proof: Observe that under \mathbb{P}_0 , the statistic $\sum_{i=1}^n X_i$ has normal $\mathcal{N}(0, n)$ distribution while for each $S \in \mathcal{C}$, under \mathbb{P}_S , it is distributed as $\mathcal{N}(\mu K, n)$. Thus, $R(f) \leq 2e^{-(\mu K)^2/(8n)}$. \square

A test based on maxima

Another natural test is based on the fact that under the alternative hypothesis for some $S \in \mathcal{C}$, $X_S = \sum_{i \in S} X_i$ is normal $(\mu K, K)$. Consider the *maximum test*

$$f(\mathbf{x}) = 1 \quad \text{if and only if} \quad \max_{S \in \mathcal{C}} X_S \geq \frac{\mu K + \mathbb{E}_0 \max_{S \in \mathcal{C}} X_S}{2} .$$

The test statistic $\max_{S \in \mathcal{C}} X_S$ is often referred to as a *scan statistic* and has been thoroughly studied for a wide range of applications, see Glaz, Naus, and Wallenstein [20]. Here we only need the following simple observation.

Proposition 2 *The risk of the maximum test f satisfies $R(f) \leq \delta$ whenever*

$$\mu \geq \frac{\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S}{K} + 2\sqrt{\frac{2}{K} \log \frac{2}{\delta}} .$$

In the analysis it is convenient to use the following simple Gaussian concentration inequality, see Tsirelson, Ibragimov, and Sudakov [34].

Lemma 3 (TSIRELSON'S INEQUALITY.) *Let $X = (X_1, \dots, X_n)$ be an vector of n independent standard normal random variables. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ denote a Lipschitz function with Lipschitz constant L (with respect to the Euclidean distance). Then for all $t > 0$,*

$$\mathbb{P}\{f(X) - \mathbb{E}f(X) \geq t\} \leq e^{-t^2/(2L^2)} .$$

Proof of Proposition 2: Simply note that under the null hypothesis, for each $S \in \mathcal{C}$, X_S is a zero-mean normally distributed random variable with variance $K = |S|$. Since $\max_{S \in \mathcal{C}} X_S$ is a Lipschitz function of \mathbf{X} with Lipschitz constant \sqrt{K} , by Tsirelson's inequality, for all $t > 0$,

$$\mathbb{P}_0 \left\{ \max_{S \in \mathcal{C}} X_S \geq \mathbb{E}_0 \max_{S \in \mathcal{C}} X_S + t \right\} \leq e^{-t^2/(2K)} .$$

On the other hand, under \mathbb{P}_S for a fixed $S \in \mathcal{C}$,

$$\max_{S' \in \mathcal{C}} X_{S'} \geq X_S \sim \mathcal{N}(\mu K, K)$$

and therefore

$$\mathbb{P}_S \left\{ \max_{S \in \mathcal{C}} X_S \leq \mu K - t \right\} \leq e^{-t^2/(2K)} ,$$

which completes the proof. \square

The maximum test is often easier to compute than the optimal test f^* , though maximization is not always possible in polynomial time. If the value of $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S$ is not exactly known, one may replace it in the definition of f by any upper bound and then the same upper bound will appear in the performance bound.

Proposition 2 shows that the maximum test is guaranteed to work whenever μ is at least $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S / K + \text{const.} / \sqrt{K}$. Thus, in order to better understand the behavior of the maximum test (and thus obtain sufficient conditions for the optimal test to have a low risk), one needs to understand the expected value of $\max_{S \in \mathcal{C}} X_S$ (under \mathbb{P}_0). As the maximum of Gaussian processes have been studied extensively, there are plenty of directly applicable results available for expected maxima. The textbook of Talagrand [33] is dedicated to this topic. Here we only recall some of the basic facts.

First note that one always has $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq \sqrt{2K \log N}$ but sharper bounds can be derived by chaining arguments, see Talagrand's [33] for an elegant and advanced treatment. The classical chaining bound of Dudley [16] works as follows. Introduce a metric on \mathcal{C} by

$$d(S, T) = \sqrt{\mathbb{E}_0 (X_S - X_T)^2} = \sqrt{d_H(S, T)}, \quad S, T \in \mathcal{C}$$

where $d_H(S, T) = \sum_{i=1}^n \mathbb{1}_{\{\mathbf{1}_{\{i \in S\}} \neq \mathbf{1}_{\{i \in T\}}\}}$ denotes the Hamming distance. For $t > 0$, let $N(t)$ denote the t -covering number of \mathcal{C} with respect to the metric d , that is, the smallest number of open balls of radius t that cover \mathcal{C} . By Dudley's theorem, there exists a numerical constant C such that

$$\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq C \int_0^{\text{diam}(\mathcal{C})} \sqrt{\log N(t)} dt$$

where $\text{diam}(\mathcal{C}) = \max_{S, T \in \mathcal{C}} d(S, T)$ denotes the diameter of the metric space \mathcal{C} . Note that since $|S| = K$ for all $S \in \mathcal{C}$, $\text{diam}(\mathcal{C}) \leq \sqrt{2K}$. Dudley's theorem is not optimal but it is relatively easy to use. Dudley's theorem has been refined, based on "majorizing measures", or "generic chaining" which gives sharp bounds, see, for example, Talagrand [33].

Remark. (THE VC DIMENSION.) In certain cases it is convenient to further bound Dudley's inequality in terms of the VC dimension [35]. Recall that the VC dimension $V(\mathcal{C})$ of \mathcal{C} is the largest positive integer m such that there exists an m -element set $\{i_1, \dots, i_m\} \subset \{1, \dots, n\}$ such that for all 2^m subsets $A \subset \{i_1, \dots, i_m\}$ there exists an $S \in \mathcal{C}$ such that $S \cap \{i_1, \dots, i_m\} = A$. Haussler [22] proved that the covering numbers of \mathcal{C} may be bounded as

$$N(t) \leq e \cdot (V(\mathcal{C}) + 1) \left(\frac{2en}{t^2} \right)^{V(\mathcal{C})}$$

so by Dudley's bound,

$$\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq C \sqrt{V(\mathcal{C}) K \log n}.$$

An interesting alternative to the maximum test, proposed and investigated by Durot and Rozenholc [17] and Arlot, Blanchard and Roquain [5], is based on the idea that under the null hypothesis the distribution of the vector \mathbf{X} does not change if the sign of each component is changed randomly, while under the alternative hypothesis the distribution changes. In [17] and [5] methods based on symmetrization and bootstrap are suggested and analyzed. Such tests are meaningful and interesting in the setup of the present paper as well and it would be interesting to analyze their behavior.

3 Lower bounds

In this section we investigate conditions under which the risk of any test is large. We start with a simple universal bound that implies that regardless of what the class \mathcal{C} is, small risk cannot be achieved unless μ is substantially large compared to $K^{-1/2}$.

A universal lower bound

An often convenient way of bounding the Bayes risk R^* is in terms of the Bhattacharyya measure of affinity (Bhattacharyya [9])

$$\rho = \rho_{\mathcal{C}}(\mu) = \frac{1}{2} \mathbb{E}_0 \sqrt{L(\mathbf{X})} .$$

It is well known (see, e.g., [13, Theorem 3.1]) that

$$1 - \sqrt{1 - 4\rho^2} \leq R^* \leq 2\rho .$$

Thus, 2ρ essentially behaves as the Bayes error in the sense that R^* is near 1 when 2ρ is near 1, and is small when 2ρ is small. Observe that, by Jensen's inequality,

$$2\rho = \mathbb{E}_0 \sqrt{L(\mathbf{X})} = \int \sqrt{\frac{1}{N} \sum_{S \in \mathcal{C}} \phi_S(\mathbf{x}) \phi_0(\mathbf{x})} d\mathbf{x} \geq \frac{1}{N} \sum_{S \in \mathcal{C}} \int \sqrt{\phi_S(\mathbf{x}) \phi_0(\mathbf{x})} d\mathbf{x} .$$

Straightforward calculation shows that for any $S \in \mathcal{C}$,

$$\int \sqrt{\phi_S(\mathbf{x}) \phi_0(\mathbf{x})} d\mathbf{x} = e^{-\mu^2 K/8}$$

and therefore we have the following.

Proposition 4 *For all classes \mathcal{C} , $R^* \geq 1/2$ whenever $\mu \leq \sqrt{(4/K) \log(4/3)}$.*

This shows that no matter what the class \mathcal{C} is, detection is hopeless if μ is of the order of $K^{-1/2}$. This classical fact goes back to Le Cam [26].

A lower bound based on overlapping pairs

The next lemma is due to Arias-Castro, Candès, Helgason, and Zeitouni [3]. For completeness we recall their proof.

Proposition 5 *Let S and S' be drawn independently, uniformly, at random from \mathcal{C} and let $Z = |S \cap S'|$. Then*

$$R^* \geq 1 - \frac{1}{2} \sqrt{\mathbb{E} e^{\mu^2 Z} - 1} .$$

Proof: As noted above, by the Cauchy–Schwarz inequality,

$$R^* = 1 - \frac{1}{2} \mathbb{E}_0 |L(\mathbf{X}) - 1| \geq 1 - \frac{1}{2} \sqrt{\mathbb{E}_0 |L(\mathbf{X}) - 1|^2}$$

Since $\mathbb{E}_0 L(\mathbf{X}) = 1$,

$$\mathbb{E}_0 |L(\mathbf{X}) - 1|^2 = \text{Var}_0(L(\mathbf{X})) = \mathbb{E}_0 [L(\mathbf{X})^2] - 1 .$$

However, by definition $L(\mathbf{X}) = \frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S - K\mu^2/2}$, so we have

$$\mathbb{E}_0 [L(\mathbf{X})^2] = \frac{1}{N^2} \sum_{S, S' \in \mathcal{C}} e^{-K\mu^2} \mathbb{E}_0 e^{\mu(X_S + X_{S'})} .$$

But

$$\begin{aligned} \mathbb{E}_0 e^{\mu(X_S + X_{S'})} &= \mathbb{E}_0 \left[e^{\mu \sum_{i \in S \setminus S'} X_i} e^{\mu \sum_{i \in S' \setminus S} X_i} e^{2\mu \sum_{i \in S \cap S'} X_i} \right] \\ &= (\mathbb{E}_0 e^{\mu X})^{2(K - |S \cap S'|)} (\mathbb{E}_0 e^{2\mu X})^{|S \cap S'|} \\ &= e^{\mu^2(K - |S \cap S'|) + 2\mu^2 |S \cap S'|} , \end{aligned}$$

and the statement follows. \square

The beauty of this proposition is that it reduces the problem to studying a purely combinatorial quantity. By deriving upper bounds for the moment generating function of the overlap $|S \cap S'|$ between two elements of \mathcal{C} drawn independently and uniformly at random, one obtains lower bounds for the critical value of μ . This simple proposition turns out to be surprisingly powerful as it will be illustrated in various applications below.

A lower bound for symmetric classes

We begin by deriving some simple consequences of Proposition 5 under some general symmetry conditions on the class \mathcal{C} . The following proposition shows that the universal bound of Proposition 4 can be improved by a factor of $\sqrt{\log(1 + n/K)}$ for all sufficiently symmetric classes.

Proposition 6 *Let $\delta \in (0, 1)$. Assume that \mathcal{C} satisfies the following conditions of symmetry. Let S, S' be drawn independently and uniformly at random from \mathcal{C} . Assume that (i) the conditional distribution of $Z = |S \cap S'|$ given S' is identical for all values of S' ; (ii) for any fixed $S_0 \in \mathcal{C}$ and $i \in S_0$, $\mathbb{P}\{i \in S\} = K/n$. Then $R^* \geq \delta$ for all μ with*

$$\mu \leq \sqrt{\frac{1}{K} \log \left(1 + \frac{4n(1-\delta)^2}{K} \right)}.$$

Proof: We apply Proposition 5. By the first symmetry assumption it suffices to derive a suitable upper bound for $\mathbb{E}[e^{\mu^2 Z}] = \mathbb{E}[e^{\mu^2 Z} | S']$ for an arbitrary $S' \in \mathcal{C}$. After a possible relabeling, we may assume that $S' = \{1, \dots, K\}$ so we can write $Z = \sum_{i=1}^K \mathbb{1}_{\{i \in S\}}$. By Hölder's inequality,

$$\begin{aligned} \mathbb{E}[e^{\mu^2 Z}] &= \mathbb{E} \left[\prod_{i=1}^K e^{\mu^2 \mathbb{1}_{\{i \in S\}}} \right] \\ &\leq \prod_{i=1}^K \left(\mathbb{E} \left[e^{K\mu^2 \mathbb{1}_{\{i \in S\}}} \right] \right)^{1/K} \\ &= \mathbb{E} \left[e^{K\mu^2 \mathbb{1}_{\{1 \in S\}}} \right] \quad (\text{by assumption (ii)}) \\ &= \left(e^{\mu^2 K} - 1 \right) \frac{K}{n} + 1. \end{aligned}$$

Proposition 5 now implies the statement. \square

Surprisingly, the lower bound of Proposition 6 is close to optimal in many cases. This is true, in particular when the class \mathcal{C} is “small,” made precise in the following statement.

Corollary 7 *Assume that \mathcal{C} is symmetric in the sense of Proposition 6 and that it contains at most n^α elements where $\alpha > 0$. Then $R^* \geq 1/2$ for all μ with*

$$\mu \leq \sqrt{\frac{1}{K} \log \left(1 + \frac{n}{K} \right)}$$

and $R^* \leq 1/2$ for all μ with

$$\mu \geq \sqrt{\frac{2\alpha}{K} \log n}.$$

Proof: The first statement follows from Proposition 6 while the second from Proposition 2 and the fact that $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq \sqrt{2K \log |\mathcal{C}|}$. \square

The proposition above shows that for any small and sufficiently symmetric class, the critical value of μ is of the order of $\sqrt{(\log n)/K}$, at least if $K \leq n^\beta$ for some $\beta \in (0, 1)$. Later we will see examples of “large” classes for which Proposition 6 also gives a bound of the correct order of magnitude.

Negative association

The bound of Proposition 6 may be improved significantly under an additional condition of negative association that is satisfied in several interesting examples (see Section 4 below). Recall that a collection Y_1, \dots, Y_n of random variables is *negatively associated* if for any pair of disjoint sets $I, J \subset \{1, \dots, n\}$ and (coordinate-wise) non-decreasing functions f and g ,

$$\mathbb{E}[f(Y_i, i \in I)g(Y_j, j \in J)] \leq \mathbb{E}[f(Y_i, i \in I)] \mathbb{E}[g(Y_j, j \in J)] .$$

Proposition 8 *Let $\delta \in (0, 1)$ and assume that the class \mathcal{C} satisfies the conditions of Proposition 6. Suppose that the labels are such that $S' = \{1, 2, \dots, K\} \in \mathcal{C}$. Let S be a randomly chosen element of \mathcal{C} . If the random variables $\mathbb{1}_{\{1 \in S\}}, \dots, \mathbb{1}_{\{K \in S\}}$ are negatively associated then $R^* \geq \delta$ for all μ with*

$$\mu \leq \sqrt{\log \left(1 + \frac{n \log(1 + 4(1 - \delta)^2)}{K^2} \right)} .$$

Proof: We proceed similarly to the proof of Proposition 6. We have

$$\begin{aligned} \mathbb{E}[e^{\mu^2 Z}] &= \mathbb{E} \left[\prod_{i=1}^K e^{\mu^2 \mathbb{1}_{\{i \in S\}}} \right] \\ &\leq \prod_{i=1}^K \mathbb{E} \left[e^{K \mu^2 \mathbb{1}_{\{i \in S\}}} \right] \quad (\text{by negative association}) \\ &= \left(\left(e^{\mu^2} - 1 \right) \frac{K}{n} + 1 \right)^K . \end{aligned}$$

Proposition 5 and the upper bound above imply that R^* at least δ for all μ such that

$$\mu \leq \sqrt{\log \left(1 + \frac{n \left((1 + 4(1 - \delta)^2)^{1/K} - 1 \right)}{K} \right)} .$$

The result follows by using $e^y \geq 1 + y$ with $y = K^{-1} \log(1 + 4(1 - \delta)^2)$. \square

4 Examples

In this section we consider various concrete examples and work out upper and lower bounds for the critical range of μ .

4.1 Disjoint sets

We start with the simplest possible case, that is, when all $S \in \mathcal{C}$ are disjoint (and therefore $KN \leq n$). Fix $\delta \in (0, 1)$. Then, under \mathbb{P}_0 , the X_S are independent normal $(0, K)$ random variables and the bound $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq \sqrt{2K \log N}$ is close to being tight. By applying the maximum test f , we see that $R^* \leq R(f) \leq \delta$ whenever

$$\mu \geq \sqrt{\frac{2 \log N}{K}} + 2\sqrt{\frac{2 \log(2/\delta)}{K}}.$$

To see that this bound gives the correct order of magnitude, we may simply apply Proposition 5. Here Z may take two values:

$$Z = K \quad \text{with probability } 1/N, \quad \text{and} \quad Z = 0 \quad \text{with probability } 1 - 1/N.$$

Thus,

$$\mathbb{E} e^{\mu^2 Z} - 1 = \frac{1}{N} (e^{\mu^2 K} - 1) \leq \frac{1}{N} e^{\mu^2 K}$$

and therefore $R^* \geq \delta$ whenever

$$\mu \leq \sqrt{\frac{\log(4N(1-\delta)^2)}{K}}.$$

So in this case the critical transition occurs when μ is of the order of $\sqrt{(1/K) \log N}$. In Section 6 we use this simple lower bound to establish lower bounds for general classes \mathcal{C} of sets. Note that in this simple case one may directly analyze the risk of the optimal test and obtain sharper bounds. In particular, the leading constant in the lower bound is suboptimal. However, in this paper our aim is to understand some general phenomena and we focus on orders of magnitude rather than on nailing down sharp constants.

Remark. (MULTIPLE HYPOTHESIS TESTING.) Taking $K = 1$ in the above example, we obtain a connection with multiple hypothesis testing. In the latter, one tests “ $m_i = 0$ ” against “ $m_i = \mu$ ” for every $1 \leq i \leq N$, and traditionally uses as test statistics $X_i, 1 \leq i \leq N$, to build a multiple testing procedure (often rejecting all the hypotheses corresponding to large X_i), see, for instance, Romano and Wolf [30] Arlot, Blanchard and Roquain [6]. Such a procedure, which is calibrated to control the family wise error rate or the false discovery rate, can be also used to test “ $\forall i \in S, m_i = 0$ ” against “ $\forall i \in S, m_i = \mu$ ” with a controlled type I error rate, by rejecting if one of the hypotheses is rejected (however note that the entire risk is not controlled because of type II errors). As a consequence, any thresholding-based multiple testing procedure falls in the framework of this example, with $S = \{\{1\}, \dots, \{n\}\}$, $K = 1$ and $N = n \geq 2$.

4.2 K -sets

Consider the example when \mathcal{C} contains all sets $S \subset \{1, \dots, n\}$ of size K . Thus, $N = \binom{n}{K}$. As mentioned in the introduction, this problem is very well understood as sharp bounds and sophisticated tests are available, see, for example, Ingster [24], Baraud [7], Donoho and Jin [14]. We include it for illustration purposes only and we warn the reader that the obtained bounds are not sharpest possible.

Let $\delta \in (0, 1)$. It is easy to see that the assumptions of Proposition 8 are satisfied (this follows, e.g., from Proposition 11 of Dubhashi and Ranjan [15]) and therefore $R^* \geq \delta$ for all

$$\mu \leq \sqrt{\log \left(1 + \frac{n \log(1 + 4(1-\delta)^2)}{K^2} \right)}.$$

This simple bound turns out to have the correct order of magnitude both when $n \gg K^2$ (in which case it is of the order of $\sqrt{\log(n/K^2)}$) and when $n \ll K^2$ (when it is of the order of $\sqrt{n/K^2}$).

This may be seen by considering the two simple tests described in Section 2 in the two different regimes. Since

$$\frac{\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S}{K} \leq \frac{\sqrt{2K \log \binom{n}{K}}}{K} \leq \sqrt{2 \log \left(\frac{ne}{K} \right)},$$

we see from Proposition 2 that when $K = O(n^{(1-\epsilon)/2})$ for some fixed $\epsilon > 0$, then the threshold value is of the order of $\sqrt{\log n}$. On the other hand, when K^2/n is bounded away from zero, then the lower bound implied by Proposition 8 above is of the order $\sqrt{n/K^2}$ and the averaging test provides a matching upper bound by Proposition 1.

Note that in this example the maximum test is easy to compute since it suffices to find the K largest values among X_1, \dots, X_n .

4.3 Perfect matchings

Let \mathcal{C} be the set of all perfect matchings of the complete bipartite graph $K_{m,m}$. Thus, we have $n = m^2$ edges and $N = m!$, and $K = m$. By Proposition 1 (i.e., the averaging test), for $\delta \in (0, 1)$, one has $R(f) \leq \delta$ whenever $\mu \geq \sqrt{8 \log(2/\delta)}$.

To show that this bound has the right order of magnitude, we may apply Proposition 8. The symmetry assumptions hold obviously and the negative association property follows from the fact that $Z = |S \cap S'|$ has the same distribution as the number of fixed points in a random permutation. The proposition implies that for all m , $R^* \geq \delta$ whenever

$$\mu \leq \sqrt{\log(1 + \log(1 + 4(1 - \delta)^2))}.$$

Note that in this case the optimal test f^* can be approximated in a computationally efficient way. To this end, observe that computing

$$\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} = \frac{1}{m!} \sum_{\sigma} \prod_{j=1}^m e^{\mu X_{(j, \sigma(j))}}$$

(where the summation is over all permutations of $\{1, \dots, m\}$) is equivalent to computing the permanent of an $m \times m$ matrix with non-negative elements. By a deep result of Jerrum, Sinclair, and Vigoda [25], this may be done by a polynomial-time randomized approximation.

4.4 Stars

Consider a network of m nodes in which each pair of nodes interacts. One wishes to test if there is a corrupted node in the network whose interactions slightly differ from the rest. This situation may be modeled by considering the class of *stars*.

A star is a subgraph of the complete graph K_m which contains all $K = m - 1$ edges containing a fixed vertex (see Figure 1). Consider the set \mathcal{C} of all stars. In this setting, $n = \binom{m}{2}$ and $N = m$.



Figure 1: A star [36].

In this case Corollary 7 is applicable and we obtain that if \mathcal{C} is the class of all stars in K_m then for any $\epsilon > 0$,

$$\lim_{m \rightarrow \infty} R^* = \begin{cases} 0 & \text{if } \mu \geq (\sqrt{2} + \epsilon) \sqrt{\frac{\log m}{m}} \\ 1 & \text{if } \mu \leq (1 - \epsilon) \sqrt{\frac{\log m}{m}}. \end{cases}$$

4.5 Spanning trees

Consider again a network of m nodes in which each pair of nodes interact. One may wish to test if there exists a corrupted connected subgraph containing each node. This leads us to considering the class of all spanning trees as follows.

Let $1, 2, \dots, n = \binom{m}{2}$ represent the edges of the complete graph K_m and let \mathcal{C} be the set of all spanning trees of K_m . Thus, we have $N = m^{m-2}$ spanning trees and $K = m - 1$ (See, e.g., [28].) By Proposition 1, the averaging test has risk $R(f) \leq \delta$ whenever $\mu \geq \sqrt{4 \log(2/\delta)}$.

This bound is indeed of the right order. To see this, we may start with Proposition 5. There are (at least) two ways of proceeding. One is based on negative association. Even though Proposition 8 is not applicable because of the lack of symmetry in \mathcal{C} , negative association still holds. In particular, by a result of Feder and Mihail [18] (see also Grimmett and Winkler [21] and Benjamini, Lyons, Peres, and Schramm [8]), if S is a random uniform spanning tree of K_m , then the indicators $\mathbb{1}_{\{1 \in S\}}, \dots, \mathbb{1}_{\{n \in S\}}$ are negatively associated. This means that, if S and S' are independent uniform spanning trees and $Z = |S \cap S'|$,

$$\begin{aligned} \mathbb{E} \left[e^{\mu^2 Z} \right] &= \mathbb{E} \mathbb{E} \left[e^{\mu^2 |S \cap S'|} \mid S' \right] \\ &= \mathbb{E} \mathbb{E} \left[e^{\mu^2 \sum_{i \in S'} \mathbb{1}_{\{i \in S\}}} \mid S' \right] \\ &\leq \mathbb{E} \prod_{i \in S'} \mathbb{E} \left[e^{\mu^2 \mathbb{1}_{\{i \in S\}}} \mid S' \right] \quad (\text{by negative association}) \\ &\leq \mathbb{E} \prod_{i \in S'} \left(\frac{2}{m} e^{\mu^2} + 1 \right) \\ &= \left(\frac{2}{m} e^{\mu^2} + 1 \right)^{m-1} \\ &\leq \exp \left(2e^{\mu^2} \right). \end{aligned}$$

This, together with Proposition 5 shows that for any $\delta \in (0, 1)$, $R^* \geq \delta$ whenever

$$\mu \leq \sqrt{\log \left(1 + \frac{1}{2} \log(1 + 4(1 - \delta)^2) \right)}.$$

We note here that the same bound can be proved by a completely different way that does not use negative association. The key is to note that we may generate the two random spanning trees based on $2(m - 1)$ independent random variables $X_1, \dots, X_{2(m-1)}$ taking values in $\{1, \dots, m - 1\}$ as in Aldous [1] (See also [11]). The key property we need is that if Z_i denotes the number of common edges in the two spanning trees when X_i is replaced by an independent copy X'_i while keeping all other X_j 's fixed, then

$$\sum_{i=1}^{2(m-1)} (Z - Z_i)_+ \leq Z$$

(the details are omitted). For random variables satisfying this last property, an inequality of Boucheron, Lugosi, and Massart [10] implies the sub-Poissonian bound

$$\mathbb{E} \exp(\mu^2 Z) \leq \exp \left(\mathbb{E} Z (e^{\mu^2} - 1) \right).$$

Clearly, $\mathbb{E} Z = 2(n - 1)/n \leq 2$, so essentially the same bound as above is obtained.

As the bounds above show, the computationally trivial average test has a close-to-optimal performance. In spite of this, one may wish to use the optimal test f^* . The ‘‘partition function’’ $(1/N) \sum_{S \in \mathcal{C}} e^{\mu X_S}$ may be computed by an algorithm of Propp and Wilson [29], who introduced a random sampling algorithm

that, given a graph with non-negative weights w_i over the edges, samples a random spanning tree from a distribution such that the probability of any spanning tree S is proportional to $\prod_{i \in S} w_i$. The expected running time of the algorithm is bounded by the cover time of an associated Markov chain that is defined as a random walk over the graph in which the transition probabilities are proportional to the edge weights. If μ is of the order of a constant (as in the critical range) then the cover time is easily shown to be polynomial (with high probability) as all edge weights $w_i = e^{\mu^2 X_i}$ are roughly of the same order both under the null and under the alternative hypotheses.

4.6 Cliques

Another natural application is the class of all cliques of a certain size in a complete graph. More precisely, the random variables X_1, \dots, X_n are associated with the edges of the complete graph K_m such that $\binom{m}{2} = n$ and let \mathcal{C} contain all cliques of size k . Thus, $K = \binom{k}{2}$ and $N = \binom{m}{k}$. This case is more difficult than the class of K -sets discussed above because negative association does not hold anymore. (This may be easily seen by considering the indicator variables of two adjacent edges both being in the randomly chosen clique.) Also, computationally the class of cliques is much more complex. A related, well-studied model starts with the subgraph K_m containing each edge independently with probability $1/2$, as null hypothesis. The alternative hypothesis is the same as the null hypothesis, except that there is a clique of size k on which each edge is independently present with probability $p > 1/2$. This is called the ‘‘hidden clique’’ problem (usually only the special case $p = 1$ is considered). Despite substantial interest in the hidden clique problem, polynomial time detection algorithms are only known when $k = \Omega(\sqrt{n})$ [2, 19]. We may obtain the hidden clique model from our model by thresholding at weight zero (retaining only edges whose normal random variable is positive), and so our model is easier for testing than the hidden clique model. However, it seems likely that designing an efficient test in the normal setting will be as difficult as it has proved for hidden cliques. It would be of interest to construct near-optimal tests that are computable in polynomial time for larger values of k .

We have the following bounds for the performance of the optimal test. It shows that when k is a most of the order of \sqrt{m} , the critical value of μ is of the order of $\sqrt{(1/k) \log(m/k)}$. The proof below may be adjusted to handle larger values of k as well but we prefer to keep the calculations more transparent.

Proposition 9 *Let \mathcal{C} represent the class of all $N = \binom{m}{k}$ cliques of a complete graph K_m and assume that $k \leq \sqrt{m(\log 2)/e}$. Then*

(i) *for all $\delta \in (0, 1)$, $R^* \leq \delta$ whenever*

$$\mu \geq 2\sqrt{\frac{1}{k-1} \log\left(\frac{me}{k}\right)} + 4\sqrt{\frac{\log(2/\delta)}{k(k-1)}} ,$$

(ii) *$R^* \geq 1/2$ whenever*

$$\mu \leq \sqrt{\frac{1}{k} \log\left(\frac{m}{2k}\right)} .$$

Proof: (i) follows simply by a straightforward application of Proposition 2 and the bound $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S \leq \sqrt{2K \log N}$.

To prove the lower bound (ii), by Proposition 5, it suffices to show that if S, S' are k -cliques drawn randomly and independently from \mathcal{C} and Z denotes the number of edges in the intersection of S and S' , then $\mathbb{E} [\exp(\mu^2 Z)] \leq 2$ for the indicated values of μ .

Because of symmetry, $\mathbb{E} [\exp(\mu^2 Z)] = \mathbb{E} [\exp(\mu^2 Z) | S']$ for all S' and therefore we might as well fix an arbitrary clique S' . If Y denotes the number of vertices in the clique $S \cap S'$ then $Z = \binom{Y}{2}$. Moreover, the distribution of Y is hypergeometrical with parameters m and k . If B is a binomial random variable with parameters k and k/m , then since $\exp(\mu^2 x^2/2)$ is a convex function of x , an inequality of Hoeffding [23] implies that

$$\mathbb{E} [e^{\mu^2 Z}] = \mathbb{E} [e^{\mu^2 Y^2/2}] \leq \mathbb{E} [e^{\mu^2 B^2/2}] .$$

Thus, it remains to derive an appropriate upper bound for the moment generating function of the squared binomial. To this end, let $c > 1$ be a parameter whose value will be specified later. Using

$$B^2 \leq B \left(k \mathbb{1}_{\{B > c \frac{k^2}{m}\}} + c \frac{k^2}{m} \right)$$

and the Cauchy–Schwarz inequality, it suffices to show that

$$\mathbb{E} \left[\exp \left(\mu^2 c \frac{k^2}{m} B \right) \right] \cdot \mathbb{E} \left[\exp \left(\mu^2 k B \mathbb{1}_{\{B > c \frac{k^2}{m}\}} \right) \right] \leq 4. \quad (2)$$

We show that, if μ satisfies the condition of (ii), for an appropriate choice of c , both terms on the left-hand side are at most 2.

The first term on the left-hand side of (2) is

$$\mathbb{E} \left[\exp \left(\mu^2 c \frac{k^2}{m} B \right) \right] = \left(1 + \frac{k}{m} \left(\exp \left(\mu^2 c \frac{k^2}{m} \right) - 1 \right) \right)^k$$

which is at most 2 if and only if

$$\frac{k}{m} \left(\exp \left(\mu^2 c \frac{k^2}{m} \right) - 1 \right) \leq 2^{1/k} - 1.$$

Since $2^{1/k} - 1 \geq (\log 2)/k$, this is implied by

$$\mu \leq \sqrt{\frac{m}{ck^2} \log \left(1 + \frac{m \log 2}{k^2} \right)}.$$

To bound the second term on the left-hand side of (2), note that

$$\begin{aligned} \mathbb{E} \left[\exp \left(\mu^2 k B \mathbb{1}_{\{B > c \frac{k^2}{m}\}} \right) \right] &\leq 1 + \mathbb{E} \left[\mathbb{1}_{\{B > c \frac{k^2}{m}\}} \exp \left(\mu^2 k B \right) \right] \\ &\leq 1 + \left(\mathbb{P} \left\{ B > c \frac{k^2}{m} \right\} \right)^{1/2} \left(\mathbb{E} \left[\exp \left(\mu^2 k B \right) \right] \right)^{1/2}, \end{aligned}$$

by the Cauchy–Schwarz inequality, so it suffices to show that

$$\mathbb{P} \left\{ B > c \frac{k^2}{m} \right\} \cdot \mathbb{E} \left[\exp \left(\mu^2 k B \right) \right] \leq 1.$$

Denoting $h(x) = (1+x) \log(1+x) - x$, Chernoff’s bound implies

$$\mathbb{P} \left\{ B > c \frac{k^2}{m} \right\} \leq \exp \left(-\frac{k^2}{m} h(c-1) \right).$$

On the other hand,

$$\mathbb{E} \left[\exp \left(\mu^2 k B \right) \right] = \left(1 + \frac{k}{m} \exp \left(\mu^2 k \right) \right)^k,$$

and therefore the second term on the left-hand side of (2) is at most 2 whenever

$$1 + \frac{k}{m} \exp \left(\mu^2 k \right) \leq \exp \left(\frac{k}{m} h(c-1) \right).$$

Using $\exp\left(\frac{k}{m}h(c-1)\right) \geq 1 + \frac{k}{m}h(c-1)$, we obtain the sufficient condition

$$\mu \leq \sqrt{\frac{1}{k} \log h(c-1)}.$$

Summarizing, we have shown that $R^* \geq 1/2$ for all μ satisfying

$$\mu \leq 2 \cdot \min \left(\sqrt{\frac{1}{k} \log h(c-1)}, \sqrt{\frac{m}{ck^2} \log \left(1 + \frac{m \log 2}{k^2}\right)} \right).$$

Choosing

$$c = \frac{m}{k} \frac{\log(m/k)}{\log(m \log 2/k^2)}$$

(which is greater than 1 for $k \leq \sqrt{m(\log 2)/e}$), the second term on the right-hand side is at most $\sqrt{(1/k) \log(m/k)}$. Now observe that since $h(c-1) = c \log c - c + 1$ is convex, for any $a > 0$, $h(c-1) \geq c \log a - a + 1$. Choosing $a = \frac{\log(m/k)}{\log(m \log 2/k^2)}$, the first term is at least

$$\sqrt{\frac{1}{k} \log \left(\frac{m}{k} - \frac{\log(m/k)}{\log(m \log 2/k^2)} \right)} \geq \sqrt{\frac{1}{k} \log \left(\frac{m}{2k} \right)}$$

where we used the condition that $m \log 2/k^2 \geq e$ and that $x \geq 2 \log x$ for all $x > 0$. \square

Remark. (A RELATED PROBLEM.) A closely related problem arising in the exploratory analysis of microarray data (see Shabalin, Weigman, Perou, and Nobel [31]) is when each member of \mathcal{C} represents the K edges of a $\sqrt{K} \times \sqrt{K}$ biclique of the complete bipartite graph $K_{m,m}$ where $m = \sqrt{n}$. (A biclique is a complete bipartite subgraph of $K_{m,m}$.) The analysis and the bounds are completely analogous to the one worked out above, the details are omitted.

5 On the monotonicity of the risk

Intuitively, one would expect that the testing problem becomes harder as the class \mathcal{C} gets larger. More precisely, one may expect that if $\mathcal{A} \subset \mathcal{C}$ are two classes of subsets of $\{1, \dots, n\}$, then $R_{\mathcal{A}}^*(\mu) \leq R_{\mathcal{C}}^*(\mu)$ holds for all μ . The purpose of this section is to show that this intuition is wrong in quite a strong sense as not only such general monotonicity property does not hold for the risk, but there are classes $\mathcal{A} \subset \mathcal{C}$ for which $R_{\mathcal{A}}^*(\mu)$ is arbitrary close to 1 and $R_{\mathcal{C}}^*(\mu)$ is arbitrary close to 0 for the same value of μ .

However, monotonicity does hold if the class \mathcal{C} is sufficiently symmetric. Call a class \mathcal{C} *symmetric* if for the optimal test $f_{\mathcal{C}}^*(\mathbf{x}) = \mathbb{1}_{\{(1/N) \sum_{S \in \mathcal{C}} \exp(\mu \sum_{i \in S} x_i) \geq \exp(K\mu^2/2)\}}$, the value of $\mathbb{P}_T\{f_{\mathcal{C}}^*(\mathbf{X}) = 0\}$ is the same for all $T \in \mathcal{C}$. Note that several of the examples discussed in Section 4 satisfy the symmetry assumption, such as the classes of K -sets, stars, perfect matchings, and cliques. However, the class of spanning trees is not symmetric in the required sense.

Theorem 10 *Let \mathcal{C} be a symmetric class of subsets of $\{1, \dots, n\}$. If \mathcal{A} is an arbitrary subclass of \mathcal{C} , then for all $\mu > 0$, $R_{\mathcal{A}}^*(\mu) \leq R_{\mathcal{C}}^*(\mu)$.*

Proof: In this proof we fix the value of $\mu > 0$ and suppress it in the notation. Recall the definition of the alternative risk measure

$$\bar{R}_{\mathcal{C}}(f) = \mathbb{P}_0\{f(\mathbf{X}) = 1\} + \max_{S \in \mathcal{C}} \mathbb{P}_S\{f(\mathbf{X}) = 0\}$$

which is to be contrasted with our main risk measure

$$R_{\mathcal{C}}(f) = \mathbb{P}_0\{f(\mathbf{X}) = 1\} + \frac{1}{N} \sum_{S \in \mathcal{C}} \mathbb{P}_S\{f(\mathbf{X}) = 0\}.$$

The risk \bar{R} is obviously monotone in the sense that if $\mathcal{A} \subset \mathcal{C}$ then for every f , $\bar{R}_{\mathcal{A}}(f) \leq \bar{R}_{\mathcal{C}}(f)$. Let $\bar{f}_{\mathcal{C}}^*$ and $f_{\mathcal{C}}^*$ denote the optimal tests with respect to both measures of risk.

First observe that if \mathcal{C} is symmetric, then $\bar{R}_{\mathcal{C}}(f_{\mathcal{C}}^*) = R_{\mathcal{C}}(f_{\mathcal{C}}^*)$. But since $R_{\mathcal{C}}(f) \leq \bar{R}_{\mathcal{C}}(f)$ for every f , we have

$$\bar{R}_{\mathcal{C}}(\bar{f}_{\mathcal{C}}^*) \leq \bar{R}_{\mathcal{C}}(f_{\mathcal{C}}^*) = R_{\mathcal{C}}(f_{\mathcal{C}}^*) \leq R_{\mathcal{C}}(\bar{f}_{\mathcal{C}}^*) \leq \bar{R}_{\mathcal{C}}(\bar{f}_{\mathcal{C}}^*) .$$

This means that all inequalities are equalities and, in particular, $\bar{f}_{\mathcal{C}}^* = f_{\mathcal{C}}^*$.

Now if \mathcal{A} is an arbitrary subclass of \mathcal{C} , then

$$R_{\mathcal{C}}^* = R_{\mathcal{C}}(f_{\mathcal{C}}^*) = \bar{R}_{\mathcal{C}}(\bar{f}_{\mathcal{C}}^*) \geq \bar{R}_{\mathcal{A}}(\bar{f}_{\mathcal{C}}^*) \geq R_{\mathcal{A}}(\bar{f}_{\mathcal{C}}^*) \geq R_{\mathcal{A}}(f_{\mathcal{A}}^*) = R_{\mathcal{A}}^* ,$$

which completes the proof. \square

Theorem 11 *For every $\epsilon \in (0, 1)$ there exist n , μ , and classes $\mathcal{A} \subset \mathcal{C} \subset \{1, \dots, n\}$ such that $R_{\mathcal{A}}^*(\mu) \geq 1 - \epsilon$ and $R_{\mathcal{C}}^*(\mu) \leq 2\epsilon$.*

Proof: We work with L_1 distances. For any class \mathcal{L} , denote $\phi_{\mathcal{L}}(\mathbf{x}) = \frac{1}{N} \sum_{S \in \mathcal{L}} \phi_S(\mathbf{x})$. Recall that

$$R_{\mathcal{L}}^*(\mu) = 1 - \frac{1}{2} \int |\phi_0(\mathbf{x}) - \phi_{\mathcal{L}}(\mathbf{x})| d\mathbf{x} .$$

Given ϵ , we fix an integer $K = K(\epsilon)$ large enough that $K + 1 \geq 1/\epsilon$ and that

$$\sqrt{\frac{\log(4(K+1)\epsilon^2)}{K+1}} \geq \sqrt{\frac{8}{K} \log\left(\frac{2}{\epsilon}\right)},$$

and let $n = n(\epsilon) = (K + 1)^2$. We let \mathcal{A} consist of M disjoint subsets of $\{1, \dots, n\}$, each of size $K + 1$. We let \mathcal{B} consist of all sets of the form $\{1, \dots, K, i\}$, where i ranges from $K + 1$ to n , and assume \mathcal{A} has been chosen so that $\mathcal{A} \cap \mathcal{B} = \emptyset$. We then let $\mathcal{C} = \mathcal{A} \cup \mathcal{B}$. We take

$$\mu = \sqrt{\frac{\log(4(K+1)\epsilon^2)}{K+1}},$$

so that, as seen in Section 4.1, we have $R_{\mathcal{A}}^*(\mu) \geq 1 - \epsilon$. We will require an upper bound on $R_{\mathcal{B}}^*(\mu)$, which we obtain by considering the averaging test on variables $1, \dots, K$,

$$f(\mathbf{x}) = \mathbb{1}_{\{\sum_{i=1}^K x_i \geq \frac{\mu K}{2}\}} .$$

Just as in Proposition 1, we have $R(f) \leq \epsilon$ whenever $\mu \geq \sqrt{\frac{8}{K} \log\left(\frac{2}{\epsilon}\right)}$, which is indeed the case by our choices of μ and K . It follows that $R_{\mathcal{B}}^*(\mu) \leq \epsilon$. We remark that

$$\int |\phi - \phi_{\mathcal{A}}| = 2 - 2R_{\mathcal{A}}^*(\mu) \leq 2\epsilon .$$

We let $M = |\mathcal{B}| = (K + 1)^2 - K$; then $N = |\mathcal{C}| = M + K + 1 = (K + 1)^2 + 1$, and note

$$\begin{aligned} \int |\phi - \phi_{\mathcal{C}}| &= \int \left| \phi - \frac{(K+1)\phi_{\mathcal{A}} + M\phi_{\mathcal{B}}}{N} \right| \\ &= \int \left| \frac{(K+1)(\phi - \phi_{\mathcal{A}}) + M(\phi - \phi_{\mathcal{B}})}{N} \right| \\ &\geq \frac{M}{N} \int |\phi - \phi_{\mathcal{B}}| - \frac{(K+1)}{N} \int |\phi - \phi_{\mathcal{A}}| \\ &\geq (1 - \epsilon) \int |\phi - \phi_{\mathcal{B}}| - 2\epsilon^2 \\ &= (1 - \epsilon)(2 - 2R_{\mathcal{B}}^*(\mu)) - 2\epsilon^2 \\ &\geq 2 - 4\epsilon . \end{aligned}$$

Thus, $R_{\mathcal{C}}^*(\mu) \leq 2\epsilon$. □

Observe that non-monotonicity of the Bhattacharyya affinity also follows from the same argument. To this end, we may express $\rho_{\mathcal{C}}(\mu) = \frac{1}{2} \int \sqrt{\phi_0(\mathbf{x})\phi_S(\mathbf{x})}d\mathbf{x}$ in function of the Hellinger distance

$$H(\phi_0, \phi_{\mathcal{C}}) = \sqrt{\int (\sqrt{\phi_0(\mathbf{x})} - \sqrt{\phi_{\mathcal{C}}(\mathbf{x})})^2 d\mathbf{x}}$$

as $\rho_{\mathcal{C}}(\mu) = \frac{1}{2} - \frac{1}{4}H(\phi_0, \phi_{\mathcal{C}})^2$. Recalling (see, e.g., Devroye and Györfi [12, p.225]) that

$$H(\phi_0, \phi_{\mathcal{C}})^2 \leq \int |\phi_0(\mathbf{x}) - \phi_{\mathcal{C}}(\mathbf{x})| \leq 2H(\phi_0, \phi_{\mathcal{C}}),$$

we see that the same example as in the proof above, for n large enough, shows the non-monotonicity of the Bhattacharyya affinity as well.

6 Lower bounds on based random subclasses and metric entropy

In this section we derive lower bounds for the Bayes risk $R^* = R_{\mathcal{C}}^*(\mu)$. The bounds are in terms of some geometric features of the class \mathcal{C} . Again, we treat \mathcal{C} as a metric space equipped with the canonical distance $d(S, T) = \sqrt{\mathbb{E}_0(X_S - X_T)^2}$ (i.e., the square root of the Hamming distance $d_H(S, T)$).

For an integer $M \leq N$ we define a real-valued parameter $t_{\mathcal{C}}(M) > 0$ of the class \mathcal{C} as follows. Let $\mathcal{A} \subset \mathcal{C}$ be obtained by choosing M elements of \mathcal{C} at random, without replacement. Let the random variable τ denote the smallest distance between elements of \mathcal{A} and let $t_{\mathcal{C}}(M)$ be a median of τ .

Theorem 12 *Let $M \leq N$ be an integer. Then for any class \mathcal{C} ,*

$$R_{\mathcal{C}}^* \geq 1/4$$

whenever

$$\mu \leq \min \left(\sqrt{\frac{\log(M/16)}{K}}, \frac{8 \log(\sqrt{3}/8)}{\sqrt{K - t_{\mathcal{C}}(M)^2/2}} \right).$$

To interpret the statement of the theorem, note that

$$K - \tau^2/2 = \max_{\substack{S, T \in \mathcal{A} \\ S \neq T}} |S \cap T|$$

is the largest overlap between any pair of elements of \mathcal{A} . Thus, just like in Proposition 5, the distribution of the overlap between random elements of \mathcal{C} plays a key role in establishing lower bounds for the optimal risk. However, while in Proposition 5 the moment generating function $\mathbb{E} \exp(\mu^2 |S \cap T|)$ of the overlap between two random elements determines an upper bound for the critical value of μ , here it is the median of the largest overlap between many random elements that counts. The latter seems to carry more information about the fine geometry of the class. In fact, invoking a simple union bound, upper bounds for $\mathbb{E} \exp(\mu^2 |S \cap T|)$ may be used together with Theorem 12.

In applications often it suffices to consider the following special case.

Corollary 13 *Let $M \leq N$ be the largest integer for which zero is a median of $\max_{\substack{S, T \in \mathcal{A} \\ S \neq T}} |S \cap T|$ where \mathcal{A} is a random subset of \mathcal{C} of size M (i.e., $t_{\mathcal{C}}(M)^2 = 2K$). Then $R_{\mathcal{C}}^*(\mu) \geq 1/4$ for all $\mu \leq \sqrt{\log(M/16)/K}$.*

Example. (SUB-SQUARES OF A GRID.) To illustrate the corollary, consider the following example which is the simplest in a family of problems investigated by Arias-Castro, Candès, and Durand [4]: assume that n and K are both perfect squares and that the indices $\{1, \dots, n\}$ are arranged in a $\sqrt{n} \times \sqrt{n}$ grid. The class \mathcal{C} contains all $\sqrt{K} \times \sqrt{K}$ sub-squares. Now if S and T are randomly chosen elements of \mathcal{C} (with or without replacement) then, if $(K + 1)^2 \leq 2\sqrt{n}$,

$$\mathbb{P}\{|S \cap T| \neq 0\} \geq \frac{(\sqrt{n} - 2K)^2}{(\sqrt{n} - K + 1)^2} \cdot \frac{K}{(\sqrt{n} - K + 1)^2} \geq \frac{K}{n}$$

and therefore

$$\mathbb{P}\left\{\max_{\substack{S, T \in \mathcal{A} \\ S \neq T}} |S \cap T| = 0\right\} = 1 - \mathbb{P}\left\{\max_{\substack{S, T \in \mathcal{A} \\ S \neq T}} |S \cap T| > 0\right\} \geq 1 - M^2 \frac{K}{n}$$

which is at least $1/2$ if $M \leq \sqrt{n/(2K)}$ in which case $t_{\mathcal{C}}(M)^2 = 2K$. Thus, by Corollary 13, $R_{\mathcal{C}}^*(\mu) \geq 1/4$ for all $\mu \leq \sqrt{\log(n/(512K))/(2K)}$. This bound is of the optimal order of magnitude as it is easily seen by an application of Proposition 2.

In some other applications a better bound is obtained if some overlap is allowed. A case in point is the example of stars from Section 4.4. In that case any two elements of \mathcal{C} overlap but by taking $M = N(= m)$, we have $K - t_{\mathcal{C}}(M)^2/2 = 1$, so Theorem 12 still implies $R_{\mathcal{C}}^*(\mu) \geq 1/4$ whenever $\mu \leq \sqrt{(1/K) \log(m/16)}$.

The main tool of the proof of Theorem 12 is Slepian's lemma which we recall here [32]. (For this version, see Ledoux and Talagrand [27, Theorem 3.11].)

Lemma 14 (SLEPIAN'S LEMMA.) *Let $\boldsymbol{\xi} = (\xi_1, \dots, \xi_N), \boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_N) \in \mathbb{R}^N$ be zero-mean Gaussian vectors such that for each $i, j = 1, \dots, N$,*

$$\mathbb{E}\xi_i^2 = \mathbb{E}\zeta_i^2 \quad \text{for each } i = 1, \dots, N \quad \text{and} \quad \mathbb{E}\xi_i \xi_j \leq \mathbb{E}\zeta_i \zeta_j \quad \text{for all } i \neq j.$$

Let $F : \mathbb{R}^N \rightarrow \mathbb{R}$ be such that for all $\mathbf{x} \in \mathbb{R}^N$ and $i \neq j$,

$$\frac{\partial^2 F}{\partial x_i \partial x_j}(\mathbf{x}) \leq 0.$$

Then $\mathbb{E}F(\boldsymbol{\xi}) \geq \mathbb{E}F(\boldsymbol{\zeta})$.

Proof of Theorem 12: Let $M \leq N$ be fixed and choose M sets from \mathcal{C} uniformly at random (without replacement). Let \mathcal{A} denote the random subclass of \mathcal{C} obtained this way. Denote the likelihood ratio associated to this class by

$$L_{\mathcal{A}}(\mathbf{X}) = \frac{\frac{1}{M} \sum_{S \in \mathcal{A}} \phi_S(\mathbf{X})}{\phi_0(\mathbf{X})} = \frac{1}{M} \sum_{S \in \mathcal{A}} V_S$$

where $V_S = e^{\mu X_S - K\mu^2/2}$. Then the optimal risk of the class \mathcal{C} may be lower bounded by

$$R_{\mathcal{C}}^*(\mu) - R_{\mathcal{A}}^*(\mu) = \frac{1}{2} (\mathbb{E}_0 |L_{\mathcal{A}}(\mathbf{X}) - 1| - \mathbb{E}_0 |L_{\mathcal{C}}(\mathbf{X}) - 1|) \geq -\frac{1}{2} \mathbb{E}_0 |L_{\mathcal{A}}(\mathbf{X}) - L_{\mathcal{C}}(\mathbf{X})|.$$

Denoting by $\widehat{\mathbb{E}}$ expectation with respect to the random choice of \mathcal{A} , we have

$$\begin{aligned}
R_{\mathcal{C}}^*(\mu) &\geq \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \frac{1}{2}\mathbb{E}_0\widehat{\mathbb{E}}\left|\frac{1}{M}\sum_{S\in\mathcal{A}}V_S - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right| \\
&\geq \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \frac{1}{2}\sqrt{\mathbb{E}_0\widehat{\mathbb{E}}\left(\frac{1}{M}\sum_{S\in\mathcal{A}}V_S - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right)^2} \\
&\geq \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \frac{1}{2}\sqrt{\mathbb{E}_0\left[\frac{1}{M}\cdot\frac{1}{N}\sum_{T\in\mathcal{C}}\left(V_T - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right)^2\right]} \\
&\quad (\text{since the variance of a sample without replacement is less than that with replacement}) \\
&= \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \frac{1}{2\sqrt{M}}\sqrt{\frac{1}{N}\sum_{T\in\mathcal{C}}\mathbb{E}_0\left(V_T - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right)^2}
\end{aligned}$$

An easy way to bound the right-hand side is by writing

$$\begin{aligned}
\mathbb{E}_0\left(V_T - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right)^2 &\leq 2\mathbb{E}_0(V_T - 1)^2 + 2\mathbb{E}_0\left(1 - \frac{1}{N}\sum_{S\in\mathcal{C}}V_S\right)^2 \\
&\leq 2\mathbb{E}_0(V_T - 1)^2 + \frac{2}{N}\sum_{S\in\mathcal{C}}\mathbb{E}_0(1 - V_S)^2 \\
&= 4\text{Var}(V_T) = 4\left(e^{\mu^2 K} - 1\right).
\end{aligned}$$

Summarizing, we have

$$R_{\mathcal{C}}^*(\mu) \geq \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \sqrt{\frac{e^{\mu^2 K} - 1}{M}} \geq \widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) - \frac{1}{4}$$

where we used the assumption that $\mu \leq \sqrt{(1/K)\log(M/16)}$. Thus, it suffices to prove that $\widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) \geq 1/2$.

We bound the optimal risk associated with \mathcal{A} in terms of the Bhattacharyya affinity

$$\rho_{\mathcal{A}}(\mu) = \frac{1}{2}\mathbb{E}_0\sqrt{\frac{(1/M)\sum_{S\in\mathcal{A}}\phi_S(\mathbf{X})}{\phi_0(\mathbf{X})}} = \frac{1}{2}\mathbb{E}_0\sqrt{\frac{1}{|\mathcal{A}|}\sum_{S\in\mathcal{A}}V_S}.$$

Recalling from Section 3 that $R_{\mathcal{A}}^*(\mu) \geq 1 - \sqrt{1 - 4\rho_{\mathcal{A}}(\mu)^2}$ and using that $\sqrt{1 - 4x^2}$ is concave, we have

$$\widehat{\mathbb{E}}R_{\mathcal{A}}^*(\mu) \geq 1 - \sqrt{1 - 4\left(\widehat{\mathbb{E}}\rho_{\mathcal{A}}(\mu)\right)^2}.$$

Therefore, it suffices to show that the expected Bhattacharyya affinity $\widehat{\mathbb{E}}\rho_{\mathcal{A}}(\mu)$ corresponding to the random class \mathcal{A} satisfies

$$\widehat{\mathbb{E}}\rho_{\mathcal{A}}(\mu) = \frac{1}{2}\widehat{\mathbb{E}}\mathbb{E}_0\sqrt{\frac{1}{|\mathcal{A}|}\sum_{S\in\mathcal{A}}V_S} \geq \frac{\sqrt{3}}{4}.$$

In the argument below we fix the random class \mathcal{A} , relabel the elements so that $\mathcal{A} = \{1, 2, \dots, |\mathcal{A}|\}$, and bound $\rho_{\mathcal{A}}(\mu)$ from below. Denote the minimum distance between any two elements of \mathcal{A} by τ . To bound $\rho_{\mathcal{A}}(\mu)$, we apply Slepian's lemma with the function

$$F(\mathbf{x}) = \sqrt{\frac{1}{|\mathcal{A}|}\sum_{i=1}^{|\mathcal{A}|}e^{\mu x_i - K\mu^2/2}},$$

where $\mathbf{x} = (x_1, \dots, x_{|\mathcal{A}|})$. Simple calculation shows that the mixed second partial derivatives of F are negative, so Slepian's lemma is indeed applicable.

Next we introduce the random vectors $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$. Let the components of $\boldsymbol{\xi}$ be indexed by elements $S \in \mathcal{A}$ and define $\xi_S = X_S = \sum_{i \in S} X_i$. Thus, under \mathbb{P}_0 , each ξ_S is normal $(0, K)$ and $\mathbb{E}F(\boldsymbol{\xi})$ is just the Bhattacharyya affinity $\rho_{\mathcal{A}}(\mu)$. To define the random vector $\boldsymbol{\zeta}$, introduce $N+1$ independent standard normal random variables: one variable G_S for each $S \in \mathcal{A}$ and an extra variable G_0 . Recall that the definition of τ guarantees that the minimal distance between any two elements of \mathcal{A} is at least τ . Now let

$$\zeta_S = G_S \frac{\tau}{\sqrt{2}} + G_0 \sqrt{K - \frac{\tau^2}{2}}.$$

Then clearly for each $S, T \in \mathcal{A}$, $\mathbb{E}\zeta_S^2 = K$ and $\mathbb{E}\zeta_S\zeta_T = K - \tau^2/2$ ($S \neq T$). On the other hand, $\mathbb{E}\xi_S^2 = K$ and

$$\mathbb{E}\xi_S\xi_T = |S \cap T| = K - \frac{d(S, T)^2}{2} \leq K - \frac{\tau^2}{2} = \mathbb{E}\zeta_S\zeta_T.$$

Therefore, by Slepian's lemma, $\rho_{\mathcal{A}}(\mu) = \mathbb{E}F(\boldsymbol{\xi}) \geq \mathbb{E}F(\boldsymbol{\zeta})$. However,

$$\begin{aligned} \mathbb{E}F(\boldsymbol{\zeta}) &= \mathbb{E} \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} e^{\mu\zeta_S - K\mu^2/2}} \\ &= \mathbb{E} \sqrt{e^{\mu\sqrt{K-\tau^2/2}G_0 - (K-\tau^2/2)\mu^2/2} \frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} e^{\mu\tau G_S/\sqrt{2} - \tau^2\mu^2/4}} \\ &= \mathbb{E} e^{\mu\sqrt{K-\tau^2/2}G_0/2 - (K-\tau^2/2)\mu^2/4} \mathbb{E} \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} e^{\mu\tau G_S/\sqrt{2} - \tau^2\mu^2/4}} \\ &= e^{-\mu^2(K-\tau^2/2)/8} \mathbb{E} \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} e^{\mu\tau G_S/\sqrt{2} - \tau^2\mu^2/4}}. \end{aligned}$$

To finish the proof, it suffices to observe that the last expression is the Bhattacharyya affinity corresponding to a class of disjoint sets, all of size $\tau^2/2$, of cardinality $|\mathcal{A}| = M$. This case has been handled in the first example of Section 4 where we showed that

$$\mathbb{E} \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} e^{\mu\tau G_S/\sqrt{2} - \tau^2\mu^2/4}} \geq R_{\mathcal{A}}^* \geq 1 - \frac{1}{2} \sqrt{\frac{1}{M} e^{\mu^2\tau^2/2}} \geq \frac{3}{4}$$

where again we used the condition $\mu \leq \sqrt{\log(M/16)/K}$ and the fact that $\tau^2/2 \leq K$.

Therefore, under this condition on μ , we have that for any fixed \mathcal{A} ,

$$\rho_{\mathcal{A}}(\mu) = \frac{1}{2} \mathbb{E}F(\boldsymbol{\zeta}) \geq \frac{3}{8} e^{-\mu^2(K-\tau^2/2)/8}$$

and therefore

$$\widehat{\mathbb{E}}\rho_{\mathcal{A}}(\mu) \geq \frac{3}{16} e^{-\mu^2(K-t_{\mathcal{C}}(M)^2/2)/8}$$

where $t_{\mathcal{C}}(M)$ is the median of τ . This concludes the proof. \square

Remark. (AN IMPROVEMENT.) At the price of losing a constant factor in the statement of Theorem 12, one may replace the parameter $t_{\mathcal{C}}(M)$ by a larger quantity. The idea is that by thinning the random subclass \mathcal{A} one may consider a subset of \mathcal{A} that has better separation properties. More precisely, for an even integer $M \leq N$ we may define a real-valued parameter $\bar{t}_{\mathcal{C}}(M) > 0$ of the class \mathcal{C} as follows. Let $\mathcal{A} \subset \mathcal{C}$ be obtained

by choosing M elements of \mathcal{C} at random, without replacement. Order the elements S_1, \dots, S_M of \mathcal{A} such that

$$\min_{i \neq 1} d(S_1, S_i) \geq \min_{i \neq 2} d(S_2, S_i) \geq \dots \geq \min_{i \neq M} d(S_M, S_i)$$

and define the subset $\widehat{\mathcal{A}} \subset \mathcal{A}$ by $\widehat{\mathcal{A}} = \{A_1, \dots, A_{M/2}\}$. Let the random variable τ denote the smallest distance between elements of $\widehat{\mathcal{A}}$ and let $\bar{t}_{\mathcal{C}}(M)$ be the median of τ . It is easy to see that the proof of Theorem 12 goes through, and one may replace $t_{\mathcal{C}}(M)$ by $\bar{t}_{\mathcal{C}}(M)$ (by adjusting the constants appropriately). One simply needs to observe that since each V_S is non-negative,

$$\rho_{\mathcal{A}}(\mu) = \frac{1}{2} \mathbb{E}_0 \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \mathcal{A}} V_S} \geq \frac{1}{2} \mathbb{E}_0 \sqrt{\frac{1}{|\mathcal{A}|} \sum_{S \in \widehat{\mathcal{A}}} V_S} = \frac{1}{\sqrt{2}} \rho_{\widehat{\mathcal{A}}}(\mu).$$

If $\bar{t}_{\mathcal{C}}(M)$ is significantly larger than $t_{\mathcal{C}}(M)$ that the gain may be substantial.

If the class \mathcal{C} is symmetric then thanks to Theorem 10, the theorem above can be improved and simplified. If the class is symmetric, instead of having to work with randomly chosen subclasses, one may optimally choose a separated subset. Then the bounds can be expressed in terms of the metric entropy of \mathcal{C} , more precisely, by its *packing numbers* with respect to the canonical distance $d(S, T) = \sqrt{\mathbb{E}_0(X_S - X_T)^2}$.

We say that $\mathcal{A} \subset \mathcal{C}$ is a t -separated set (or t -packing) if for any $S, T \in \mathcal{A}$, $d(S, T) \geq t$. For $t < \sqrt{2K}$, define the *packing number* $M(t)$ as the size of a maximal t -separated subset \mathcal{A} of \mathcal{C} . It is a simple well-known fact that packing numbers are closely related to the covering numbers introduced in Section 2 by the inequalities $N(t) \leq M(t) \leq N(t/2)$.

Theorem 15 *Let \mathcal{C} be symmetric in the sense of Theorem 10 and let $t \leq \sqrt{2K}$. Then*

$$R_{\mathcal{C}}^* \geq 1/2$$

whenever

$$\mu \leq \min \left(\sqrt{\frac{\log(M(t)/16)}{K}}, \frac{8 \log(\sqrt{3}/2)}{\sqrt{K - t^2/2}} \right).$$

Proof: Let $\mathcal{A} \subset \mathcal{C}$ be a maximal t -separated subclass. Since \mathcal{C} is symmetric, by Theorem 10, $R_{\mathcal{C}}^* \geq R_{\mathcal{A}}^*$ so it suffices to show that $R_{\mathcal{A}}^* \geq 1/2$ for the indicated values of μ . The rest of the proof is identical to that of Theorem 12. \square

To interpret this result, take $t = \sqrt{2K(1 - \epsilon)}$ for some $\epsilon \in (0, 1/2)$. Then, by the theorem, $R^* \geq 1/2$ if

$$\mu \leq \frac{1}{\sqrt{K}} \min \left(\frac{8 \log(\sqrt{3}/2)}{\sqrt{\epsilon}}, \sqrt{\log(M(\sqrt{2K(1 - \epsilon)})/16)} \right).$$

As an example, suppose that the class \mathcal{C} is such that there exists a constant $V > 0$ such that $M(t) \sim (n/t^2)^V$. (Recall that all classes with VC dimension V have an upper bound of this form for the packing numbers, see Remark p. 5.) In this case one may choose $\epsilon \sim V \log(n/K)$ and obtain a sufficient condition of the form $\mu \geq c\sqrt{(V/K) \log(n/K)}$ (for some constant c), closely matching the bound obtained for the maximum test by Dudley's chaining bound.

7 Optimal versus maximum test: an analysis of the type I error

In all examples considered above, upper bounds for the optimal risk R^* are derived by analyzing either the maximum test or the averaging test. As the examples show, very often these simple tests have a near-optimal

performance. The optimal test f^* is generally more difficult to study. In this section we analyze directly the performance of the optimal test. More precisely, we derive general upper bounds for the type I error (i.e., the probability that the null hypothesis is rejected under \mathbb{P}_0) of f^* . The upper bound involves the expected value of the maximum of a Gaussian process indexed by a sparse subset of \mathcal{C} and can be significantly smaller than the maximum over the whole class that appears in the performance bound of the maximum test in Proposition 2. Unfortunately we do not have an analogous bound for the type II error.

We consider the type I error of the optimal test f^*

$$\mathbb{P}_0\{f^*(\mathbf{X}) = 1\} = \mathbb{P}_0\{L(\mathbf{X}) > 1\} = \mathbb{P}_0\left\{\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} > e^{K\mu^2/2}\right\}.$$

An easy bound is $\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} \leq e^{\mu \max_{S \in \mathcal{C}} X_S}$ so

$$\mathbb{P}_0\{L(\mathbf{X}) > 1\} \leq \mathbb{P}_0\left\{\max_S X_S > K\mu/2\right\}$$

Thus, $\mathbb{P}_0\{L(\mathbf{X}) > 1\} \leq \delta$ whenever $\mu \geq (1/K)\mathbb{E}_0 \max_S X_S + \sqrt{(2/K) \log(1/\delta)}$. Of course, we already know this from Proposition 2 where this bound was derived for the (suboptimal) test based on maxima.

In order to understand the difference between the performance of the optimal test f^* and the maximum test, one needs to compare the random variables $(1/\mu) \log \frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S}$ and $\max_{S \in \mathcal{C}} X_S$.

Proposition 16 *For any $\delta \in (0, 1)$, the type I error of the optimal test f^* satisfies*

$$\mathbb{P}_0\{f^*(\mathbf{X}) = 1\} \leq \delta$$

whenever

$$\mu \geq \frac{2}{K} \mathbb{E}_0 \max_{S \in \mathcal{A}} X_S + \sqrt{\frac{32 \log(2/\delta)}{K}}.$$

where \mathcal{A} is any $\sqrt{K}/2$ -cover of \mathcal{C} .

If \mathcal{A} is a minimal $\sqrt{K}/2$ -cover of \mathcal{C} then $(1/K)\mathbb{E}_0 \max_{S \in \mathcal{A}} X_S \leq \sqrt{\frac{2 \log N(\sqrt{K}/2)}{K}}$. By ‘‘Sudakov’s minoration’’ (see Ledoux and Talagrand [27, Theorem 3.18]) this upper bound is sharp up to a constant factor.

It is instructive to compare this bound with that of Proposition 2 for the performance of the maximum test. In Proposition 16 we were able to replace the expected maximum $\mathbb{E}_0 \max_{S \in \mathcal{C}} X_S$ by $\mathbb{E}_0 \max_{S \in \mathcal{A}} X_S$ where now the maximum is taken over a potentially much smaller subset $\mathcal{A} \subset \mathcal{C}$. It is not difficult to construct examples when there is a substantial difference, even in the order of magnitude, between the two expected maxima so we have a genuine gain over the simple upper bound of Proposition 2. Unfortunately, we do not know if an analog upper bound holds for the type II error $(1/N) \sum_{S \in \mathcal{C}} \mathbb{P}_S\{f^*(\mathbf{X}) = 0\}$ of the optimal test f^* . In cases when $\mathbb{E}_0 \max_{S \in \mathcal{A}} X_S \ll \mathbb{E}_0 \max_{S \in \mathcal{C}} X_S$, we suspect that the maximum test is far from optimal. However, to verify this conjecture, one would need a similar analysis for the type II error as well.

Proof: Introduce the notation

$$M_{\mathcal{C}}(\mu) = \mathbb{E}_0 \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} \right).$$

Then

$$\begin{aligned} \mathbb{P}_0 \left\{ \frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} > e^{K\mu^2/2} \right\} &= \mathbb{P}_0 \left\{ \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} \right) > \frac{K\mu}{2} \right\} \\ &= \mathbb{P}_0 \left\{ \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu X_S} \right) - M_{\mathcal{C}}(\mu) > \frac{K\mu}{2} - M_{\mathcal{C}}(\mu) \right\}. \end{aligned}$$

We use Tsirelson's inequality (Lemma 3) to bound this probability. To this end, we need to show that the function $h : \mathbb{R}^N \rightarrow \mathbb{R}$ defined by

$$h(\mathbf{x}) = \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu \sum_{i \in S} x_i} \right)$$

is Lipschitz (where $\mathbf{x} = (x_1, \dots, x_N)$). Observing that

$$\frac{\partial h}{\partial x_j}(\mathbf{x}) = \frac{\frac{1}{N} \sum_{S \in \mathcal{C}} \mathbb{1}_{\{j \in S\}} e^{\mu x_S}}{\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu x_S}} \in (0, 1),$$

we have

$$\|\nabla h(\mathbf{x})\|^2 = \sum_{j=1}^n \left(\frac{\partial h}{\partial x_j}(\mathbf{x}) \right)^2 \leq \sum_{j=1}^n \frac{\partial h}{\partial x_j}(\mathbf{x}) = K$$

and therefore h is indeed Lipschitz \sqrt{K} . By Tsirelson's inequality, we have

$$\mathbb{P}_0\{f^*(\mathbf{X}) = 1\} \leq \exp \left(-\frac{(K\mu/2 - M_{\mathcal{C}}(\mu))^2}{2K} \right).$$

Thus, the type I error is bounded by δ if

$$\mu \geq \frac{2M_{\mathcal{C}}(\mu)}{K} + \sqrt{\frac{8}{K} \log \frac{1}{\delta}}.$$

It remains to bound $M_{\mathcal{C}}(\mu)$.

Let $t \leq \sqrt{2K}$ be a positive integer and consider a minimal t -cover of the set \mathcal{C} , that is, a set $\mathcal{A} \subset \mathcal{C}$ with cardinality $|\mathcal{A}| = N(t)$ such that, if $\pi(S)$ denotes an element in \mathcal{A} whose distance to $S \in \mathcal{C}$ is minimal then $d(S, \pi(S)) \leq t$ for all $S \in \mathcal{C}$. Then clearly,

$$M_{\mathcal{C}}(\mu) \leq \mathbb{E}_0 \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu(X_S - X_{\pi(S)})} \right) + \mathbb{E}_0 \max_{S \in \mathcal{A}} X_S.$$

To bound the first term on the right-hand side, note that, by Jensen's inequality,

$$\mathbb{E}_0 \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} e^{\mu(X_S - X_{\pi(S)})} \right) \leq \frac{1}{\mu} \log \left(\frac{1}{N} \sum_{S \in \mathcal{C}} \mathbb{E}_0 e^{\mu(X_S - X_{\pi(S)})} \right) \leq \frac{\mu t^2}{2}$$

since for each S , $d_H(X_S, X_{\pi(S)}) \leq t^2$ and therefore $X_S - X_{\pi(S)}$ is a centered normal random variable with variance $d_H(X_S, X_{\pi(S)})$. For the second term we have

$$\mathbb{E}_0 \max_{S \in \mathcal{A}} X_S \leq \sqrt{2K \log N(t)}.$$

Choosing $t^2 = K/4$, we obtain the proposition. □

Acknowledgments. We thank Ery Arias-Castro and Emmanuel Candès for discussions on the topic of the paper. We also thank the referees for their valuable remarks. Parts of this work were done at the Bellairs Research Institute of McGill University and the Saturna West Island Mathematics Center.

References

- [1] D.J. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM Journal on Discrete Mathematics*, 3:450–465, 1990.
- [2] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *Random Structures and Algorithms*, 13:457–466, 1999.
- [3] E. Arias-Castro, E.J. Candès, H. Helgason, and O. Zeitouni. Searching for a trail of evidence in a maze. *The Annals of Statistics*, 36:1726–1757, 2008.
- [4] E. Arias-Castro, E. Candès, and A. Durand. Detection of abnormal clusters in a network. Technical report, UCSD, 2009.
- [5] S. Arlot, G. Blanchard, and E. Roquain. Some non-asymptotic results on resampling in high dimension, I: Confidence regions. *Annals of Statistics*, to appear, 2009.
- [6] S. Arlot, G. Blanchard, and E. Roquain. Some non-asymptotic results on resampling in high dimension, II: Multiple tests. *Annals of Statistics*, to appear, 2009.
- [7] Y. Baraud. Non asymptotic minimax rates of testing in signal detection. *Bernoulli*, 8:577–606, 2002.
- [8] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Uniform spanning forests. *The Annals of Probability*, 29:1–65, 2001.
- [9] A. Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhya, Series A*, 7:401–406, 1946.
- [10] S. Boucheron, G. Lugosi, and P. Massart. A sharp concentration inequality with applications. *Random Structures and Algorithms*, 16:277–292, 2000.
- [11] A. Broder. Generating random spanning trees. In *30th Annual Symposium on Foundations of Computer Science*, pages 442–447, 1989.
- [12] L. Devroye and L. Györfi. *Nonparametric Density Estimation: The L_1 View*. John Wiley, New York, 1985.
- [13] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag, New York, 1996.
- [14] D. Donoho and J. Jin. Higher criticism for detecting sparse heterogeneous mixtures. *Annals of Statistics*, 32:962–994, 2004.
- [15] D. Dubdashi and D. Ranjan. Balls and bins: a study in negative dependence. *Random Structures and Algorithms*, pages 99–124, 1998.
- [16] R.M. Dudley. Central limit theorems for empirical measures. *Annals of Probability*, 6:899–929, 1979. Correction in 7:909–911, 1979.
- [17] C. Durot and Y. Rozenholc. An adaptive test for zero mean. *Mathematical Methods of Statistics*, 15: 26–60, 2006.
- [18] T. Feder and M. Mihail. Balanced matroids. In *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 26–38, New York, NY, USA, 1992. ACM.
- [19] U. Feige and R. Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures and Algorithms*, 16(2):195–208, 2000.
- [20] J. Glaz, J. Naus, and S. Wallenstein. *Scan Statistics*. Springer, New York, 2001.

- [21] G.R. Grimmett and S.N. Winkler. Negative association in uniform forests and connected graphs. *Random Structures & Algorithms*, 24:444–460, 2004.
- [22] D. Haussler. Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69:217–232, 1995.
- [23] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- [24] Y.I. Ingster. Minimax detection of a signal for l_p^n -balls. *Mathematical Methods of Statistics*, 7:401–428, 1999.
- [25] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM*, 51:671–697, 2004.
- [26] L. Le Cam. On the assumptions used to prove asymptotic normality of maximum likelihood estimates. *The Annals of Mathematical Statistics*, 41:802–828, 1970.
- [27] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Springer-Verlag, New York, 1991.
- [28] J.W. Moon. *Counting Labelled Trees*. Number 1 in Canadian Mathematical Monographs. Canadian Mathematical Congress, Montreal, 1970.
- [29] J.G. Propp and D.B. Wilson. How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *Journal of Algorithms*, 27:170–217, 1998.
- [30] J.P. Romano and M. Wolf. Exact and approximate stepdown methods for multiple hypothesis testing. *Journal of the American Statistical Association*, 100:94–108, 2005.
- [31] A.A. Shabalin, V.J. Weigman, C.M. Perou, and A.B. Nobel. Finding large average submatrices in high dimensional data. *Annals of Applied Statistics*, page to appear, 2009.
- [32] D. Slepian. The one-sided barrier problem for Gaussian noise. *Bell System Tech. J.*, 41:463–501, 1962.
- [33] M. Talagrand. *The generic chaining*. Springer, New York, 2005.
- [34] B.S. Tsirelson, I.A. Ibragimov, and V.N. Sudakov. Norm of gaussian sample function. In *Proceedings of the 3rd Japan-U.S.S.R. Symposium on Probability Theory*, volume 550 of *Lecture Notes in Mathematics*, pages 20–41. Springer-Verlag, Berlin, 1976.
- [35] V.N. Vapnik and A.Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- [36] K. Vonnegut. *Breakfast of champions*. Delacorte Press, 1973.